

Ink-Bleed Reduction Using Functional Minimization

Grani A. Hanasusanto Zheng Wu Michael S. Brown
School of Computing, National University of Singapore
brown@comp.nus.edu.sg

Abstract

Ink-bleed interference is undesirable as it reduces the legibility and aesthetics of affected documents. We present a novel approach to reduce ink-bleed interference using functional minimization. In particular, we show how to modify the Chan-Vese active contour model to incorporate information from the front and back sides of the ink-bleed document. This contour model is particularly useful as it does not require edge extraction or explicit thresholding of the document. In addition, we show how functional minimization can again be used to restore broken foreground strokes that arise when strong ink-bleed overlaps with foreground strokes. The experimental results show that our functional minimization method produces better results than recent ink-bleed reduction techniques. To provide a complete framework, we also show how simple user assistance can be further exploited to improve the results.

1. Introduction and Contribution

As recto-verso documents (i.e. double-sided text documents) age, several intrinsic and extrinsic factors cause the ink to seep (or bleed) through the paper to interfere with the opposite side. While conservators in the archival and library community have physical restoration techniques to reduce ink-bleed, the physical restoration is costly and time-consuming and carries the risk of irreversible damage. The goal of this paper is to reduce ink-bleed in digital facsimiles of the original, a solution welcomed in lieu of physical restoration.

The contribution of our work is an ink-bleed reduction technique based on functional minimization. Unlike previous work, we treat the problem as one of active contour segmentation and minimize variational energy within an image and between the recto and verso images. To ameliorate the effects of broken strokes that may arise when strong ink-bleed overlaps with foreground ink, we again use functional minimization to complete the strokes. Comparisons with recent user-assisted ink-bleed techniques show that our approach produces consistently better results. Moreover, our

segmentation is fast and can be used in a semi-interactive software framework. Following a recent trend of exploiting user interactions, we also demonstrate how simple user markup can be used to help when a manuscript page has varying ink-bleed characteristics.

The remainder of this paper is organized as follows. Section 2 describes related work; section 3 overviews our framework; section 4 presents results and comparisons with recent techniques. A discussion and summary are provided in section 5.

2. Related Work

Related work is discussed for both ink-bleed reduction and active contours. For further information readers are directed to [24] and [3] respectively.

Ink-Bleed Reduction Image-based ink-bleed reduction methods can be discussed by whether they operate from a single image or two images (i.e. aligned recto and verso images). Single image methods typically employ various types of thresholding [7, 12, 16], but these simple methods can only handle relatively easy ink-bleed. In the work by Wang et al. [23], images are preprocessed in the wavelet domain before thresholding. The method is parameter intensive and targets the images where the handwriting is significantly slanted. Recent work by Wolf [24] avoids thresholding and instead uses a dual-layer MRF with a single observation field based on a Gaussian degradation model governing the MRF smoothness term. This work uses k-means clustering to generate the initial MRF label assignments which requires the ink-bleed and foreground-ink intensities to be distinguishable.

Two-image approaches are generally more effective given the additional information. These approaches include Tonazzini et al.'s method [21] which used blind signal separation to restore original images by assuming a linear mixture model. These authors also introduced a two-layer-MRF approach based on the same linear mixture model [20]. However, the linear model cannot be always justified for ink-bleed as shown in [24]. Tan et al. [19] used a two-image wavelet-based method, which has the similar limits as its single-image counterpart in [23]. Moghaddam et

al. [13] simultaneously reduce ink-bleed in both the recto and verso images using three diffusion models for recto foreground-ink, verso foreground-ink and background respectively. This method relies on five user-tuned parameters and uses scans of real paper without text to obtain the background information.

Recent work by Huang et al. [9] introduces a user-assisted approach that avoids prior assumptions about the document by having the user provide training examples of ink-bleed, foreground-ink, and background using simple scribble. An MRF is used to help enforce smoothness in the final classification. This approach requires sufficient markup to produce a satisfactory result. Lu et al. [11] extend this work by directing the user to pixel regions classified with low confidence to obtain this sufficient markup. Because these methods do not rely on any assumptions or user-tuned parameters, we use [9] and [11] as the primary benchmark compared against our technique. In the cases of difficult ink-bleed, we can also incorporate simple user input to denote the area to be locally processed. We show that our approach is adaptive to local area statistics and can produce better overall results than [9, 11] and other techniques [7, 19].

