MULTI-SCALE SEMI-TRANSPARENT BLOTCH REMOVAL ON ARCHIVED PHOTOGRAPHS USING BAYESIAN MATTING TECHNIQUES AND VISIBILITY LAWS

A.J. Crawford[1,2]*, V. Bruni*[1]*, A.C. Kokaram*[3]*, and D. Vitulano*[1]

[1] Istituto per le Applicazioni del Calcolo - C.N.R., Viale del Policlinico 137, 00161 Rome Italy
[2]Dip. di Modelli e Metodi Matematici per le Scienze Applicate Universitá di Roma "La Sapienza"
[3]Electronic and Electrical Engineering Department University of Dublin, Trinity College, Ireland
ajcrawford@gmail.com

ABSTRACT

This paper presents an automatic technique to remove semi-transparent blotches (due to moisture) from archived photographs and documents. Blotches are processed in the HSV space. While chroma components are processed using a simple texture synthesis method, the intensity component is split into an over-complete wavelet representation. In the approximation band, the blotch is modelled as an alpha matte which reduces the intensity of the image in a non-uniform yet smooth manner. The alpha matte is estimated using a Bayesian approach and its effect reversed. Wavelet details are left unchanged in the case of perfect semi-transparency or attenuated using visibility laws whenever dirt and dust cause spurious edges. Experimental results achieved on many historical photographs show the effectiveness of the proposed approach.

Index Terms— Image restoration, Bayes procedures, Minimization methods, Wavelet transforms

1. INTRODUCTION

Archives such as Fratelli Alinari in Florence, contain huge numbers of old photographs. Other examples of archived material are old documents, film and books. With the growth of the internet, it is now possible to provide efficient and wide access to the archived material. However, most of this material suffers from degradation and damage. Therefore, restoration may be required before they are made available.

In recent years, digital restoration techniques have been developed to deal with many kinds of degradation. Digital techniques have the advantage that they will not damage the original material after the initial scanning process. Also, as techniques are improved, they can be re-applied to the digital scan whereas this may not be the case with physical methods.

This paper describes an user independent technique to remove semi-transparent blotches from archived photographs. Blotches of this type are caused by water penetration into the paper. As already investigated in literature [1], the automatic detection and restoration of this kind of defect is a difficult problem because of the variability in shape, colour and intensity of this degradation along with its semi-transparency. As is the case for the detection [2], restoration takes place in the HSV colour space, as this space corresponds closely to human perception [3]. Since the photographs in question are sepia images, and therefore almost constant in colour, a simple texture synthesis method is sufficient to remove the defect from the chroma channels. The luminance channel, however, also contains image detail in the degraded regions due to semi-transparency. The proposed model differs from classical inpainting, where models are based on missing information [4, 5], because the preservation of the whole original information is fundamental from an historical point of view. The luminance component is split into an over-complete wavelet representation until a suitable scale level J. This latter is estimated in the detection phase and accounts for the maximum visibility of the blotch on the degraded image [2]. Thus, the approximation band is processed by a Bayesian restoration technique (Section 2.1). Wavelet details are left unchanged if the blotch satisfies the semi-transparency hypothesis or attenuated according to perception laws in case of dust and dirt causing spurious edges (Section 2.2). The inverse wavelet transform provides the restored luminance. This is finally combined with the restored chroma components (Section 3) to give the final result.

2. THE PROPOSED MODEL

The physical formation of a water blotch on paper (photo, book, etc.) can be modelled by the spreading and the penetration of water droplets into porous material [6, 7]. If the droplet is modeled as a semi-sphere, during the spreading its radius grows to an equilibrium value, that is determined by the contact angle θ. From this point on, the liquid is only absorbed depending on the porosity of the considered medium. In ideal conditions, the central pores absorb more than the ex-

Thanks to the Italian Ministry of Education for funding (FIRB project no.RBNE039LLC, "A knowledge-based model for digital restoration and enhancement of images concerning archaeological and monumental heritage of the Mediterranean coast") and to F.lli Alinari SpA for providing the images.

ternal pores, since they come in contact with the liquid earlier (see Fig. 1). As a matter of fact, the evolution of a drop involves different parameters, such as the geometry of the original drop and the regularity of the surface of the paper, that are unknown in real applications. The problem can be simplified by modeling the water drop as a semi sphere of radius R and assuming the contact angle to be $\leq \frac{\pi}{2}$. In cases where the spreading and absorption processes have been completed, one can expect a small contact angle at the equilibrium, while it is larger if one of the two processes has been disturbed. In the second case, the blotch presents an evident edge.

