

Degraded Document Bleed-Through Removal

Róisín Rowley-Brooke, Anil Kokaram
Department of Electronic & Electrical Engineering, Sigmedia Group
Trinity College Dublin
Ireland
Email: {rowleybr, anil.kokaram}@tcd.ie

Abstract—This paper presents a Bayesian approach for bleed-through reduction in degraded document images based on a simple linear degradation model. A variation of ICM optimisation is used whereby samples are drawn for the bleed-through reduced images, whilst the remaining variables are estimated via the mode of their conditional probabilities. The proposed method is tested on various samples of scanned manuscript images with different degrees of degradation, and the results show some convincing removal of bleed-through.

I. INTRODUCTION

Reduction in legibility due to progressive degradation is often encountered in the study of documents, particularly handwritten. Many libraries host large collections of manuscripts which are vulnerable to such degradations due to the fragile nature of the writing media. Physical restoration of degraded documents is an expensive and time intensive process that could also negatively affect the integrity of the original. Restoration methods using automatic image processing techniques therefore have become increasingly popular as they have the advantage of being able to make any number of alterations to the degraded document appearance, whilst leaving the original intact.

Loss of textual information in degraded documents may be classified into four categories:

- (i) Fading of text due to light exposure or flaking ink.
- (ii) Obscured or missing text due to degradation of the writing medium. This may be caused by damp, mould, parasites, or inherent brittleness in the medium.
- (iii) Bleed-through interference, where ink has seeped through from one side of a page to the other.
- (iv) Digitisation of documents may introduce noise artifacts and degrade the textual information. For example non stationary noise due to variable illumination [1], show-through caused by the scanning of double sided documents [2], and images that appear warped as a result of inherent curvature in the document, for example due to binding.

This paper focuses solely on degradation caused by bleed-through interference, and proposes a method for bleed-through removal based on a linear degradation model. A previous version of this work can be found in [3], however, the iteration scheme has been altered, and new priors have been added.

II. RELATED WORK

Digital restoration of degraded document images, and more specifically the bleed-through problem, has become increas-

ingly popular recently. Different approaches may be categorised into one of two groups; blind and non-blind. Blind methods work with one side of the document only, whereas non-blind methods work with registered recto and verso sides, on the assumption that images of both are available. Due to the relative lack of image data, blind methods often involve an intensity based thresholding step, such as hysteresis thresholding in [4], and the recursive unsupervised classification method of [5]. Tonazzini et al. [6] separate the RGB colour channels of a single image into foreground, background, and bleed-through classes using independent component analysis (ICA). More recently Tonazzini in [7] applies this method to different colour space representations of manuscript images for bleed-through removal and information content maximisation. Wolf in [1] addresses the problem via a dual-layer Markov Random Field (MRF), using two hidden label fields and one observation field. Non-blind methods make use of a greater amount of data, however they are essentially two-stage processes. Firstly recto and verso sides must be registered so that they are aligned and of the same resolution (this is not a trivial problem), then the second stage comprises the bleed-through removal. Even given the increased data available, many non-blind methods are based around intensity differences rather than textual structure. Tonazzini et al. extended the ICA method to double sided documents in [8], using grayscale images and the flipped verso image as one of the sources. The methods contained in [9] and [10] focus on early music documents. Castro et al. [9] use Sauvola's thresholding algorithm [11] combined with fuzzy classification, while [10] proposes extensions to two binarisation methods, namely symmetric/non-symmetric Kullback-Leibler (KL) thresholding algorithms and the binarisation algorithm of Gatos et al. [12]; adding in a second threshold level for the bleed-through interference. Huang et al. in [13] and [14], following an increased interest in 'interactive computer vision', propose a method for bleed-through reduction that takes a small set of user input training data for the background, foreground text, and bleed-through text of both recto and verso pages, and uses this to locate and remove bleed-through interference. Moghaddam et al. [15] use diffusion models for the recto and verso texts, and the background medium, and apply a reverse diffusion model to remove interference. More recently, they have apply this diffusion method in a unified framework, [16], using variational models for blind, non-blind, and severe bleed-through removal. The linear model presented here is similar to that of Tonazzini et al in [8]. However,

constant mixing parameters across the images are not assumed, nor is it assumed that the extent of bleed-through is the same on both sides. In addition, this model simpler and unified under a Bayesian framework.

