ELSEVIER

Signal Processing: *Image Communication* 18 (2003) 91–101

SIGNAL PROCESSING: IMAGE COMMUNICATION

www.elsevier.com/locate/image

Postprocessing of decoded color images by adaptive linear filtering

Ulug Bayazit*

Department of Electronics Engineering, Isik University, Buyukdere Cad., Maslak, Istanbul 80670, Turkey

Received 22 March 2002; received in revised form 10 July 2002; accepted 22 July 2002

Abstract

This paper presents an image adaptive linear filtering method for the reconstruction of the *RGB* (red, blue, green) color coordinates of a pixel from the lossy compressed luminance/chrominance color coordinates. In the absence of quantization noise, the *RGB* coordinates of a pixel can be perfectly reconstructed by employing a standard, fixed filter whose support includes only the luminance/chrominance coordinates at the spatial location of the pixel. However, in the presence of quantization noise, a filter with a larger support, that also spatially extends over the luminance/chrominance coordinate planes, is capable of exploiting the statistical dependence among the luminance/chrominance coordinate planes, and thereby yields more accurate reconstruction than the standard, fixed filter. We propose the optimal (in the minimum mean squared error sense) determination of the coefficients of this adaptive linear filter at the image encoder by solving a system of regression equations. When transmitted as side information to the image decoder, the filter coefficients need not incur significant overhead if they are quantized and compressed intelligently. Our simulation results demonstrate that the distortion of the decompressed color coordinate planes can be reduced by several tenths of a dB with negligible overhead rate by the application of our image adaptive linear filtering method.

Keywords: Color image coding; Color image processing; Adaptive linear filter

1. Introduction

In order to achieve high compression performance in color image coding, statistical dependencies among color coordinate planes must be exploited. Towards this end, the highly correlated *RGB* color coordinates can either be jointly coded by means of vector quantization, or can be decorrelated by means of a linear coordinate transformation and the resulting less correlated coordinate planes can be independently coded. Vector quantization is performed by grouping either the color coordinates [1,2,7], or the wavelet coefficients of the color coordinates [3,4] into vectors.

Karhunen Loeve Transform (KLT) is the optimal linear coordinate transformation which yields uncorrelated transform coefficients. However, *KLT* is rarely used for this purpose, since it is data dependent, and the transformation matrix needs to be transmitted to the decoder as side information. Instead, fixed linear coordinate transformations, yielding highly uncorrelated one

*Tel.: +90-212-286-2960 x1803; fax: +90-212-285-9453.
E-mail address: bayazit@isikun.edu.tr (U. Bayazit).

PII: S0923-5965(02)00083-8

luminance and two chrominance coordinates (*YUV*, *YIQ* or *YCrCb*) for each *RGB* coordinate triple, are commonly used. In [14], independent coding of C.I.E. $l \times a \times b$ color coordinates is shown to yield higher coding performance than independent coding of *RGB* color coordinates.

It has been observed that the luminance and chrominance coordinates are practically uncorrelated, but not statistically independent [5]. This is especially true for large transition regions of the image such as the edges, where the magnitude of the luminance and the chrominance coordinates are correlated. Therefore, independent coding of the luminance and chrominance (e.g. *YUV*) planes followed by the application of the standard, fixed filter of [13, p. 67] having the transformation matrix

$$G = \begin{bmatrix} 1.000 & 0.000 & 1.400 \\ 1.000 & -0.395 & -0.581 \\ 1.000 & 2.032 & 0.000 \end{bmatrix} \qquad (1)$$

for converting the quantized luminance/chrominance (*YUV*) coordinates back to the *RGB* coordinates yields suboptimal reconstruction since such an approach cannot exploit the statistical dependencies of spatial detail among the luminance and chrominance planes. The support of the standard, fixed filter, G, does not cover the luminance and chrominance coordinates of the neighbors of the filtered pixel. On the other hand, the quantization noise of the high performance transform based grayscale coding methods of [9,10] appears predominantly at the high frequency edges and one can considerably reduce it if these dependencies are exploited. In [11,12], the spatial orientation tree of [10] is extended to link the chrominance and luminance planes as well as the frequency bands towards this end.

