COLOR IMAGE DENOISING BASED ON MULTICHANNEL NON-LOCAL MEANS FUSION

Jingjing Dai, Oscar C. Au, Feng Zou, Chao Pang, and Lu Fang

Hong Kong University of Science and Technology
Department of Electronic and Computer Engineering
Clear Water Bay, Kowloon, Hong Kong

ABSTRACT

In this paper, we investigate the problem of color image denoising, and propose a novel algorithm called multichannel non-local means fusion (MNLMF), building on the grayscale denoiser non-local means filter. By analyzing and modeling the inter-channel correlation in color images, we formulate the color noise reduction as a minimization problem with a specifically-designed penalty function which fully takes advantages of the inter-channel prior information. The optimal solution is derived consisting of constructing multiple non-local means spanning all three channels and fusing them together. The weights in the fusion are optimized to minimize the overall denoising error. Simulation results under various noise levels demonstrate that when compared to other state-of-the-art algorithms, the proposed MNLMF achieves competitive performance both in terms of the color peak signal-to-noise ratio (cPSNR) and in perceptual quality.

Index Terms— color image denoising, inter-channel correlation, non-local means

1. INTRODUCTION

During the past decades, image denoising has seen a huge amount of work, and many directions have been successfully visited. Some famous work includes total-variation (TV) based restoration [1], non-local means filter (NLM) [2], wavelet-domain thresholding [3], etc. Despite the richness of work on the denoising of grayscale images, color image denoising has received much less attention. As color image production and application are getting extremely popular nowadays, color noise reduction is becoming an essential task in today's image processing systems. The key issue in color image denoising is how to exploit the inter-channel correlation in the RGB signal, and there are two feasible strategies in the literature [4].

The first strategy is to avoid the correlation issue by converting the RGB signal into a more 'decorrelated' color space and apply the single-component denoiser to every component individually in the new space. [5] addresses color noise reduction in the hue-saturation-value (HSV) space through geometrical "good" continuation of hue while treating the noise in the other two dimensions (saturation and value) through independent scalar anisotropic diffusion. The luminance-chrominance space has been found to be a good choice in the wavelet-based methods [6].

The second strategy is to introduce coupling between RGB components and devise a non-separable denoising formula. [7] takes into account the inter-channel dependency by modeling the color-difference signal, and achieves denoising by alternate projections onto intra-channel and inter-channel constraint sets. The color adaption of the popular NLM is presented in [2] and [8], of which the core idea is to enhance the robustness of weights by computing the resemblance of color patches. [4] devises a vector extension of the wavelet-domain monochannel denoiser [3] invented by the same authors to effectively process multichannel images.

In this paper, we propose a novel color image denoising algorithm called multichannel non-local means fusion (MNLMF), building on the grayscale denoiser NLM. Color noise reduction is formulated as the minimization of a specifically-designed penalty function which makes full use of the inter-channel correlation. The optimal solution for every channel is derived consisting of constructing multiple non-local means spanning all three channels and fusing them together. The optimal weights in the fusion are derived to minimize the overall denoising error.

The remainder of this paper is organized as follows. In the next section, we describe the prior art of NLM for grayscale and color image denoising. The proposed MNLMF is developed and presented in details in Section III. Section IV is dedicated to the quantitative and qualitative performance evaluation of MNLMF as well as comparison with other state-of-the-art techniques and finally Section V concludes this work.

2. PRIOR ART

In this section, we consider an N-pixel image, and denote it by the vector $\mathbf{y}$ with column-stacked pixel values. The image is corrupted by additive white Gaussian noise (AWGN) $\mathbf{n}$ with the noise standard deviation σ_n, and the resulting noisy image is denoted as $\mathbf{z} = \mathbf{y} + \mathbf{n}$. The NLM [2] achieves the denoised estimate $\hat{\mathbf{y}}(k)$ as follows

$$\hat{\mathbf{y}}(k) = \frac{\sum_{l \in \mathcal{S}_k} w(k,l)\mathbf{z}(l)}{\sum_{l \in \mathcal{S}_k} w(k,l)}, \tag{1}$$