III. ALIGNMENT/REGISTRATION

As mentioned previously, the recto and verso sides must be registered so that the bleed-through interference and originating text match up on both sides. To do this the registration method proposed by Dubois et al. in [17] has been modified by taking a set of user selected points (minimum 3) to initialise the process. The user selects corresponding control points on both recto and verso images that indicate locations of the same textual features. The initialisation is then an affine model derived from a least squares fit to the displacements between these locations. In the experiments performed this relatively simple step improves the computation time of the algorithm and yields better alignment. In what follows *verso* refers to the flipped and registered image of the original un-aligned verso side.

IV. DEGRADATION MODEL

It is assumed that the intensity of an observed (degraded) recto pixel $I_r(h, k)$, at location (h, k) in the image, is a linear combination of the clean recto image pixel $Y_r(h, k)$ and some proportion of the clean verso pixel $Y_v(h, k)$. The combination is controlled by $\alpha_v(h, k)$, a mixing parameter and masks defined on both sides, $M_r(h, k), M_v(h, k)$. Due to the symmetrical nature of the problem, the model for an observed verso pixel, $I_v(h, k)$, is similar. Therefore the models for each side are as follows (where we discard pixel co-ordinates for brevity).

$$\begin{aligned} I_r &= Y_r + M_r(1 - M_v)\alpha_v Y_v + \rho \\ I_v &= Y_v + M_v(1 - M_r)\alpha_r Y_r + \nu \end{aligned} \quad (1)$$

α_v and α_r are the proportion of the clean image intensity that has bled through from one side to the other. M_r and M_v here are binary masks that have value 0 where the corresponding image is foreground text, and 1 everywhere else. The combination of these two mask terms ensures that bleed-through cannot occur where foreground text is present. This assumption that bleed-through is not visible through foreground text is reasonable, as it is only in very extreme cases of bleed-through degradation that the interference is visible through the foreground characters. Finally, the noise terms ρ, ν are Gaussian $\mathcal{N}(0, \sigma_{\rho\rho}^2), \mathcal{N}(0, \sigma_{\nu\nu}^2)$.

V. BAYESIAN FRAMEWORK

Under a Bayesian framework, the estimation of the parameters $\theta = [\alpha_v, \alpha_r, M_v, M_r, Y_r, Y_v]$ proceeds using a Maximum A Posteriori (MAP) scheme. The p.d.f. of the variables given the observed data I_r, I_v , at a single location is then as follows.

$$p(\theta|I_r, I_v, \tilde{M}, \tilde{\alpha}) \propto p(I_r, I_v|\theta, \tilde{M}, \tilde{\alpha})p(\theta|\tilde{M}, \tilde{\alpha}) \quad (2)$$

where $\tilde{M}, \tilde{\alpha}$ represent the existing state of the mask and linear mixing parameter (on the recto or verso side as appropriate) in

the neighbourhood of the pixel site currently being considered. The likelihood and prior distributions are described in what follows.

A. Likelihood

Following the degradation model, and assuming that the two sides are independent, the likelihood combines the influence of both the recto and verso sides to yield another Gaussian distribution.

$$\begin{aligned} p(I_r, I_v|\theta, \tilde{M}, \tilde{\alpha}) &\propto \\ &\exp - \left\{ \frac{1}{2\sigma_{\rho\rho}^2} [I_r - Y_r - M_r(1 - M_v)\alpha_v Y_v]^2 \right. \\ &\quad \left. + \frac{1}{2\sigma_{\nu\nu}^2} [I_v - Y_v - M_v(1 - M_r)\alpha_r Y_r]^2 \right\} \end{aligned} \quad (3)$$

