

Bleed-through cancellation in non-rigidly misaligned recto-verso archival manuscripts based on local registration

Pasquale Savino, Anna Tonazzini, Luigi Bedini

the date of receipt and acceptance should be inserted later

Abstract Ancient manuscripts written on both pages of the sheet are frequently affected by ink bleeding from the reverse side. This phenomenon produces a significant degradation of both the foreground text and the general appearance of the manuscript. Effective digital image restoration techniques may require the use of the content of both document sides, thus needing their perfect alignment. Although often available, recto and verso are usually not aligned, either for rigid misalignments occurring during acquisition, or for non-rigid deformations of the sheet. In this paper we propose a novel method to restore color recto-verso manuscript images in a piecewise manner, without the need of a preliminary, global registration of the two sides. We assume that at the local level any deformation can be approximated by a displacement, and subdivide the two images into small patches of same size. For each pair of patches at the same location, their relative shift is estimated by cross-correlation, thus allowing their straightforward alignment. A bleed-through removal algorithm is then applied to the registered patches. By spanning the entire images, this procedure returns free of interferences versions of the images in their original acquisition layout. The experiments show that the restoration results so obtained are better than those obtained with the classical approach that first registers the whole recto-verso pair and then performs restoration. Further advantages are a much lower computational cost, the possibility to manage non-global and non-rigid deformations, and the unaltered geometry and color appearance of the two restored images.

The authors are with the Istituto di Scienza e Tecnologie dell'Informazione, Consiglio Nazionale delle Ricerche, Via G. Moruzzi, 1, I-56124 PISA, Italy. e-mails: pasquale.savino@isti.cnr.it, luigi.bedini@isti.cnr.it, anna.tonazzini@isti.cnr.it

Address(es) of author(s) should be given

1 Introduction

The seeping of ink from the reverse side (bleed-through) is a severe distortion affecting most ancient archival manuscripts. This degradation greatly impairs the legibility and fruition of the manuscript contents, besides being highly unpleasant. Physical restoration of these manuscripts cannot be performed, as the chemical substances used to remove the unwanted ink could also damage or even destroy the foreground ink. Thus, digitization and digital image processing is the only viable way to manage this problem. On one hand, digital acquisitions, especially the multispectral ones, could be a solution *per se*. Indeed, it might happen that, under some specific wavelengths, the seeping ink tends to fade. More frequently, image processing techniques applied to the digital images offer a number of different strategies to approach the problem. When the aim is to simply improve readability, thresholding techniques, either global or local, could be the first attempt to be made. However, the usually high variability of the intensity of the bleed-through pattern, which can be sometimes as dark as the foreground text, can make thresholding ineffective. On another hand, thresholding unavoidably removes also other features of the manuscript, such as paper watermarks and texture, and other marks (e.g. stamps or annotations), which could be useful to the scholar to study the origin of the document.

Since our aim is not only to improve readability, but also to remove the unwanted, uninformative interferences while preserving the original appearance and color of the manuscript, more sophisticated image processing techniques must be employed. Most methods working on a single side are based on classification of the pixels in background, foreground and bleed-through, and then on the inpainting of the identified bleed-through areas with a simulated background [1]. However, the techniques that uses both recto and verso are often more effective, since they

exploit all the available information [2–13]. As a counterpart, using recto-verso images entails their preliminary, very accurate alignment. Nowadays many libraries and archives have specialized digitization devices, e.g. high resolution CCD cameras, mounted on mechanical equipments that guarantee a stable setup. Nonetheless, misalignments frequently occur, due to the human intervention for repositioning the sheet when acquiring the verso image, or to accidental movements of the camera. Furthermore, the sheet might be intrinsically distorted in a non-rigid way.

Registration of recto-verso images is a challenging task, due to sparsity of bleed-through, different intensity of the same stroke in the two sides, and different degree of curvature of recto and verso in books. Most methods consider a global affine transformation, estimated by minimizing the intensity differences between the two sides [14], by using the Fourier-Mellin transform [15], or through a block-by-block strategy, where for each block in the recto the corresponding verso block is searched for in a larger window by spatial cross-correlation [16]. Projective transformations have also been proposed, e.g. in [17], where a feature-based method employs the corners detected from extracted character contours, and in [18], where the features are corresponding points located from the relative shifts of pairs of small recto-verso patches, computed via cross power spectrum. Non-rigid registration is proposed in [19], combining a global affine transformation with a free-form, hierarchical transformation based on B-splines and detected recto-verso corresponding points. A thin plate spline smoothness constraint is then applied for minimizing the residual complexity between the two sides, as done in [20]. In [21], a similar method is proposed, where the two sides are first globally aligned using the page outline, and then a local grid point warp is applied, combining the norm of differences between both intensity and gradient with a content preserving smoothness penalty.

In any case, the registration of the recto and verso images of a manuscript is a time-consuming task that can even become impossible if local deformations are present, as may happen when the documents are acquired from bound books, or when humidity produced bubbles on the paper. On the other hand, registration is not necessary *per se*, but it is just a way to make possible the spatial match of the information contained in the two sides, for analysis tasks that require its joint use. Thus, recto-verso registration can be considered as part of the whole bleed-through removal process.

The bleed-through cancellation method proposed in this paper overcomes the above discussed limitations and difficulties, since it does not require the global registration of the two entire document sides. Indeed, it is based on local registration of small recto-verso patches, performed jointly with their restoration. We assume that the deformation between the two sides, though not describable by a unique ge-

ometrical transformation, is locally rigid. This implies that, at the very local level, it can be approximated by a translation only. Thus, for every individual pair of image patches at the same location and of a given small size, we compute their relative shift through cross-correlation of their gradients. The verso patch exactly aligned with the recto patch is then found by simple translation. Any bleed-through cancellation algorithm, which exploits information from both sides but acts locally, can then be applied to the patches so aligned. By spanning the images with the sequence of all the adjacent, non-overlapping patches, this procedure returns the restored versions of the entire images.

We compare this patch-by-patch alignment plus restoration modality with the classical one, where the two images are first globally registered and then globally restored. The quality of the final restored images is taken as a measure of the quality of the registration as well. For a quantitative quality measure, we applied our method to the 25 pairs of the bleed-through database described in [22], and available at [23]. The pairs of this database are already registered, so that we artificially distorted each verso side through a projective transformation that is typical of acquisitions of recto-verso flat manuscripts. We also provide qualitative results obtained for real documents, affected by an unknown and probably non-rigid misalignment.

The paper is organized as follows. In Section 2, the whole procedure is described, and the method for computing the relative shift of the recto-verso patches is formalized and discussed. Section 3 regards the quantitative evaluation of the method on the bleed-through database of [22]. Section 4 is devoted to the discussion of the application of the method to real manuscripts pairs. Finally, Section 5 concludes the paper with the description of possible improvements and extensions.

2 Patch-by-patch recto-verso alignment and bleed-through cancellation

As anticipated in the Introduction, the method proposed here assumes that, for each side, joint registration and bleed-through removal is performed sequentially on each pair of the small, non-overlapping patches in which the images can be subdivided. We assume that, at the very local level, the reciprocal deformation between the two sides amounts to a translation only. For each small patch of one side, the homologous (i.e. of same size and at the same location) patch in the opposite side is first selected. The relative displacement between the two patches is estimated from the cross-correlation of their gradients, relying on the shift theorem of the Fourier Transform. The estimated shift is then used to locate the best matching opposite patch. We assume herein a perfect registration of the three color planes in the two sides, which means that the geometric deformation between

