BAYESIAN RESTORATION OF COLOR IMAGES USING A NON-HOMOGENOUS CROSS-CHANNEL PRIOR

Manu Parmar, Stanley J. Reeves, and Thomas S. Denney, Jr.

Department of ECE, Auburn University, Auburn, AL-36849, USA.

ABSTRACT

Gauss-Markov Random Field (GMRF) image models are commonly used in many Bayesian-based imaging techniques to define priors that lead to computationally tractable solutions for the image restoration problem. Recently, the color reconstruction literature has demonstrated, and effectively employed, the high correlation among the bands of a color image for color reconstruction. In this paper, we formulate a compound GMRF prior based on cross-channel spatial derivatives that reflects the smoothness in the color-difference space in addition to the often used intra-channel smoothness assumption. The proposed model is used to develop an effective method for restoring sparsely sampled color images in the presence of noise. The value of the proposed method is demonstrated on the problem of color reconstruction for single-sensor cameras.

Index Terms— Color reconstruction, Bayesian restoration

1. INTRODUCTION

The problem of restoration of images from sparse and noisy data is generally ill-posed and stable solutions are obtained only by the use of prior information in the estimates for true images. Bayesian methods are well suited for the problem of image restoration since they allow for flexible and effective means to incorporate prior information into the solution. Within the Bayesian paradigm, a common choice for the estimator is the maximum a posteriori (MAP) solution which provides an estimate for the true image x from the noisy data y as the maximum of the posterior probability density function $\mathrm{p}(x|y)$, which is in effect the most likely image given the occurrence of the data y. The choice of $\mathrm{p}(x)$, the prior density function, greatly affects the quality of the solution.

Markov random field (MRF) image models effectively describe the smoothness and local-nature of features in the general natural scene and have been extensively used to define image priors. MRFs are derived from potential functions that may be thought of as representing a quantity akin to energy. The desirability of configurations of local intensities depends on the value of the resulting potential function. Common potential functions of note are functions of the type $\psi(x_i - x_j)$, where x_i and x_j are intensity values at the i^{th} and j^{th} pixel respectively. A common prior model is the Gauss-Markov random field (GMRF) where ψ is a quadratic. A major issue with the GMRF model is its behavior across edges. Cost functions based on GMRF models will penalize large intensity differences between pixels and characteristically oversmooth across edges. Typically, local features and edges are accommodated by augmenting GMRF models with line processes [1] that lend a degree of adaptation to the estimation procedure. The line process l acts to inhibit smoothing across edges and encourages smoothing across pixels that do not lie across edges. The MAP estimate in this case is found as

$$\hat{x}, \hat{l} = \arg\max_{x,l} \mathrm{p}(x, l|y). \tag{1}$$

Bayesian restoration of multichannel images has previously been addressed in the literature. The authors [2] use multiple line processes to define a potential function that has cross-channel line-process terms. The cost function derived from the resulting GMRF prior reduces the penalty on intensity differences in a particular channel if line processes in other channels indicate the presence of an edge. This approach is useful since edge features in scenes typically appear at boundaries of objects and are reflected across all color channels. The Bayesian approach has also been used to address the specific problem of color reconstruction from single-sensor data [3]. The authors use a GMRF prior that includes only one line process that reflects the probability of the presence of an edge in all three color channels of a color image.

In this paper we propose a novel prior model for color images that incorporates cross-channel edge information. The proposed model is the result of a GMRF model augmented with line processes that attenuate the penalty on large intensity differences to prevent smoothing across edges. In addition, the prior has cross channel terms that control spatial smoothing in the color-difference channels. Section 2 details the reasoning behind this approach. The image model and restoration algorithm are developed in Section 3. In Section 4 the efficacy of the proposed algorithm is demonstrated on the problem of reconstruction of color images acquired by single-sensor digital cameras.

2. COLOR IMAGE MODEL

Digital color images are obtained either by direct acquisition with digital cameras or by scanning prints or slides obtained by film-based cameras. In both cases, the digital imaging device typically acquires three color bands in the red, green, and blue (RGB) regions of the spectrum. The forward model in the discrete form for the signal acquired at a pixel location is given by

$$c_i = \mathscr{F}\left(\sum_{k=1}^{N} r_k s_{i_k} + \eta_i\right), \qquad i = \mathrm{R}, \mathrm{G}, \mathrm{B} \tag{2}$$

where the visible range of wavelengths is sampled N times, c_i is the intensity of the i^{th} color, r_k and s_{i_k} are samples of scene irradiance and sensitivity of the i^{th} sensor respectively;