B. Priors

Priors for each of the important variables can be designed using some intuition of the visual effects that they embody. The Mask variables M_r, M_v have the effect of delineating the regions covered by foreground text. Clearly therefore, these variables should be smooth in local patches and hence the usual Gibbs energy prior for spatial smoothness makes sense. However, in regions where there is no bleed-through, it is clear that these variables should be 1. Given that it is possible to estimate roughly (using K-means clustering on the original degraded document) the regions of text and regions of background, it is sensible then to include a prior constraining these variables in the non-text, that is, background region. Therefore the prior for the masks is as follows (the prior for M_r only is shown here, M_v is similar).

$$p(M_r|\tilde{M}) \propto \exp - \left\{ \sum_{\tilde{M}} (M_r - \tilde{M})^2 \lambda_M + \beta_r (1 - M_r) \right\} \quad (4)$$

Here \tilde{M} represents the current state of M_r in the 8-connected neighbourhood of the current site, and λ_M is a smoothness hyperparameter to encourage spatial smoothness (in essence an Ising model in this case). In the experiments to follow $\lambda_M = 1$. β_r is a penalty for setting $M_r = 0$ (foreground) in the regions of non-text, and is estimated in the initialisation step discussed below.

A prior for the mixing parameters, α_v, α_r follows similar logic. Firstly, smoothness is again encouraged with the Gibbs energy term. It is then important to constrain α_v, α_r near some reasonable value that would yield picture material that is close in intensity to the observed background (non-textual) regions where α_v, α_r are non-zero. Without this constraint there is nothing limiting α_v, α_r to create useful images since smoothness of itself does not constrain the absolute mixing amount. Hence the prior (for α_v only, as α_r is similar) is as follows.

$$p(\alpha_v|\tilde{\alpha}) \propto \exp - \left\{ \sum_{\tilde{\alpha}} (\alpha_v - \tilde{\alpha})^2 \lambda_{\alpha} + (\alpha_v - \tilde{\alpha}_v)^2 \lambda_c \right\} \quad (5)$$

Here, $\tilde{\alpha}$ are values of α_v in the neighbourhood (again, 8-connected) of the current site, and λ_α , λ_c are smoothness and constraining hyperparameters respectively, set to $\lambda_\alpha = \lambda_c = 55$ in what follows. As can be seen, the constraint on the magnitude of α_v is Gaussian (the second term in Eq.5), and $\bar{\alpha}_v$ is a *rough* estimate of α_v , obtained from the initialisation step.

Finally, a prior is used for the underlying clean image data to encourage the average brightness of the restored document bleed-through regions to match the average brightness of the background. This is a simple Gaussian prior and again the parameters are derived from the initialisation step (again the case for the recto side Y_r only is shown).

$$p(Y_r|\bar{Y}_b) \propto \exp\left\{-\frac{1}{2\sigma_{\rho\rho}^2}M_r(1-M_v)(Y_r-\bar{Y}_b)^2\right\} \quad (6)$$

Here, \bar{Y}_b is an estimate of the average background intensity of the clean image, obtained in the initialisation, and $\sigma_{\rho\rho}^2$ is the variance of the background noise as defined above in Section IV.

VI. SOLUTION

To solve for all the variables, ICM optimisation is used. However a slightly modified version is adopted here in that *samples* are drawn for the underlying clean images Y_r , Y_v , while selecting the mode of the conditionals for the remaining variables. The process is clearly iterative and the variables in θ are solved for in turn. The solution for each of the variables are presented below in the case of the recto side only, as, again, the verso equations are similar due to the symmetry of the problem.