As an alternative to the joint coding of the luminance and chrominance planes, the application of a linear filter with a support that extends spatially as well as across the luminance and chrominance coordinate planes, to transform the (possibly independently) quantized luminance and chrominance coordinates back to the *RGB* coordinates, is proposed in this paper. Such a filter can exploit the statistical dependencies among the (quantized) luminance and chrominance coordinate planes. Given a region of support for the filter in the luminance and chrominance planes, the *optimal* adaptive linear filter should be designed to yield a reconstruction error for the *RGB* coordinates that is statistically orthogonal to the luminance and chrominance values inside the support. The approach here is similar to the approach outlined in [6] for the design of custom subband synthesis filter banks.

The optimal adaptive linear filter coefficients are determined at the encoder following the compression and subsequent decompression[1] of the luminance and chrominance coordinate planes and are transmitted to the decoder as side information. The decoder can apply the filter to the decompressed luminance/chrominance planes to reconstruct the *RGB* planes. The set of coefficients for a region of support of reasonable size need not incur substantial overhead if the most significant bits in the representation of the coefficients are entropy coded.

As depicted in Fig. 1, the optimal adaptive linear filter for luminance/chrominance to *RGB* coordinate transformation can be used in a postprocessor at the back-end of existing color image decompression systems. At the decoder, the postprocessor transforms the reconstructed *RGB* planes by means of the standard, fixed filter of [13, p. 66] to the luminance/chrominance planes and transforms the luminance/chrominance planes by means of the optimal adaptive linear filter to get the improved versions of the reconstructed *RGB* planes. The postprocessor does not need to have any knowledge about the color image compression and decompression systems in use.

The organization of this paper is as follows. In Section 2, the optimal adaptive linear filter for luminance/chrominance to *RGB* coordinate transformation is described in detail. Section 3 outlines the lossy coefficient compression technique based on adaptive arithmetic coding which substantially reduces the overhead rate requirement for the transmission of the filter coefficients. Experimental results are provided in Section 4 followed by

[1] We assume the integration of the decompression function into the image encoder. This resembles scalable coding since the enhancement layer can only be coded after the base layer is decompressed and subtracted from original image.

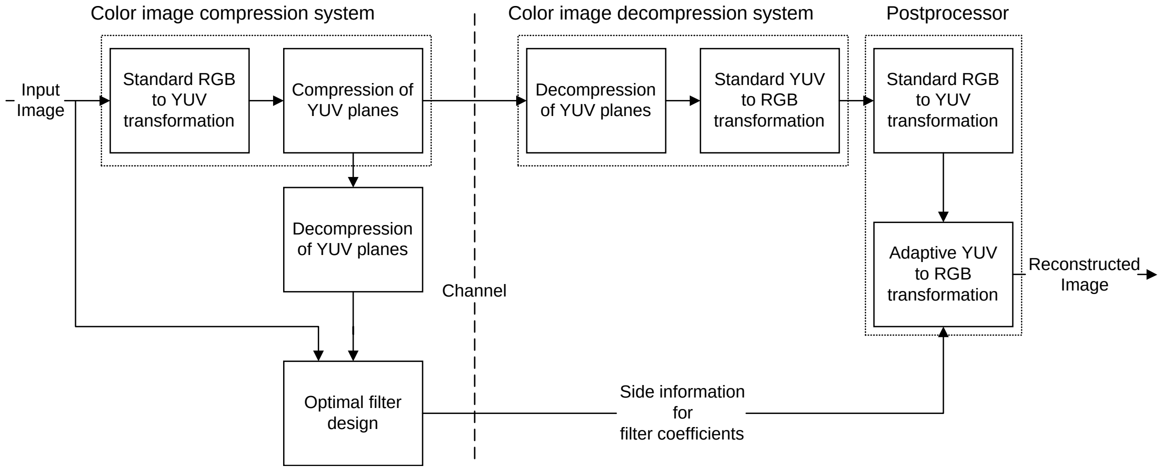

Fig. 1. *YUV* to *RGB* transformation by image adaptive linear filtering can be implemented in a postprocessor at the back end of a generic color image compression system. The filter coefficients that are determined at the encoder need to be conveyed to the decoder as side information.

concluding remarks in Section 5. Appendix A provides an analysis of the estimation gain with the optimal adaptive linear filter and shows that a positive gain can be expected only when there is correlation between adjacent pixels.