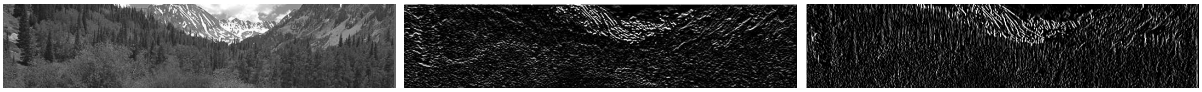

(a) From left to right: The red channel shown in grayscale, ∇_HR; and ∇_VR

(b) From left to right: The R-G channel difference image; ∇_H(R-G); and ∇_V(R-G)

(c) From left to right: The R-B channel difference image; ∇_H(R-B); and ∇_V(R-B)

Fig. 1. A representation of horizontal and vertical gradients obtained as the first differences in the respective directions.

η_i is noise, and $\mathscr{F}(.)$ is a non-linear function that describes the characteristics of the imaging system.

Channel intensities in (2) are functions of the inner products of the scene spectral content described by the irradiance and the respective sensor sensitivities. The irradiance incident on a sensor due to the general natural scene is a consequence of the reflectance of the scene and the radiance of the illuminant. Typically, the illuminant (sunlight, camera flash, fluorescent light, etc.) has broad support across the spectrum. Also, sensor-sensitivities have considerable overlapping support, especially among adjoining bands (red and green, green and blue). As a result, the channel intensities detected at a spatial location are well correlated.

Color images can be decomposed into the luminance channel, which describes brightness, and two chrominance channels, which convey information about color [4]. It is well known that most significant structures in color images manifest predominantly in the luminance channel. Color image enhancement algorithms take advantage of this feature, for instance, unsharp masking as applied to color images is performed only on the luminance channel. Luminance is commonly defined as a linear combination of the color channels, although the weights for each channel differ among various treatments. The appearance of major features in the luminance channel suggests that feature edges in typical color images are very well correlated. This phenomenon was demonstrated by the authors [5] by decomposing each color channel of a set of images into 4 bands by filtering with directional (0, $\pi/4$, $\pi/2$, and $3\pi/4$) high-pass filters. It was shown that corresponding high frequency components across the color channels are highly correlated. An equivalent assumption is that the color-difference channels are band-limited. The smoothness of the R-G and R-B color-difference channels in a natural scene is illustrated in the images in the left column in Figs. 1(b) and 1(c) respectively.

It follows from the high correlation among high frequency components of the color channels that the gradients of the color channels will be highly correlated. This is illustrated in Fig. 1. Figure 1(a) shows the red channel of an image from Kodak's PhotoCD PCD0992 [6] and the horizontal and vertical gradients found as the respective first differences of the red channel image (denoted hereafter with the symbols ∇_H and ∇_V respectively). Figs. 1(b) and 1(c) show the color-difference (∇_HR - ∇_HG, and ∇_VR - ∇_VG) images and the corresponding images for the R-B channel respectively. The color-difference images are characteristically smooth and the difference images of channel gradients illustrate their small magnitude. Table 1 shows the correlation coefficients among color channels and color channel gradients for an image obtained by stitching together all images in Eastman Kodak's PhotoCD PCD0992 [6]. Correlation coefficients for channel gradients in each direction are very high in each case, and highest for adjacent bands ($\nabla_\theta R \nabla_\theta G$ and $\nabla_\theta B \nabla_\theta R$).

Table 1. Correlation coefficients found for the database of images in Eastman Kodak's PhotoCD PCD0992 [6].

ρ	i = R, j = G	i = G, j = B	i = B, j = R
i, j	0.8534	0.9230	0.7560
$\nabla_H i, \nabla_H j$	0.9777	0.9765	0.9561
$\nabla_V i, \nabla_V j$	0.9751	0.9776	0.9565
$\nabla_{DL} i, \nabla_{DL} j$	0.9760	0.9778	0.9546
$\nabla_{DR} i, \nabla_{DR} j$	0.9751	0.9774	0.9532

2.1. Prior model

From the discussion in Section 2, it follows that a suitable prior model for color images will be formed from potential functions of the type

$$\psi\left(\left(x_i^k - x_j^k\right) - \left(x_i^{k'} - x_j^{k'}\right)\right), \tag{3}$$

where i, j are pixel locations, k, k' are R,G,B, and $k \neq k'$. The cost functions derived from such priors will penalize the difference between the first-difference of pixel intensities between color channels.