A. Mixing Parameter Estimate

Each mixing parameter is present in only one of the observation terms, therefore the conditional probability $p(\alpha_v|\cdot)$ at a site is as follows.

$$p(\alpha_v|M_v, M_r, Y_r, Y_v, I_r, I_v) \propto \exp\left\{-\frac{1}{2\sigma_{\rho\rho}^2}(I_r - Y_r - M_r(1 - M_v)\alpha_v Y_v)^2 + \sum_{\tilde{\alpha}} (\alpha_v - \tilde{\alpha})^2 \lambda_\alpha + (\alpha_v - \bar{\alpha}_v)^2 \lambda_c\right\} \quad (7)$$

Using ICM, an estimate of α_v at that site is obtained analytically since the expression is quadratic in α_v . Hence

$$\hat{\alpha}_v = \frac{2\sigma_{\rho\rho}^2 [\sum_{\tilde{\alpha}} \tilde{\alpha} \lambda_\alpha + \bar{\alpha}_v \lambda_c] + (I_r - Y_r)M_r(1 - M_v)Y_v}{2\sigma_{\rho\rho}^2 [\sum_{\lambda_\alpha} \lambda_\alpha + \lambda_c] + (M_r(1 - M_v)Y_v)^2} \quad (8)$$

B. Mask Estimate

The estimates for the masks are generated using the following conditional at a pixel site

$$p(M_r|\alpha_r, \alpha_v, M_v, Y_r, Y_v, I_r, I_v) \propto \exp\left\{-\frac{1}{2\sigma_{\rho\rho}^2}(I_r - Y_r - M_r(1 - M_v)\alpha_v Y_v)^2 + \frac{1}{2\sigma_{\nu\nu}^2}(I_v - Y_v - M_v(1 - M_r)\alpha_r Y_r)^2 + \sum_{\tilde{M}} (M_r - \tilde{M})^2 \lambda_M + \beta_r(1 - M_r)\right\} \quad (9)$$

In this case the estimation is performed numerically since M_r is binary. Hence both $M_r = 0, 1$ are substituted in the expression above and whichever yields the greater probability is selected.

C. Clean Image Estimate

The estimates for the clean images are generated with the conditional $p(Y_r|\cdot)$ as follows.

$$p(Y_r|\alpha_r, \alpha_v, M_r, M_v, Y_v, I_r, I_v) \propto \exp\left\{-\frac{1}{2\sigma_{\rho\rho}^2}(I_r - Y_r - M_r(1 - M_v)\alpha_v Y_v)^2 + \frac{1}{2\sigma_{\nu\nu}^2}(I_v - Y_v - M_v(1 - M_r)\alpha_r Y_r)^2 + \frac{1}{2\sigma_{\rho\rho}^2}M_r(1 - M_v)(Y_r - \bar{Y}_b)^2\right\} \quad (10)$$

This distribution is clearly Gaussian, however instead of maximising the conditional as part of the ICM process, a sample is drawn from this distribution within $\pm T$ standard deviations of the mean. This strategy has been employed by other authors working in video and audio restoration [18]. The reasoning behind this is that the mean tends to generate oversmooth images, while using an unconstrained random draw is visibly chaotic. Therefore by drawing samples within some distance of the mean, a textural component in the underlying signal is allowed for and the iterative process performs better. The required draw is therefore $Y_r \sim \mathcal{N}(\bar{Y}_r, \sigma_Y^2)$. By completing the square in the conditional above the mean, \bar{Y}_r , and variance, σ_Y^2 , are extracted as follows.

$$\bar{Y}_r = \frac{I_r + M_r(1 - M_v)(\bar{Y}_b - \alpha_v Y_v) + \sigma^2(I_v - Y_v)M_v(1 - M_r)\alpha_r}{1 + \sigma^2(M_v(1 - M_r)\alpha_r)^2 + M_r(1 - M_v)} \quad (11)$$

where $\sigma^2 = \sigma_{\rho\rho}^2/\sigma_{\nu\nu}^2$

$$\sigma_Y^2 = \frac{\sigma_{\rho\rho}^2 \sigma_{\nu\nu}^2}{\sigma_{\nu\nu}^2(1 + M_r(1 - M_v)) + M_v(1 - M_r)\alpha_r \sigma_{\rho\rho}^2} \quad (12)$$

The estimate for σ_Y^2 clearly depends on estimates for $\sigma_{\rho\rho}^2, \sigma_{\nu\nu}^2$ and in the proposed algorithm these are generated by measurement from the data.