where $\mathcal{S}_k$ is the search region centered at k and $w(k,l)$ is the weight for $\mathbf{z}(l)$ when computing $\hat{\mathbf{y}}(k)$. The motivation behind NLM is that there exist many similar patterns in natural images, and we call this feature self-similarity. The main idea of NLM is to make use of this self-similarity to suppress the noise. The weight $w(k,l)$ is determined by $d(k,l)$, the distance between the patches centered respectively at k and l, which is calculated as

$$d(k,l) = \frac{\sum_{b \in \mathcal{B}} (\mathbf{z}(k+b) - \mathbf{z}(l+b))^2}{|\mathcal{B}|}, \tag{2}$$

with $\mathcal{B}$ defining the patch, and $|\mathcal{B}|$ being its total size. An exponential kernel is used to relate $w(k,l)$ with $d(k,l)$:

$$w(k,l) = e^{-\frac{d(k,l)/\sigma_n^2}{h^2}}, \tag{3}$$

where h is the filtering parameter. As such, a small distance (high similarity) leads to a large weight, and vice versa.

The theoretical explanation of NLM is given in [9] by introducing the following penalty function:

$$\eta^2(\mathbf{y}) = \|\mathbf{y}-\mathbf{z}\|_2^2 + \lambda \sum_{k\in\mathcal{I}} \sum_{l\in\mathcal{S}_k} w(k,l)\cdot(\mathbf{y}(k)-\mathbf{y}(l))^2, \quad (4)$$

where $\mathcal{I}$ defines the whole image region. The first term corresponds to the log-likelihood function under AWGN. The second term stands for the prior constraint of self-similarity that requires each pixel in the denoised image to resemble other pixels in its vicinity, provided they have similar surrounding. The λ is the regularization parameter balancing the weight between these two terms. The optimal solution minimizing η can be solved iteratively. Using one iteration with the initialization of $\mathbf{y}^0 = \mathbf{z}$ basically means that the prior term requires pixels in the output denoised image to have similar values to pixels in the corresponding search region in $\mathbf{z}$. Putting this idea directly into the penalty term gives

$$\eta_a^2(\mathbf{y}) = \sum_{k\in\mathcal{I}} \sum_{l\in\mathcal{S}_k} w(k,l)\cdot(\mathbf{y}(k)-\mathbf{z}(l))^2. \quad (5)$$

Then the optimal solution that minimizes η_a^2 is exactly NLM.

When the color signal is concerned, the observed image is corrupted by an additive channel-wise white Gaussian noise, and the observation model turns into

$$\mathbf{z}_R = \mathbf{y}_R + \mathbf{n}_R, \mathbf{z}_G = \mathbf{y}_G + \mathbf{n}_G, \mathbf{z}_B = \mathbf{y}_B + \mathbf{n}_B, \quad (6)$$

with the subscript R, G, and B denoting the red, green, and blue channel respectively. The non-local means for color (NLMC) works as follows [2, 8]:

$$\hat{\mathbf{y}}_c(k) = \frac{\sum_{l\in\mathcal{S}_k} w(k,l)\mathbf{z}_c(l)}{\sum_{l\in\mathcal{S}_k} w(k,l)}, \quad (7)$$

with $c \in \{R,G,B\}$. The inter-channel correlation in the RGB signal is enforced by computing the distance between the three-channel color patches, to enhance the robustness of the weights towards noise. Correspondingly, the term $\frac{d(k,l)}{\sigma_n^2}$ in (3) is substituted by

$$\frac{\sum_{b\in\mathcal{B}} \delta_{k+b,l+b}^T C_n^{-1} \delta_{k+b,l+b}}{3|\mathcal{B}|}, \quad (8)$$

where $\delta_{k+b,l+b} = [\mathbf{z}_R(k+b)-\mathbf{z}_R(l+b), \mathbf{z}_G(k+b)-\mathbf{z}_G(l+b), \mathbf{z}_B(k+b)-\mathbf{z}_B(l+b)]^T$, and C_n is the 3×3 covariance matrix of noise. In this paper, C_n is assumed to be known in advance.

3. ALGORITHM

3.1. Motivation

It has been well known that the RGB channels are highly correlated with each other. Therefore, when color images are concerned, denoising algorithms should exploit this inter-channel correlation to expect a better denoising performance. In NLMC, the inter-channel correlation is enforced through improved weight calculation. The motivation behind this is that the red, green, and blue channels manifest almost identical image patterns. Correspondingly, if the patches centered at two pixels are similar (or dissimilar) in one channel, then they are also similar (or dissimilar) in the other two channels. Therefore, the filtering of RGB channels can employ the same set of weights, and the computed weights are more noise-robust when all three channels are involved in the distance calculation. The NLMC is successful in improving the weight calculation, but the inter-channel correlation has not been fully exploited.