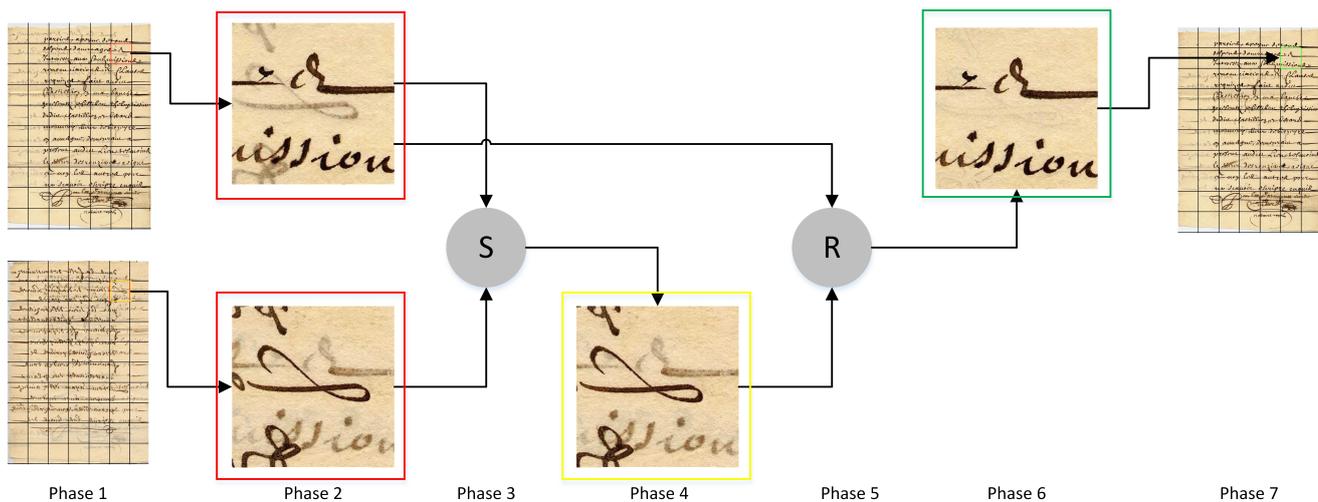


Fig. 1 Diagram illustrating the single-step restoration process of one recto patch

the recto and verso side is unique for the three pairs of homologous channels. As a consequence, we can compute the relative shifts between homologous patches in a single pair of channels only.

The bleed-through cancellation algorithm that we chose is an improvement of the one detailed in [12]. We would like however to highlight that any of the existing recto-verso restoration algorithms can be applied to the aligned patches, provided that it works locally. The specific algorithm used here was originally designed to simultaneously restore registered whole recto-verso pairs, either grayscale or RGB. This algorithm first identifies the bleed-through areas, based on a pixel-by-pixel criterium, and then fill-in them with estimated background values. In this work we apply the algorithm in a patch-by-patch modality, and improve the way to estimate locally the background values, now computed at the patch level.

We will distinguish between single-step restoration, to indicate the joint patch-by-patch alignment plus restoration method proposed here, and two-step restoration, to indicate that a preliminary, global registration of the two sides is performed prior patch-by-patch restoration.

The whole registration-restoration procedure is applied two times. The first time, we consider the recto as the reference image, and detect the pairs of patches by simultaneously moving a square window across the two sides in such a way to span the whole image domain. For each pair, we align the verso patch on the recto patch by simple translation, and then restore the only recto patch. Since the reference recto patches being restored are adjacent and non-overlapping, once all of them have been processed, the whole restored recto is readily obtained, without any need for interpolation across the restored patches themselves. Note that the restored recto remains geometrically unaltered. The second time the procedure is repeated by inverting the role of recto

and verso. Thus, this time, we obtain the geometrically unaltered restored verso.

In case of RGB images, the restoration of every patch is performed independently for the three color channels.

The diagram in Figure 1 illustrates the single-step restoration process of one recto patch. Phase 1 shows the grid of all the patches in which the reference recto image (top) and the verso image (bottom) are subdivided. Note that the number and size of the patches are not the actual ones used for this manuscript; we depicted larger patches for make clearer the interpretation of the diagram. The two homologous (at same position) patches at hand are highlighted with a red box, and individually shown, enlarged, at Phase 2. Their misalignment is apparent. The two patches are used as input for the computation of their cross-correlation, and then of their mutual shift (Phase 3). Based on the estimated shift, a different verso patch that best matches with the given recto patch can be now located in the verso image. For illustration purposes, in the whole verso image (Phase 1) we highlighted with a yellow box the location of this matching patch. The aligned verso patch is shown enlarged at Phase 4, and given as input, along with the original recto patch, to the restoration algorithm of Phase 5. This produces their free-of-interferences versions, shown at Phase 6 (only the recto is shown). Finally, the restored recto patch is put back in the whole recto image being restored (the patch highlighted with the green box at Phase 7). The procedure illustrated in the diagram is repeated for each pair of homologous patches, selected in any order. For clarity safe, in the diagram we refer to patches picked up in sequence from top to bottom and from left to right. As a matter of fact, in the recto of Phase 7 the patches that precede the one at hand are shown in their restored version, while the following ones are still degraded.

2.1 Alignment of the patches

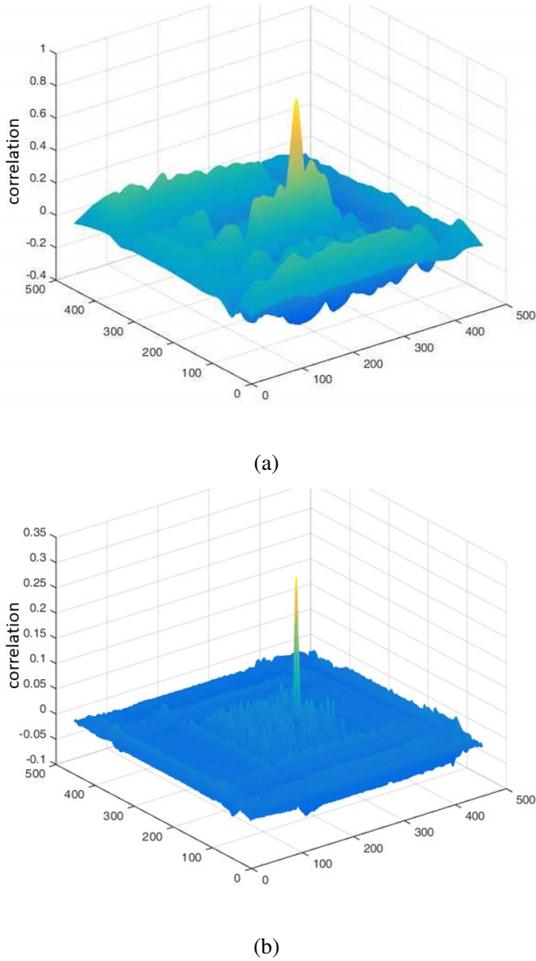


Fig. 2 Correlation matrix computed on the intensities (a) and on the gradients (b).

With reference to the physical manuscript, given two patches f_r and f_v such that, in a common support, it is $f_r(x + \Delta x_0, y + \Delta y_0) = f_v(x, y)$, the following relationship between their FTs holds true:

$$F_v(\omega_x, \omega_y) = F_r(\omega_x, \omega_y) e^{j(\omega_x \Delta x_0 + \omega_y \Delta y_0)} \quad (1)$$

from which:

$$\frac{F_v(\omega_x, \omega_y)}{F_r(\omega_x, \omega_y)} = e^{j(\omega_x \Delta x_0 + \omega_y \Delta y_0)} \quad (2)$$

Theoretically, the inverse FT of the ratio in eq. (2) would return a delta of Dirac impulse located in $(\Delta x_0, \Delta y_0)$. Nevertheless, as always happens when inverse filtering is used, due to the presence of random noise, dissimilar parts and gain changes, this operation produces a very noisy map, where a trustable peak cannot be located. A more robust estimate of the cross correlation between f_r and f_v is given by

the inverse FT of the cross power spectrum:

$$\frac{F_v(\omega_x, \omega_y) \cdot F_r^*(\omega_x, \omega_y)}{|F_v(\omega_x, \omega_y)| \cdot |F_r^*(\omega_x, \omega_y)|} = e^{j(\omega_x \Delta x_0 + \omega_y \Delta y_0)} \quad (3)$$

where $*$ denotes the complex conjugate. The location of the now well emerging peak of the cross correlation function so computed defines the relative displacement between the two patches.

To further improve the estimation of the shifts, we compute cross-correlation of the gradients of the patches, rather than of their intensities. Indeed, though depicting the same scene, recto-verso images have different intensities. Specifically, dark strokes in one side (foreground text) are lighter in the other (bleed-through pattern), and viceversa. As for images of a same scene taken with different sensors (e.g., RGB color channels), we thus inferred that recto and verso mainly correlate in correspondence of the object borders or textures, and that a measure of correlation is more reliable when performed on the gradients of the patches.