A secondary estimation method for Y_r , Y_v that is used every N iterations ($N = 10$ in this case) is to maximise conditionals that ignore the mask terms and prior, hence

$$p(Y_r | \alpha_r, \alpha_v, Y_v, I_r, I_v) \propto \exp - \left\{ \frac{1}{2\sigma_{\rho\rho}^2} (I_r - Y_r - \alpha_v Y_v)^2 + \frac{1}{2\sigma_{vv}^2} (I_v - Y_v - \alpha_r Y_r)^2 \right\} \quad (13)$$

Due to the quadratic nature of Eq.13 in Y_r , an estimate of Y_r at a site is obtained analytically as follows

$$\hat{Y}_r = \frac{I_r - \alpha_v Y_v + \sigma^2 (I_v - Y_v) \alpha_r}{1 + \sigma^2 \alpha_r^2} \quad (14)$$

Where σ^2 is defined as above.

D. Initialisation

Good initial estimates are required for ICM to converge usefully, see Figure 1. In this work the masks are initialised using k-means clustering on the intensity channel of the observed images, using 2 or 3 clusters, dependent on the severity of the bleed-through. Assuming 3 clusters are used, then the lightest cluster is considered to be background region where there is no bleed-through. The darkest cluster is then the foreground text, with the remaining cluster an initial estimate of the bleed-through regions. The mask penalties β_r, β_v are calculated from the initial mask estimates (for β_r) by creating another binary mask B_r coincident with the darkest two clusters from the recto k-means process. β_r is then configured as $100(1 - B_r)$. The masks and observed images are also used to estimate the noise variance for both recto and verso sides, based on the variance of regions where $B_r, B_v = 0$, also the estimates for mean background intensity \bar{Y}_b (for each side as appropriate) are taken to be the means of these regions. The initial mixing parameters are then obtained from the mask estimates and the observed images as follows (for α_v). Finally initial clean recto

At each pixel location

$\alpha_v = 0$

if recto is background ($M_r = 1$) **then**

if verso is foreground ($M_v = 0$) **then**

if $Y_r > \bar{Y}_b$ **then**

$$\alpha_v = \frac{Y_r - \bar{Y}_b}{Y_v}$$

end if

end if

end if

and verso estimates are obtained by substituting the relevant initial estimates into Eq.1.

E. Algorithm

The final algorithm may be enumerated as follows

- 1) Register the recto and verso sides of the image using the process outlined in Section III.
- 2) Initialise all variables as described above

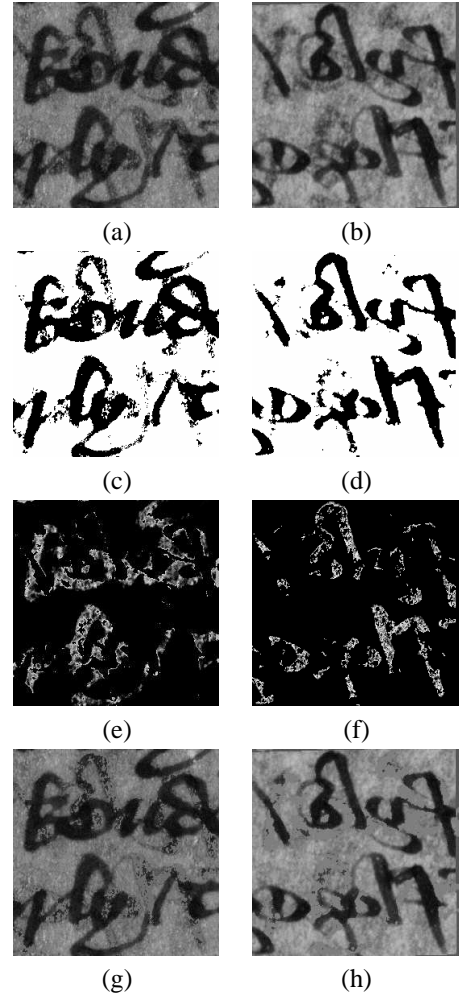


Fig. 1. (a),(b) degraded recto and verso images of the Piers example, (c),(d) initial mask estimates using k-means clustering ($k = 3$), (e),(f) initial estimates of α_r and α_v (white represents 1), (g),(h) initial clean recto and verso estimates.