2. Optimal filtering of luminance and chrominance coordinates for reconstruction of RGB coordinates

Let the original *RGB* coordinate planes and the reconstructed *YUV* luminance and chrominance coordinate planes of a color image be given. Without loss of generality to other coordinate systems such as *YCrCb*, *YIQ* or *KLT*, we shall be using the *YUV* coordinate system throughout the remainder of this paper unless otherwise mentioned. The original value of the *RGB* coordinate $\zeta \in \{R, G, B\}$, and the decoded value of the *YUV* coordinate $\gamma \in \{Y, U, V\}$ at pixel location (i,j) are denoted by $x_\zeta(i,j)$ and $r_\gamma(i,j)$, respectively. The optimal linear filter g_ζ^*, used to reconstruct the coordinate plane $\zeta \in \{R, G, B\}$, should satisfy

$$g_\zeta^* = \arg\min_{g_\zeta} \sum_{i,j} \|x_\zeta(i,j) - g_\zeta^{\mathrm{T}} u(i,j)\|^2, \tag{2}$$

where $u(i,j)$ is the $3(2L+1)(2K+1)$ dimensional vector of decoded luminance and chrominance values of pixels in the $(2L+1)\times(2K+1)$ spatial support centered at location (i,j). Specifically, $u(i,j) = [u_Y^{\mathrm{T}}(i,j) u_U^{\mathrm{T}}(i,j) u_V^{\mathrm{T}}(i,j)]^{\mathrm{T}}$.

In the above, each $(2L+1)(2K+1)$ dimensional subvector $u_\gamma^{\mathrm{T}}(i,j)$, $\gamma \in \{Y, U, V\}$ is defined as a stack of columns $\{v_\gamma^{\mathrm{T}}(i,j-l)\colon l = -L, -L+1, \ldots, L-1, L\}$ so that

$$u_\gamma^{\mathrm{T}}(i,j) = [v_\gamma^{\mathrm{T}}(i,j-L) v_\gamma^{\mathrm{T}}(i,j-L+1) \ldots \\ v_\gamma^{\mathrm{T}}(i,j+L-1) v_\gamma^{\mathrm{T}}(i,j+L)]^{\mathrm{T}},$$

where

$$v_\gamma^{\mathrm{T}}(m,n) = [r_\gamma(m-K,n) r_\gamma(m-K+1,n) \ldots \\ r_\gamma(m+K-1,n) r_\gamma(m+K,n)]^{\mathrm{T}}$$

is a column of decoded values of the *YUV* coordinate γ centered at location (m,n).

For each *RGB* coordinate ζ, the minimization in Eq. (2) is achieved by requiring the error $x_\zeta(i,j) - g_\zeta^{\mathrm{T}} u(i,j)$ to be orthogonal to $u(i,j)$ in the statistical sense. This yields $3(2L+1)(2K+1)$ regression equations in $3(2L+1)(2K+1)$ unknowns expressed succinctly as

$$R_{uu} g_\zeta^* = r_{x_\zeta u}, \tag{3}$$

where, for an image of height N pixels and width M pixels,

$$R_{uu} = \frac{1}{NM} \sum_{i,j} u(i,j) u^{\mathrm{T}}(i,j)$$

is the $3(2L+1)(2K+1) \times 3(2L+1)(2K+1)$ correlation matrix estimate and

$$r_{x_\zeta u} = \frac{1}{NM} \sum_{i,j} x_\zeta(i,j) u(i,j)$$

is the $3(2L+1)(2K+1)$ dimensional cross-correlation vector estimate.

3. Lossy compression of filter coefficients

The filter coefficients determined by solving Eq. (3) are conveyed to the decoder as side information. Simple uniform scalar quantization of the coefficients allows adequate representation of the most significant bits of the coefficients to yield a desirably low overhead rate. Judicious arithmetic coding of the most significant bits in the representation can further cut down the overhead rate.

For each color plane ζ, the number of filter coefficients is $3(2L+1)(2K+1)$. Assuming 4 byte floating point precision, the overhead rate requirement for an image with height $N = 512$, width $M = 512$, filter support height $2L+1 = 3$, and filter support width $2K+1 = 3$ is 0.0099 bits/pixel. Assuming the largest coefficient magnitude to be bounded by 2, if the magnitude of each coefficient is represented by uniform quantization to 11 (most significant) bits, and transmitted along with one sign bit, the overhead rate is brought down to 0.00372 bits/pixel.