Figure 2 shows the correlation matrices computed on the intensities and on the gradients, respectively, for a typical pair of patches. Note how the peak of the gradient correlation matrix is much better defined than that of the intensity correlation matrix.

Whether using intensities or gradients, for an effective estimation of the relative shift the two patches must share some common strokes, either see-through in one side and foreground in the opposite side, or mixed see-through and foreground in both sides. Indeed, it is apparent that when the common portion of text is very little, the estimate might be inaccurate. Thus, the patches must be large enough to share a significant portion of text, while they must be small enough to assume that their misalignment can be approximated by a translation only. To simultaneously satisfy these opposite conditions, a good compromise is to make the size of the patch depending on the character size. For recto-verso pairs as large as around 3000×4500 , which is typical of manuscript letters of the XVI–XVII centuries acquired at very high resolution, we experimentally found that patches whose size is between 150×150 and 200×200 give satisfactory results. When it is apparent that in some regions the misalignment is high, so that the two patches could share a too small portion of common text, it may also be convenient to compute cross-correlation between the recto patch and an enlarged window containing the homologous verso patch.

As a further drawback, it is clear that when one of the two patches or both are pure background, the shift computed with the method above is meaningless.

To cope with these inaccurate or meaningless shift estimates, the shifts for all the pair of patches are computed off-line, possible outliers are identified on the basis of their deviation from local means, and then corrected by averaging them with their four neighborhoods. In Figure 4, the

```

function [RectoRestored] = SSR(Recto,Verso,PatchSize)

[RowRectoSize, ColRectoSize] = size(Recto);
% N and M are the number of patches in vertical and horizontal direction
N = floor(RowRectoSize/PatchSize);
M = floor(ColRectoSize/PatchSize);
for i = 1:N
    for j = 1:M
        RectoPatch = selectPatch(Recto,i,j,PatchSize);
        VersoPatch = selectPatch(Verso,i,j,PatchSize);
        % Shift contains the horizontal and vertical shifts needed to
        % align the recto and verso patches
        Shift(i,j) = calculateShift(RectoPatch,VersoPatch);
    end
end
Shift = correctOutliers(Shift);

for i = 1: N
    for j = 1: M
        RectoPatch = selectPatch(Recto,i,j,PatchSize);
        VersoPatchShifted = shiftPatch(Verso,i,j,Shift(i,j),PatchSize);
        % RectoBg and VersoBg are the average background values for
        % the recto patch and the aligned verso patch.
        RectoBg = computeBackground(RectoPatch);
        VersoBg = computeBackground(VersoPatchShifted);
        [RectoPatchRestored] = restorePatch(RectoPatch,VersoPatchShifted,RectoBg,VersoBg);
        RectoRestored = composeImage(RectoPatchRestored,i,j,PatchSize);
    end
end

```

Fig. 3 Pseudocode of the function SSR implementing the single-step restoration algorithm.

maps of the (x,y) values of the shifts of a few, adjacent recto-verso patches are shown before and after the correction of the outliers.

2.2 Implementation of the single-step restoration algorithm

The pseudo code of the function SSR, which implements the single-step restoration algorithm, is shown in Figure 3. The input parameters of the function are the `Recto` and `Verso` images and the size in pixel of the patches (`PatchSize`). The function returns the restored recto image (`RectoRestored`). For easier understanding, in the reported pseudo code we made some simplifying assumptions: the images are considered to be graylevel, and the number of rows and columns in the images are assumed to be an exact multiple of the patch size. The pseudo code can be easily modified to account for color images of any size, as the actual implemented algorithm does.

Initially, the function determines the number of horizontal and vertical patches (N and M). The subsequent loop on N and M selects the various recto-verso patch pairs in the images, calculates their mutual shifts through the function `calculateShift`, which implements the method described in Section 2.1. The shift of each verso patch with respect to the homologous recto patch is stored in the matrix

`Shift`, whose size is $N \times M$. The shift matrix is then corrected for possible outliers, as explained at the end of Section 2.1, through the function `correctOutliers`.

The second loop selects again the recto patches and, by means of function `shiftPatch`, the shifted verso patches. Then, for each pair of patches, it estimates the average background values `RectoBg` and `VersoBg`, respectively, through the `computeBackground` function, and performs the restoration of the recto, through the function `restorePatch`. This function implements the bleed-through cancellation algorithm detailed in [12]. Each restored recto patch is then inserted in the proper position into the restored recto image `RectoRestored`, through function `composeImage`.

3 Quantitative analysis

We quantitatively measured the results of our single-step method on the 25 pairs of the entire database in [23], and compared them with the results of the two-step method, based on the registration algorithm proposed in [18]. The pairs of this database are already registered, so that, assuming the alignment to be optimal, we also measured on them the quality of the restoration algorithm alone, and used these measurements as a baseline.

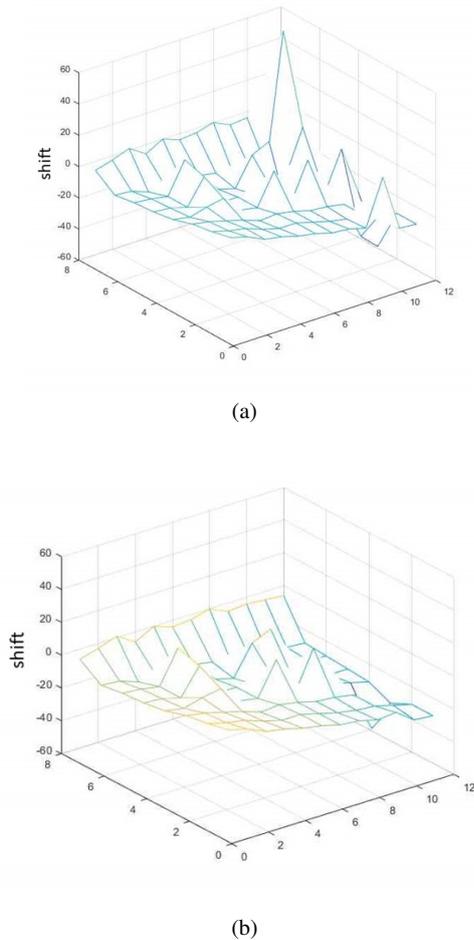


Fig. 4 Typical map of the relative shifts estimated for a few adjacent recto-verso patches: (a) with outliers; (b) after correction of the outliers.

In order to test the single-step and the two-step methods on this database, we artificially distorted the verso side of each image pair through a typical projective transformation previously estimated on a real pair of manuscripts. The matrix P of the used projective transformation is the following:

$$P = \begin{bmatrix} 0.969 & -0.016 & -1.071e-05 \\ -0.002 & 0.983 & 3.621e-07 \\ 16.181 & 19.539 & 0.999 \end{bmatrix} \quad (4)$$

where the various coefficients account for translations, scale factors and projective deformations.

To illustrate the complete procedure described in the previous section we refer to the third pair of the database, shown in Figures 5 (a) and 5 (b).

We first considered the undistorted, pre-registered pair available in the database, and compared the results of the patch-by-patch restoration algorithm (Figures 5 (e) and 5 (f)) with those furnished by the algorithm in [11] (Figures 5 (b) and 5 (c)). Note that, in the verso side, a few bleed-

through strokes left in the result of [11] (highlighted with the red box) are fully removed by our algorithm, while some foreground strokes at the bottom of the recto side that are lost in the result of [11] are preserved by our algorithm.

We then show, in Figure 6 (b), the verso distorted according to the projective transform of eq. (4). Figures 6 (e) and (f) show the result of the application of the patch-by-patch, single-step alignment/restoration procedure proposed (only the recto is shown). For comparison, Figures 6 (c) and (d) show the result of the two-step restoration, i.e. restoration applied after the preliminary global registration of the pair in Figures 6 (a) and 6 (b).

Note that the quality of the results provided by the two methods is similar. This was largely expected, since, in this case, the artificial deformation is globally rigid. However, in the result of the two-step method some bleed-through borders remain unremoved, and some strokes of the foreground text are not preserved. This seems to indicate that, due to computational approximations when applying the deformation to the verso, and the necessary interpolation to render the digital image, the actual deformation is not truly global. Thus, acting locally results to be beneficial anyway.