For an $M \times N$ color image, the true image x is defined as a realization of a random process defined on a 3-D rectangular lattice S with $3MN$ points. In addition we introduce two sets of line processes l_θ and c_θ, $\theta = H, V, DL, DR$, for the horizontal, vertical, and left and right diagonal directions respectively. The line processes l_θ model intra-channel intensity transitions and c_θ model the transitions in the color difference channels. Fig. 2 shows all line processes associated with the RGB intensities at a spatial location in the image.

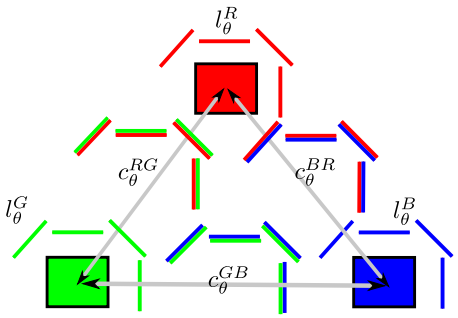

Fig. 2. Representation of a point in the 3-D lattice with associated line processes. Red, green and blue pixels are shown surrounded by the respective line processes that denote intra-channel edges (l_θ^k). Line processes for the cross-channel terms ($c_\theta^{kk'}$) are appropriately labeled.

The prior joint density for x, l_θ, and c_θ is defined as a Gibbs density to ensure that the resulting field is a MRF (Hammersley-Clifford theorem). A Gibbs field has a density function of the form

$$p(x) = \frac{\exp\left(-\sum_{i\in\mathscr{C}} \tilde{v}_i(x)\right)}{Z_x}, \tag{4}$$

where Z_x in the denominator normalizes the density function, $\tilde{v}_i(x)$ is the potential function defined over the set of cliques $\mathscr{C}$. A subset $\mathscr{C}$ of S is a clique if every pair of distinct sites in $\mathscr{C}$ are neighbors. Figure 3 shows the set of cliques used to define the proposed prior model

$$\begin{aligned} p(x, l_\theta, c_\theta) = \frac{1}{Z_x} \exp\Big(- \frac{1}{2\sigma_x^2} \sum_i \sum_{k,k'} \sum_\theta \zeta (x_i^k)^2 \\ + \left((x_i^k - x_{i:+\theta}^k) - (x_i^{k'} - x_{i:+\theta}^{k'})\right)^2 \tilde{c}_{\theta_i}^{kk'} \\ + (x_i^k - x_{i:+\theta}^k)^2 \, \tilde{l}_{\theta_i}^k \Big) \end{aligned} \tag{5}$$

where $1 \le i \le MN$; $\theta = H, V, DL, DR$; $\tilde{l}_{\theta_i} = 1 - l_{\theta_i}$, $\tilde{c}_{\theta_i} = 1 - c_{\theta_i}$; $k, k' = R, G, B$, $k \ne k'$ (x^k are random processes with intensities of the three color channels); and the term ζx_{ij}^2 keeps the density in (5) from being improper. ζ is set to a number that is small enough for this term to have little effect on the solution. We used $\zeta = 10^{-3}$ for all reconstructions in this paper. The index $i : +\theta$ refers to the pixel location adjacent to i in the direction of θ (illustrated in Fig. 3), and $i : -\theta$ will refer to the spatial location $i : +\theta + \pi$.

2.2. Degradation Model

The degradation model is presented in the matrix vector form as:

$$y = Ax + w, \tag{6}$$

where $y \in \Re^{3MN\times 1}$ is formed by stacking the three column-ordered color channels. The degradation to the true image $x \in \Re^{3MN\times 1}$ is described by A, which is a block matrix that has expressions for intra-channel blur as its diagonal blocks and cross-channel blur as the off-diagonal blocks. The additive noise w is zero mean, white, Gaussian noise with variance σ_w^2. The color channels are ordered column-wise and stacked such that $x = [x^{\mathrm{R}T}, x^{\mathrm{G}T}, x^{\mathrm{B}T}]^T$ (y is similarly arranged).