The proposed model splits the sepia degraded image in the color space HSV. The two chroma components H and S are directly restored in the physical space while, as argued in the Introduction, the V component is split into an overcomplete wavelet basis until the optimal scale level, J, computed in [2]. The approximation band A and the wavelet detail bands $\{D_j\}_{1\leq j\leq J}$ are separately restored according to the transparency model and perception laws yielding $\{\tilde{D}_j\}_{1\leq j\leq J}$ and $\tilde{A}$. Finally, the inverse wavelet transform is performed to achieve $\tilde{V}$. The latter along with $\tilde{H}$ and $\tilde{S}$ provide the final restored image. Empirical studies showed that the regularity and support length of the Daubechies wavelet with 4 vanishing moments, *db4*, agree with the common structure of a significant number of blotches.

2.1. Restoration of the luminance approximation band

As the blotch does not completely obscure the clean image, the luminance approximation band can be modeled as a multilayer image similar to [8], i.e. the luminance approximation band is modeled as a mixture between the clean image layer and the blotch layer [9]. The layers mix is based on the following relationship:

$$A(\mathbf{x}) = \alpha(\mathbf{x})I(\mathbf{x}) + \epsilon(\mathbf{x}) \tag{1}$$

where $A(\mathbf{x})$ is the observed luminance approximation band at point $\mathbf{x}$, $\alpha(\mathbf{x})$ the distortion layer and $I(\mathbf{x})$ the clean luminance approximation band. Noise is represented by $\epsilon(\mathbf{x}) \sim N(0, \sigma_\epsilon^2)$.

The correct values of I and α are those which maximise $p(I, \alpha|A, \sigma_\epsilon^2)$. Bayes' law gives the following relationship

$$p(I, \alpha|A, \sigma_\epsilon^2) \propto p(A|I, \alpha, \sigma_\epsilon^2)p(\alpha|\overline{\alpha})p(I|\overline{I}) \tag{2}$$

where $\overline{\alpha}$ and $\overline{I}$ are α and I in the neighbourhood of $\mathbf{x}$ respectively. It is now easier to compute the likelihoods on the right hand side of the equation in place of $p(I, \alpha|A, \sigma_\epsilon^2)$. The first two likelihoods on the right hand side ensure that alpha matches the behaviour of the blotch described, i.e. *i)* α must mix to give the observed data; *ii)* α must be smooth. The third term ensures that I values are similar. The probabilities from expression (2) can be represented as follows:

$$p(A|I, \alpha, \sigma_\epsilon^2) \propto exp(\frac{-(A(\mathbf{x}) - \alpha(\mathbf{x})I(\mathbf{x}))^2}{2\sigma_\epsilon^2}) \tag{3}$$

$$p(\alpha|\overline{\alpha}) \propto exp(-\sum_{k=0}^{n} \lambda_k(\alpha(\mathbf{x}) - \alpha(\mathbf{x}+q_k))^2) \tag{4}$$

$$p(I|\overline{I}) \propto exp(-\sum_{k=0}^{n} \lambda_k(I(\mathbf{x}) - I(\mathbf{x}+q_k))^2) \tag{5}$$

where $\mathbf{x}+q_k$ is a neighbouring sample and λ_k is a weight based the distance to this sample. These expressions show that maximising $p(I, \alpha|A, \sigma_\epsilon^2)$ is equivalent to minimising the following energy:

$$\begin{aligned} E &= W_1 \frac{(A(\mathbf{x}) - \alpha(\mathbf{x})I(\mathbf{x}))^2}{2\sigma_\epsilon^2} \\ &+ W_2 \sum_{k=0}^{n} \lambda_k(\alpha(\mathbf{x}) - \alpha(\mathbf{x}+q_k))^2 \\ &+ W_3 \sum_{k=0}^{n} \lambda_k(I(\mathbf{x}) - I(\mathbf{x}+q_k))^2 \end{aligned} \tag{6}$$

Weights W_1,W_2 and W_3 regulate the emphasis on the different constraints modeled by the 3 terms of (6).