The most significant bits of the coefficient magnitude up to the leading "1" in the representation may be adaptive arithmetic coded since the distribution for these bits is biased towards the "0" bit. However, the magnitudes of the coefficients at the centers of spatial supports are observed to be large and the leading 1 occurring in the most significant bits has a greater likelihood for these coefficients. Therefore, these coefficients are coded with a different probability model than the one used to code the other coefficients. The bits following the leading 1 in the representation are not arithmetic coded since the distribution for these bits is not necessarily biased towards the 0 bit.

4. Experimental setup and results

The *YUV* to *RGB* coordinate transformation by optimal adaptive linear filtering can be implemented in conjunction with any compression/decompression engine whose internal mechanisms are unknown. In our simulations, we have utilized the wavelet transform based SPIHT (Set Partitioning in Hierarchical Trees) grayscale image coding method of [9] to compress and decompress the *YUV* coordinate planes. SPIHT is not only competitive with the latest high performance image coding methods, but also the bitstreams generated by SPIHT are embedded and the resulting size of the compressed grayscale image file can be scaled to any desired precision. This facilitates rate-distortion optimal rate allocation to *YUV* coordinate planes by a method such as [8], followed by independent coding of each of the *YUV* coordinate planes with SPIHT.

Specifically, a bitplane of wavelet coefficients of a luminance/chrominance plane is coded by SPIHT only if the ratio of the decrease in distortion to the increase in rate is more than a threshold. This strategy closely parallels the BFOS rate allocation strategy of [8]. We take the granularity of allocation to be a bitplane for simplicity of implementation. The performance of the proposed optimal linear filtering method does not depend on the rate allocation strategy used.

In our experiments, we have used the 512×512 color still images Lenna, Peppers and Sailboat taken from the USC image database. The *RGB* planes were converted to the *YUV* planes by the standard, fixed filter of [13, p. 66] and the *YUV* planes were separately coded by SPIHT. Table 1 reports the number of bitplanes coded by the rule stated in the previous paragraph and the rate allocated to each of the luminance and chrominance planes. For each image, the rate-distortion threshold used for rate allocation is also tabulated in the first column. The reconstructed *YUV* planes are transformed to *RGB* planes by the application of the optimal adaptive *YUV* to *RGB* transformation filter and the application of the standard, fixed *YUV* to *RGB* transformation filter of Eq. (1). We have also considered the performance of the method of independent application of an

Table 1
Number of bitplanes coded and rate allocated to each YUV component

		Y	U	V
Lenna $\lambda = 25$	No. levels	10	8	8
	Rate	0.71014	0.18854	0.32660
Peppers $\lambda = 25$	No. levels	10	8	10
	Rate	0.86270	0.27945	0.66116
Sailboat $\lambda = 50$	No. levels	10	8	9
	Rate	1.48477	0.51282	1.14926

Table 2
PSNR of color planes reconstructed by the application of the standard, fixed, optimal adaptive interplane, and adaptive intraplane filters to decompressed YUV planes

	Lenna			Peppers			Sailboat		
	R	G	B	R	G	B	R	G	B
Standard	37.6507	36.4945	33.4416	34.3874	35.9378	33.1241	34.3996	35.4310	32.0051
Adaptive interplane $K = L = 0$ (compressed coef.)	37.6516	36.4926	33.4254	34.3873	35.9397	33.1181	34.4091	35.4320	32.0116
Adaptive interplane $K = L = 1$ (compressed coef.)	38.0251	36.5568	33.7123	34.4456	35.9685	33.4033	34.6727	35.4897	32.2351
Adaptive interplane $K = L = 1$ (floating precision)	38.0430	36.5819	33.7119	34.4475	35.9687	33.4072	34.6757	35.4970	32.2397
Adaptive interplane $K = L = 2$ (compressed coef.)	38.1080	36.6532	33.8265	34.6381	36.0310	33.4649	34.7826	35.5780	32.3322
Adaptive intraplane $K = L = 1$ (compressed coef.)	37.7312	36.4860	33.4468	34.3905	35.9039	33.1774	34.4682	35.3993	32.0610

intraplane adaptive spatial filter to each of the reconstructed YUV planes prior to the application of the standard, fixed YUV to RGB transformation filter. The intraplane adaptive spatial filter for each YUV plane is similarly designed by solving a set of regression equations involving only the original and reconstructed values of that YUV plane.