- 3) Using a checkerboard visitation pattern for sites
 - a) Generate $\hat{M}_v, \hat{M}_r, \hat{\alpha}_r, \hat{\alpha}_v$ using the expressions above, (across all sites in separate image passes) updating in place.
 - b) i) If $\text{mod}(\text{iteration no.}, 10) = 0$ generate Y_r, Y_v without the masks, as in Eq.14
 - ii) Else draw samples for Y_r, Y_v , as above.
- 4) Terminate for iterations = 30.
- 5) Goto 3

VII. RESULTS

Small (255x255) recto/verso patches were used for testing, extracted from larger high resolution (600dpi) images. The manuscript images used were from *Piers Plowman* 'B' text from the late 14th century, fol.22, MS.201, Corpus Christi Library, Oxford (*Piers*), and a Welsh dictionary from 16th/17th Century, fol.1, MS.16 Jesus College Library, Oxford (*Welsh*). A further larger, lower resolution example was also tested

from fol.13 of the same *Piers Plowman* manuscript (*Piers2*). Figure 1(a,b) show the recto and verso sides respectively for *Piers*, and Figure 3(a,b) show the Welsh samples. These images have been pre-registered using the technique outlined in Section III. Initial parameter estimates are also shown in Figure 1. The initialization for the clean images shows some bleed-through removal already but the initialization for the mixing parameters are far from optimal. In the absence of ground truth, it is difficult to quantify results obtained as all evaluations will involve some degree of subjectivity. The results presented here are evaluated visually. Figure 2 shows output over 30 iterations on *Piers*. As can be seen the bleed-through is fairly well removed and the mask estimates match well to the area of the foreground text on the recto and verso sides. The mixing parameters also coincide well with the non-binary nature of the bleed-through. Figure 3 shows the results of the algorithm on Welsh. Only 5 iterations were required for this example as the degradation was relatively light. The bleed-through is, again, reasonably removed, however it is noticeable that in regions where the bleed-through area is larger than the originating text (as can be seen on the verso side) the algorithm performs poorly. In Figure 4 the results of the algorithm after 30 iterations on *Piers2* are shown. It is noticeable in this example that the image contains much more detail, and the bleed-through is very severe. However the restoration results were reasonable, with only the most extreme bleed-through regions remaining.

VIII. CONCLUSION

A relatively simple algorithm for bleed-through removal has been presented that relies on estimates generated from local pixel neighbourhoods via a linear degradation model in an ICM-based scheme. The process converges and yields good bleed-through removal in most areas. However, the model has been shown to be lacking in a few areas. Firstly, in *Piers* the verso image certainly has a more blurry appearance than the recto; the degradation model does not encompass this sharpness difference. This means that the spatial smoothness on both sides is not the same and that could influence estimation detrimentally. Secondly, it is noted that the algorithm does not perform optimally in cases where the bleed-through regions are much larger than the originating text due to the porosity of the medium (as in the Welsh verso example). Finally, as is the case with many bleed-through removal algorithms, this model does not perform optimally where *large* areas of bleed-through text are present that are comparable in intensity, or darker than the foreground text. These issues are being addressed currently

ACKNOWLEDGMENT

This research has been funded by the Irish Research Council for Science, Engineering, and Technology ‘Embark’ Initiative.

REFERENCES

- [1] C. Wolf, “Document ink bleed-through removal with two hidden markov random fields and a single observation field,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 32, no. 3, pp. 431–447, 2010.