Active Contour The main objective of an active contour algorithm is to evolve a curve to segment an image based on one or several constraints. Typically the constraints are incorporated into an energy functional such that they are achieved when the energy is minimized. The classical approach in active contour segmentation [10, 5, 22] is to stop the curve evolution at the edges, i.e. the curve will try to locate the pixels where absolute gradient $|\nabla u(x, y)|$ is maximum.

Region based active contour approaches avoid using solely edge information (e.g. [6, 15, 14, 25, 4, 18]). A well-known region-based method is the Chan-Vese model [4] in which the curve is obtained by minimizing the intensity variance inside and outside the curve. Our work extends the Chan-Vese model, but unique to our application is the incorporation of the relationship between the recto-verso image pair. By incorporating this information we produce significantly better results than using the Chan-Vese model alone.

3. System Overview

Our overall framework including foreground stroke segmentation, stroke completion, and incorporation of user input is discussed here. Our input is assumed to be a pair of aligned recto and verso images. We use the alignment approach in [9], where the back image is mirrored and a global translation alignment is first applied. The images are then divided into patches in order to compute local displacement between each front image patch and its corresponding patch on the back image. From the local displacements, thin-plate-spline (TPS) interpolation is used to warp the back

image to align with the front¹. We note that recto-verso image alignment is a non-trivial problem, however, further details are outside the scope of this paper, instead, we refer the reader to [9] and other recent related techniques (e.g. [1]).

3.1. Ink-Bleed Removal with Active Contours

We begin by describing the implementation of the Chan-Vese model and our extension to the energy functional. Here, the curve is represented by the zero level curve

$$C = \{(x, y) \in \Omega | \phi(x, y) = 0\}, \quad (1)$$

where $\Omega \subset \mathbb{R}^2$ is the given image domain. The curve C is evolved via the function $\phi(x, y, t)$, a surface in \mathbb{R}^3 . In the following, we also define $\phi_1 = \{(x, y) \in \Omega | \phi > 0\}$ (i.e. pixels inside the contour) and $\phi_2 = \{(x, y) \in \Omega | \phi < 0\}$ (i.e. pixels outside the contour) as the targeted foreground and non-foreground (including both ink-bleed and background) segments respectively.

The Chan-Vese model [4] minimizes the energy functional:

$$\begin{aligned} E = & \mu \times \text{length}(C) \\ & + \lambda_1 \int_{\phi_1} |u_0(x, y) - c_1|^2 dx dy \\ & + \lambda_2 \int_{\phi_2} |u_0(x, y) - c_2|^2 dx dy, \end{aligned} \quad (2)$$

where $\text{length}(C)$ is the length of the curve, u_0 is the intensity of the recto image to be segmented, c_1 and c_2 are the average intensity in the domain ϕ_1 and ϕ_2 , respectively, and μ , λ_1 , and λ_2 are weighting factors set to 1 in this work. Essentially the idea of this model is to stop the curve evolution when the intensity variance in each domain separated by the curve C is minimized.

The ingenuity of the Chan-Vese model, when applied to our segmentation problem, is that it does not rely on sharp edges to segment the ink from its background. This is particularly useful since the foreground strokes are often diffuse due to aging of the paper and ink. However, using the basic Chan-Vese formulation, the obtained foreground domain ϕ_1 will also contain ink-bleed pixels. Based on the assumption that the intensity of the foreground is typically lower than that of the corresponding ink-bleed or background on the reverse side, we propose a new energy functional:

$$E = \text{length}(C) + E_1 + \lambda E_2, \quad (3)$$

¹This approach performs well on most documents from our partners at the National Archive of Singapore. It tends to fail for documents having weak ink-bleed, but for these cases thresholding approaches are applicable.