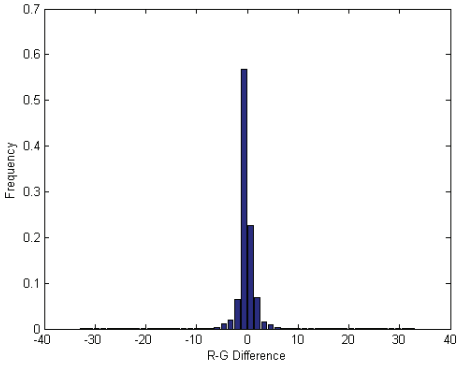

Fig. 1. Histogram of R-G difference after low frequency compensation.

The inter-channel correlation can be exploited by cross-channel weighted prediction [10]. Take the red and green channels as an example and consider any red and green pixel y_R and y_G, the predictor $p(y_R)$ for the red channel generated from the green channel is constructed as

$$p(y_R) = wy_G + o, \quad (9)$$

where w is the weighting factor and o is the additive offset. Taking into account the high similarity between the spatial gradients in RGB channels [10], (9) can be simplified by setting $w = 1$. Further, by optimizing the offset to minimize the prediction error, (9) turns into

$$p(y_R) = y_G + \bar{y}_R - \bar{y}_G, \quad (10)$$

where $\bar{y}_R$ and $\bar{y}_G$ denote the mean of y_R and y_G respectively. The popularity of the predictor in (10) indicates that after mean compensation, the pixel intensities of different channels resemble each other. With $\mathbf{y}_R^{lp}$, $\mathbf{y}_G^{lp}$, and $\mathbf{y}_B^{lp}$ we denote the low-passed red, green, and blue channels used to approximate the mean (the superscript lp is used to denote the low-passed signal in the remaining part of this paper if not specified), the inter-channel correlation manifests itself in the form of the similarity between $(\mathbf{y}_R - \mathbf{y}_R^{lp})$, $(\mathbf{y}_G - \mathbf{y}_G^{lp})$, and $(\mathbf{y}_B - \mathbf{y}_B^{lp})$. This argument is further supported by the wavelet-domain analysis in [6], which claims that the high-frequency wavelet coefficients of red, green, and blue channels are not only highly correlated with but also approximately equal to each other.

Here, we conduct a simple experiment to verify the inter-channel similarity on color images. We apply the 3×3 average filter as the low-pass filter, and study the difference between different color pairs after low frequency compensation. The histogram of $((\mathbf{y}_R - \mathbf{y}_R^{lp}) - (\mathbf{y}_G - \mathbf{y}_G^{lp}))$ taken from one image is plotted in Fig. 1. As observed, most difference values cluster around zero, and the frequency of a large difference is pretty low. Similar histograms are observed in other channel pairs and in other test images. Motivated by this observation, we propose MNLMF to improve over the NLMC, which is the topic of the next section.

3.2. Multichannel Non-local Means Fusion

We introduce the inter-channel similarity as shown in the previous section into η_a^2 in (5), and construct the following penalty function

for color image denoising:

$$\begin{aligned}\eta_c^2(\mathbf{y}_R,\mathbf{y}_G,\mathbf{y}_B) &= \lambda_R\sum_{k\in\mathcal{I}}\sum_{l\in\mathcal{S}_k} w(k,l)\cdot(\mathbf{y}_R(k)-\mathbf{z}_R(l))^2 \\ &+\lambda_G\sum_{k\in\mathcal{I}}\sum_{l\in\mathcal{S}_k} w(k,l)\cdot(\mathbf{y}_G(k)-\mathbf{z}_G(l))^2 \\ &+\lambda_B\sum_{k\in\mathcal{I}}\sum_{l\in\mathcal{S}_k} w(k,l)\cdot(\mathbf{y}_B(k)-\mathbf{z}_B(l))^2 \\ &+\lambda_{RG}\cdot\|\mathbf{y}_R-\mathbf{L}\mathbf{y}_R-\mathbf{y}_G+\mathbf{L}\mathbf{y}_G\|_2^2 \\ &+\lambda_{RB}\cdot\|\mathbf{y}_R-\mathbf{L}\mathbf{y}_R-\mathbf{y}_B+\mathbf{L}\mathbf{y}_B\|_2^2 \\ &+\lambda_{GB}\cdot\|\mathbf{y}_G-\mathbf{L}\mathbf{y}_G-\mathbf{y}_B+\mathbf{L}\mathbf{y}_B\|_2^2 \quad (11)\end{aligned}$$