2.1.1. Algorithm

There are two main steps in the restoration process. The first step assigns a "first guess" to each pixel inside the blotch. The value selected is chosen from the "clean" area close to the pixel. The next step uses the Iterative Conditional Mode (ICM) algorithm [10] to minimise the energy E from equation (6). The "first guess" step is a simple image synthesis method. For each pixel, its value is taken as a sample drawn from local "clean" pixels similar to the method used in [9]. The local region is composed of all clean pixels contained within a circle, centered on the current pixel. The radius of the circle is proportional to the distance between the pixel and the nearest edge of the blotch. Specifically the radius is defined as: $log(d(\mathbf{x}) + 1) + S_{\psi_J}$,where $d(\mathbf{x})$ is the distance to the edge and S_{ψ_J} is the wavelet support at the considered scale level J.

The first guess provides a reasonable solution for the blotch. However, it is void of the details clearly visible due to the semitransparent properties of the blotch. To recover the details, E must be minimised using the original image and the first guess acting as the initial conditions, I_0 and α_0. The minimisation is carried out using the ICM algorithm recursively improving estimates for I and α as follows:

$$I_1 \sim p(I|A, \alpha_0, \sigma_\epsilon^2) \quad \alpha_1 \sim p(\alpha|A, I_1, \sigma_\epsilon^2)$$

$$I_2 \sim p(I|A, \alpha_1, \sigma_\epsilon^2) \quad \alpha_2 \sim \ldots$$

In practice, the value of α is fixed, and the E is calculated for a large range of I ($[0, 0.01, 0.02, 0.03, \ldots, 0.99, 1]$ for a

normalised image). I is selected as that which gives the minimum value of E. The process is repeated fixing I and calculating E for a range of α. This is repeated until the whole blotch converges, i.e. the restored approximation band $\tilde{A}$ is achieved. The blotch is processed from the outside-in on the premise that values drawn from closer neighbourhoods are more likely to be accurate. The blotch is divided into an onion-like orbital rings calculated using morphological operators. Initially, the outer ring is processed. As more iterations are preformed, inner-rings are also processed until the whole blotch is included. Outer rings are reprocessed to take account of changes on neighbouring rings. Within each ring, adjacent pixels are not processed consecutively in order to avoid any accumulated error.

2.2. Restoration of the luminance wavelet details

The choice of the wavelet basis made at the beginning of this section, also gives the minimum admissible contact angle θ_{min} yielding a visible boundary. It is the one realizing the minimum of the error of the $(n-1)^{th}$ Taylor expansion around the point $x = \overline{R}$ of the function $y(x) = \sqrt{R^2 - x^2}, \quad |x| \leq \overline{R}$, where $\overline{R} = Rsin(2\theta)$ and n is the number of vanishing moments of the adopted wavelet. The function $y(x)$ models the arc of the blotch shape, according to Fig. 1. For the *db4* wavelet ($n = 4$) it holds

$$\theta_{min} = \frac{1}{2}\, arcsin(\frac{1}{\sqrt{6}}). \tag{7}$$

The difference between the ideal contact angle θ_{min} and the measured angle $\theta = arctg\left(\frac{h}{\overline{R}}\right)$, where h and $\overline{R}$ respectively are the height and half of the width of the analysed blotch[1], can be used for defining suitable contrast measures in the wavelet domain. h and $\overline{R}$ define the amount of shrinking to perform on detail coefficients in order to have a blotch with an invisible boundary.

It is worth noticing that: *i)* θ also gives the slope $\alpha = \tan(2\theta)$ of the straight line approximating the arc near the point $x = \overline{R}$; *ii)* in order to separate the left and the right boundary of the blotch, it is necessary that $2\overline{R} > S_{\psi_j} + 1$.

