

A Non-Parametric Framework for Document Bleed-Through Removal

Róisín Rowley-Brooke François Pitié Anil Kokaram
 Department of Electronic and Electrical Engineering
 Trinity College Dublin, Ireland

{rowleybr@, fpitie@mee., anil.kokaram}@tcd.ie

Abstract

This paper presents recent work on a new framework for non-blind document bleed-through removal. The framework includes image preprocessing to remove local intensity variations, pixel region classification based on a segmentation of the joint recto-verso intensity histogram and connected component analysis on the subsequent image labelling. Finally restoration of the degraded regions is performed using exemplar-based image inpainting. The proposed method is evaluated visually and numerically on a freely available database of 25 scanned manuscript image pairs with ground truth, and is shown to outperform recent non-blind bleed-through removal techniques.

1. Introduction

Ink bleed-through degradation poses one of the most difficult problems in document restoration. It occurs when ink has seeped through from one side of the page and interferes with text on the other side. This can significantly reduce the legibility of affected documents. Physical restoration of degraded documents is an invasive, expensive, and time consuming process that may affect the integrity of the original. It is therefore preferable to perform document restoration on a digital copy, where any number of changes may be made whilst leaving the original document intact.

Previous approaches to bleed-through removal struggle with severe bleed-through, where the intensity ranges of bleed-through and foreground regions overlap significantly. Furthermore in previous non-blind approaches [10, 8, 12], though intensity and spatial information from both recto and verso sides of the page are used to locate bleed-through regions, processing is performed separately on each side. The aim of this paper is to present a fully automated, non-parametric approach to non-blind bleed-through removal that can deal with a wider degree of degradation than other approaches, whilst producing results which preserve the characteristics of the original document. In order to achieve this we break the process into discreet stages. Firstly pre-

processing is performed on recto and verso images separately to enforce uniform global intensity characteristics. Secondly a two stage classification is performed on both sides of the document simultaneously to locate regions of bleed-through degradation. Thirdly clean background plate images are created using texture synthesis, and finally restored recto and verso images are obtained by blending the original degraded images and the clean background plates in regions classified as bleed-through.

The rest of the paper is set out as follows. Section 2 examines previous works on bleed-through removal. Section 3.1 describes the preprocessing stage. In Section 3.2 the two stage classification process is presented, and in Section 3.3 the details of the restoration stage are described. Visual and numerical comparisons between the proposed method and three recent non-blind removal methods, using the database and methodology proposed in [13] are made in Section 4, and finally the conclusions are presented in Section 5.

2. Previous Work

Approaches to bleed-through reduction generally fall into one of two groups; blind or non-blind, depending on whether they operate on one or both sides of the document. The image intensity is the main source of information used, with spatial information included in some approaches.

Blind methods are normally based on the assumption that there is a clear distinction between foreground text and bleed-through intensities, and hence use intensity based classification algorithms, such as hysteresis thresholding in [5], iterative K-means clustering and principal component analysis (PCA) in [4], and independent component analysis (ICA) in [17, 16]. This assumption does not hold for severe cases where the bleed-through intensity can be equivalent to or darker than the foreground text, and so intensity information alone is not enough to remove bleed-through successfully. Thus Wolf [19] introduces spatial smoothness priors in the form of a dual-layer MRF.

Non-blind methods make use of intensity information from both sides of the page, however the sides must first be registered so that they are aligned and of the same reso-

lution. Some non-blind methods use comparative intensity information from both sides to improve the performance of well known binarisation algorithms. For example the Kullback-Leibler thresholding algorithm and the binarisation algorithm of Gatos et al. [6] are improved for bleed-through interference by adding in secondary threshold levels in [1], and the Sauvola and Pietikainen adaptive binarisation algorithm [14], improved by fuzzy classification, is used in [2]. The ICA method is extended to double-sided documents in [18], using the recto and verso images as the sources for a blind-source separation. A model based approach is used by Moghaddam and Cheriet in [9], where a function of the difference in intensities between the two sides is used to indicate bleed-through regions. Physical diffusion-based models are defined for the foreground, bleed-through, and background, and then a reverse diffusion model is applied to bleed-through regions to remove interference. The same authors incorporate this diffusion model into a unified framework [10], using variational models for both blind and non-blind bleed-through removal with spatial smoothness enforced in the wavelet domain. Huang et al. in [7] and [8] proposed a user assisted method that classifies each pixel based on the ratio of intensities between the two sides, with spatial smoothness is enforced in a dual-layer MRF framework. The data cost energy is defined from a small set of user input training data, in the form of coloured strokes drawn by the user in foreground, background, and bleed-through regions on both sides. More recently Rowley-Brooke and Kokaram [12] proposed to represent the degradation via linear mixing models combined with foreground text masks, and to estimate restored image intensities explicitly, thus preserving the background texture of the document. A dual MRF framework is used to enforce spatial smoothness.




