Mohamed Kamel
Aurélio Campilho (Eds.)

LNCS 4633

Image Analysis and Recognition

4th International Conference, ICIAR 2007
Montreal, Canada, August 2007
Proceedings

Springer

Methods for Written Ancient Music Restoration

Pedro Castro and J.R. Caldas Pinto

IDMEC/IST, Technical University of Lisbon
Av. Rovisco Pais, 1049-001 Lisboa, Portugal
prmc@rnl.ist.utl.pt,
jcpinto@dem.ist.utl.pt

Abstract. Degradation in old documents has been a matter of concern for a long time. With the easy access to information provided by technologies such as the Internet, new ways have arisen for consulting those documents without exposing them to yet more dangers of degradation. While restoration methods are present in the literature in relation to text documents and artworks, little attention has been given to the restoration of ancient music. This paper describes and compares different methods to restore images of ancient music documents degraded over time. Six different methods were tested, including global and adaptive thresholding, color clustering and edge detection. In this paper we conclude that those based on the Sauvola's thresholding algorithm are the better suited for our proposed goal of ancient music restoration.

Keywords: Ancient Music Restoration, Image Processing, Document Degradation, Thresholding, Clustering, Edge Detection.

1 Introduction

Documents of ancient music frequently appear with multiple signs of degradation. Most of these signs result from the aging process, especially if poor care was taken to store and preserve the documents over time. It is becoming a common task in institutions like libraries to provide easy access to different types of information. With ancient documents, the process usually evolves from the creation of high resolution images, using image scanners, to the generation of different versions of these images, differing in aspects like image quality and size. One common feature to all of these images is that they continue to reveal the deterioration that was present in the original documents.

Restoration can be thought of as a transformation process that gives the original aspect to ancient document images that show a certain state of degradation. Degradation, on the other hand, can be seen as "every sort of less-than-ideal properties of real document images, e.g. coarsening due to low digitizing resolution, ink/toner, drop-outs and smears, thinning and thickening, geometric deformations, etc" [1]. Restoration is needed not only to improve the appearance of a document but also to improve the results of further segmentation and recognition operations. By clearing artifacts from the images there will be less room in the future for misinterpretation.

M. Kamel and A. Campilho (Eds.): ICIAR 2007, LNCS 4633, pp. 1194–1205, 2007.

Degradation can be divided into three types [2], according to the areas of a document that are subject to interference: 1) background degradation; 2) foreground degradation; and 3) global degradation. Background degradation includes blotches due to humidity, marks resulting from ink that traverses the paper (bleed-through) or from the scanning process (show-through), underlines, strokes of pen, annotations, and the superimposition of other symbols. Degradation on the foreground generally leads to broken or touching foreground objects, for instance, characters or musical symbols. Age effects can affect the ink components of a document leading to ink disappearance and some gaps can even appear in the document image causing significant loss of data and therefore affecting the document's content. Global degradation affects documents in their entirety, relating to a transformation that can be observed in a document as a whole, i.e., without affecting uniquely the foreground or the background. It includes, for instance, geometrical transformations and torn pages.

Background degradations account for the majority of degradations found in documents of ancient music, and therefore constitute the main focus of this paper. Examples of these degradations are depicted in Fig. 1. Water blotches are characterized by having a mainly convex shape (due to the diffusion of water molecules in the paper), a color that is darker than the neighborhood (due to the dust which is attracted in the paper texture), and an even darker border area where the dust accumulates [3]. Bleed-through refers to the sipping of ink from one side of a page to the other. It can be quite damaging, showing intensity levels that can be even darker than the true valid musical symbols in the foreground. Underlines are more frequent in text documents than they are in music documents, but strokes, annotations and superimposed symbols do appear quite often.

The restoration of ancient music images constitutes a topic that has been given little attention in the literature. Many papers exist that focus on optical music recognition, even for ancient music, but no emphasis is given to the restoration of those documents. On the other hand, a variety of methods do exist with application to text documents, and also some with application to images of maps, paintings and photographic prints.

