

Document Bleed-through Removal using Sparse Image Inpainting

Muhammad Hanif, Anna Tonazzini
Pasquale Savino, and Emanuele Salerno
*Institute of Information Science and Technologies
Italian National Research Council, Pisa, Italy
Email: muhammad.hanif@isti.cnr.it*

Gregory Tsagakatakis
*Institute of Computer Science,
Foundation of Research and Technology - Hellas
Heraklion, Crete Greece*

Abstract—Bleed-through is a pervasive degradation in ancient documents, caused by the ink of the opposite side of the sheet that has seeped through the paper fiber, and appears as an extra, interfering text. Bleed-through severely impairs document readability and makes it difficult to decipher the contents. Digital image restoration techniques have been successfully employed to remove or significantly reduce this distortion. The main theme is to identify the bleed-through pixels and estimate an appropriate replacement for them, in accordance to their surrounding. This paper proposes a two-step image restoration method, exploiting information from the recto and verso images. First, based on a non-stationary linear model of the two texts overlapped in the recto-verso pair, the bleed-through pixels are identified. In the second step, a sparse representation based image inpainting technique, with a non-negative sparsity constraint, is used to find an appropriate replacement for the bleed-through pixels. Thanks to the power of dictionary learning and sparse image reconstruction methods, the natural texture of the background is well reproduced in the bleed-through areas, and even a their possible overestimation is effectively corrected, so that the original appearance of the document is preserved. The experiments are conducted on the images of a popular database of ancient documents, and the results validate the performance of the proposed method compared to the state of the art.

Keywords—Ancient document restoration; image inpainting; bleed-through removal; sparse representation;

I. INTRODUCTION

Ancient documents constitute the primary source of information from the middle ages, and provide insight into the culture, civilization, events and lifestyles of the distant past. In the years, these documents have been exposed to different type of degradations, and most of these ancient classics had a very narrow escape from total annihilation. In particular, in the manuscripts written on both sides of the sheet, often the ink had seeped through and appears as a degradation on the other side. Ink penetration through the paper is mainly due to aging, humidity, ink chemical property or paper porosity [1]. This kind of degradation is termed as bleed-through, and impairs legibility and interpretation of the document contents [2].

Nawadays, digital images of ancient documents and manuscripts are widely used for preservation, distribution, retrieval, and analysis. On these images, digital image processing techniques can be applied to perform any alteration to the document appearance, while preserving the original intact. Specifically, digital image processing techniques have been attempted for the virtual restoration



Figure 1: An example of bleed-through removal.

of documents affected by bleed-through, with some impressive results. Besides improving document readability, bleed-through removal is also a critical preprocessing step for tasks such as feature extraction, optical character recognition, segmentation, and automatic transcription.

In the literature, bleed-through removal is addressed as a classification problem, where the document image is subdivided into three components: background (the paper support), foreground (the main text), and bleed-through [1]. The existing methods can be divided into two main categories: blind, where the image of a single side is used, and non-blind, where the information of both the recto and verso sides of the document is required. Most of the earlier blind methods involve an intensity based thresholding step. However, thresholding may also destroy other useful document features, such as stamps, annotations, or paper watermarks. Thus, thresholding is not suitable when the aim is to preserve as much as possible the original appearance of the document. In order to remove bleed-through selectively, in [3] an independent component analysis (ICA) method is proposed to separate the foreground, background, and bleed-through layers from an RGB image. For grayscale images, in [4] a conditional random field (CRF) method is proposed. For the non-blind case an ICA method is outlined in [5], whereas the work in [6] proposes a variational approach, based on non-linear diffusion and wavelet transforms.

In addition to bleed-through identification, finding a suitable replacement for the affected pixels is equally important. Indeed, most of the methods either produce a binary image or replace the bleed-through areas with a smooth or random pattern that alters the document original look. In some method, e.g. [7], a preliminary “clean” background for the entire image is estimated, but this is usually a very laborious task.