where λ_R, λ_G, λ_B, λ_{RG}, λ_{RB}, and λ_{GB} are the regularization parameters. The first three prior constraints correspond to the self-similarity in each channel, and the last three prior constraints account for the inter-channel similarity, with $\mathbf{L}$ being an $N\times N$ matrix representing the low-pass operation. The optimal solution minimizing η_c^2 can be achieved by nulling the derivative of η_c^2 with respect to $\mathbf{y}_R$, $\mathbf{y}_G$, and $\mathbf{y}_B$. While directly solving it is possible, it requires an inversion of a very large matrix. Instead, we try to seek a simplified formulation and solution.

We then take a different angle to look at the constraint of the inter-channel similarity. Take the term $\|\mathbf{y}_R-\mathbf{L}\mathbf{y}_R-\mathbf{y}_G+\mathbf{L}\mathbf{y}_G\|_2^2$ as an example, this constraint requires $\mathbf{y}_R(k)$ to be similar to $(\mathbf{y}_G(k)-\mathbf{y}_G^{lp}(k)+\mathbf{y}_R^{lp}(k))$, as well as $\mathbf{y}_G(k)$ to be similar to $(\mathbf{y}_R(k)-\mathbf{y}_R^{lp}(k)+\mathbf{y}_G^{lp}(k))$. Recall that the self-similarity constraints require $\mathbf{y}_R(k)$ and $\mathbf{y}_G(k)$ to be similar to pixels in the corresponding search region in $\mathbf{z}_R$ and $\mathbf{z}_G$ respectively. Combined with these self-similarity constraints, $\|\mathbf{y}_R-\mathbf{L}\mathbf{y}_R-\mathbf{y}_G+\mathbf{L}\mathbf{y}_G\|_2^2$ can be reformulated as

$$\begin{aligned}&\sum_{k\in\mathcal{I}}\sum_{l\in\mathcal{S}_k} w(k,l)\cdot[(\mathbf{y}_R(k)-(\mathbf{z}_G(l)-\mathbf{z}_G^{lp}(l)+\mathbf{z}_R^{lp}(l)))^2 \\ &+(\mathbf{y}_G(k)-(\mathbf{z}_R(l)-\mathbf{z}_R^{lp}(l)+\mathbf{z}_G^{lp}(l)))^2].\end{aligned}$$

The other two terms of inter-channel similarity can be revised in the same way, yielding the new penalty function, which turns to be separable with respect to $\mathbf{y}_R$, $\mathbf{y}_G$, and $\mathbf{y}_B$. Thus, the optimal solutions for the three channels can be solved independently. Take the red channel as an example, the part of the penalty function related to it is

$$\begin{aligned}\eta_R^2(\mathbf{y}_R) &= \sum_{k\in\mathcal{I}}\sum_{l\in\mathcal{S}_k} w(k,l)\cdot[\lambda_R\cdot(\mathbf{y}_R(k)-\mathbf{z}_R(l))^2 \\ &+\lambda_{RG}\cdot(\mathbf{y}_R(k)-(\mathbf{z}_G(l)-\mathbf{z}_G^{lp}(l)+\mathbf{z}_R^{lp}(l)))^2 \\ &+\lambda_{RB}\cdot(\mathbf{y}_R(k)-(\mathbf{z}_B(l)-\mathbf{z}_B^{lp}(l)+\mathbf{z}_R^{lp}(l)))^2]\end{aligned}$$