Many segmentation and binarization approaches that aim to extract clear text from either noisy or textured backgrounds have been reported and compared in the literature [4,5,6,7,8], some with direct relation to old documents [9,10,11,12,13,14]. Edge-based segmentation has also been reported to work with the observation that edges of the valid foreground are sharper than those of the interference [15,16]. Other approaches exist that attempt to recover from defects found in artworks [3,17,18,19], typically following a semi-automatic approach that requires a user to interactively select the locations of the defects.

The rest of the paper is organized as follows. Section 2 describes the methods to be used throughout the experiments. Section 3 details the methodology used during the evaluation, and presents and analyzes the results. Section 4 concludes the paper and indicates future work directions.

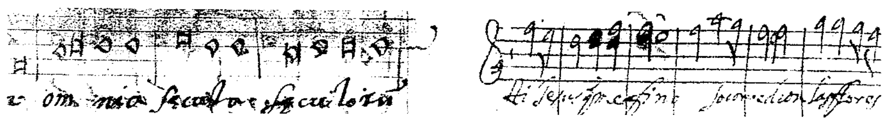

(a) (b)

Fig. 1. Images of ancient music showing background degradation

2 Methods

Interference in the background of a document can be generally detected using segmentation in the feature space. Because the background interference usually presents intensity levels that are a mid-term between those the background and the foreground, thresholding and clustering methods can be used to separate these defects from the rest of the image. However, there are also cases of extreme degradation in which other methods have to be used. In this sense, edge-based segmentation has been used, as it can be observed that edges surrounding degradation objects are usually smeared while those of the foreground are sharp.

Of the existing methods, only a subset can be applied to restore documents of ancient music. The selected methods to be used throughout the experiments are described next.

2.1 Niblack Thresholding

Niblack's method [20] performs adaptive thresholding and was selected because it is frequently cited and has been thoroughly compared with other types of documents. Comparisons have revealed that it outperforms other thresholding algorithms under different application domains [7,8,13].

This algorithm calculates a threshold value for each pixel based on the mean and standard deviation of all the pixels in a local neighborhood. A window of size $N \times N$ is moved over the image and the threshold value, for a pixel (x, y), is calculated as:

$$t(x, y) = m(x, y) + K \cdot s(x, y) \ , \tag{1}$$

where $m(x, y)$ and $s(x, y)$ are the mean and standard deviation values, respectively, in a local neighborhood of size $N \times N$ of pixel (x, y).

The parameters N and K highly influence the segmentation result. The window size N must be small enough to keep the details, but also large enough to remove noise. The value of K is used to decide how much of the total print object boundary is retained. A difficulty naturally arises on how to determine the best N and K parameters. We tested N and K ranging from 9 to 71 and from -0.2 to -1, respectively. The best parameter set did vary from image to image, but the values of $N = 51$ and $K = -0.8$ were found to average better.

2.2 Sauvola Thresholding

Sauvola's method [21] performs adaptive thresholding and was chosen because it is a modification of Niblack's method aimed at dealing better with the cases in which the background contains light texture, big intensity variations and uneven illumination.

In this modification, the threshold values are computed with the dynamic range of standard deviation, as:

$$t(x, y) = m(x, y) \cdot \left[1 + K \cdot \left(\frac{s(x, y)}{R} - 1\right)\right] \ , \tag{2}$$

where $m(x, y)$ and $s(x, y)$ are still the mean and standard deviation values, respectively, in a local neighborhood of size $N \times N$, as they were in Niblack's formula. R is the dynamic range of standard deviation and K now assumes positive values. Our tests revealed that the value of N was of limited influence to the quality of the segmentation, except with images containing bigger music note heads, in which case lower N values would not detect their centers. The value of R was fixed to 128, as used by Sauvola [21]. From tests performed with N and K ranging from 5 to 31 and from 0.2 to 0.8, respectively, the values of $N = 15$ and $K = 0.2$ provided the best results.