For each image in the database a binary ground-truth mask of the foreground text is provided. Although these ground-truth images are synthetic, i.e. created manually, some authors have used them for a quantitative analysis of the results, comparing them with the binarized versions of the restored images. Since the ground-truth masks are available for the undistorted images, we can compare them with our binarized results only for the recto side. Figure 7 shows the available ground-truth for the original recto of Figure 6 (a) compared with the binarized restored recto of Figure 6 (c).

As binarization algorithm we used the adaptive Sauvola algorithm [24]. As quality indices we computed the probability $FgError$ that a pixel in the foreground text was classified as background, the probability $BgError$ that a background or bleed-through pixel was classified as foreground, and the $WTotError$, that is, the weighted mean of $FgError$ and $BgError$, with the weights being the numbers of the foreground pixels and the background pixels as they result from the corresponding ground-truth images. The weighted total error $WTotError$ indicates the probability that any pixel in the image was misclassified. According to [22], these quality indices are defined as:

$$\begin{aligned} FgError &= \frac{1}{N_{Fg}} \sum_{t \in GT(Fg)} |GT(t) - B(t)| \\ BgError &= \frac{1}{N_{Bg}} \sum_{t \in GT(Bg)} |GT(t) - B(t)| \\ WTotError &= \frac{N_{Fg} FgError + N_{Bg} BgError}{N} \end{aligned} \quad (5)$$

where GT is the ground truth, B is the binarized restoration result, $GT(Fg)$ is the foreground region of the ground-truth

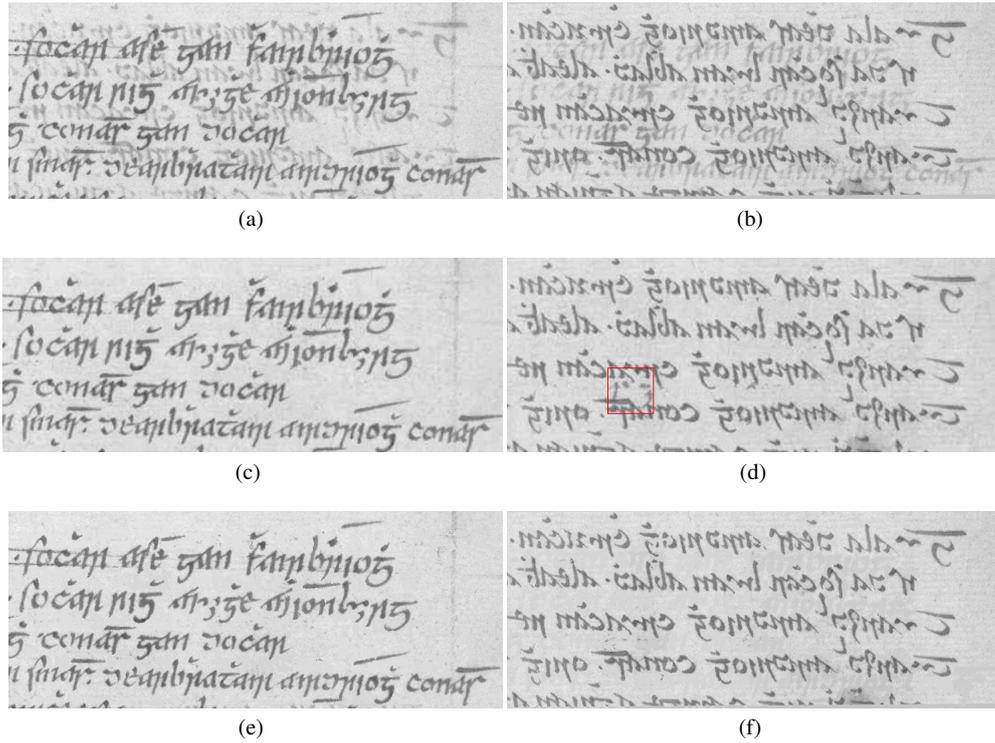


Fig. 5 Application of the patch-by-patch restoration to a real pre-registered graylevel recto-verso pair: (a) original degraded recto; (b) original degraded verso; (c) recto restored with the algorithm in [11]; (d) verso restored with the algorithm in [11]; (e) recto restored with our patch-by-patch restoration algorithm; (f) verso restored with our patch-by-patch restoration algorithm. Original images (a) and (b) digitized by Irish Script On Screen (www.isos.dias.ie).

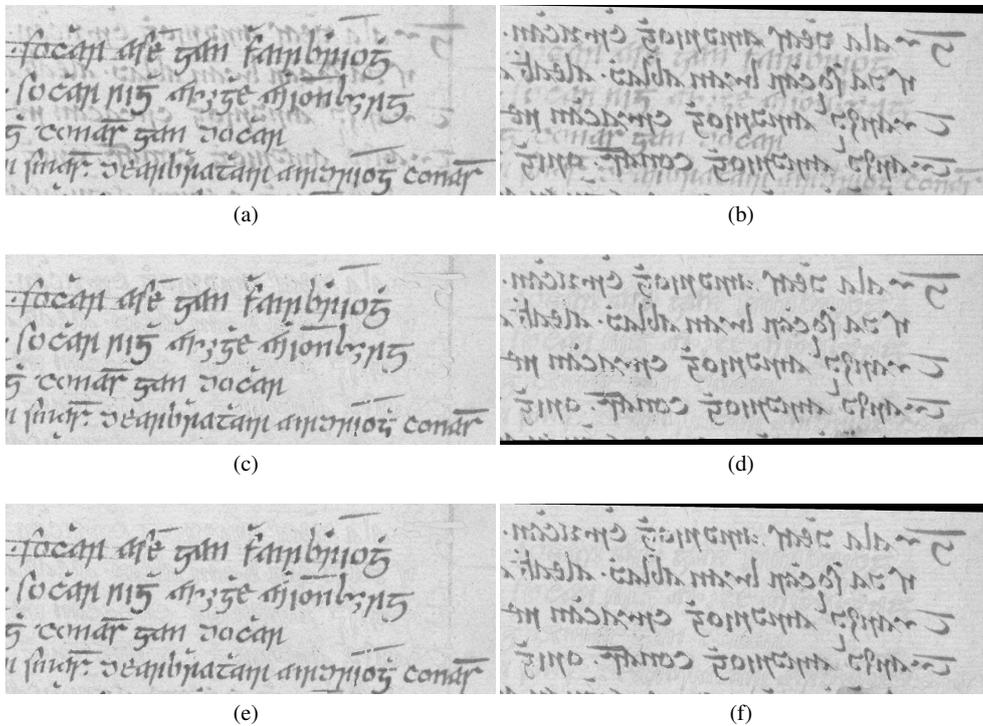


Fig. 6 Application of the single-step alignment/restoration to the pair of Figures 5 (a) and 5 (b), where the verso side has been geometrically distorted by the projective transformation of eq. (4) : (a) original degraded recto; (b) original degraded and distorted verso; (c) two-step restored recto; (d) two-step restored verso; (e) single-step restored recto; (f) single-step restored verso.

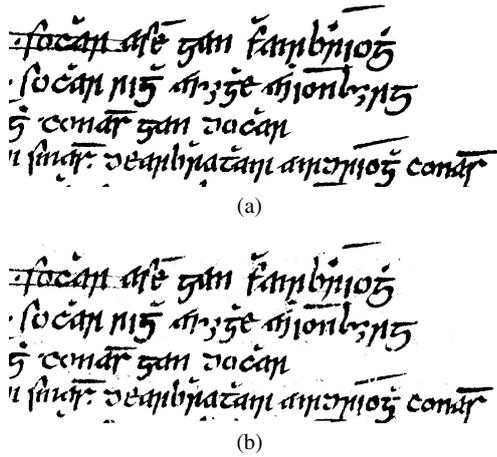


Fig. 7 Comparison between the binarized version of the restored recto of Figure 6 (e) and the available ground-truth: (a) manually generated ground-truth of the undegraded recto; (b) binarized restored recto. Original image (a) provided by Irish Script On Screen (www.isos.dias.ie).