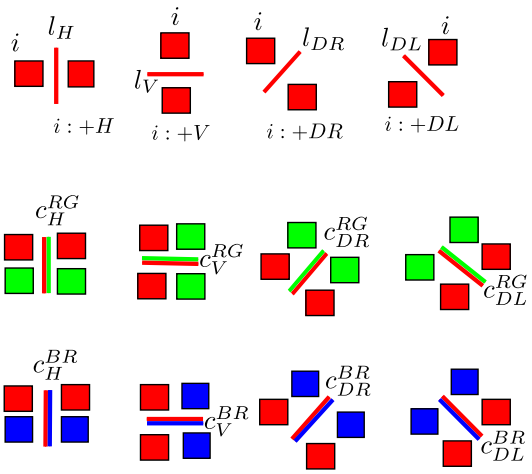

Fig. 3. The set of cliques associated with a red pixel at location i. Locations of $i : +\theta$, $\theta = H, V, DL, DR$ are labeled.

3. ALGORITHM DERIVATION

The maximum a posteriori (MAP) estimate of x, and the line processes given the prior described in (5) is obtained by maximizing

$$p(x, l_\theta, c_\theta) = \frac{p(y \mid x)\, p(x, l_\theta, c_\theta)}{p(y)}. \tag{7}$$

As $p(y)$ is constant with respect to x, l_θ, and c_θ, the optimal values of x, l_θ, and c_θ are the solution to the following optimization problem

$$\hat{x}, \hat{l}_\theta, \hat{c}_\theta = \arg\max_{x, l_\theta, c_\theta} p(y \mid x)\, p(x, l_\theta, c_\theta), \tag{8}$$

where $p(y \mid x)$ is the likelihood function in the presence of additive, uncorrelated Gaussian noise. Optimizing the cost function in (8) simultaneously over x, l_θ and c_θ is non-convex and computationally prohibitive. Instead, we iteratively update the estimate of x and then update the estimate of the edge variables. Results derived in [7] based on the the iterated conditional mode (ICM) algorithm [8] and the iterated conditional average (ICA) technique are used to update x, l_θ and c_θ.

ICM is a greedy iterative algorithm that sequentially updates values of pixels by maximizing their conditional posterior probability. Specifically, the value of x_i^k that maximizes the conditional probability of x_i^k, given all the remaining pixels x_j^k, $x_j^{k'}$, $x_i^{k'}$; $i \ne j$, and $k \ne k'$, and the associated edge variables is found. The required conditional probability is given by

$$p(x_i^k \mid l_\theta^k, c_\theta^{kk'}, x_{\setminus ik}) = p(y \mid x)\, p(x_i^k \mid x_{\mathrm{nb}}, l_\theta^k, c_\theta^{kk'}), \tag{9}$$

where $x_{\setminus ik}$ is the set of all pixels in the image except the i^{th} pixel of color k, and x_{nb} are neighboring pixels of x_i^k.

It can be shown [7] that the conditional posterior density of a pixel contingent on its neighboring elements x_{nb} is

$$p\left(x_i^k \mid y, x_{\mathrm{nb}}, l_{\theta_{\mathrm{nb}}}, c_{\theta_{\mathrm{nb}}}\right) = \frac{1}{\sqrt{2\pi}\,\sigma_i} e^{\left(-(x_i^k - \mu_i)^2 / 2\,\sigma_i^2\right)}, \tag{10}$$

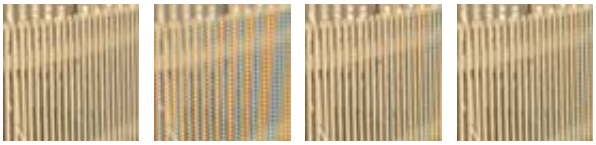

(a) Original (b) Bilinear (c) POCS (d) Proposed

Fig. 4. Experimental results for an image cropped from Image 19 (lighthouse) from [6].

where

$$\mu_i = \frac{a_i^T\left(y - Ax_{-i}\right) + (\sigma_w^2/\sigma_x^2)\rho_i^k}{a_i^H a_i + (\sigma_w^2/\sigma_x^2)\gamma_i^k}, \tag{11}$$

and a_i is the column of A corresponding to pixel x_i, x_{-i} is x with a zero in the i^{th} pixel, and

$$\begin{aligned}\rho_i^k = \sum_{\theta,k,k'} \Big(& x_{i:+\theta}^k \tilde{l}_{\theta_{i:+}}^k + \left(x_{i:+\theta}^k + x_i^{k'} - x_{i:+\theta}^{k'}\right)\tilde{c}_{\theta_{i:+}}^{kk'} \\ & + x_{i:-\theta}^k \tilde{l}_{\theta_i}^k + \left(x_{i:-\theta}^k + x_i^{k'} - x_{i:-\theta}^{k'}\right)\tilde{c}_{\theta_i}^{kk'}\Big) \\ \gamma_i^k = \zeta + \sum_{\theta,k,k'} & \left(\tilde{l}_{\theta_{i:+}}^k + \tilde{c}_{\theta_{i+}}^{kk'} + \tilde{l}_{\theta_i}^k + \tilde{c}_{\theta_i}^{kk'}\right)\end{aligned} \tag{12}$$