This paper presents a two-step method to address the restoration of documents affected by bleed-through using pre-registered recto and verso images. First, the bleed-

through pattern is selectively identified on both sides, then image inpainting is used to suitably fill-in the areas to be corrected. In general, any off-the-shelf bleed-through identification method can be used in the first step. Here we adopt the algorithm described in [8], which is simple and very fast. Although efficient in locating the bleed-through patterns, this algorithm lacks a proper strategy to replace the unwanted bleed-through pixels. The simple replacement with the predominant background value causes unpleasant imprints of the bleed-through pattern to be visible in the restored image. An interpolation based inpainting technique is presented in [9], but the filled-in areas are mostly smooth. We use a sparse image representation based inpainting, with a non-negativity constraint, to find a befitting fill-in for the bleed-through pixels, in accordance to their neighbourhood. This sparse inpainting step, which constitutes the main contribution of the paper, enhances the quality of the restored image and preserves the natural paper texture.

The paper is organized as follows. Next section briefly introduces sparse image representation and dictionary learning. The bleed-through identification method is presented in Section 3. Section 4 describes the proposed sparse inpainting technique and the comparative results are discussed in Section 5. Section 6 concludes the paper.

II. SPARSE IMAGE REPRESENTATION

In the last few years, sparse signal representation has emerged as a powerful tool for efficiently representing high dimensional data. The underlying assumption is that signals such as audio and images are naturally generated by a multivariate linear model, driven by a small number of bases or regressors. The basis set, called dictionary, is either fixed and predefined, i.e., Fourier, Wavelet, Cosine, etc., or adaptively learned from a training set [10]. Sparse representation methods have proven successful in solving inverse problems arising in a variety of signal and image processing applications, such as denoising, reconstruction, classification, and compression. While the underlying key constraint is that the observed signal is sparse, explicitly meaning that it can be adequately represented using a small set of dictionary atoms, their particularity is that the dictionary is also learned to find the one that best describes the observed signal.

Given a dataset $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N] \in \mathbb{R}^{n \times N}$, its sparse representation consists of learning an overcomplete dictionary, $\mathbf{D} \in \mathbb{R}^{n \times K}$, with $K > n$, and a sparse coefficient matrix, $\mathbf{X} \in \mathbb{R}^{K \times N}$, such that $\mathbf{y}_i \approx \mathbf{D}\mathbf{x}_i$, by solving the optimization problem given as

$$\min_{\mathbf{D}, \mathbf{X}} \|\mathbf{Y} - \mathbf{D}\mathbf{X}\|_F^2 \text{ s.t. } \|\mathbf{x}_i\|_p \leq s$$

where the \mathbf{x}_i 's are the column vectors of \mathbf{X} , s is the desired sparsity level, and $\|\cdot\|_p$ is the ℓ_p norm, with $0 \leq p \leq 1$. **The choice of an overcomplete dictionary ($K > n$) allows more flexibility and richer data representation.** Most of these methods consist of a two stage optimization scheme: sparse coding and dictionary update. In the first stage the sparsity constraint produces a sparse linear approximation

for the observed data. Finding the exact sparse approximation is a NP-hard problem [11], but using approximate solutions has proven to be a good compromise. Commonly used sparse approximation algorithms are Matching Pursuit (MP) [12], Basis Pursuits (BP)[13], Focal Underdetermined System Solver (FOCUSS)[14], and Orthogonal Matching Pursuit (OMP)[15].

As per the dictionary that leads to sparse decomposition, although working with pre-defined dictionaries may be simple and fast, their performance might be not good for every task, due to their global-adaptivity nature [16]. Instead, learned dictionaries are adaptive to both the signals and the processing task at hand, thus resulting in a far better performance [17].