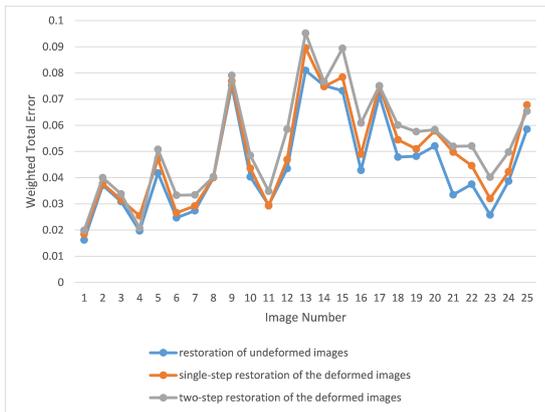


Fig. 8 Plots of the Weighted Total Errors for the 25 images of the database in [23]. Blue line: restoration of the undistorted images; red line: single-step restoration of the distorted images; gray line: two-step restoration of the distorted images.

image constituted of N_{Fg} pixels, $GT(Bg)$ is the complementary background region of the ground-truth image constituted of N_{Bg} pixels, and N is the total number of pixels in the image.

The plots of Figure 8 compare the $WTotError$ quality measure obtained, at each image, with the results of three methods: i) the patch-by-patch restoration applied to the original, undistorted pair (blue line), which we consider as our reference; ii) the single-step alignment/restoration applied to the distorted pair (red line); iii) the two-step registration/restoration applied to the distorted pair (gray line).

From the plots, it is apparent that the combined alignment plus restoration method proposed performs almost identically to the restoration algorithm applied to the original aligned pairs, whereas it is much better than the two-step method where the pairs are registered off-line. We already

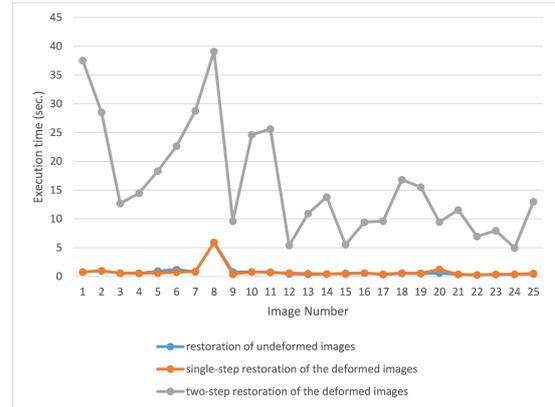


Fig. 9 Plots of the execution times (in seconds) for the 25 images of the database in [23]. Blue line: restoration of the undistorted images; red line: single-step restoration of the distorted images; gray line: two-step restoration of the distorted images.

commented the probable reasons of this behaviour from the qualitative analysis of the results of Figure 6.

Finally, it is worth highlighting again the simplicity of our method, which leads to a very fast algorithm. The execution times for each of the 25 images are reported in the plots of Figure 9, compared with those of the two-step algorithm.

To conclude the analysis of our method on artificially distorted images of the database in [23], we attempted bleed-through cancellation on the same pair shown in Figure 5, where the verso has been now distorted through an elastic warping operator available as a tool of Photoshop. This warping operator was manually applied in different ways to various areas of the image.

In a first example, we tried to simulate the typical situation occurring when the acquisition is made from a book. We left the recto side flat, and reproduced the effect due to the curvature of the verso side in correspondence of the binding regions of the book. Figures 10 (a) and 10 (b) show the distorted verso and its superposition in transparency with the original, flat verso. Figures 10 (c) and 10 (d) show the restored recto and verso sides obtained through the single-step method.

In a second example we applied to the verso a stronger local, elastic deformation, not necessarily corresponding to a real situation. Our aim was just to qualitatively evaluate the robustness of our shift-based patch-by-patch alignment against locally elastic deformations. Figures 11 (a) and 11 (b) show the distorted verso and its superposition in transparency with the original, flat verso. It is apparent the non-stationarity and non-rigidity of the deformation, which produces significant scale changes in some image areas.

Despite the fact that our method, at present, does not treat scale changes, the mechanism of working on small patches is partially able to overcome this inconvenience. The results of the single-step method, shown in Figures 11 (c)

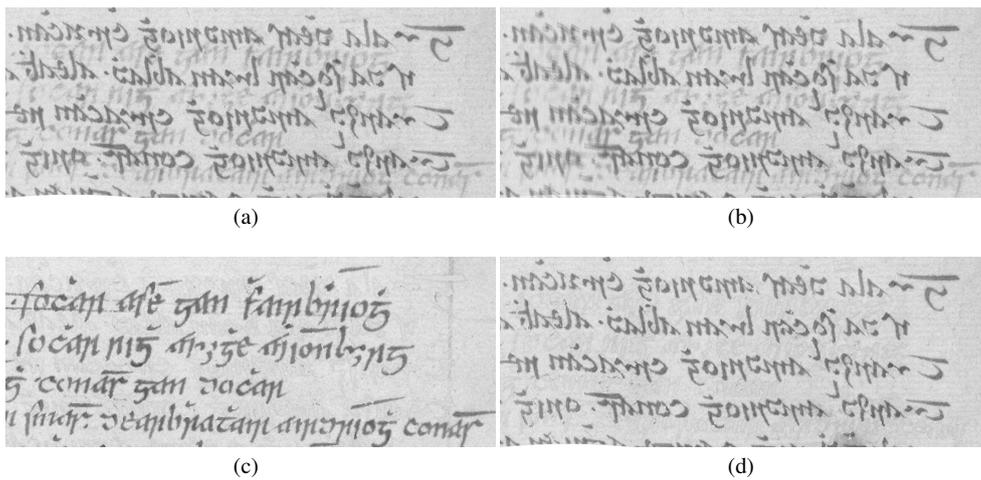


Fig. 10 Application of a local warping to the verso of Figure 5 (b): (a) distorted verso; (b) superposition in transparency of the original verso and the distorted verso; (c) restored recto; (d) restored verso.

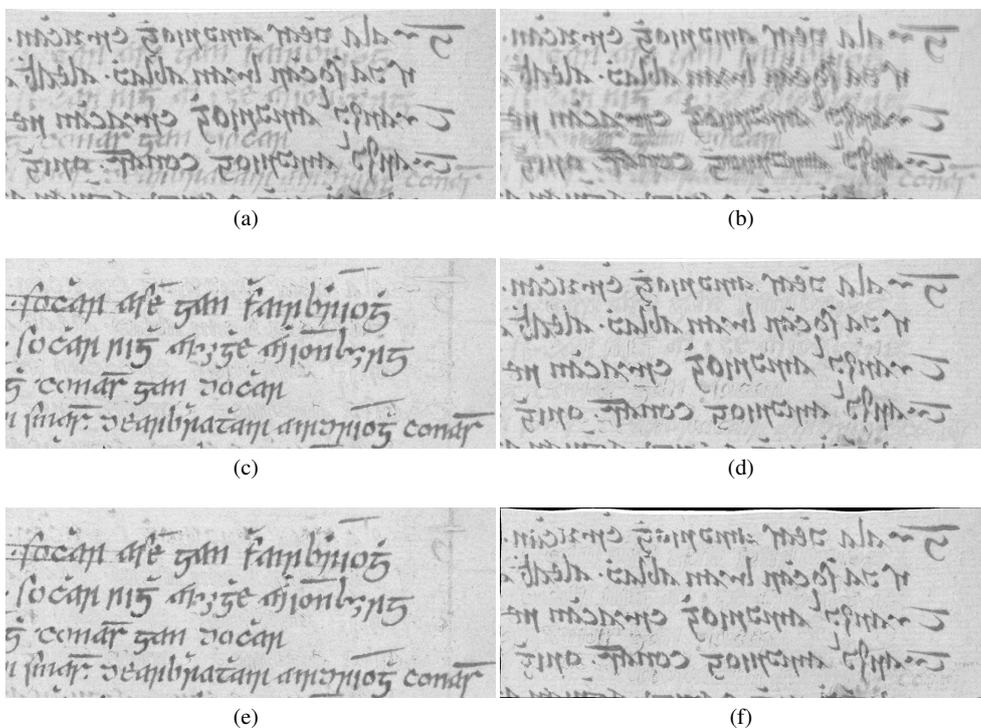


Fig. 11 Application of local elastic warpings to the verso of Figure 5 (b): (a) distorted verso; (b) superposition in transparency of the original verso and the distorted verso; (c) recto restored by the single-step method; (d) verso restored by the single-step method; (e) recto restored after global registration through the Elastix toolbox; (f) verso restored after global registration through the Elastix toolbox.

and 11 (d), can be considered satisfactory, though the bleed-through is not perfectly removed. It is worth to highlight that no rigid global registration technique could be able to make flat the distorted image, so that we did not apply the two-step method in this situation. Instead, we attempted to flatten the distorted verso (i.e. to align it on the flat recto), by using a recently proposed image registration software toolbox, namely Elastix [25], designed for the non-rigid registration of elastically distorted medical images.