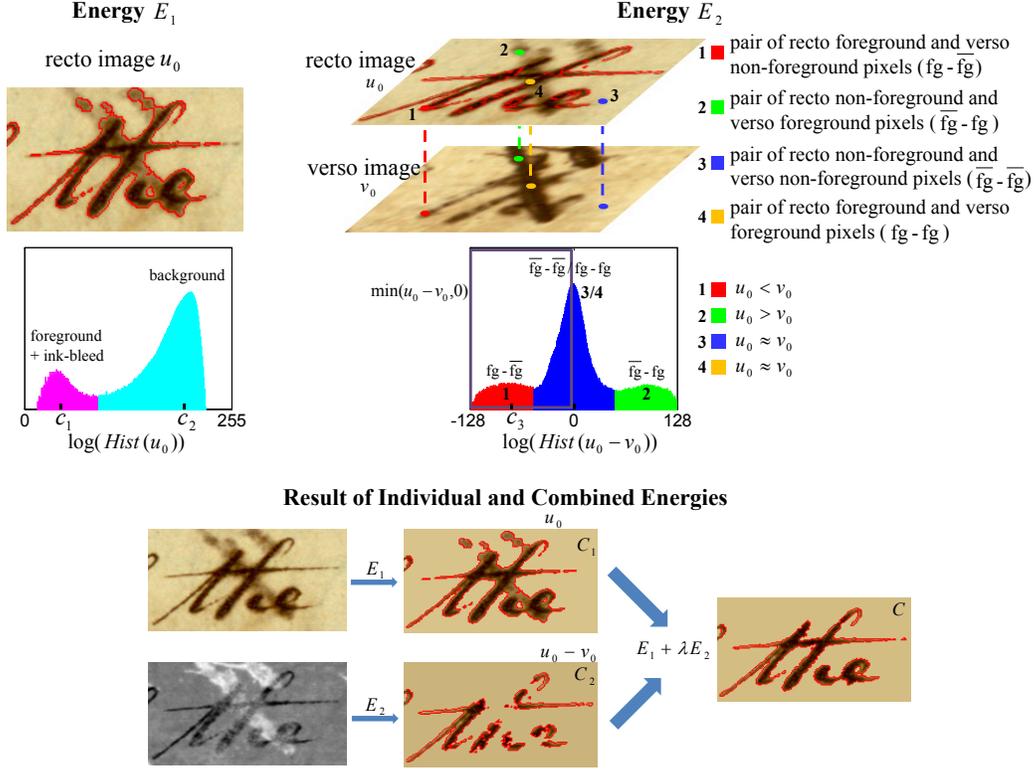


Figure 1: This figure shows the associated terms of the two energies E_1 and E_2 and the combined effect of $E_1 + \lambda E_2$. A detailed description of this diagram is in the text (directly below).

with

$$E_1 = \int_{\phi_1} |u_0(x, y) - c_1|^2 dx dy + \int_{\phi_2} |u_0(x, y) - c_2|^2 dx dy \quad (4)$$

$$E_2 = \int_{\phi_1} |u_0(x, y) - v_0(x, y) - c_3|^2 dx dy + \int_{\phi_2} |u_0(x, y) - v_0(x, y)|^2 dx dy \quad (5)$$

where $v_0(x, y)$ is the intensity of the aligned verso page, λ is a multiplicative constant, c_3 is the average of $\min(u_0 - v_0, 0)$ inside the curve.

The combination of E_1 and E_2 requires a bit of examination to fully appreciate their mutual benefits and we use Fig. 1 to help explain the two terms' effects. The top row of Fig. 1 diagrams the E_1 and E_2 energy terms separately. As previously stated, the E_1 energy finds a contour, call it C_1 , that separates the image into two domains that have minimal variances. Examining the histogram of the recto image shown for E_1 , we see that c_1 and c_2 are centered at the two peaks in the histogram, representing 1) the foreground and ink-bleed intensities and 2) the background intensities.

The E_2 energy term operates from the difference of the verso and recto image, $u_0 - v_0$, as shown in Fig. 1. Examining the histogram of the difference image, $u_0 - v_0$, we see the majority of pixels are about $u_0 - v_0 = 0$ with two smaller peaks on either side. Since we know that the foreground-ink on the recto image should be darker in intensity than the corresponding ink-bleed or background in the verso image, we can restrict the domain for these pixels to be $u_0 - v_0 < 0$, using the $\min(u_0 - v_0, 0)$ operation. This forces the mean of pixels inside the curve, c_3 , to only consider foreground pixels. The result is a functional energy that finds a contour, C_2 , whose inside pixels are $u_0 - v_0 < 0$ and the outside pixels have $u_0 - v_0$ closer to 0 (i.e. the overall mean of the histogram). These outside pixels correspond to the second term in the Eq. (5). This allows us to pull out the foreground only as shown in the bottom image of Fig. 1. However, this tends to over segment the foreground-ink and will remove pixels where recto foreground-ink and verso foreground-ink overlap. The best result is obtained when combining these two terms as shown in Fig. 1. This balance between the results of E_1 and E_2 is controlled by the parameter λ in Eq. (3).