For a given set of signals \mathbf{Y} , dictionary learning algorithms generate a representation of signal \mathbf{y}_i as a sparse linear combination of the atoms \mathbf{d}_k for $k = 1, \dots, K$,

$$\hat{\mathbf{y}}_i = \mathbf{D}\mathbf{x}_i \quad (1)$$

Dictionary learning algorithms distinguish themselves from traditional model-based methods by the fact that, in addition to \mathbf{x}_i , they also train the dictionary \mathbf{D} to better fit the dataset \mathbf{Y} . The solution is generated by iteratively alternating between the sparse coding stage,

$$\hat{\mathbf{x}}_i = \arg \min_{\mathbf{x}_i} \|\mathbf{y}_i - \mathbf{D}\mathbf{x}_i\|^2; \text{ subject to } \|\mathbf{x}_i\|_0 \leq s \quad (2)$$

for $i = 1, \dots, N$, and the dictionary update stage for the \mathbf{X} obtained from the sparse coding stage

$$\mathbf{D} = \arg \min_{\mathbf{D}} \|\mathbf{Y} - \mathbf{D}\mathbf{X}\|_F^2. \quad (3)$$

where $\|\cdot\|_F$ represents the Frobenius norm and $\|\cdot\|_0$ is the ℓ_0 -norm, which counts the non-zero elements in \mathbf{x} . If \mathbf{x} is sparse enough, ℓ_0 -norm can be replaced by ℓ_1 -norm.

Dictionary learning algorithms are often sensitive to the choice of s . The update step can either be sequential, as in [18], or parallel, as in [19]. In sequential dictionary learning, the dictionary update minimization problem (3) is split into K sequential minimizations, by optimizing the cost function (3) for each individual atom while keeping fixed the remaining ones. Most proposed algorithms have kept the two stages optimization procedure, the difference appearing mainly in the dictionary update stage, with some exceptions having a difference in the sparse coding stage as well [10]. In the method proposed in [18], which has become a benchmark in dictionary learning, each column \mathbf{d}_k of \mathbf{D} and its corresponding row of coefficients \mathbf{x}_k^{row} are updated based on a rank-1 matrix approximation of the error for all the signals when \mathbf{d}_k is removed

$$\begin{aligned} \{\mathbf{d}_k, \mathbf{x}_k^{row}\} &= \arg \min_{\mathbf{d}_k, \mathbf{x}_k^{row}} \|\mathbf{Y} - \mathbf{D}\mathbf{X}\|_F^2 \\ &= \arg \min_{\mathbf{d}_k, \mathbf{x}_k^{row}} \|\mathbf{E}_k - \mathbf{d}_k \mathbf{x}_k^{row}\|_F^2. \end{aligned} \quad (4)$$

where $\mathbf{E}_k = \mathbf{Y} - \sum_{i=1, i \neq k}^K \mathbf{d}_i \mathbf{x}_i^{row}$. The singular value decomposition (SVD) of $\mathbf{E}_k = \mathbf{U}\mathbf{\Delta}\mathbf{V}^T$ is used to find the closest rank-1 matrix approximation of \mathbf{E}_k . The \mathbf{d}_k update is taken as the first column of \mathbf{U} , and the \mathbf{x}_k^{row} update is taken as the first column of \mathbf{V} multiplied by the

first element of Δ . To avoid the loss of sparsity in \mathbf{x}_k^{row} that would be created by the direct application of the SVD on \mathbf{E}_k , in [18] it was proposed to modify only the non-zero entries of \mathbf{x}_k^{row} resulting from the sparse coding stage. This is achieved by taking into account only the signals \mathbf{y}_i that use the atom \mathbf{d}_k in (4), or by taking the SVD of $\mathbf{E}_k^R = \mathbf{E}_k \mathbf{I}_{w_k}$, where $w_k = \{i | 1 \leq i \leq N; \mathbf{x}_k^{row}(i) \neq 0\}$ and \mathbf{I}_{w_k} is the $N \times |w_k|$ submatrix of the $N \times N$ identity matrix obtained by retaining only those columns whose index numbers are in w_k , instead of the SVD of \mathbf{E}_k .

III. BLEED-THROUGH IDENTIFICATION

The method we adopt to identify the bleed-through pixels is based on a comparison of the intensities of each pair of corresponding pixels in the recto and verso sides, provided that these are very accurately aligned. The rationale behind this simple approach is given by two basic assumptions, which we observed holding true in the majority of the manuscripts examined: i) the two sides are usually written by the same writer with the same ink; ii) due to paper porosity, the ink seeps and diffuses through the paper fiber. Hence, in general, the bleed-through pattern is a smeared and lighter version of the opposite text that has generated it. Nonetheless, this does not mean that, in a same side, bleed-through is lighter than the foreground text. On the contrary, in each side the intensity of bleed-through is usually very variable, that is, highly non-stationary, and sometimes can be as dark as the foreground. As already said, this makes hard to find thresholds to selectively identify it, especially if we want to protect the background texture and other slight document marks that genuinely belong to the given side.