Elastix exploits local similarity measures (local correlation and local mutual information), which are well suited to cope with image intensity inhomogeneity and non-stationarity, to iteratively update the transformation using gradient descent. Elastix provides a broad range of intensity-based registration options and settings, so that it has many parameters to be set. We run the code several times to empirically find the parameter values that produced the best results. The aligned recto-verso pair was given as input to

the patch-by-patch restoration algorithm, whose results are shown in Figures 11 (e) and 11 (f). As it can be appreciated, the quality of the restored images obtained with the two procedures is similar.

3.1 Robustness of patch alignment

In Section 2.1 we discussed the necessary conditions for a reliable estimation of the mutual shifts between pairs of recto-verso patches through cross-correlation of the intensities or the edges. These conditions mainly regard the size of the patches, which must be large enough to make the two recto-verso patches share a significant portion of text, and small enough to assume that their misalignment can be approximated by a translation only. We also provided practical reliefs to fix possible inaccurate or meaningless shift estimates. However, it is apparent that also the measurement of correlation between two signals can be unreliable itself, for instance in case of large differences in the amplitude of the two signals. In our specific application, this happens when the degree of ink seeping is low, so that there is a large difference between the intensity (or gradient) of the bleed-through pattern and that of the foreground text that generated it.

To test the robustness of the patch alignment strategy against different degrees of bleed-through, we performed the following synthetic experiment. We used the first pair of images of the public dataset in [22] [23], since for those images binary ground-truths of the clean recto and verso sides are available. We generated a synthetic, typical document background, one for the recto and another one for the verso, and then placed on them the two clean foreground texts. In this way, we obtained clean recto and verso images. We then mixed the two clean images according to the density model described in [12], by using different percentages of ink seeping. We considered values of bleed-through percentage ranging from 5% to 50%. For simplicity sake, and without loss of generality, the mixing model was considered space-invariant in this experiment. As the generated recto and verso images are perfectly aligned, we then applied to the verso image a fixed, global shift, whose values are comparable to the typical shifts that we found in real recto-verso document images. We applied a deformation constituted by a translation only since this allows us to precisely evaluate the quality of the patch alignment process. Figure 12, first and second row, shows the unshifted recto-verso images for some of the degrees of bleed-through analyzed.

At each percentage of bleed-through, we computed cross-correlation of each pair of all the adjacent and non-overlapping recto-verso patches in the two images. The estimated mutual shift values were compared with the true shift applied to the verso image, which is known by construction.

We computed the average of the shifts estimated for the entire set of patches, and their histogram, i.e. the percentages of patches returning the same shift value. In case of perfect shift estimation for all patches, the average should be equal to the true shift, and the histogram should consist of a single impulse, located in correspondence of the true shift value.

We have experimentally found that for percentage values of bleed-through higher than 20% the average shift values in the x and y directions are always almost equal to the true shift, whereas the histogram peak is exactly located in the true shift, for both x and y directions (see Figure 12, third and fourth row). For each percentage of bleed-through, we also compared the quality of restoration in the cases of shifted and unshifted verso image. For percentage values of bleed-through higher than 20%, the weighted total errors for shifted and unshifted verso images are comparable.

When the percentage values of bleed-through are lower than 20%, shift estimation returns random values, so that the application of the restoration algorithm is meaningless. On the other hand, for those percentages bleed-through is practically imperceptible, as it can be appreciated from Figure 12, and it can be removed by simple thresholding.

4 Experiments on real color manuscripts

In this section we will discuss the results of an experiment performed on a recto-verso RGB manuscript. Figure 13 shows the original recto and reflected verso images, together with their superposition in transparency.

From Figure 13 (c) it clearly appears an unidentifiable misalignment between the two sides. We first attempted the two-step restoration modality, assuming a projective deformation between the two sides. This produced the results in Figures 14 (a) and 14 (c), where some bleed-through strokes are still visible. We inferred that the relative deformation between the two sides is likely to be non globally rigid. The application of the single-step method proposed in this paper gives the much better results of Figures 14 (b) and 14 (c).

The observation of enlarged details allows a more comprehensive discussion of the differences between the two methods. Figures 15 (a) and 15 (b) show two homologous areas in the original recto and verso sides. Figure 15 (c) show the verso area after the global registration of the two whole images. It is apparent that, while the horizontal shift component of the mutual deformation has been corrected, the vertical shift component is only little reduced. As a consequence, the restoration algorithm was not able to remove some bleed-through strokes (see Figure 15 (e)). The improved result obtained with the single-step algorithm is shown in Figure 15 (d). Note also that, besides exhibiting an imperfect correction of the geometrical deformation, the detail of Figure 15 (c) appears manifestly smoother than its

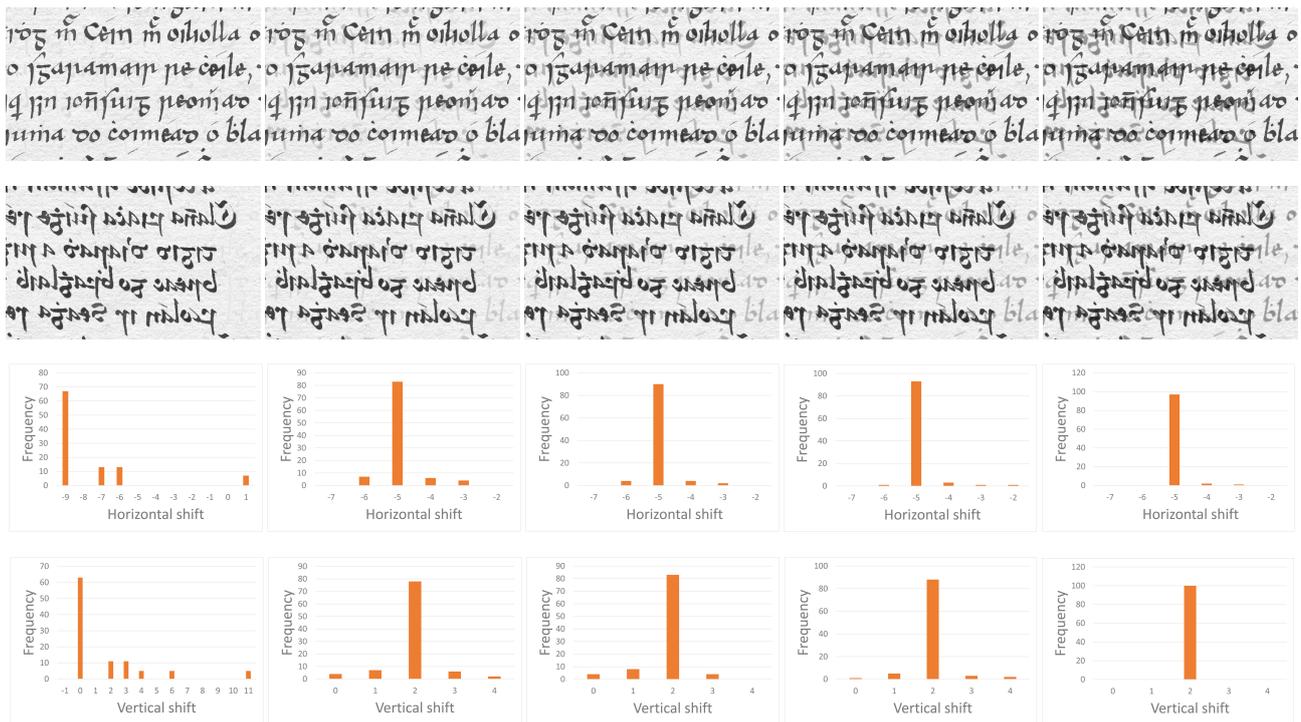


Fig. 12 Evaluation of the robustness of patch alignment for different degrees of bleed-through. First and second row: recto and undeformed verso with (from left to right) 10%, 20%, 30%, 40%, 50% of bleed-through; third and fourth row: corresponding histograms of the horizontal and vertical shifts computed on the entire set of patches.