The conditional mean is also the conditional mode and maximizes the probability in Equation (9). The ICM algorithm consists of iteratively replacing pixel x_i^k with its conditional mean μ_i^k, i.e.

$$x_i^{k^{p+1}} = \frac{a_i^T\left(y - Ax_{-i}^p\right) + (\sigma_w^2/\sigma_x^2)\rho_i^p}{c + (\sigma_w^2/\sigma_x^2)\gamma_i^p}, \tag{13}$$

where $[.]^p$ denotes the value of a variable after the p^{th} iteration. The (ICA) algorithm [7] is used to update the line processes by iteratively updating a single line variable with its mean conditioned on x and all the other edge variables. The restoration algorithm iterates alternately between the pixel-updates and line variable updates until convergence. Updated values are used in subsequent iterations.

4. EXPERIMENTS

We demonstrate the proposed algorithm on the problem of reconstruction of full-color images acquired with single-sensor digital cameras. Such cameras use a sensor-array that acquires only one color at a particular spatial location. The full color image is reconstructed in an operation commonly referred to as demosaicking. The system model in this is case is

$$y = SHx + w \tag{14}$$

where S is a sampling matrix that and H represents blur. The proposed algorithm was used to reconstruct images sampled by a Bayer CFA. For lack of space only results of one experiment are shown here while more results may be found online at: *www.eng.auburn.edu/˜reevesj/research/ICIP07*. The blur matrix H for this particular experiment was set at the identity to facilitate comparison with demosaicking methods that do not include a deblurring operation. The values for σ_x, and σ_w were set at 10^{-2} and 10^{-1} respectively. Experimental results are shown in Fig. 4. Table 2 gives the RMSE values for the color channels of the reconstructed images.

Table 2. Channel RMS errors for images in Fig. 4

PSNR	Red	Green	Blue
Bilinear	22.6339	9.3327	23.1383
POCS [5]	7.4252	3.1876	8.1445
Proposed	6.4902	4.0993	6.8904

In the digital camera pipeline, deblurring is typically carried out as a post-processing operation. The demosaicking operation that occurs earlier is inherently a nonlinear blurring operation. Deblurring carried out at a late stage in the imaging pipeline can not take advantage of linear models of blurring. In contrast, the proposed algorithm can be used to simultaneously deblur and demosaic single-senor camera images.

5. CONCLUSIONS

In this paper we have proposed a novel cross-channel GMRF prior model for color images. The proposed prior model takes advantage of the high correlation between high-frequency edge content across the color channels of an image. We also propose a Bayesian algorithm that uses the prior model to obtain MAP estimates for the restoration of color images. The proposed algorithm can be used to jointly demosaic and deblur images acquired with single-sensor digital cameras. The efficacy of the algorithm is demonstrated in an experiment in which we demosaic a sparsely sampled color image to arrive at a good estimate of the full-color image.

6. REFERENCES

[1] S.Geman and D.Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721–741, 1984.

[2] R.Molina, *et al.*, "Bayesian multichannel image restoration using compound gauss-markov random fields," *Image Processing, IEEE Transactions on*, vol. 12, no. 12, pp. 1642–1654, Dec. 2003.

[3] M.Parmar, *et al.*, "Bayesian edge-preserving color image reconstruction from color filter array data," in *Computational Imaging III*, C. A.Bouman and E. L.Miller, Eds., vol. 5674, no. 1. SPIE, 2005, pp. 259–268.

[4] G.Sharma, Ed., *Digital Color Imaging Handbook*. Boca Raton, FL, USA: CRC Press, Inc., 2002, ch. one: Color fundamentals for digital imaging.

[5] B.Gunturk, *et al.*, "Color plane interpolation using alternating projections," *Image Processing, IEEE Transactions on*, vol. 11, no. 9, pp. 997–1013, Sept. 2002.

[6] Eastman Kodak Company, "PhotoCD PCD0992," (http://r0k.us/graphics/kodak/).

[7] T. S.Denney Jr. and S. J.Reeves, "Bayesian image reconstruction from fourier-domain samples using prior edge information," *Journal of Electronic Imaging*, vol. 14, no. 4, p. 043009, 2005.

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