Fig. 2. Iterations 5, 15, 30 for the *Piers* example. (a-c) Y_r , (d-f) Y_v , (g-i) M_r , (j-l) M_v , (m-o) α_r , (p-r) α_v . For iteration 15, the estimate for Y_r was obtained by setting $M_r = 1$ and $M_v = 0$ to increase smoothness on the final result, similarly for Y_v .

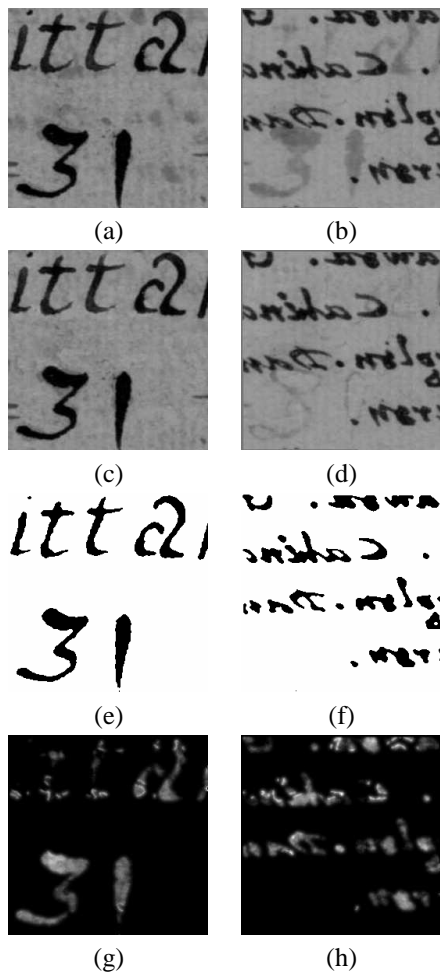


Fig. 3. Results for the Welsh example (a) I_r , (b) I_v , (c) Y_r , (d) Y_v , (e) M_v , (f) α_r , (g) α_v .

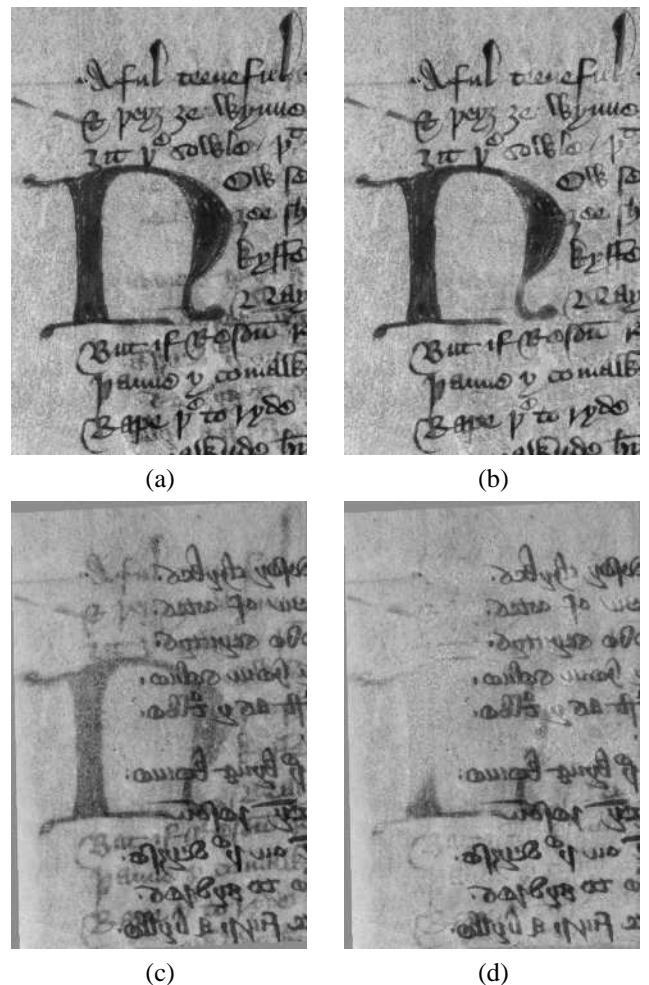


Fig. 4. Result after 30 iterations for the Piers2 example (a) I_r , (b) Y_r , (c) I_v , (d) Y_v