2.2.1. Algorithm

1. Compute the contact angle $\theta = arctg\left(\frac{h}{\overline{R}}\right)$.
2. $If \quad \theta_{min} < \theta$, then leave the details unaltered and set $\tilde{D}_j = D_j, \quad 1 \leq j \leq J$;
   $else$ the blotch has a visible boundary. For each row (column) of the considered detail band D_j, apply the following attenuation:

$$\tilde{D}_j(\mathbf{x}) = min(1, w_j(\mathbf{x}))\; D_j(\mathbf{x}), \quad \forall\,(\mathbf{x}) \in \Omega_j \tag{8}$$

with

$$w_j(\mathbf{x}) = \frac{1}{\frac{|D_j(\mathbf{x})|}{c_1\sigma_{ext}}\;\frac{|D_j(\mathbf{x})|}{c_2 H_{loc}}\;\frac{|D_j(\mathbf{x})|}{|D_j(\mathbf{x}) - D_j(N(\mathbf{x}))|}} \tag{9}$$

where Ω_j is the region of the blotch boundary at scale level j ($|\Omega_j| = S_{\psi_j}$), σ_{ext} is the standard deviation of the external part of the blotch; $H_{loc} = \overline{R}\; tg(\theta_{min})$ is the minimum height measured by the adopted wavelet; $c_1 = 1.02$ and $c_2 = 0.98$ are the Weber coefficients; $\frac{|D_j(\mathbf{x})|}{|D_j(\mathbf{x}) - D_j(N(\mathbf{x}))|}$ is the local contrast computed using the Weber's law — $N(\mathbf{x})$ indicates the local neighbourhood of the analysed pixel.

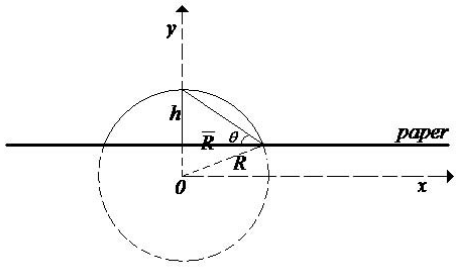

Fig. 1. Sketch of the drop causing the blotch.

Summing up, each pixel of the boundary is attenuated according to the contrast masking $\frac{|D_j(\mathbf{x})|}{c_1\sigma_{ext}}$, the contrast sensitivity $\frac{|D_j(\mathbf{x})|}{c_2 H_{loc}}$ and the local contrast $\frac{|D_j(\mathbf{x})|}{|D_j(\mathbf{x}) - D_j(N(\mathbf{x}))|}$ [11].

2.3. Restoration of chroma components

Although colour images are being processed, the clean image is almost constant in colour. In the areas affected by the blotch, Hue and Saturation values are increased. However, there is no underlying colour detail as in the luminance channel. Therefore, the simple texture synthesis method adopted in the first step for the luminance process, is sufficient to remove the effects of the blotches from the chroma channels.

3. RESULTS

The proposed model has been tested on various cases. However, in this paper just three representative examples are presented. A more complete set of results can be found at `firb.alinari.it`. The examples shown here are of historical photographs from the Alinari Archive in Florence and are depicted in Fig. 2 (a), (c) and (e). The first two are not extreme cases: these blotches do not contain dirt which covers original information, therefore the absorption process reached its

[1] The computation of the angle in real conditions would require the knowledge of some parameters like the resolution of the acquisition, the illumination and so on. We neglect all these parameters and we consider the intensity value of the image as the one indicating the height of the blotch inside a pore: the darker the blotch the deeper its penetration.

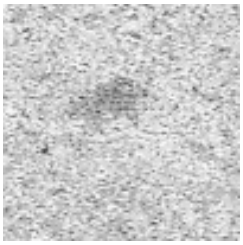

(a) Piramide Original

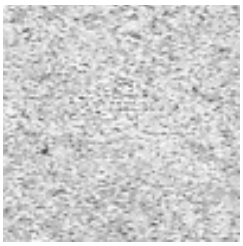

(b) Piramide Restored

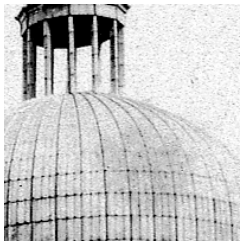

(c) Cupola Original

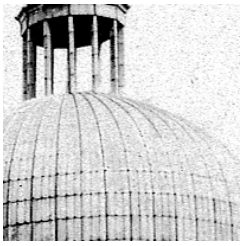

(d) Cupola Restored

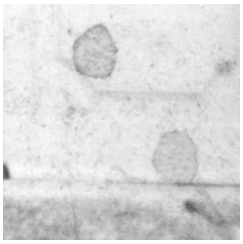

(e) Porta Original

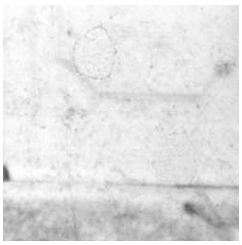

(f) Porta Restored

Fig. 2. Original luminance channel of the Pyramide (a), Cupola (c) and Porta (d) images. The restoration of the respective images are shown in (b), (d) and (f)