To minimize the combined $E_1 + \lambda E_2$ energy, we use the strategy in [17] that directly optimizes Eq. (3) instead



Figure 2: Examples of our stroke completion method for restoring broken strokes. Shown from left to right are: input images, initial results with over segmentation, determined completion domains and completion results.

of solving the typical Euler-Lagrange equations to produce a curve evolution with respect to artificial time. This approach typically converges in less than 10 iterations. It usually takes under 5s on a standard 2.66 GHz CPU to process the input image of 1M pixels, which is the typical size of documents in the resolution intended for viewing on the screen. For a letter-size document page scanned in higher resolution of 300dpi (typically 8M pixels), the processing time is usually 30-40s.

3.2. Broken-Stroke Completion

In severely degraded documents, the ink-bleed intensity can be very similar to the foreground-ink intensity. In such cases, the parameter λ in Eq. (3) needs to be increased higher to remove the ink-bleed. As we increase λ (i.e. weighting E_2 more), broken stroke where the ink-bleed and foreground overlap is more likely. Such broken stroke is often unavoidable and a problem for other methods described in Section 2 as well.

To complete the broken strokes, we propose to apply another functional minimization algorithm. To determine the completion domain, i.e. the region where the broken stroke needs to be restored, we apply a statistical significance testing by assuming that for the given recto image, pixel intensity u_0 in the background domain ϕ_2 follows a normal distribution. We use one-tailed test at a 99% confidence level to get the dark pixels in ϕ_2 as the completion domain:

$$D = \{(x, y) \in \Omega | z(x, y) < -2.325\} \quad (6)$$

where $z(x, y) = (u_0(x, y) - \bar{u}_0)/\sigma$ is the standard score of $u_0(x, y)$ with respect to sample mean \bar{u}_0 and standard deviation σ over the ϕ_2 domain.

For the functional minimization algorithm, we implement the modified Cahn-Hilliard inpainting model introduced in [2, 8] for binary images. This is well suited to our application as the segmentation equation (3) involves only two regions ϕ_1 and ϕ_2 .

We construct an image f_0 corresponding to the output of

the segmentation algorithm in Section 3.1 such that:

$$f_0(x, y) = \begin{cases} 0 & \text{if } (x, y) \in \phi_1; \\ 1 & \text{if } (x, y) \in \phi_2 \setminus D; \\ 0.5 & \text{if } (x, y) \in D. \end{cases} \quad (7)$$

An energy functional is then minimized to find a new image f , that fills the completion domain D based on the modified Cahn-Hilliard equation expressed as:

$$E = \int_{\Omega} \left[\frac{\epsilon}{2} |\nabla f|^2 + \frac{1}{\epsilon} W(f) \right] dx dy + \beta \int_{\Omega \setminus D} (f - f_0)^2 dx dy, \quad (8)$$

where $W(f) = f^2(f - 1)^2$ has two minima at $f = 0$ and $f = 1$, and ϵ is a parameter that measures the transition region between $f = 0$ and $f = 1$. Here, β measures the fitness level of the solution image f to the given image f_0 outside the completion domain D . The discretized $\frac{\delta f}{\delta t}$ that minimizes the energy functional is solved for a new f at the next artificial time step by the way of 2-dimensional FFT (see [8] for more details).

For our experiments, we empirically set β to 10^{10} . Energy minimization is then performed in two stages with large ϵ (0.8) for the first stage and subsequently smaller ϵ (0.01) for the second stage. A detailed explanation on the two stages and the suitable parameter ϵ can be found in [2].