Figure 1. An example of an image with local intensity variations before (top), and after (bottom) detrending.

As highlighted in Section 1, the method proposed here seeks to emulate the non-parametric approach of [8], in that no assumptions about the document properties need to be made, whilst maintaining the restoration goal of [12], that

is to preserve the intrinsic characteristics of the document.

3. Proposed Method

3.1. Preprocessing

Registration of the recto and verso images is an essential preprocessing step for non-blind bleed-through reduction as it ensures that bleed-through pixels are aligned with their originating text pixels from the opposite side. For the purposes of this paper, it is assumed that the input recto and verso images are already registered - those in the database used for testing were registered manually.

Prior to classification it is necessary to compensate for any variations in the intensity profile over the document image, for example due to page binding or water stains. These effects can interfere with bleed-through restoration methods that rely on global intensity properties. Since many document imaging projects perform little or no image enhancement it can not be assumed that the resultant images have uniform global intensity properties. Any such trends are removed by exploiting firstly the fact that the only constant in most document images is the background medium; the presence and strength of any text or other foreground information may change across a page, and secondly that the largest proportion of pixels corresponds to background. Therefore the recto and verso images are adjusted separately by applying local intensity offsets such that the peaks of the local intensity histograms, corresponding to mean local background intensities, are aligned. This is performed by examining intensity histograms of overlapping blocks in the original image and storing the corresponding peak intensities. The mean of these peaks is set as the target intensity, and an offset matrix is defined from the difference between the target and stored local peak values. The block size used on the test images was set empirically to 200×200 pixels with an overlap of 50.

3.2. Classification

The proposed method aims to create a joint labelling of recto and verso images, from a set of four ‘pair’ labels: background on both sides, *bghg*, recto foreground and verso bleed-through, *fgbl*, recto bleed-through and verso foreground, *blfg*, or foreground on both sides, *fgfg*. Thus the recto and verso images r, v are treated as a joint image p , and each pixel pair $r(i, j), v(i, j)$ is treated as a single pixel $p(i, j)$ with intensity pair \mathbf{x} in the range $[0, 255]$, where 0 corresponds to white, and 255 to black. The motivation for considering pair rather than single intensities is to reduce the instances of overlap between labels. Further processing can then be performed in the spatial domain to correct any misclassifications due to the remaining overlap regions. There are therefore two stages to classification, firstly a joint histogram of intensity pairs is segmented into four regions




Figure 2. Contour plot of a joint histogram of recto-verso intensities.

corresponding to the four labels using a MRF-based clustering approach. This histogram labelling is then used as a map to obtain an initial image labelling. Secondly a set of rules governing connected label components in the image labelling is applied to produce the final labels for the recto-verso image p . Details of the two stages are contained in what follows.

3.2.1 Histogram Segmentation

Examining the joint histogram in Fig. 2, it is clear from the large peak in the points with lighter intensity that the largest proportion of pixels in p will correspond to regions where both recto and verso are background ($bgbg$). There are then three smaller peaks corresponding to the three remaining labels. We propose to assign labels based on the distance from these maxima. However, we want to ensure that the histogram segmentation maps to an image segmentation that is as spatially smooth as possible. So the labelling is formulated as a MRF framework with a spatial smoothness prior based in the recto-verso image domain rather than the joint histogram domain. Therefore the energy E corresponding to a particular labelling l is defined as follows:

$$E(l) = \sum_{\mathbf{x} \in \mathcal{X}} [U_{\mathbf{x}}(l_{\mathbf{x}}) + \alpha \sum_{\mathbf{y} \in \mathcal{N}_{\mathbf{x}}} V(l_{\mathbf{x}}, l_{\mathbf{y}})] \quad (1)$$

Where $l_{\mathbf{x}} \in \{bgbg, fgbl, blfg, fgfg\}$ is the label at point \mathbf{x} in the histogram, and \mathcal{X} corresponds to the set of non-zero histogram entries. The remaining terms in Eq. 1 are discussed below.