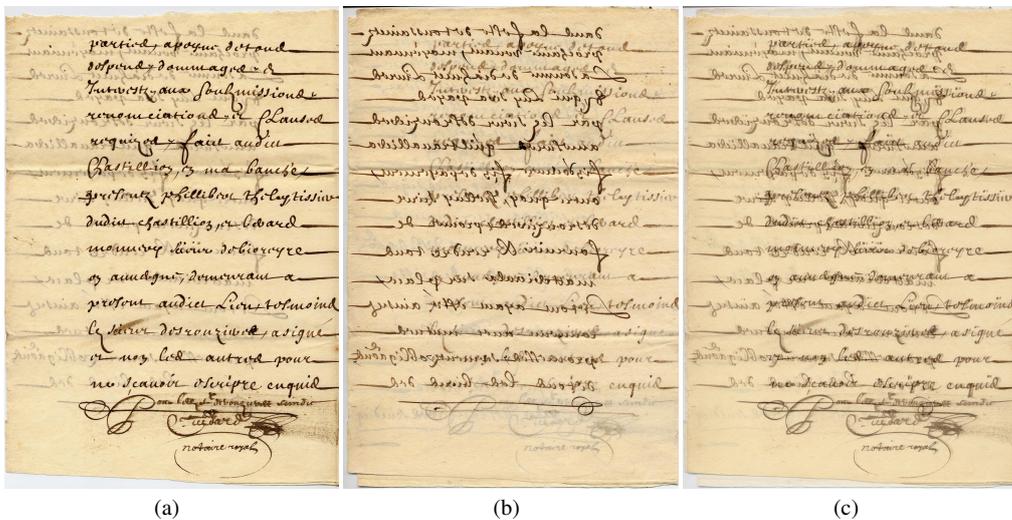


Fig. 13 A real RGB manuscript: (a) original degraded recto; (b) original degraded (mirrored) verso; (c) superposition in transparency of the original recto and the mirrored verso.

original. This degradation persists in its poorly restored version of Figure 15 (e). Oversmoothing is a typical effect of applying an off-line registration to the image pair, since interpolation of the verso pixels is necessary in order to correct the geometric deformation. This unpleasant effect is absent in the result of the single-step method, since we only admit translations between the patches. Note also that the restored area has maintained its geometrical, original asset.

Another detail is shown in the subsequent Figure 16. Here the two-step method performs better than in the previous detail (see Figure 16 (c)), meaning that the local misalignment has been satisfactorily corrected by the global registration process. Consequently, the bleed-through strokes are almost completely removed. However, due to the little misalignment still left, their borders remain unre-moved, making the bleed-through still visible. The result of

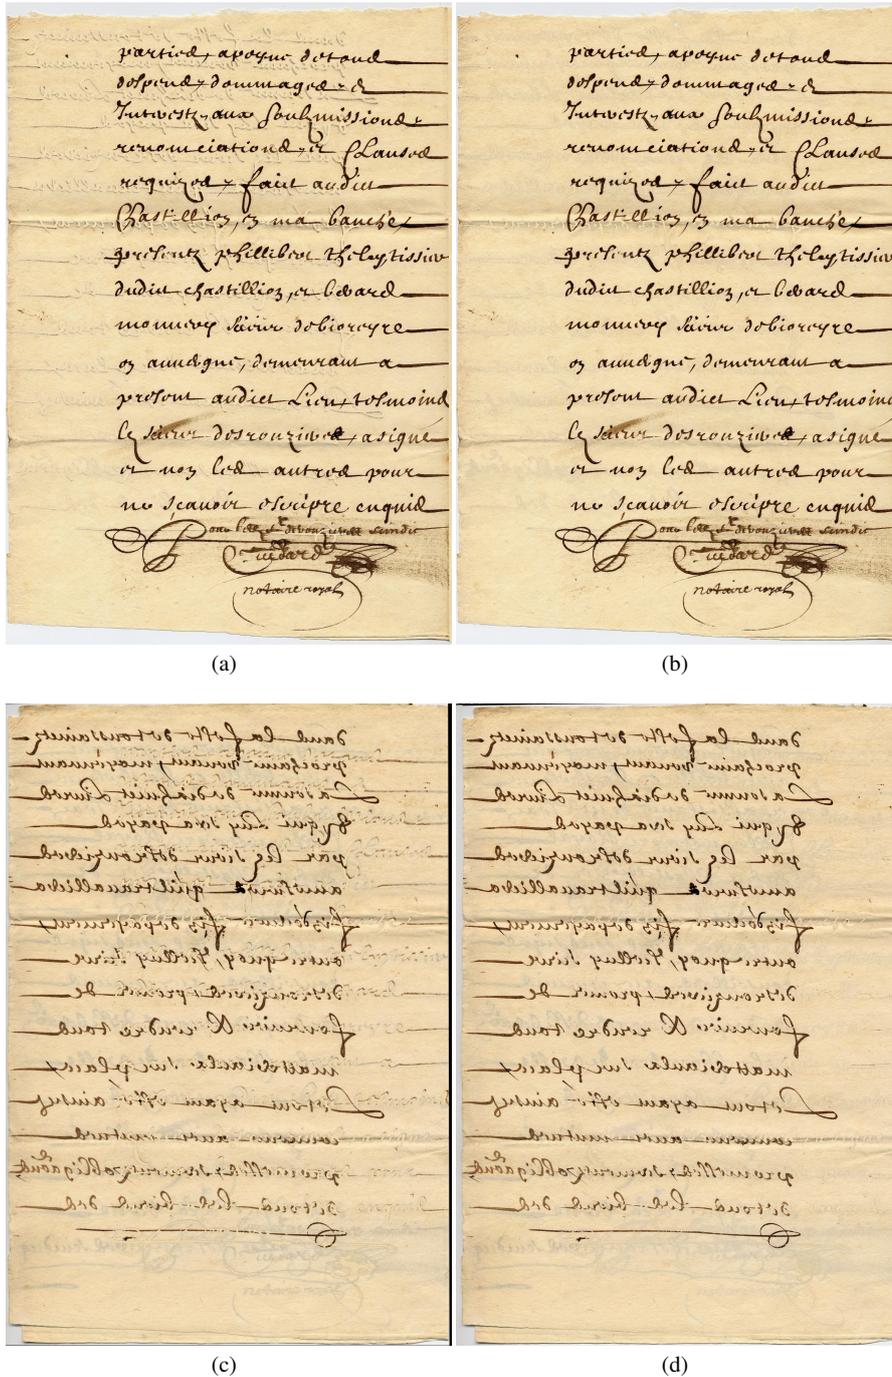


Fig. 14 Restoration to the misaligned pair of Figures 13 (a) and 13 (b): (a) recto restored with the two-step method; (b) recto restored by the proposed single-step method; (c) verso restored with the two-step method; (d) verso restored by the proposed single-step method.

the single-step method is superior, as shown in Figure 16 (b).

5 Conclusion and future work

We proposed a fully automatic procedure for the joint registration and restoration of recto-verso misaligned

manuscripts affected by bleed-through, assuming the relative deformation to be locally rigid. For each pair of small homologous patches in the two sides, the verso is registered on the recto by a simple translation, estimated by cross-correlation of image gradients, rather than image intensities. Then, a pixel-by-pixel identification of the bleed-through pattern, followed by inpainting with locally estimated back-

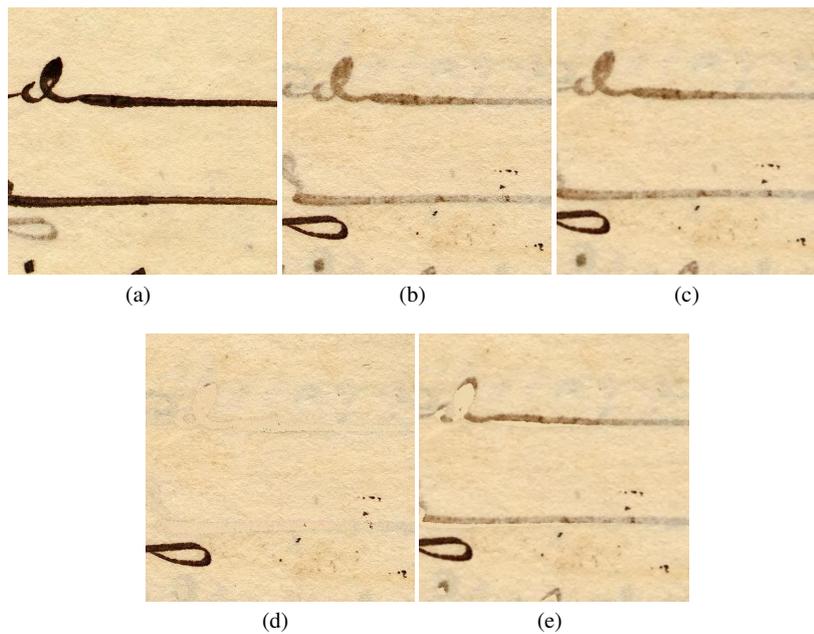


Fig. 15 An enlarged detail of the manuscript in Figures 13 and 14: (a) original recto; (b) original verso; (c) verso aligned on the recto through global image registration; (d) verso restored with the single-step proposed method; (e) verso restored after preliminary, global image registration.