Targeting now the minimization of η_R^2, we null its derivative with respect to $\mathbf{y}_R$, and derive the optimal solution as

$$\hat{\mathbf{y}}_R(k)=w_{Rr}\hat{\mathbf{y}}_{Rr}(k)+w_{Rg}\hat{\mathbf{y}}_{Rg}(k)+w_{Rb}\hat{\mathbf{y}}_{Rb}(k), \quad (12)$$

where

$$\hat{\mathbf{y}}_{Rr}(k)=\frac{\sum_{l\in\mathcal{S}_k} w(k,l)\mathbf{z}_R(l)}{\sum_{l\in\mathcal{S}_k} w(k,l)}, \quad (13)$$

$$\hat{\mathbf{y}}_{Rg}(k)=\frac{\sum_{l\in\mathcal{S}_k} w(k,l)(\mathbf{z}_G(l)-\mathbf{z}_G^{lp}(l)+\mathbf{z}_R^{lp}(l))}{\sum_{l\in\mathcal{S}_k} w(k,l)}, \quad (14)$$

$$\hat{\mathbf{y}}_{Rb}(k)=\frac{\sum_{l\in\mathcal{S}_k} w(k,l)(\mathbf{z}_B(l)-\mathbf{z}_B^{lp}(l)+\mathbf{z}_R^{lp}(l))}{\sum_{l\in\mathcal{S}_k} w(k,l)}, \quad (15)$$

$$w_{Rr}=\frac{\lambda_R}{\lambda_R+\lambda_{RG}+\lambda_{RB}}, \quad (16)$$

$$w_{Rg}=\frac{\lambda_{RG}}{\lambda_R+\lambda_{RG}+\lambda_{RB}}, \quad (17)$$

$$w_{Rb}=\frac{\lambda_{RB}}{\lambda_R+\lambda_{RG}+\lambda_{RB}}. \quad (18)$$

This optimal solution for the red channel in essence constructs the weighted average in three planes, z_R, $(z_G-z_G^{lp}+z_R^{lp})$, and $(z_B-z_B^{lp}+z_R^{lp})$, and correspondingly generates three denoised estimates, $\hat{\mathbf{y}}_{Rr}$, $\hat{\mathbf{y}}_{Rg}$, and $\hat{\mathbf{y}}_{Rb}$, which are finally fused together.

Recall that the main idea of NLM is to make use of similar patterns in one image to suppress the noise. Naturally, if more similar patterns occur, better denoising performance is expected. Note that the weighted average of NLMC in (7) is restricted to being in the same channel, and thus only similar patterns in the same channel are exploited. In fact, the inter-channel resemblance suggests that there exist similar patterns cross channels. Our optimal solution then allows the weighted average to span all three channels. In this case, more similar patches can be made use of, and thus the denoising performance can be improved.

The remaining issue is to choose w_{Rr}, w_{Rg}, and w_{Rb}. The goal of optimizing them is to minimize the denoising error $\epsilon_R=\frac{1}{N}\|\hat{\mathbf{y}}_R-\mathbf{y}_R\|_2^2$. With $\mathbf{w}_R=[w_{Rr},w_{Rg},w_{Rb}]^T$, $\mathbf{e}_{Rr}=\hat{\mathbf{y}}_{Rr}-\mathbf{y}_R$, $\mathbf{e}_{Rg}=\hat{\mathbf{y}}_{Rg}-\mathbf{y}_R$, and $\mathbf{e}_{Rb}=\hat{\mathbf{y}}_{Rb}-\mathbf{y}_R$, ϵ_R can be written as

$$\epsilon_R = \frac{1}{N}\|[\hat{\mathbf{y}}_{Rr},\hat{\mathbf{y}}_{Rg},\hat{\mathbf{y}}_{Rb}]\mathbf{w}_R-\mathbf{y}_R\|_2^2=\mathbf{w}_R^T C_R\mathbf{w}_R, \quad (19)$$

where C_R is the matrix of denoising error, defined as:

$$C_R=\frac{1}{N}\begin{bmatrix}\mathbf{e}_{Rr}^T\mathbf{e}_{Rr} & \mathbf{e}_{Rr}^T\mathbf{e}_{Rg} & \mathbf{e}_{Rr}^T\mathbf{e}_{Rb}\\ \mathbf{e}_{Rg}^T\mathbf{e}_{Rr} & \mathbf{e}_{Rg}^T\mathbf{e}_{Rg} & \mathbf{e}_{Rg}^T\mathbf{e}_{Rb}\\ \mathbf{e}_{Rb}^T\mathbf{e}_{Rr} & \mathbf{e}_{Rb}^T\mathbf{e}_{Rg} & \mathbf{e}_{Rb}^T\mathbf{e}_{Rb}\end{bmatrix}.$$