Unary Terms: The unary energy $U_{\mathbf{x}}(l_{\mathbf{x}})$ represents the cost of a point \mathbf{x} being assigned label $l_{\mathbf{x}}$. Document background regions generally have a lower range of intensities than foreground, so to prevent over classification of points as $bgbg$, $U_{\mathbf{x}}(l_{\mathbf{x}})$ is defined as the mahalanobis distance between point \mathbf{x} and the centre of the label cluster corresponding to $l_{\mathbf{x}}$.

Table 1. Likelihood of co-occurrence of labels

$\begin{matrix} l_{\mathbf{y}} \\ l_{\mathbf{x}} \end{matrix}$	$bgbg$	$fgbl$	$blfg$	$fgfg$
$bgbg$	0.66	0.00065	0.0069	0.00013
$fgbl$	0.0065	0.13	0.0001	0.0022
$blfg$	0.0069	0.0001	0.13	0.0021
$fgfg$	0.00013	0.0022	0.0021	0.046

Binary connections: The neighbourhood structure $\mathcal{N}_{\mathbf{x}}$ of a point \mathbf{x} is defined in the image domain as follows.

$$\mathcal{N}_{\mathbf{x}} = \{\mathbf{y} | \mathbf{y} = p(i', j'), \mathbf{x} = p(i, j), (i', j') \in \mathcal{N}_{i, j}\} \quad (2)$$

So each instance of an intensity pair \mathbf{x} is located in the recto-verso image p , then the corresponding points in the joint histogram of the 4-connect neighbours in p of these instances are added to the neighbourhood of \mathbf{x} .

Binary Terms: The pairwise energy $V(l_{\mathbf{x}}, l_{\mathbf{y}})$ represents the cost of neighbouring points in the histogram being assigned labels $l_{\mathbf{x}}$ and $l_{\mathbf{y}}$ respectively. Using a subset of the image pairs in the database, we set pairwise energies empirically as the negative log likelihood of co-occurrence of labels in that subset when labelled using the ground truth provided (see Table 1).

Smoothness Weight: A smoothness weight is applied to $V(l_{\mathbf{x}}, l_{\mathbf{y}})$ to balance the influence of the binary and unary energies, and depends on the range of intensities in the recto-verso image. When the range of intensities is small, there is a greater overlap between labels, and so there is less information available from the recto-verso intensities. In these cases the smoothness weight needs to be higher to reflect the increased reliance on spatial information. By optimising α over a subset of the images in the database with respect to ground truth histogram segmentations obtained from the ground truth provided, we observed that there is a linear relationship between the number of distinct intensities present in the recto-verso image and the optimal alpha value. Therefore we fit a linear model to the optimised results, and a unique α for each image pair is calculated based on the number of distinct intensities, I , as follows.

$$\alpha = I \times (5.8845 \times 10^{-7}) - 0.0024522 \quad (3)$$

Optimisation: Quadratic Pseudo-Boolean Optimization (QPBO) [11] is used to solve Eq. 1, repeatedly iterating through each of the four label options until convergence.

3.2.2 Image Segmentation

Following colour segmentation, the image labelling is initialised by using the histogram labelling as a look up table for pixels in the recto-verso image p . A subset of pixels will inevitably be misclassified due to the overlapping nature of the histogram label boundaries, however as the pairwise energy used in the histogram segmentation is derived from neighbourhoods in the image domain, spatial smoothness




Figure 3. This image illustrates the implementation of the connected component boundary rules. Left to right: recto extract, verso extract, image labelling before rules applied, and after rules applied. Row 1: Misclassified *bbg* components (dark blue) are corrected. Row 2: *fgfg* components (pink) are replaced with *fgbl* (green). Row 3: *fgbl* components (green) connected to *blfg* (light blue), but not *fgfg* (pink) are replaced with *blfg*. Row 4: A *blfg* component is connected to *fgfg*, but not *bbg* so is replaced with *fgfg*.