Thus, based on assumption ii), we propose to selectively identify bleed-through and discriminate it from other desired features by looking for those pixels that are lighter than the corresponding pixels of an artificially blurred version of the opposite side [8]. To do so, at each pixel we compute the “seeping level” from a side to the other. From verso to recto the seeping level is given by:

$$q_v(i, j) = \frac{1 - \frac{s_r(i, j)}{b_r}}{h_v(i, j) \otimes \left(1 - \frac{s_v(i, j)}{b_v}\right) + \epsilon} \quad (5)$$

where the subscripts r, v indicate recto and verso, respectively. The seeping level q_r from recto to verso is given by the same equation, by inverting the role of the two sides. In eq. (5) $s(i, j)$ is the intensity at pixel (i, j) , b is an estimate of the predominant value in the background, h is a Point Spread Function (PSF) of unit volume, describing the smearing of the seeping ink, and \otimes indicates convolution. We assume h as a Gaussian function, with standard deviation approximately estimated from the extent of the character smearing. We remark that the use of a PSF allows a better matching of a same stroke in the two sides. At each pixel, we retain the smallest between the two seeping levels computed with eq. (5), to indicate a bleed-through pixel in the related side, and set to zero the other, to indicate a foreground pixel in the

opposite side. In [8], the restoration algorithm is completed by substituting the pixels having positive seeping level with the background predominant value b , and leaving unchanged the areas where $q = 0$. The experimental results proved the efficiency of this algorithm in removing bleed-through while leaving unaltered other salient marks, such as stamps, pencil annotations, and paper textures. These marks can be of paramount importance for the scholar studying the document contents and its origin.

Nevertheless, this method tends to overestimate the bleed-through areas. Indeed, we would like the two seeping levels to be both zero in the occlusions, *i.e.* those areas where the two texts overlap, and in the areas of background in both sides. Of course, based on eq. (5), this requirement cannot be fulfilled in the occlusions, and might fail to be met in some background pixels as well. At present, we attempt to correct these situations by [thresholding the absolute difference between the two images via the Otsu algorithm \[20\], given that the normalized intensities of text-text or background-background pixels are almost equal](#). Although with this strategy we obtain improvements, some pixels remain misclassified as bleed-through. However, this is not a true drawback, since the sparse inpainting algorithm that we propose only requires that “all” the genuine bleed-through pixels are recognized, and is able to cope with their possible overestimation, by properly replacing those pixels with the correct values.

IV. SPARSE BLEED-THROUGH INPAINTING

After successful identification of the bleed-through pixels, the next task is to find a suitable replacement for them. In this paper, we treat the bleed-through pixels as missing or corrupt image regions, and use sparse image inpainting to estimate pertinent fill-in values, which are consistent with the known uncorrupted surrounding regions. In recent years, image inpainting techniques have been widely used in image restoration, target removal, and compression. Generally, image inpainting techniques can be divided into two groups: diffusion based and exemplar based inpainting [21]. The diffusion based methods use a parametric model or partial differential equations, which extend the local structure from the surrounding to the internal of the region to be repaired [22]. In the exemplar based methods, an image block is selected as a unit, and the information is replicated from the known part of the image to the unknown region [23]. Comparatively, the exemplar based methods are faster and exhibit better performance, but use only a single best matching block to estimate the unknown pixels. This greedy kind of approach often introduces artifacts and also consumes more time in finding the best match for each image patch [21].

Recently, sparse representation based image inpainting algorithms have been reported with impressive results [24]. As sparse representation works on image patches, the main idea is to find the optimal sparse representation for each image patch and then estimate the missing pixels in a patch using the sparse coefficients of the known pixels.