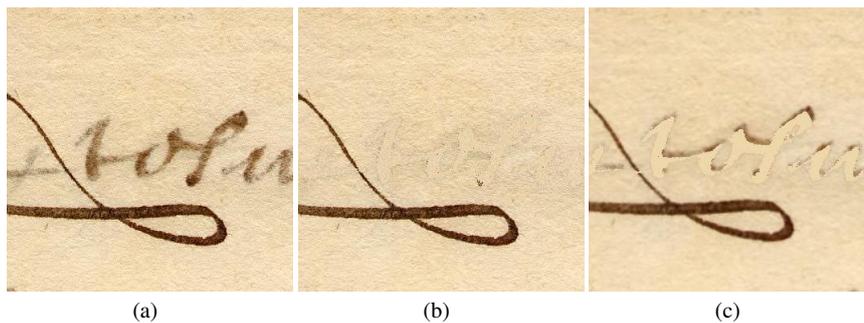


Fig. 16 An enlarged detail of the manuscript in Figures 13 and 14: (a) original verso; (b) verso restored with the single-step proposed method; (c) verso restored with the two-step method.

ground values, is performed. The procedure is repeated by inverting the two sides, thus furnishing their restored versions while leaving unaltered the original geometric appearance. The results show a significant improvement of the local registration performance, with a consequent much more effective removal of the bleed-through, when compared with the classical two-step procedure that globally aligns the two sides off-line. Furthermore, the single-step procedure is much faster. We experimented the procedure also in the case of mild elastic deformations, obtaining a satisfactory performance as well.

Straightforward future studies could consist in experimenting this patch-by-patch joint alignment and restoration with other bleed-through removal algorithms available in the literature. Furthermore, the presently adopted restoration algorithm could be improved, with respect to both the identification step and the inpainting step. For instance, in [26]

we proposed an inpainting technique based on image sparse representation and dictionary learning.

As per the registration, future investigations could regard the extension of the method to stronger and locally varying elastic deformations of the sheet. At the patch level, such deformations could be approximated by affine transformations accounting also for scale changes, and estimated, e.g., by a Fourier-Mellin transform.

References

1. 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.
2. A. Tonazzini, E. Salerno, and L. Bedini, "Fast correction of bleed-through distortion in grayscale documents by a blind source separation technique," *Int. Journal on Document Analysis and Recognition*, vol. 10, pp. 17–25, June 2007.

3. G. A. Hanasusanto, Z. Wu, and M. S. Brown, "Ink-bleed reduction using functional minimization," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2010, pp. 825–832.
4. Y. Huang, M. S. Brown, and D. Xu, "User assisted ink-bleed reduction," *IEEE Transactions on Image Processing*, vol. 19, no. 10, pp. 2646–2658, 2010.
5. A. Tonazzini, I. Gerace, and F. Martinelli, "Multichannel blind separation and deconvolution of images for document analysis," *IEEE Trans Image Processing*, vol. 19, no. 4, pp. 912–925, April 2010.
6. F. Merrikh-Bayat, M. Babaie-Zadeh, and C. Jutten, "Using non-negative matrix factorization for removing show-through," in *Proc. LVA/ICA*, 2010, pp. 482–489.
7. R. F. Moghaddam and M. Cheriet, "A variational approach to degraded document enhancement," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 32, no. 8, pp. 1347–1361, 2010.
8. M. S. C. Almeida and L. B. Almeida, "Nonlinear separation of show-through image mixtures using a physical model trained with ica," *Signal Processing*, vol. 92, p. 872884, 2012.
9. F. Martinelli, E. Salerno, I. Gerace, and A. Tonazzini, "Non-linear model and constrained ML for removing back-to-front interferences from recto-verso documents," *Pattern Recognition*, vol. 45, pp. 596–605, 2012.
10. R. Rowley-Brooke and A. Kokaram, "Bleed-through removal in degraded documents," *Proc. SPIE 8297 Document Recognition and Retrieval XIX, 82970T-10*, 2012.
11. R. Rowley-Brooke, F. Pitié, and A. C. Kokaram, "A non-parametric framework for document bleed-through removal," in *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*, 2013, pp. 2954–2960. [Online]. Available: <http://dx.doi.org/10.1109/CVPR.2013.380>
12. 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.
13. Y. Ai, W. Li, T.-H. Chan, and Q. Liao, "Removal of bleed-through effect via nonnegative least-correlation," in *Proceedings of International Conference on Acoustic, Speech and Signal Processing (ICASSP 2014)*. IEEE, 2014, pp. 2774–2778.
14. E. Dubois and A. Pathak, "Reduction of bleed-through in scanned manuscript documents," in *Proc. IS&T Image Processing, Image Quality, Image Capture Systems Conference*, 2001, pp. 177–180.
15. A. Tonazzini, G. Bianco, and E. Salerno, "Registration and enhancement of double-sided degraded manuscripts acquired in multispectral modality," in *Proc. 10th International Conference on Document Analysis and Recognition ICDAR 2009*, 2009, pp. 546–550.
16. J. Wang, M. S. Brown, and C. L. Tan, "Accurate alignment of double-sided manuscripts for bleed-through removal," in *Proc. 8th IAPR Workshop on Document Analysis Systems*, 2008, pp. 69–75.
17. B. Li, W. Wang, and H. Ye, "Multi-sensor image registration based on algebraic projective invariants," *Optics express*, vol. 21, pp. 9824–9838, 2013.
18. P. Savino and A. Tonazzini, "Digital restoration of ancient color manuscripts from geometrically misaligned recto-verso pairs," *Journal of Cultural Heritage*, vol. 19, pp. 511–521, 2016.
19. J. Wang and C. L. Tan, "Non-rigid registration and restoration of double-sided historical manuscripts," in *Proc. Int. Conf. on Document Analysis and Recognition (ICDAR)*, 2011, p. 13741378.
20. A. Myronenko and X. Song, "Intensity-based image registration by minimizing residual complexity," *IEEE Transactions on Medical Imaging*, vol. 29, pp. 1882–1891, 2010.
21. R. Rowley-Brooke, F. Pitié, and A. C. Kokaram, "Nonrigid recto-verso registration using page outline structure and content preserving warps," in *Proceedings of the 2nd International Workshop on Historical Document Imaging and Processing*. ACM, 2013, pp. 8–13.
22. R. Rowley-Brooke, F. Pitié, and A. Kokaram, "A ground truth bleed-through document image database," in *Theory and Practice of Digital Libraries*, ser. Lecture Notes in Computer Science, P. Z. adn G. Buchanan, E. Rasmussen, and F. Loizides, Eds., vol. 7489. Springer, 2012, pp. 185–196.
23. Irish Script On Screen Project, "<http://www.isos.dias.ie>," 2012.
24. J. Sauvola and M. Pietikäinen, "Adaptive document image binarization," *Pattern Recognition*, vol. 33, pp. 225–236, 2000.
25. S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. Pluim, "Elastix: A toolbox for intensity based medical image registration," *IEEE Transactions on Medical Imaging*, vol. 29, pp. 196–205, 2010.
26. M. Hanif, A. Tonazzini, P. Savino, and E. Salerno, "Non-local sparse image inpainting for document bleed-through removal," *Journal of Imaging*, vol. 4, 2018.