Fig. 2 shows two examples that require a large λ to ensure that the dark ink-bleed is completely removed. The side-effect is that the regions where foreground and ink-bleed overlap are segmented as ϕ_2 . The completion domain is found using Eq. (6) and inpainted using the Cahn-Hilliard Eq. (8). The original image pixels under this completed binary mask are used in the final result shown as the last column of Fig. 2. Processing time for stroke-completion is slower than segmentation, taking roughly 2mins per 1M image. This, however, should be considered a post-processing step that is necessary only in difficult cases.

3.3. Extension Towards Interactive Approach

For some documents, the foreground and ink-bleed intensities vary spatially as shown in Fig. 3. Here, the foreground-ink is much lighter in the bottom region of the document as compared to the ink-bleed intensity in the top portion. One nice benefit of our approach is that it adapts to the characteristics of the input domain. This means we simply need to denote the region we would like to apply the functional minimization to locally gain a new segmentation that is tuned to the local characteristics. In this manner, the curve C is evolved by taking into account the statistics c_1 , c_2 , and c_3 in the locality, resulting in an optimized segmentation of ϕ_1 and ϕ_2 that otherwise would be non-trivial under global minimization. Additionally, the user could visually adjust a local λ to obtain the best segmentation in the selected region.



Figure 3: The results of applying our functional minimization to a selected region to improve the results when ink-bleed characteristics vary spatially. Our results are compared with those of the iterative user-assisted approach [11].

Fig. 3 shows our initial result for the global minimization. The algorithm is able to remove the difficult ink-bleed and still maintain the consistency of the foreground strokes (see zoomed region 3) for most of the document. However, for the bottom region of the document, the foreground-ink has significantly different characteristics from that of the

overall document. This causes the global minimization to segment foreground-ink here into the ϕ_2 domain. Fig. 3 shows the result of local minimization which is performed over the region in red, with the same λ as that used globally. The foreground strokes have been successfully recovered in this local region (see zoomed regions 1 and 2) because the



Figure 4: Experimental results and comparison to classification based approaches (iterative markup approach [11] and single-pass markup approach with MRF [9]).

local region statistics are used when minimizing Eq. (3). Fig. 3 also shows that our algorithm performs better than the recent user-assisted approach that helps direct the user to troubled locations [11].

4. Results

We first show the subjective results of our method in Fig. 4, compared against the user-assisted approaches by Huang et al. [9] and Lu et al. [11]. The results from the user-assisted approaches are shown in the middle column of Fig. 4. The top three rows are implemented by the iterative markup approach [11] while the bottom three rows are obtained using the single-pass user-markup that incorporates additional MRF for smoothing [9]. Our approach is successful in removing ink-bleed for all the cases. When compared against the iterative markup [11] (rows 1-3), our approach is more effective in preserving the complete foreground strokes. For the image in row 2, the result from [11] has several broken (or erased) foreground strokes. For the images in rows 1 and 3, the results of [11] remove some very light foreground strokes. For the examples compared against [9] (rows 4-6), our algorithm is able to better separate the foreground and ink-bleed and also maintain the consistency of foreground stroke, while the work by [9] tends

to fail in removing the very dark ink-bleed.

For a quantitative evaluation, we compare our method against three approaches, including the adaptive thresholding by Drira et al. [7], the wavelet method by Tan et al. [19] and Huang et al.’s markup with MRF method [9]. As Lu et al.’s iterative markup method [11] and Huang et al.’s method share the same underlying approach, Lu et al.’s method is not included. The four methods are applied to ten documents which cover a wide spectrum of ink-bleed. For a fair comparison, no stroke completion (Section 3.2) and local processing (Section 3.3) are performed by our method.

For each result image, we count the numbers of the erroneously segmented words. The errors can be divided into four types: the foreground words totally or partly *Missed* in the output, the foreground words still *Interfered* by the ink-bleed in the output, the foreground words suffering from both the *Missed & Interfered* problems and the *Added* words which are the retained ink-bleed and locate away from the true foreground words. Fig. 5 shows the examples of these errors. The error counting results are given in Table 1. For each result, we compute the recall and precision scores as

$$\text{recall} = \frac{W - W^E}{W} \quad \text{and} \quad \text{precision} = \frac{W - W^E}{W - W^M + W^A}, \quad (9)$$

where W^E is the sum amount of the *Missed*, *Interfered*



Figure 5: Examples of *Missed*, *Interfered*, *Missed & Interfered* and *Added* errors in result images. Shown from top to bottom are: input images, outputs that would be considered as having no error and outputs with errors.

and *Missed & Interfered* words, W^M is the amount of the *Missed* words and W^A is the amount of the *Added* words². For each method, we show the recall and precision scores averaged on the ten trials in the table. Overall, our method consistently generates fewer errors than the competing approaches. Because this quantitative evaluation still has elements of subjectivity, the original images together with the results have been placed on the contact author’s webpage.