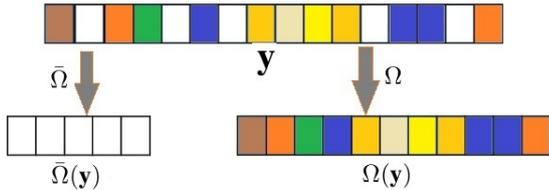


Figure 2: Extraction of known and bleed-through pixels from a patch.

Mathematically, the image inpainting problem is formulated as the reconstruction of the underlying complete image (in a column vector form) $\mathbf{C} \in \mathbb{R}^W$ from its observed incomplete version $\mathbf{I} \in \mathbb{R}^L$, where $L < W$. We assume a sparse representation of \mathbf{C} over a dictionary $\mathbf{D} \in \mathbb{R}^{n \times K}$: $\mathbf{C} \approx \mathbf{D}\mathbf{X}$. The incomplete image \mathbf{I} and the complete image \mathbf{C} are related through:

$$\mathbf{I} = \mathbf{M}\mathbf{C} \approx \mathbf{M}\mathbf{D}\mathbf{X} \quad (6)$$

where $\mathbf{M} \in \mathbb{R}^{L \times W}$ represents the layout of the missing pixels. Assuming that a well trained dictionary \mathbf{D} is available, the problem boils down to the estimation of sparse coefficients $\hat{\mathbf{X}}$ such that the underlying complete image $\hat{\mathbf{C}}$ is given by $\hat{\mathbf{C}} = \mathbf{D}\hat{\mathbf{X}}$. To learn the dictionary \mathbf{D} , a training set \mathbf{Y} is created by extracting overlapping patches of size $\sqrt{p_s} \times \sqrt{p_s}$ from the image at location $j = 1, 2, \dots, P$, where P is the total number of patches and p_s represents the patch size, i.e., the total number of pixels in a patch. Then we have $\mathbf{y}_j = \mathbf{R}_j(\mathbf{I})$, where $\mathbf{R}_j(\cdot)$ is an operator that extracts patch \mathbf{y}_j from the image \mathbf{I} , and its transpose, denoted by $\mathbf{R}_j^T(\cdot)$, is able to put back a patch into the j^{th} position in the reconstructed image. Since patches are overlapped, \mathbf{C} can be recovered from $\{\hat{\mathbf{y}}_j\}$ by averaging all the overlapping patches according to

$$\hat{\mathbf{C}} = \sum_{j=1}^P \mathbf{R}_j^T(\hat{\mathbf{y}}_j) / \sum_{j=1}^P \mathbf{R}_j^T(\mathbf{1}_{p_s}) \quad (7)$$

where $\hat{\mathbf{y}}_j$ is the inpainted patch, given by eq. (1).

A. Bleed-through Patch Inpainting

In each patch \mathbf{y} with bleed-through pixels, we have a known part and a missing or bleed-through part. Let Ω be an operator that extracts the known pixels in a patch and $\bar{\Omega}$ extracts the missing pixels. Then $\Omega(\mathbf{y})$ represents the known pixels and $\bar{\Omega}(\mathbf{y})$ represents the missing pixels in a patch \mathbf{y} , as illustrated in Fig.2. Given a well trained dictionary \mathbf{D} , the sparse reconstruction of patches with bleed-through pixels can be formulated as

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\Omega(\mathbf{y}) - \Omega(\mathbf{D}\mathbf{x})\|^2; \text{ s.t. } \|\mathbf{x}\|_1 \leq s \quad (8)$$

$\Omega(\mathbf{D}\mathbf{x})$ represents the estimated pixels values of the reconstructed patch $\mathbf{D}\mathbf{x}$ at location $\Omega(\mathbf{y})$. After obtaining the sparse coefficients $\hat{\mathbf{x}}$ using (8), an estimate for the bleed-through pixels can be obtained using

$$\bar{\Omega}(\mathbf{y}) = \bar{\Omega}(\mathbf{D}\hat{\mathbf{x}}) \quad (9)$$