equilibrium in ideal conditions. In fact, the measured critical angle is $\theta = 9^o$. The detection algorithm [2] automatically selects the second scale level, as maximum level of the overcomplete wavelet representation. In this case the approximation band has to be restored while no restoration is required on the details. The final results are shown in Fig.2 (b) and (d). Fig. 2(e) shows a third example and a common case: non ideal conditions for the absorption process. Dirt causes a darker boundary due to the water diffusion on the paper. The critical angle is now $\theta = 36^o$ and both approximation and details have to be recovered. Artifacts in the details are eliminated accounting for visibility laws, as explained in Section 2. The final result is shown in Fig. 2(f). In all presented examples, the blotches are almost invisible in the restored images, even if the restoration of one the blotches in Fig. 2 (e) is not perfect because of the remaining slight boundary due to dirt.

Finally, it is worth noticing that the whole framework does not require any user interaction. In fact, the physical modeling of the defect along with visibility laws enable it to automatically adapt to the analysed image. Errors in the restoration process may arise in situations where the blotch is located close to a significant image discontinuity. In this case, samples from an inappropriate region of the image could be used for the "first guess", affecting the final result. In the case of a blotch edge corresponding to a clean image edge, information could be incorrectly removed in the wavelet detail restoration.

4. CONCLUSIONS

In this paper we have presented a model for recovering old documents affected by moisture. It is based on a Bayesian attenuation applied on the approximation band of an overcomplete wavelet representation. Wavelet details are recovered through a strategy based on visibility laws, just in case that degradation causes spurious edges due to dirt. First results are encouraging since the proposed approach exploits semi-transparency of this type of degradation. Nonetheless there are some cases where dirt causes a critical loss of original information. Future research will be then oriented to also study this case and improve the performances of the proposed approach.

5. REFERENCES

[1] F. Stanco, L. Tenze, and G. Ramponi, "Virtual restoration of vintage photographic prints affected by foxing and water blotches," *Journal of Electronic Imaging*, vol. 14, no. 4, Dec. 2005.

[2] V. Bruni, A.J. Crawford, and D. Vitulano, "Visibility based detection of complicated objects: A case study," in *Conference on Visual Media Production (CVMP)*, Nov. 2006.

[3] Rafael C. Gonzalez and Richard E. Woods, *Digital Image Processing*, Prentice Hall, 2nd edition, 2002.

[4] M. Beltramio, G. Sapiro, V. Caselles, and B. Bellester, "Image inpainting," *SIGGRAPH 2000*, pp. 417–424, 2000.

[5] A. Criminisi, P. Perez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1200–1212, Sept. 2004.

[6] A. Clarke, T.D. Blake, K. Carruthers, and A. Woodward, "Spreading and imbibition of liquid droplets on porous surfaces," *Langmuir Letters 2002 American Chemical Society*, vol. 18, no. 8, pp. 2980–2984, 2002.

[7] D. Seveno, V. Ledauphine, G. Martic, and M. Voué, "Spreading drop dynamics on porous surfaces," *Langmuir 2002 American Chemical Society*, vol. 18, no. 20, pp. 7496–7502, 2002.

[8] John Y. A. Wang and Edward H. Adelson, "Representing moving images with layers," *IEEE Transactions on Image Processing*, vol. 3, no. 5, pp. 625–638, Sept. 1994.

[9] P.R. White, W.B. Collis, S. Robinson, and A.C. Kokaram, "Inference matting," in *Conference on Visual Media Production (CVMP)*, Nov. 2005.

[10] J. R. Besag, "On the analysis of dirty pictures," *Journal of the Royal Statistical Society B*, vol. 48, pp. 259–302, 1986.

[11] Stefan Winkler, *Digital Video Quality - Vision Models and Metrics*, John Wiley and Sons, 2005.