The optimization of $\mathbf{w}_R$ can be formulated as

$$\begin{aligned}\min_{\mathbf{w}_R} \quad & \mathbf{w}_R^T C_R\mathbf{w}_R \\ s.t. \quad & \mathbf{w}_R^T\cdot\mathbf{1}=1.\end{aligned}$$

Using Lagrange multipliers, the optimal $\mathbf{w}_R$ is

$$\mathbf{w}_R=\frac{C_R^{-1}\mathbf{1}}{\mathbf{1}^T C_R^{-1}\mathbf{1}}. \quad (20)$$

In order to obtain the optimal $\mathbf{w}_R$ as defined in (20), we need to know C_R. The difficulty lies in the fact that in practice we cannot compute C_R directly from the definition. Fortunately, we can extend and adapt Stein's unbiased risk estimator (SURE) [11] to achieve an accurate estimate.

4. EXPERIMENTAL RESULTS

In this section, experiments are conducted to evaluate MNLMF. As a performance measure we adopt the color peak signal-to-noise ratio (cPSNR) defined as $10log_{10}(\frac{255^2}{\text{MSE}})$, where the MSE is the mean squared error averaged over all pixels and all three channels.

To evaluate the improvements brought by multichannel nonlocal means compared with the traditional single channel nonlocal means, MNLMF is compared with NLMC. We also compare MNLMF with other state-of-the-art methods: OCP [6], LET [4], and AP [7]. The cPSNR values are displayed in Table I. Note that AP is not specially optimized for the nonuniform noise, and we

Table 1. cPSNR comparisons between MNLMF and other state-of-the-arts under various noise cases. For uniform noise, the noise standard deviations in three channels are all 25.5. For uncorrelated nonuniform noise, the noise standard deviations in RGB channel $(\sigma_{nR}, \sigma_{nG}, \sigma_{nB})$ are $(38.25, 25.5, 12.75)$. For correlated nonuniform noise, the noise standard deviations in RGB channel are $(12.75, 38.25, 25.5)$, and the correlation coefficients of cross-channel noise for (R&G, R&B, G&B) is $(0.3, 0.3, 0.3)$.

Image	Uncorrelated noise									Correlated noise			
	Uniform					Nonuniform				Nonuniform			
	OCP	AP	LET	NLMC	MNLMF	OCP	LET	NLMC	MNLMF	OCP	LET	NLMC	MNLMF
1	27.42	27.89	27.92	27.01	**28.16**	28.73	29.33	27.85	**29.44**	27.48	28.53	27.58	**28.86**
2	30.22	30.35	**31.03**	30.28	30.99	30.48	31.39	30.44	**31.46**	30.14	31.46	30.83	**31.84**
3	30.54	31.24	31.54	31.65	**32.34**	30.76	32.26	31.97	**33.22**	30.97	31.92	31.91	**32.99**
4	30.41	30.77	31.13	30.73	**31.28**	30.85	31.80	31.08	**32.24**	30.39	31.63	30.96	**31.94**
5	27.11	27.01	27.91	27.55	**28.70**	28.11	28.97	28.16	**29.80**	27.34	28.43	27.98	**29.31**
6	28.01	28.53	28.88	28.09	**29.14**	28.66	30.19	28.86	**30.28**	28.59	29.66	28.70	**29.97**
7	29.74	30.55	30.57	30.54	**31.34**	30.46	31.44	30.93	**32.37**	30.01	31.00	30.75	**32.00**
8	27.04	26.98	27.84	28.11	**28.92**	28.13	29.13	28.27	**30.05**	27.39	28.48	28.13	**29.55**
Ave.	28.81	29.17	29.60	29.25	**30.11**	29.52	30.56	29.70	**31.11**	29.04	30.14	29.61	**30.81**

have found that the performance of AP deteriorates a lot under the nonuniform noise. Therefore, we do not include AP in the comparison for the case of nonuniform noise. As observed, the integration of multichannel non-local means significantly improves the denoising performance, as MNLMF consistently outperforms NLMC, by up to 1.78 dB. In addition, MNLMF clearly outperforms OCP and AP. When compared with LET, MNLMF produces slightly worse performance only in one case, and produces higher cPSNR values most of the time, with the cPSNR gain up to 1.08 dB.