has already been enforced, which increases the likelihood that misclassified pixels will be grouped together. Therefore a full per-pixel analysis is not performed on the image labelling, and instead connected components of each label are examined, and rules governing permitted neighbouring components iteratively applied to correct misclassifications until convergence. The rules for each label are as follows.

bbg: This label covers the greatest proportion of the image, and so connected components will mostly be larger than the average character size. Smaller components correspond either to valid within character spaces, such as in ‘a’ and ‘o’, or to misclassifications. To avoid relabelling valid within character spaces, only the connected components that are less than 10% of the average character component size are analysed. Presumed to be mislabelled these components are relabelled with the label corresponding to the largest proportion of their neighbours.

fgfg: Conversely, this label covers the smallest proportion of the image, and as very dark bleed-through can often be mislabelled as *fgfg*, no assumptions can be made about the size of components and all are examined. The outer edges of components with this label must contain both *fgbl* and *blfg* labels, as overlapping text regions will originate from text alone on both sides. If this is not the case the component is relabelled *fgbl* or *blfg* according to which is present in the outer edge, or as *bbg* if neither.

fgbl: For this label, again only components less than 10% of the average character size are examined. The outer edges of these components must contain either *fgfg* and *bbg*, or *bbg* only. If the outer edge of a component contains *fgfg*, but not *bbg* also, then the component is relabelled as *fgfg*. If the outer edge of a component contains the label *blfg*, but



Figure 4. An example of an image where the gradient needs to be examined to improve the background plate. Top row left: degraded recto with faint ruled lines, right: corresponding verso. Second row left: image labelling (dark blue=texture source), right: visible artefacts in the recto background plate. Bottom left: labelling with 10% of source gradients removed (yellow), right: the improved background plate.

not *fgfg*, then it is relabelled *blfg*.

blfg: Components labelled *blfg* are processed in exactly the same way as *fgbl*, with the two labels interchanged. Example implementations of these rules are illustrated in Fig 3.

3.3. Restoration

The aim of this method is to preserve as much of the document as possible; the background texture is preserved to ensure that the experience of studying the document *image* remains close to that of studying the *physical* document. The restored recto and verso images $\hat{r}(x, y)$, $\hat{v}(x, y)$ are obtained by replacing identified bleed-through regions, where $l_i = fgbl$ for $\hat{r}(x, y)$, and $l_i = blfg$ for $\hat{v}(x, y)$, with background texture from clean background images $r_b(x, y)$, $v_b(x, y)$.

3.3.1 Clean Background

The background images are generated using the exemplar based inpainting algorithm of Criminisi et al. [3], that uses small sampled patches of nearby texture to fill image holes. The images $r_b(x, y)$, $v_b(x, y)$ for recto and verso sides are generated using regions labelled as *bbg* as the texture source, and inpainting all other label regions. Problems may be encountered with this approach in regions where faint foreground information might not have been identified during classification, with the result that foreground patterns are replicated in the background images. To mitigate this the gradients of the regions labelled as *bbg* are examined and the highest 10% of gradients removed from the inpainting source (see Fig 4).

3.3.2 Blending

Using a per-pixel replacement of bleed-through pixels with corresponding clean background pixels creates restored im-



Figure 5. An example of blending the background image with the degraded image in bleed-through boundary regions. Left to right: Degraded image, restored image with no blending, and after blending.

ages with visible seams at the edges of replaced regions. These artefacts are prevented by identifying the seam locations as the inner and outer edges of components labelled *fgbl* or *blfg*, and then performing a weighted sum of the original and background images across these edges to produce a smooth transition between original and replaced regions in the final restored results. The effect of this stage is illustrated in Fig. 5

4. Results & Discussion

The proposed method was tested on the database of 25 manuscript recto-verso image pairs with manually created binary foreground text images, presented in [13]. The results are evaluated first subjectively via a visual comparison, and then objectively, via a numerical comparison, against three recent non-blind bleed-through removal methods: (i) The dual-layer MRF approach with user trained likelihood proposed by Huang et al. [8], referred to as ‘H’ in the results, (ii) The double sided wavelet model of Moghadam and Cheriet [10], referred to as ‘M’, and (iii) The linear model based approach of Rowley-Brooke and Kokaram [12], referred to as ‘R’.