2.3 Otsu Thresholding

Otsu's method [22] was included for completeness, to perform global thresholding. It has been many times cited, evaluated [7,8] and used.

This method selects a threshold that maximizes between-class variance after creating the histogram of an intensity image. This threshold is then applied to all image pixels. It has the advantage of not requiring the input of parameters, but assumes that histograms are bimodal and illumination is uniform.

2.4 Canny Edge Detection with Niblack Post-processing

This method is a composition of algorithms, similar in part to what was used by Tan et al [15]. As outlined in Fig. 2, it starts by using a Canny edge detector [23], which is followed by foreground recovery and an adaptive thresholding algorithm, in this case Niblack's. Fig. 3 depicts the image resulting from the multiple stages of this method, which will be described next.

Canny edge detection is used to detect the edges within an image based on the observation that the edges of the valid foreground are usually sharper than

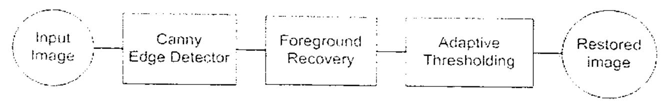

Fig. 2. Outline of the Canny edge detection with post-processing method

those of the background interference. This edge detector was selected due to the mentioned previous work with images of text documents, and also because it is known to many as the optimal edge detector, following the assumptions: 1) low error rate; 2) good localization on edge marking; and 3) minimal number of responses for a single edge. This algorithm starts by applying a Gaussian filter to smooth the image. Next, the gradient magnitude and orientation are computed using finite-difference approximations for the partial derivatives. Non-maximal suppression is applied to the gradient magnitude and double thresholding is used to detect and link edges. The two thresholds are used to eliminate weak edges. Initial seeds relating to strong edges are found using the high threshold. These seeds are grown to the pixels connected to them showing a gradient value above the low threshold. This inherently imposes the difficulty of selecting good threshold values. We opted to select the thresholds automatically based on the histogram of gradient magnitudes. The high threshold is set to the level that allows separating 30% of the highest gradient magnitudes from the rest. The low threshold is set to 40% of the high threshold, according to Canny's recommendations. Fig. 3(b) shows an example of detected edges.

Following the detection of the edges, recovery is needed to restore the original foreground. The dilation mathematical morphology operator is used with a structuring element that corresponds to a square of size $D \times D$. Its use has the purpose of filling the regions around the edges. The neighboring pixels within an $D \times D$ window centered on each edge pixel are recovered. The pixels outside the recovered region are set to the average color of the region they fill. Fig. 3(c) shows the result of applying the dilation operator to the edges that were previously detected. This image is used to recover the original pixels, resulting in the image shown in Fig. 3(d).

As can be observed in Fig. 3(d), the dilation operator retains not only valid parts of the foreground but also background and degradation that were surrounding the edges. Therefore, adaptive thresholding is used to remove the remaining interference. Niblack was once again selected for this task. Fig. 3(e) presents the result of the binarization process. This binary image can be used together with the original image to retain the original color of the foreground, as shown in Fig. 3(f). The background and interference pixels were set to their average color.

We tested D, N and K ranging from 5 to 9, from 9 to 71, and from -0.2 to -1, respectively, being N and K the parameters for Niblack's thresholding. In what concerns the D parameter used for the dilation, the value of 9 has to be used for the foreground objects to be completely filled. Contrary to what happens with text documents, where a value of 7 has been used [15], musical

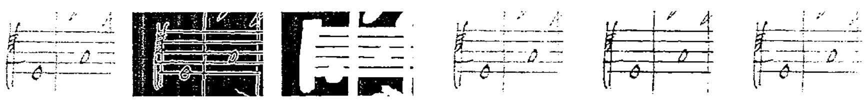
(a) (b) (c) (d) (e) (f)

Fig. 3. Images resulting from the multiple stages of the Canny edge detection method with post-processing

objects can many times be thicker thus requiring a higher D value. The values $N = 51$ and $K = -0.8$ provided the best results for the thresholding phase.