In this paper we learned a dictionary \mathbf{D} from the training set \mathbf{Y} created from the overlapping patches of an image

with bleed-through, using the method described in Section II. For optimization, we used only complete patches from \mathbf{Y} , i.e. the patches with no bleed-through pixels, selected from both background areas and foreground text. This choice of ‘clean patches’ speeds up the training process and excludes the ‘non-informative’ bleed-through pixels. After dictionary training, the sparse coefficients in (8) are estimated using the non-negative OMP (NNOMP) presented in [25]. NNOMP guarantees the selection of only positive correlated atoms in the sparse representation of each image patch. Due to the large size of bleed-through areas, the inpainting order of bleed-through patches also has a significant impact on the final restored image. Similar to [23], high priority is given to patches with structure information in the known part. Patch priority is calculated based on the continuation of edges and textures within them. For each patch, the priority is defined as product of two terms, *confidence* and *data* [23]. The *confidence* is a measure of the available reliable information, i.e., those patches which have more known pixels are filled first. The *data* term encourages propagation of linear structures into the interior part of missing areas. This patch priority scheme enables a smooth transition of linear structure and texture information from the known part to the unknown (bleed-through) part of the patch.

V. EXPERIMENTAL RESULTS

We compare the performance of the proposed method with other state-of-the-art methods including [4], [7], and [26]. For evaluation purposes, we used images from the well known database of ancient documents presented in [27][28]. This database contains 25 pairs of recto-verso images of ancient manuscripts affected by bleed-through, along with the ground-truth images, where the foreground text is manually labeled. For the proposed method, the input images are first processed for bleed-through detection as discussed in Section III.

For bleed-through inpainting, the dictionary training dataset \mathbf{Y} is constructed by selecting from the input image the overlapping patches of size 8×8 with no bleed-through pixels. We learned an overcomplete dictionary \mathbf{D} of size 64×256 from \mathbf{Y} , with sparsity level $s = 5$ and DCT as initial dictionary. These parameters are empirically adjusted considering execution time and restoration result. For each patch to be inpainted, its non-negative sparse coefficients, denoted by \mathbf{x}_j , where j indicates the number of the patch, are estimated using the learned dictionary and NNOMP, and then used to estimate the fill-in values for the bleed-through areas.

In bleed-through restoration, the efficacy is generally evaluated qualitatively, as in most cases the original clean image is not available. A visual comparison of the proposed method with other state-of-the-art methods is presented in Fig. 3 and 5. The reported results for [26] and [7] are obtained from the online available ancient manuscripts database¹. The main goal here is to remove bleed-through while leaving intact the foreground text and

¹<https://www.isos.dias.ie/>

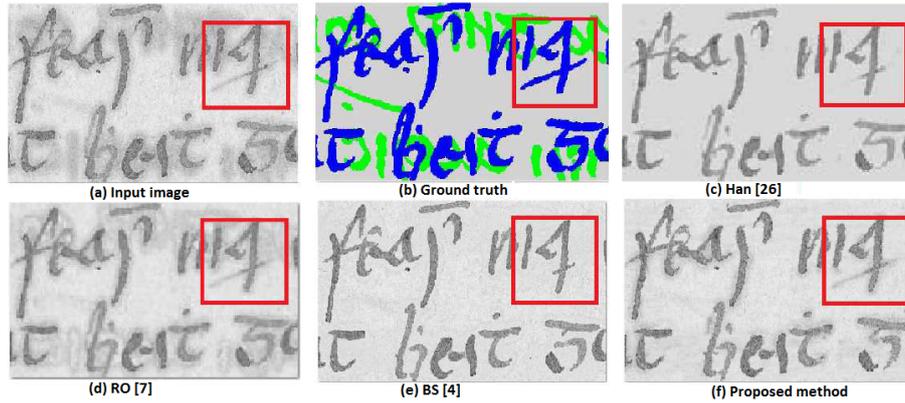


Figure 3: Visual comparison of bleed-through restoration.