5. Discussion and Summary

Our results show that the proposed algorithm performs better on a wide range of documents than previous approaches. While the idea in the user-assisted approaches [9, 11] is to avoid assumptions and parameter-tuning (endemic of other ink-bleed reduction algorithms), the trade-off is reliance on sufficient markup to perform the classification. While the “directed assistance” approach in [11] helps guide the user to where to perform the markup to improve results, our approach requires no stroke-level markup at all. Instead, if wish to apply our method locally, we only need to coarsely select the region of interest on the image itself.

Moreover, we only have a single parameter to tune (λ). Given that our segmentation algorithm can run in under 5s for images in viewing resolution, tuning this parameter can be performed in near real-time using a slider UI feature. Lastly, we are also able to complete strokes that were broken due to the overlap of ink-bleed and foreground stroke. This allows us to produce consistently good results, even on difficult examples.

We do note that our method (along with others) still suffers when the ink-bleed is very dark, or when there is significant ink-bleed corruption creating large gaps in the strokes that cannot be completed (Fig. 6 shows these cases). Another limitation of our work is that we do not segment the foreground-ink in both the recto and verso images simultaneously. This is because the optimal λ may be different

²The definitions follow those in [19, 23], but with a difference in the denominator of precision formula. By our definition, the 100% precision score is obtained when no ink-bleed is retained (but some true foreground words may be removed as well).

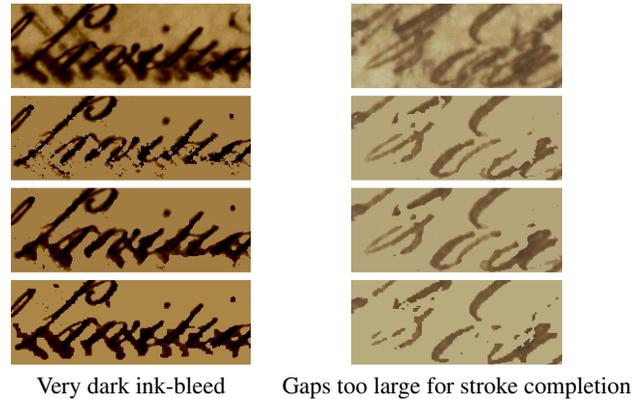


Figure 6: Examples of failure cases. Shown from top to bottom are: input images, segmentation by functional minimization, segmentation with stroke completion, and results by markup with dual MRF method [9].

for each side. Future work is to address this issue by re-examining the energy terms. Similarly, we are exploring new mechanism of stroke completion that can operate on vector-valued images to allow simultaneous stroke repair for both sides of the document.

6. Acknowledgement

This work was supported by the Singapore Agency for Science, Technology and Research (A*STAR) project (No. 0520180).

References

- [1] R. Baumann and W. B. Seales. Robust registration of manuscript images. In *ACM/IEEE Joint Conference on Digital Libraries*, 2009.
- [2] A. Bertozzi, S. Esedoglu, and A. Gillette. Inpainting of binary images using the Cahn-Hilliard equation. *IEEE Trans. Image Processing*, 16(1):285–291, 2007.
- [3] T. F. Chan, M. Moelich, and B. Sandberg. Some recent developments in variational image segmentation. In X.-C. Tai, K.-A. Lie, T. F. Chan, and S. Osher, editors, *Image Process-*