2.5 Canny Edge Detection with Sauvola Post-processing

This method is similar to the previous one, described in Sect. 2.4, with the difference that Sauvola thresholding is used in the post-processing phase, instead of Niblack's. From tests with D, N and K ranging from 7 to 9, from 5 to 31, and from 0.2 to 0.8, the values of $D = 9$, $N = 15$ and $K = 0.2$ provided the best results, where N and K are the parameters of Sauvola's thresholding.

2.6 Fuzzy C-Means Clustering

This method shares similarities with that of Pinto et al [24], which uses fuzzy clustering in the restoration of old text documents. Clustering is used in the RGB color space to collect pixels into groups according to their similarity. Groups, or clusters, are created based on the principle of maximizing the intraclass similarity and minimizing the interclass similarity. Objects within a cluster should have high similarity between each other, but should also be very dissimilar to objects in other clusters. Fuzzy clustering distinguishes itself from hard clustering by not assigning an object into a single cluster. Each object can belong to more than one cluster, with certain membership levels.

The fuzzy C-Means [25] is well known and revealed to present good results [24]. It is used to partition an image into N clusters, where each pixel has a membership degree to each cluster. The cluster centers are used to determine which cluster relates to the darkest color. All pixels that share a degree of membership greater than a value M are selected as valid foreground pixels. A difficulty arises in the selection of the ideal number of clusters, as well as on the minimum required membership value M. From tests with N and M ranging from 2 to 5 and from 0.65 to 0.95, the best results were achieved with $N = 3$ and $M = 0.75$.

3 Experiments

Experiments were performed in order to determine which of the method achieves a better segmentation. The emphasis was put on the quality of the restoration as observed in the resultant images. Computational issues were therefore not considered.

3.1 Methodology

Images of written ancient music were provided by the Portuguese Biblioteca Nacional and the Biblioteca Geral da Universidade de Coimbra, Portugal. A total of 10 images, scanned at a resolution of 150 *dpi*, were used throughout the experiments. These images contain a variety of degradations, as well as different musical notations and illumination characteristics. They are representative of the majority of images presented in the studied collections from the two referred libraries. All images were firstly manually restored, using graphics editing software, in order to be used as a standpoint for comparison.

Except for the case of Fuzzy C-Means, which clustered the images in the RGB color space, all other methods had the original images converted to gray scale before processing. This was performed using the known weighted conversion:

$$Gray = 0.3 \cdot R + 0.59 \cdot G + 0.11 \cdot B, \tag{3}$$

where $Gray$, R, G and B represent the gray, red, green and blue intensity levels, respectively. The result of all methods is a binary image containing the restored foreground and background. After binarization, this image can be applied to the original images as a mask, in order to retain the original colors of the foreground. In what concerns the background, it can be set to a fixed color or the average of its pixels, just to name a few simple alternatives.

The test images were processed by the chosen methods and compared to the manually restored images. The comparison was evaluated by the standard measures of *precision* and *recall* [26], with a slight modification. To evaluate text segmentation, these measures are typically used with the precision of a character or word. When relating to ancient music, it does not apply that well as the musical notation is very varied, including notes, clefs, key and time signatures, rests, bar and staff lines, as well as text, among other symbols. All of these symbols could be treated as characters, for instance, but that would ignore the great differences they have in shapes and sizes. Therefore, we opted to perform a bitwise comparison, and as such, the *precision* (P) and *recall* (R) measures were used as:

$$P = \frac{\text{Correctly Detected Pixels}}{\text{Total Detected Pixels}} , \tag{4}$$

$$R = \frac{\text{Correctly Detected Pixels}}{\text{Total Pixels}} , \tag{5}$$

where "Correctly Detected Pixels" refers to the foreground pixels that were correctly binarized (i.e., that are equal to those of the manually restored images) by a specific method, "Total Detected Pixels" refers to the total foreground pixels that were binarized by that method, and "Total Pixels" refers to the total foreground pixels that are present in the manually restored image.