replacing the bleed-through pixels with suitable values, so that the original look of the document is preserved. The proper inpainting of bleed-through areas is critical in preserving the paper texture in the background. As can be seen, the proposed method (Fig. 3 (f)) produces comparatively better results considering the given ground-truth image, especially in the area bounded by the red rectangle. It efficiently removes the bleed-through degradation, leaving intact the foreground text, and preserves the original look of the document. In [26] (Fig. 3 (c)) a modified Chan-Vese active contour model along with function minimization is used to remove the bleed-through using a recto-verso pair, but it flattens the background in the restored image, thus the original look of the document is compromised. The non-parametric method of [7] (Fig. 3 (d)) retains foreground text and background texture, but the bleed-through imprints are clearly visible. The recently proposed method in [4] (Fig. 3 (e)) produces better results, but some strokes of the foreground text are missing and the foreground text is more smooth compared to the original in the input image, as can be seen in the area bounded by the red rectangle. A bleed-through free colour image, obtained by using the proposed method, is shown in Fig.4. Another visual comparison is presented in Fig. 5, with the proposed method producing better results by preserving the background texture and keeps the foreground text intact. The proposed method copes very well with bleed-through removal and the dictionary based inpainting preserves the original appearance of the document.

VI. CONCLUSION

We extended the bleed-through identification method in [8] with a sparse image inpainting to estimate the appropriate fill-in for the identified bleed-through regions. Finding a befitting fill-in for the degraded pixels is a crucial task, because the imprints due to assigned values that are not in accordance with the neighborhood have unpleasant visual effects, and destroy the original look of the restored document. The proposed image inpainting technique, based on an efficient sparse image representation, calculates plausible fill-ins for the bleed-through

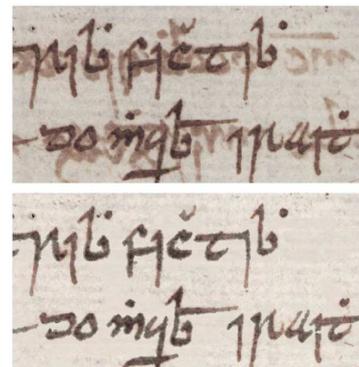


Figure 4: Ink bleed-through removal in a colour image: input image (top) and restored image (bottom).

regions using the surrounding clean pixels. In order to process only the positive correlated dictionary atoms, we used the non-negative orthogonal matching pursuit (NNOMP) for estimating the sparse coefficients. The performance of the proposed method is compared with other state-of-the-art methods on a database of recto-verso documents with bleed-through degradation.

ACKNOWLEDGMENT

This work has been partially supported by the European Research Consortium for Informatics and Mathematics (ERCIM), within the “Alain Bensoussan” Fellowship Programme.

REFERENCES

- [1] D. Fadoua, F. L. Bourgeois, and H. Emptoz, “Restoring ink bleed-through degraded document images using a recursive unsupervised classification technique,” *Document Analysis Systems VII, Lecture Notes in Computer Science*, vol. 3872. Springer, 2006.
- [2] C. Tan, R. Cao, and P. Shen, “Restoration of archival documents using a wavelet technique,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, pp. 1399–1404, 2002.
- [3] A. Tonazzini, L. Bedini, and E. Salerno, “Independent component analysis for document restoration,” *Int. J. Doc. Anal. Recogn.*, vol. 7, pp. 17–27, 2004.

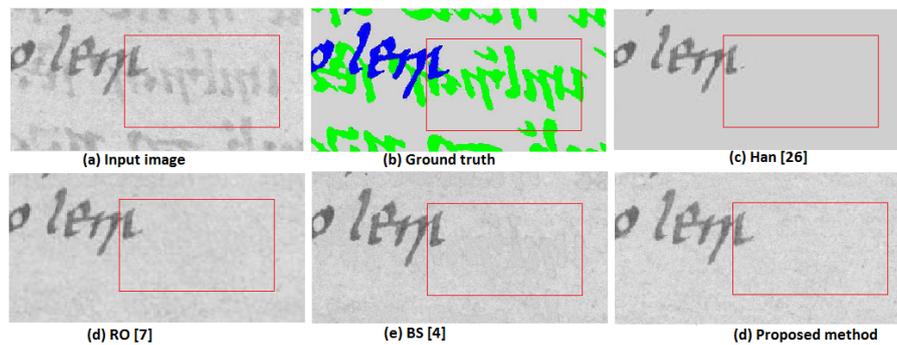


Figure 5: Visual comparison of bleed-through restoration.