For comparisons of the visual appearance, some typical examples are depicted in Fig. 2. Compared with other methods, the denoised images produced by MNLMF demonstrate higher visual quality, with less remaining noise, and fewer artifacts.

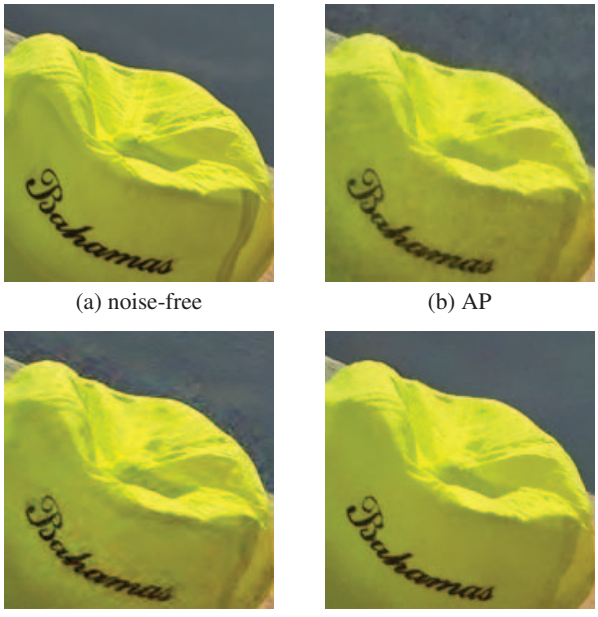

(a) noise-free (b) AP

(c) LET (d) MNLMF

Fig. 2. Visual comparisons on Image 3 corrupted by uncorrelated uniform noise.

5. CONCLUSION

In this paper, we propose a novel algorithm MNLMF for color image denoising. The color noise reduction is formulated as a minimization problem, and the optimal solution is derived. MNLMF allows the weighted average filter to span all three channels, and thus can fully exploit the inter-channel correlation in the RGB signal. Experimental results show that MNLMF produces favorable performance compared with several state-of-the-art algorithms in terms of both cPSNR and subjective quality.

6. ACKNOWLEDGEMENT

This work has been supported in part by the Research Grants Council (RGC) of the Hong Kong Special Administrative Region, China.

7. REFERENCES

[1] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.

[2] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling and Simulation*, vol. 4, no. 2, pp. 490–530, 2006.

[3] F. Luisier, T. Blu, and M. Unser, "A new sure approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Transactions on Image Processing*, vol. 16, no. 3, pp. 593–606, 2007.

[4] F. Luisier and T. Blu, "SURE-LET multichannel image denoising: Interscale orthonormal wavelet thresholding," *IEEE Transactions on Image Processing*, vol. 17, no. 4, pp. 482–492, April 2008.

[5] O. Ben-Shahar and S. W. Zucker, "Hue fields and color curvatures: A perceptual organization approach to color image denoising," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2003, vol. 2, pp. 713–720.

[6] N. X. Lian, V. Zagorodnov, and Y. P. Tan, "Edge-preserving image denoising via optimal color space projection," *IEEE Transactions on Image Processing*, vol. 15, no. 9, pp. 2575–2587, September 2006.

[7] X. Li, "On modeling interchannel dependency for color image denoising," *International Journal of Imaging Systems and Technology*, vol. 17, no. 3, pp. 163–173, 2007.

[8] B. Goossens, H. Q. Luong, J. Aelterman, A. Pizurica, and W. Philips, "A GPU-accelerated real-time NLMeans algorithm for denoising color video sequences," in *Lecture Notes in Computer Science (Advanced Concepts For Intelligent Vision Systems)*, 2010, pp. 46–57.

[9] M. Protter, M. Elad, H. Takeda, and P. Milanfar, "Generalizing the nonlocal-means to super-resolution reconstruction," *IEEE Transactions on Image Processing*, vol. 18, no. 1, pp. 36–51, 2009.

[10] S. Benierbah and M. Khamadja, "Compression of colour images by inter-band compensated prediction," in *IEE Proceedings on Vision, Image and Signal Processing*, April 2006, vol. 153, pp. 237–243.

[11] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *The Annals of Statistics*, vol. 9, pp. 1135–1151, 1981.