Precision and *recall* reflect the performance of removing interfering strokes and restoring foreground strokes, respectively. The two measures have to be used together in order to compare the methods. To relate the two measures, the $F1$ score statistical measure was used, which is defined as:

$$F1 = \frac{2 \cdot P \cdot R}{P + R} . \tag{6}$$

The is the actual value that is worth maximizing.

3.2 Results

The results of evaluating the selected methods on the 10 selected images are presented in Table 1. Sauvola's thresholding was the method to obtain the highest average F1 score, but only slightly better than Canny edge detection with Sauvola post-processing. Niblack's thresholding was better when used as a post-processor for the Canny method, but only with a small difference to when it was used alone. Fuzzy C-Means performed somewhat close to Niblack's thresholding. Otsu's global thresholding method gave the worst results, as expected.

Table 1. Detailed *precision* (P), recall (R) and $F1$ results obtained by applying six different methods to 10 images of ancient music

Image		1	2	3	4	5	6	7	8	9	10	Average
Niblack	P	88.02	83.63	48.64	83.83	75.24	69.09	91.03	81.95	91.12	95.28	80.78
	R	91.99	80.18	82.83	84.42	82.32	88.77	71.86	80.30	81.68	89.81	83.42
	F1	89.97	81.87	61.29	84.12	78.62	77.70	80.31	81.12	86.14	92.47	81.36
Sauvola	P	89.42	84.90	94.51	91.74	74.97	94.00	90.60	84.04	89.56	96.35	89.01
	R	91.17	81.37	93.17	85.56	79.32	93.00	80.33	95.66	88.80	94.79	88.32
	F1	90.28	83.10	93.84	88.54	77.08	93.50	85.16	89.47	89.18	95.57	88.57
Otsu	P	70.16	73.01	24.12	66.64	67.25	69.87	83.93	62.60	79.74	85.92	68.32
	R	90.72	85.53	36.16	81.91	85.11	96.10	73.92	78.65	80.70	87.06	79.59
	F1	79.13	78.78	28.94	73.49	75.13	80.92	78.61	69.71	80.22	86.49	73.14
Canny Niblack	P	86.22	80.12	40.81	85.19	72.33	71.79	89.84	78.41	88.96	97.05	79.07
	R	94.02	83.20	86.02	84.86	86.96	92.06	77.39	84.92	83.09	93.00	86.55
	F1	89.95	81.63	55.36	85.02	78.97	80.67	83.15	81.53	85.92	94.98	81.72
Canny Sauvola	P	88.05	84.79	93.13	91.91	78.00	92.75	90.69	84.00	89.36	96.05	88.87
	R	90.59	81.69	93.33	85.81	82.00	91.38	80.01	95.00	88.25	94.99	88.31
	F1	89.31	83.21	93.23	88.76	79.95	92.06	85.02	89.16	88.80	95.52	88.50
Fuzzy C-Means	P	76.02	85.79	33.63	87.20	59.39	89.14	95.36	74.73	90.71	98.40	79.04
	R	87.83	76.90	95.39	89.44	90.40	86.46	68.29	84.83	88.37	91.47	85.94
	F1	81.50	81.10	49.72	88.31	71.69	87.78	79.58	79.46	89.52	94.81	80.35

The two methods that used Sauvola's thresholding obtained the highest values of precision and recall, with a small difference between them. The other methods presented bigger difference between precision and recall, with recall being greater most of the times. This means that these methods demonstrated a preference to restore foreground strokes over to remove interference, causing a great part of degradation to remain intact.

Fuzzy C-Means was more relevant with images in which the degradation showed a strong color. When converted to gray scale, as performed with the other methods, the degradation color may become a gray level in some parts similar to that of the foreground. In this case, clustering in the color space may be able to group the degradation pixels together, distinguishing them from the remaining pixels.