4.1. Visual Comparison

Results of the proposed and comparison methods, focusing on severe bleed-through regions, are shown in Fig. 6. In all the examples the proposed method produces a better result compared with other methods. The user assisted method (H) copes well with dark bleed-through when it is isolated, but tends to remove foreground text in overlapping *fgfg* regions, reducing legibility. The Wavelet method (M) preserves the foreground information, but does not cope well with dark bleed-through regions, leaving visible artefacts. The linear model based approach (R) also preserves foreground information well in most cases, but again does not cope well with dark-bleed through regions. The proposed method removes most of the bleed-through in all the examples whilst preserving the foreground well. Full results on the database images are available to view at www.isos.dias.ie under ‘Research’.

Table 2. Mean Error Probabilities and RP Ranks

	H	R	M	Proposed
FgError	0.2308	0.0795	0.0746	0.0696
RP FgRank	4	3	2	1
BgError	0.0012	0.0096	0.0148	0.0085
RP BgRank	1	2	4	3
TotError	0.0413	0.0220	0.0244	0.0196
RP TotRank	4	2	3	1

4.2. Numerical Comparison

As stated in Section 1, the methodology of [13] is used for numerical comparison. We therefore convert the results of different methods to a similar format in order to compare them objectively with the ground truth provided. The binarisation method of Gatos et. al [6] is applied to restored image results of the four approaches, and then these binary images compared to the ground truth to obtain three error metrics: *FgError*, the probability that a foreground pixel is binarised as background, *BgError*, the probability that a background or bleed-through pixel is binarised as foreground, and *TotError*, the probability that any pixel in the image is binarised incorrectly. The errors of the four methods are then ranked for each image, and an overall ranking for each error metric obtained using Ranked Pairs Voting (RP) [15]. The mean error probabilities and RP ranks are shown in Table 2.

Fig. 7 provides a visual representation of the quantitative performance of each method on the database. The *BgError* is plotted against *FgError* for all 50 images restored with each method, with trend ellipses centred at the mean *BgError* and *FgError* values (see Table 2), and with shape defined from the covariances between the two metrics. The restoration target is to have each point as close to the origin as possible. Table 3 shows the results of comparing pairs of methods against each other with each entry corresponding to the percentage of images where the method listed vertically performed better than that listed horizontally.

The numerical results correlate well with the visual observations already made. Results from the user assisted method (H) [8] successfully remove most of the bleed-through; it has a very low mean *BgError*. However this is achieved as the cost of foreground information in *fgfg* regions and so this method performs worst in terms of *FgError*. The Wavelet based method (M) [10] preserves the foreground well, however does not cope well with severe bleed-through so has a high average *BgError* and is ranked fourth for this metric. The linear model-based method (R) [12] outperforms the proposed method in terms of *BgError*. This is due to the fact that the mixing parameters in the model are assigned a very high smoothness such that at each successive estimation iteration the bleed-through removed regions increase in size and regions misclassified in the initial stages are gradually blended into the background. The



Figure 6. Four examples of results of different non-blind bleed-through removal techniques. In each example from top to bottom: Degraded recto and verso images, results from the user assisted method (H) [8], results from the Wavelet method (M) [10], results from the linear-based method (R) [12], results from the proposed method.

disadvantage of such a high smoothness is that valid foreground characters connected to bleed-through regions may also be increasingly blended into the background (as can be seen in the top left example of Fig. 6). The pairwise comparison results (Table 3) and RP metric rankings highlight that the proposed method outperforms the other three in terms of foreground preservation, and overall error.

5. Conclusion

This paper has presented a new framework for non-blind bleed-through removal including preprocessing, classification, and restoration. The preprocessing stage removes intensity trends in the input images. The classification stage

has the advantage over other methods that both recto and verso images are processed simultaneously, first by performing a joint histogram segmentation, then by applying rules to label connected components in the corresponding image segmentation. The restoration is performed using exemplar based image inpainting to preserve the character of the original document image. Visual and numerical comparisons with three recent non-blind methods show that the proposed method performs better overall. It was noted that a slight improvement in the histogram segmentation can produce significantly better results in the image segmentation, so it may be useful to examine the formulation of the pairwise energy in to see if the results can be improved further.