- [4] B. Sun, S. Li, X. P. Zhang, and J. Sun, "Blind bleed-through removal for scanned historical document image with conditional random fields," *IEEE Trans. Image Process.*, pp. 5702–5712, 2016.
- [5] 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.*, vol. 10, pp. 17–27, 2007.
- [6] R. Moghaddam and M. Cheriet, "A variational approach to degraded document enhancement," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, pp. 1347–1361, 2010.
- [7] R. Rowley-Brooke, F. Pitié, and A. Kokaram, "A non-parametric framework for document bleed-through removal," *Proc. CVPR*, pp. 2954–2960, 2013.
- [8] A. Tonazzini, P. Savino, and E. Salerno, "A non-stationary density model to separate overlapped texts in degraded documents," *Signal, Image and Video Processing*, vol. 9, pp. 155–164, 2015.
- [9] I. Gerace, C. Palomba, and A. Tonazzini, "An inpainting technique based on regularization to remove bleed-through from ancient documents," *Proc. IWCIM*, pp. 1–5, 2016.
- [10] I. Tosić and P. Frossard, "Dictionary learning," *IEEE Signal Processing Magazine*, vol. 28, pp. 27–38, 2011.
- [11] J. A. Tropp and S. J. Wright, "Computational methods for sparse solution of linear inverse problems," *Proc. IEEE*, vol. 98, pp. 948–958, 2010.
- [12] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, pp. 3397–3415, 1993.
- [13] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, 1999.
- [14] I. Gorodnitsky and B. Rao, "Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm," *IEEE Trans. Signal Process.*, vol. 45, pp. 600–616, 1997.
- [15] J. Tropp, "Greed is Good: Algorithmic results for sparse approximation," *IEEE Trans. Information Theory*, vol. 50, pp. 2231–2242, 2004.
- [16] K. Kreutz-Delgado, J. Murray, B. Rao, K. Engan, T. Lee, and T. Sejnowski, "Dictionary learning algorithms for sparse representation," *Neural Comput.*, vol. 15(2), p. 349396, 2003.
- [17] B. Olshausen and D. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?" *J. Vision Research*, vol. 37, pp. 3311–3325, 1997.
- [18] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing of overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, pp. 4311–4322, 2006.
- [19] K. Engan, S. O. Aase, and J. Hakon-Husoy, "Method of optimal directions for frame design," *Proc. ICASSP*, pp. 2443–2446, 1999.
- [20] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Sys., Man., Cyber.*, vol. 9, no. 1, pp. 62–66, 1979.
- [21] C. Guillemot and O. Le Meur, "Image inpainting: Overview and recent advances," *IEEE Signal Processing Magazine*, vol. 31, pp. 127–144, 2014.
- [22] M. Bertalmio, A. Bertozzi, and G. Sapiro, "Navier-stokes, fluid dynamics, and image and video inpainting," *IEEE conference on Computer Vision and Pattern Recognition*, 2001.
- [23] P. Criminisi and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Trans. Image Process*, vol. 13, pp. 1200–1212, 2004.
- [24] T. Ogawa and M. Haseyama, "Image inpainting based on sparse representations with a perceptual metric," *EURASIP J. Adv. Signal Process.*, vol. 179, 2013.
- [25] M. Yaghoobi, D. Wu, and M. Davies, "Fast non-negative orthogonal matching pursuit," *Signal Process. Lett.*, vol. 22, pp. 1229–1233, 2015.
- [26] G. Hanasusanto, Z. Wu, and M. Brown, "Ink-bleed reduction using functional minimization," *IEEE conference on Computer Vision and Pattern Recognition*, 2010.
- [27] Irish Script On Screen Project, "http://www.isos.dias.ie," 2012.
- [28] R. Rowley-Brooke, F. Pitié, and A. Kokaram, "A ground truth bleed-through document image database," *Theory and Practice of Digital Libraries, Lecture Notes in Computer Science*, vol. 7489. Springer, pp. 185–196, 2012.