Sauvola's method was able to outperform Niblacks' with a much smaller window of size 15×15, compared to the window of size 51×51 used with Niblack's

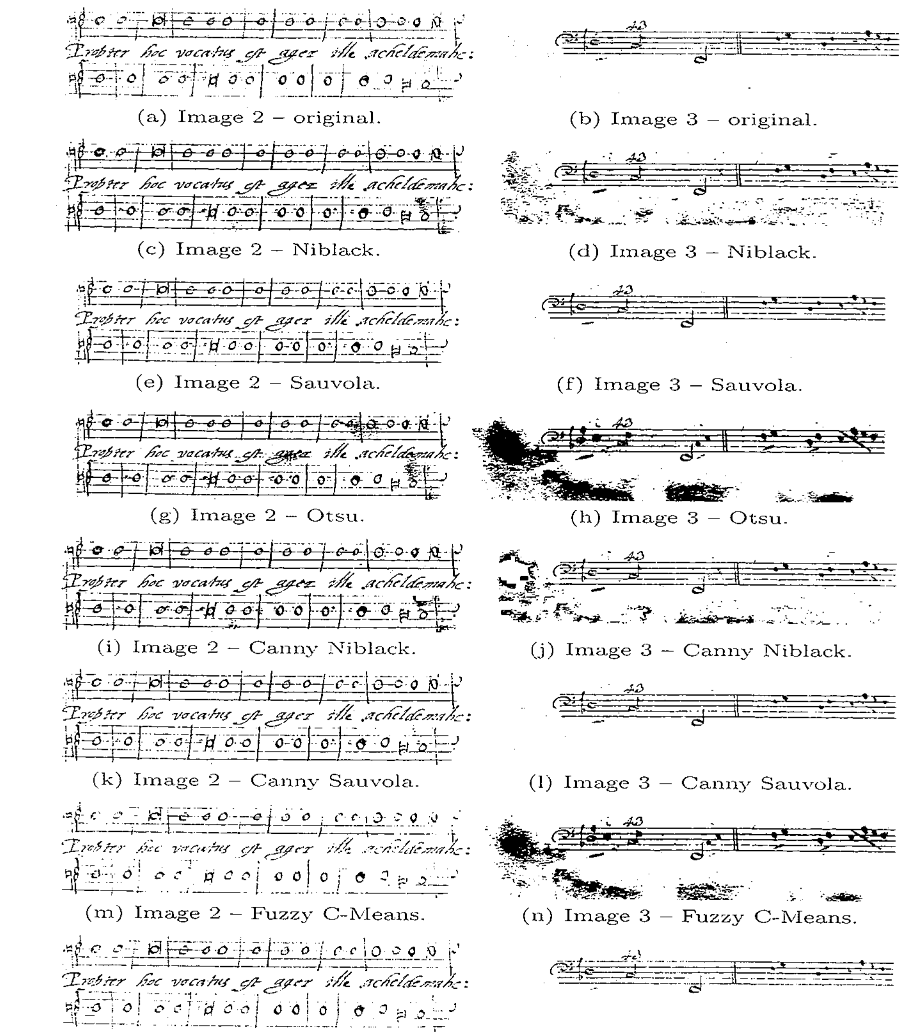

(a) Image 2 – original. (b) Image 3 – original.
(c) Image 2 – Niblack. (d) Image 3 – Niblack.
(e) Image 2 – Sauvola. (f) Image 3 – Sauvola.
(g) Image 2 – Otsu. (h) Image 3 – Otsu.
(i) Image 2 – Canny Niblack. (j) Image 3 – Canny Niblack.
(k) Image 2 – Canny Sauvola. (l) Image 3 – Canny Sauvola.
(m) Image 2 – Fuzzy C-Means. (n) Image 3 – Fuzzy C-Means.
(o) Image 2 – restored (Sauvola). (p) Image 3 – restored (Sauvola).