Table 3. Pairwise Method Rank Comparison(%)

FgError	H	R	M	Proposed
H	0	0	2	0
R	100	0	48	8
M	98	52	0	46
Proposed	100	92	54	0

BgError	H	R	M	Proposed
H	0	100	98	100
R	0	0	82	70
M	2	18	0	26
Proposed	0	30	74	0

TotError	H	R	M	Proposed
H	0	4	14	2
R	96	0	72	24
M	86	28	0	32
Proposed	98	76	68	0




Figure 7. BgError vs FgError over whole dataset

Finally, as the connected component rules do not examine all of the components in the labelling, it would be interesting to investigate replacing the rule based connected component analysis with a graph based optimization using the label components as nodes in the graph, and to see if better results are achieved.

Acknowledgement

This research has been funded by the Irish Research Council.

References

[1] J. A. Burgoyne, J. Devaney, L. Pugin, and I. Fujinaga. Enhanced bleedthrough correction for early music documents with recto-verso registration. In *Int. Conf. Music Inform. Retrieval*, pages 407–412, Philadelphia, PA, 2008.

[2] P. Castro, R. J. Almeida, and J. R. C. Pinto. Restoration of double-sided ancient music documents with bleed-through. In L. Rueda, D. Mery, and J. Kittler, editors, *CIARP*, volume

4756 of *LNCS*, pages 940–949. Springer Berlin/Heidelberg, 2007.

[3] A. Criminisi, P. Perez, and K. Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Trans. Image Process.*, 13(9):1200–1212, 2004.

[4] F. Drira, F. Le Bourgeois, and H. Emptoz. Restoring ink bleed-through degraded document images using a recursive unsupervised classification technique. In H. Bunke and A. Spitz, editors, *DAS*, volume 3872 of *LNCS*, pages 38–49. Springer Berlin / Heidelberg, 2006.

[5] R. Estrada and C. Tomasi. Manuscript bleed-through removal via hysteresis thresholding. In *10th Int. Conf. Doc. Anal. and Recogn.*, pages 753–757, Barcelona, Spain, 2009.

[6] B. Gatos, I. Pratikakis, and S. J. Perantonis. Adaptive degraded document image binarization. *J. Pattern Recogn.*, 39(3):317–327, 2006.

[7] Y. Huang, M. S. Brown, and D. Xu. A framework for reducing ink-bleed in old documents. In *IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 1–7, Anchorage, AK, 2008.

[8] Y. Huang, M. S. Brown, and D. Xu. User-assisted ink-bleed reduction. *IEEE Trans. Image Process.*, 19(10):2646–2658, 2010.

[9] R. F. Moghaddam and M. Cheriet. Low quality document image modeling and enhancement. *Int. J. Doc. Anal. Recogn.*, 11(4):183–201, 2009.

[10] R. F. Moghaddam and M. Cheriet. A variational approach to degraded document enhancement. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(8):1347–1361, 2010.

[11] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer. Optimizing binary mrf's via extended roof duality. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, 2007.

[12] R. Rowley-Brooke and A. Kokaram. Bleed-through removal in degraded documents. In *SPIE: Doc. Recogn. Retrieval Conf.*, San Francisco, CA, 2012.

[13] R. Rowley-Brooke, F. Pitié, and A. Kokaram. A ground truth bleed-through document image database. In P. Zaphiris, G. Buchanan, E. Rasmussen, and F. Loizides, editors, *TPDL*, volume 7489 of *LNCS*, Paphos, Cyprus, 2012. Springer.

[14] J. Sauvola and M. Pietikäinen. Adaptive document image binarization. *J. Pattern Recogn.*, 33(2):225–236, 2000.

[15] T. N. Tideman. Independence of clones as a criterion for voting rules. *J. Soc. Choice Welf.*, 4(3):185–206, 1987.

[16] A. Tonazzini. Color space transformations for analysis and enhancement of ancient degraded manuscripts. *J. Pattern Recogn. Image Anal.*, 20(3):404–417, 2010.

[17] A. Tonazzini, L. Bedini, and E. Salerno. Independent component analysis for document restoration. *Int. J. Doc. Anal. Recogn.*, 7(1):17–27, 2004.

[18] A. Tonazzini, E. Salerno, and L. Bedini. Fast correction of bleed-through distortion in grayscale documents by a blind source separation technique. *Int. J. Doc. Anal. Recogn.*, 10(1):17–25, 2007.

[19] C. Wolf. Document ink bleed-through removal with two hidden markov random fields and a single observation field. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(3):431–447, 2010.