- [2] G. Sharma, "Cancellation of show-through in duplex scanning," in *Image Processing, 2000. Proceedings. 2000 International Conference on*, vol. 2, 2000, pp. 609–612 vol.2.
- [3] R. Rowley-Brooke and A. Kokaram, "Bleed-through removal in degraded manuscripts," in *IET Irish Signals and Systems Conference*, Trinity College, Dublin, Ireland, 2011.
- [4] R. Estrada and C. Tomasi, "Manuscript bleed-through removal via hysteresis thresholding," in *Document Analysis and Recognition, 2009. ICDAR '09. 10th International Conference on*, 2009, pp. 753–757.
- [5] F. Drira, F. Le Bourgeois, and H. Emptoz, "Restoring ink bleed-through degraded document images using a recursive unsupervised classification technique," in *Document Analysis Systems VII*, ser. Lecture Notes in Computer Science, H. Bunke and A. Spitz, Eds. Springer Berlin / Heidelberg, 2006, vol. 3872, pp. 38–49.
- [6] A. Tonazzini, L. Bedini, and E. Salerno, "Independent component analysis for document restoration," *International Journal on Document Analysis and Recognition*, vol. 7, no. 1, pp. 17–27, 2004.
- [7] A. Tonazzini, "Color space transformations for analysis and enhancement of ancient degraded manuscripts," *Pattern Recognition and Image Analysis*, vol. 20, no. 3, pp. 404–417, 2010.
- [8] A. Tonazzini, E. Salerno, and L. Bedini, "Fast correction of bleed-through distortion in grayscale documents by a blind source separation technique," *International Journal on Document Analysis and Recognition*, vol. 10, no. 1, pp. 17–25, 2007.
- [9] P. Castro, R. J. Almeida, and J. R. C. Pinto, "Restoration of double-sided ancient music documents with bleed-through," in *Progress in Pattern Recognition, Image Analysis and Applications*, ser. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2007, vol. 4756, pp. 940–949.
- [10] J. A. Burgoyne, J. Devaney, L. Pugin, and I. Fujinaga, "Enhanced bleedthrough correction for early music documents with recto-verso registration," in *International Conference on Music Information Retrieval*, Philadelphia, USA, 2008, pp. 407–412.
- [11] J. Sauvola and M. Pietikinen, "Adaptive document image binarization," *Pattern Recognition*, vol. 33, no. 2, pp. 225–236, 2000.
- [12] B. Gatos, I. Pratikakis, and S. J. Perantonis, "Adaptive degraded document image binarization," *Pattern Recognition*, vol. 39, no. 3, pp. 317–327, 2006.
- [13] Y. Huang, M. S. Brown, and D. Xu, "A framework for reducing ink-bleed in old documents," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, 2008, pp. 1–7.
- [14] —, "User-assisted ink-bleed reduction," *Image Processing, IEEE Transactions on*, vol. 19, no. 10, pp. 2646–2658, 2010.
- [15] R. F. Moghaddam and M. Cheriet, "Low quality document image modeling and enhancement," *International Journal on Document Analysis and Recognition*, vol. 11, no. 4, pp. 183–201, 2009.
- [16] —, "A variational approach to degraded document enhancement," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 32, no. 8, pp. 1347–1361, 2010.
- [17] E. Dubois and A. Pathak, "Reduction of bleed-through in scanned manuscript documents," in *IS&T Image Processing, Image Quality, Image Capture Systems Conference (PICS2001)*, vol. 4, Montreal, Canada, 2001, pp. 177–180.
- [18] J. O'Ruanaidh and W. Fitzgerald, *Numerical Bayesian Methods Applied to Signal Processing*. Springer, 1996.