Algorithm		Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Trial 6	Trial 7	Trial 8	Trial 9	Trial 10	Average
	# of words	171	142	117	107	121	111	107	104	111	118	
Adaptive	Missed	41	47	2	4	11	20	1	4	3	16	Recall
	Interfered	30	26	27	25	22	48	24	18	59	26	55.58%
	Missed & interfered	45	45	0	1	0	12	0	0	0	6	Precision
	Added	6	16	12	31	7	25	12	11	30	11	53.51%
Wavelet	Missed	1	0	5	37	17	48	17	14	5	8	Recall
	Interfered	112	100	3	4	1	5	0	0	15	10	68.48%
	Missed & interfered	1	2	0	1	0	3	0	0	2	1	Precision
	Added	10	23	1	3	1	4	0	0	5	9	89.09%
Dual MRF	Missed	11	37	6	15	1	20	14	13	18	9	Recall
	Interfered	41	25	8	2	12	2	1	0	25	13	77.02%
	Missed & interfered	3	10	0	0	0	1	0	1	3	0	Precision
	Added	4	9	3	2	5	1	0	1	6	9	87.88%
Proposed	Missed	4	5	0	1	0	1	0	0	4	1	Recall
	Interfered	11	34	2	6	3	4	3	2	4	2	93.41%
	Missed & interfered	0	0	0	0	0	0	0	0	0	0	Precision
	Added	1	9	1	4	3	3	2	3	4	7	94.02%

Table 1: The number of the *Missed*, *Interfered*, *Missed & Interfered* and *Added* words in the results of four approaches, including the single-image adaptive thresholding method [7], the two-image wavelet-based method [19], the dual MRF method [9] and the proposed functional minimization method.

ing Based on Partial Differential Equations, pages 175–210. Springer, 2007.

- [4] T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, 2001.
- [5] L. Cohen. On active contour models and balloons. *Computer Vision Graphics Image Processing*, 53:211–218, 1991.
- [6] L. Cohen, E. Babinet, and N. Ayache. Surface reconstruction using active contour models. In *SPIE Conference on Geometric Methods in Computer Vision*, 1993.
- [7] F. Drira, F. L. Bourgeois, and H. Emptoz. Restoring ink bleed-through degraded document images using a recursive unsupervised classification technique. In *International Workshop on Document Analysis Systems*, 2006.
- [8] A. Gillette. *Image inpainting using a modified Cahn-Hilliard equation*. PhD thesis, UCLA, 2006.
- [9] Y. Huang, M. S. Brown, and D. Xu. A framework for reducing ink-bleed in old documents. In *CVPR*, 2008.
- [10] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1988.
- [11] Z. Lu, Z. Wu, and M. S. Brown. Directed assistance for ink-bleed reduction in old documents. In *CVPR*, 2009.
- [12] F. C. Mintzer et al. Toward on-line, worldwide access to vatican library materials. *IBM Journal of Research and Development*, 40(2):139–162, 1996.
- [13] R. F. Moghaddam and M. Cheriet. Low quality document image modeling and enhancement. *International Journal on Document Analysis and Recognition*, 11(4):183–201, 2009.
- [14] N. Paragios and R. Deriche. Geodesic active regions for texture segmentation. In *INRIA Technical Report 3440*, 1998.
- [15] R. Ronfard. Region-based strategies for active contour models. *International Journal of Computer Vision*, 13:229–251, 1994.
- [16] Z. Shi and V. Govindaraju. Historical document image enhancement using background light intensity normalization. In *ICPR*, 2004.
- [17] B. Song and T. Chan. A fast algorithm for level set based optimization. In *UCLA Cam Report 02-68*, 2002.
- [18] Z. L. Szapka and J. R. Tapamo. Further optimization for the Chan-Vese active contour model. In *High Performance Computing and Simulation Conference*, 2008.
- [19] C. L. Tan, R. Cao, and P. Shen. Restoration of archival documents using a wavelet technique. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(10):1399–1404, 2002.
- [20] A. Tonazzini, L. Bedini, and E. Salerno. A markov model for blind image separation by a mean-field em algorithm. *IEEE Trans. Image Processing*, 15(2):473–482, 2006.
- [21] 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*, 10(1):17–25, 2007.
- [22] V. Vasselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22:61–79, 1997.
- [23] Q. Wang, T. Xia, L. Li, and C. L. Tan. Document image enhancement using directional wavelet. In *CVPR*, 2003.
- [24] C. Wolf. Document ink bleed-through removal with two hidden markov random fields and a single observation field. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 32(3):431–447, 2010.
- [25] A. Yezzi, A. Tsai, and A. Willsky. A statistical approach to snakes for bimodal and trimodal imagery. In *ICCV*, 1999.