Fig. 4. Result of processing two of the original images with the chosen methods. Columns 1 and 2 correspond to the processing of images 2 and 3, respectively. Presented from top to bottom are: the original images; the result of binarizing the original images with Niblack thresholding, Sauvola thresholding, Otsu thresholding, Canny edge detection with Niblack post-processing, Canny edge detection with Sauvola post-processing, and Fuzzy C-Means; and the restored images using Sauvola thresholding, the method that performed the best.

method. In fact, using a 15×15 window with Niblack's method resulted in a particularly bad segmentation, with most of the symbols being barely recognizable as most of the interference would be detected as foreground too.

Two of the images used within the experiments are presented in Fig. 4, along with the results of processing them with the chosen methods. Image 2 presents less varying results, but still noticeable interference is observable in the results created by the methods of Otsu, Niblack and Canny with Niblack post-processing. The Fuzzy C-Means method was particularly aggressive with this image, with part of the staff lines getting erased. More varying results are observable in image 3. The Otsu and Fuzzy C-Means methods performed very poorly, showing big areas of degradation. Niblack's method, either when used alone or as post-processing for the Canny edge detection, also performed poorly, showing sparse but still large interference regions. Sauvola's method, either after edge detection or used alone, was able to remove almost all degradation. No method was able to remove the bleed-through pixels, which appear to the right in image 3. This is due to the high color intensity that bleed-through objects can appear with.

4 Conclusion

Ancient music images often present multiple types of degradation. Typical cases of degradation were presented and analyzed. Different written ancient music restoration methods have been described. A methodology for method evaluation was established and the selected methods compared. A discussion was performed on the obtained results.

Indeed, the use of edge detection did not seem to be of great use when compared to the use of the thresholding algorithms themselves. This seems to have less to do with the edge detection itself, but more with the way those edges are used afterwards. By observing the edge detection result, as exemplified in Fig. 3(b), it is noticeable that it typically correctly detects the edges, i.e., most of the foreground edges and almost none of the interference. The problem is that, after edge detection, the recovery phase retains all the pixels that surround the edges, within a local window, therefore including not only foreground but also interference pixels. A solution to be tested would consist on growing the edge pixels in the direction of the darkest colors, in such a way that would only preserve the true foreground.

Another problem evidenced by our experiments relates to the removal of bleed-through. As can be observed in the binarization results of image number 3, in Fig. 4, even Sauvola's method was not able to remove the sipped ink. One thing that is noticeable is that edge detection does, in fact, distinguish between bleed-through and valid foreground strokes. As the bleed-through strokes are smeared, they are typically not detected as edges. Therefore, once again, the need for a different method of reusing the detected edges, like the previously indicated solution of pixel growing, could be more successful in dealing with this problem.

Finally, some research is still necessary in order to fully automate the application of these methodologies to the mass production of restored ancient music documents. That will be the concern of future work.

Acknowledgments. This work was partly supported by: "Programa de Financiamento Plurianual de Unidades de I&D (POCTI), do Quadro Comunitário de Apoio III"; the FCT project POSC/EIA/60434/2004,(CLIMA), Ministério do Ensino Superior da Ciência e Tecnologia, Portugal; program FEDER; and "Programa Operacional Ciência e Inovação POCI2010".

References

1. Baird, H.: The state of the art of document image degradation modeling (2000)
2. Drira, F.: Towards restoring historic documents degraded over time. In: Document Image Analysis for Libraries, pp. 350–357 (2006)
3. Stanco, F., Ramponi, G.: Detection of Water Blotches in Antique Documents. In: Proc. 8th COST 276 Workshop, Trondheim, Norway (May 2005)
4. Liu, Y., Srihari, S.N.: Document image binarization based on texture features. IEEE Trans. Pattern Anal. Mach. Intell 19(5), 540–544 (1997)
5. Liang, S., Ahmadi, M., Shridhar, M.: A morphological approach to text string extraction from regular periodic overlapping text/background images. ICIP (1), 144–148 (1994)
6. Yang, Y., Yan, H.: An adaptive logical method for binarization of degraded document images. Pattern Recognition 33, 787–807 (2000)
7. Trier, .D., Jain, A.K.: Goal-directed evaluation of binarization methods. IEEE Trans. Pattern Anal. Mach. Intell 17(12), 1191–1201 (1995)
8. Taxt, T., Trier, O.D.: Evaluation of binarization methods for document images. IEEE Trans. Pattern Analysis and Machine Intelligence 17(6), 640–640 (1995)
9. Negishi, H., Kato, J., Hase, H., Watanabe, T.: Character extraction from noisy background for an automatic reference system. In: ICDAR, pp. 143–146 (1999)
10. Gatos, B., Pratikakis, I., Perantonis, S.J.: An adaptive binarization technique for low quality historical documents. In: Workshop on Document Analysis Systems, pp. 102–113 (2004)
11. Leedham, G., Varma, S., Patankar, A., Govindaraju, V.: Separating text and background in degraded document images: A comparison of global thresholding techniques for multi-stage thresholding. In: Frontiers in Handwriting Recognition, pp. 244–249 (2002)
12. Garain, U., Paquet, T., Heutte, L.: On foreground – background separation in low quality document images. International Journal on Document Analysis and Recognition 8(1), 47–63 (2000)
13. He, J., Do, Q.D.M., Downton, A.C., Kim, J.H.: A comparison of binarization methods for historical archive documents. In: ICDAR, pp. 538–542. IEEE Computer Society, Los Alamitos (2005)
14. Leedham, G., Yan, C., Takru, K., Tan, J.H.N., Mian, L.: Comparison of some thresholding algorithms for text/background segmentation in difficult document images. In: International Conference on Document Analysis and Recognition, pp. 859–864 (2003)
15. Tan, C.L., Cao, R., Wang, Q., Shen, P.: Text extraction from historical handwritten documents by edge detection. In: ICARCV2000. 6th International Conference on Control, Automation, Robotics and Vision, Singapore (December 5-8, 2000)
16. Tan, Cao, Shen: Restoration of archival documents using a wavelet technique. IEEETPAMI: IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (2002)

17. Barni, M., Bartolini, F., Cappellini, V.: Image processing for virtual restoration of artworks. IEEE MultiMedia 7(2), 34–37 (2000)
18. de Rosa, A., Bonacchi, A.M., Cappellini, V., Barni, M.: Image segmentation and region filling for virtual restoration of artworks. In: International Conference on Image Processing, vol. 1, pp. 562–565 (2001)
19. Stanco, F., Ramponi, G., Tenze, L.: Removal of Semi-Transparent Blotches in Old Photographic Prints. In: Proc. 5th COST 276 Workshop, Prague, Czech Republic (2003)
20. Niblack, W.: An Introduction to Digital Image Processing. Prentice-Hall, Englewood Cliffs (1986)
21. Sauvola, J., Pietikainen, M.: Adaptive document image binarization. Pattern Recognition 33(2), 225–236 (2000)
22. Otsu, N.: A threshold selection method from gray level histograms. IEEE Trans. Systems, Man and Cybernetics 9, 62–66 (1979)
23. Canny, J.: A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 8, 679–698 (1986)
24. Pinto, J.R.C., Bandeira, L., Sousa, J.M.C., Pina, P.: Combining fuzzy clustering and morphological methods for old documents recovery. In: Marques, J.S., Pérez de la Blanca, N., Pina, P. (eds.) IbPRIA 2005. LNCS, vol. 3523, p. 387. Springer, Heidelberg (2005)
25. Bezdek, J.C. (ed.): Pattern Recognition with Fuzzy Objective Function Algorithms. Plenum, NY (1981)
26. Junker, M., Dengel, A., Hoch, R.: On the evaluation of document analysis components by recall, precision, and accuracy. In: ICDAR, pp. 713–716 (1999)