Table 2 reports the PSNRs of the RGB planes reconstructed with the optimal adaptive linear filter, those reconstructed with the standard, fixed filter, as well as those reconstructed with the combination of the intraplane adaptive spatial filter and the standard, fixed filter. Three sets of filter support dimensions were considered for the optimal adaptive linear filter: $K = L = 0$, $K = L = 1$, $K = L = 2$. For $K = L = 0$ and $K = L = 2$, we only report the PSNRs obtained by employing filter coefficients which are lossy compressed by representing each coefficient by one sign bit and

11 most significant bits of the magnitude where the most significant bits up to the leading 1 were arithmetic coded. For $K = L = 1$, we also report the PSNRs obtained by employing filter coefficients at full, 32 bit floating precision.

The coefficients of the optimal adaptive *YUV* to *RGB* transformation filter with filter support dimensions $K = L = 0$ turns out to be almost identical to the standard, fixed *YUV* to *RGB* transformation filter. This is manifested by the PSNR results for the reconstructed *RGB* planes reported in the third and fourth rows of Table 2. Likewise, the adaptive intraplane filters applied independently to the *YUV* planes do not yield notable performance gain. On the other hand, the PSNR gain of the optimal adaptive linear filter is modest for $K = L = 1$. This gain comes at the expense of an increase in computational complexity for the accumulation of the statistics, $r_{x_\xi u}$ and R_{uu}, and an increase in the overhead rate for the transmission of the filter coefficients. Further PSNR gain, with $K = L = 2$ as the support dimensions of the optimal adaptive linear filter, is accompanied by nearly 8 fold increase in computational complexity with respect to the $K = L = 1$ case due to accumulation of the statistics, and further increase in the overhead rate. Finally, it is observed that filter coefficient compression yields negligible loss in PSNR gain.

Table 3 reports the overhead rates that correspond to the entries of Table 2. It is observed that arithmetic coding works better for compressing a large number of coefficients such as in the $K = L = 2$ case for the optimal adaptive linear filter. This is desirable, since the full precision overhead rate required for transmitting filters with a large number of coefficients is considerable.

The ratio of decrease in mean squared error to the overhead rate incurred due to the application of the optimal adaptive linear filter is far larger than λ, the rate-distortion threshold used in rate allocation during compression of the coordinate planes. For instance, this ratio is 962.94 for the $K = L = 1$ optimal adaptive linear filter with lossy compressed coefficients applied to the *YUV* coordinate planes of Lenna image which were allocated rate with threshold $\lambda = 25$.

In Fig. 2, we show a cropped section of the original 24 bit color Lenna image and its reconstructions after the compression and decompression of the *YUV* planes with SPIHT, and the subsequent *YUV* to *RGB* transformation with the standard, fixed filter and the optimal adaptive linear filter ($K = L = 2$). Since transform based high performance image coding methods like SPIHT coarsely quantize the high frequency coefficients, the quantization noise manifests itself at high frequency details like the edges. Therefore, the optimal adaptive linear filter, that exploits dependencies of spatial detail among luminance and chrominance planes, can yield several tenths of a dB gain in the reconstruction of the *RGB* planes over the standard, fixed filter for a modest overhead rate. In Fig. 3, we show the difference plane between the reconstruction error amplitude planes obtained with the optimal adaptive linear filter and the reconstruction error amplitude plane obtained with the standard fixed filter for each of the *RGB* planes of Lenna image. For the most part, the optimal adaptive linear filter yields better reconstruction than the standard fixed filter at the edges. We note, however, that the independent application of adaptive intraplane filters (not shown in this figure) to the *YUV* coordinate

Table 3
Overhead rate for the transmission of filter coefficients for the *YUV* to *RGB* transformation methods

	Lenna	Peppers	Sailboat
Standard	0	0	0
Adaptive interplane $K = L = 0$	0.0005379	0.0005341	0.0005379
Adaptive interplane $K = L = 1$ (compressed coef.)	0.003021	0.002781	0.002758
Adaptive interplane $K = L = 1$ (floating precision)	0.009888	0.009888	0.009888
Adaptive interplane $K = L = 2$	0.006939	0.006916	0.006569
Adaptive intraplane $K = L = 1$	0.001022	0.0009918	0.001003

Fig. 2. (a) Original cropped section of Lenna. (b) Cropped section reconstructed with standard *YUV* to *RGB* transformation filter. (c) Cropped section reconstructed with optimal adaptive *YUV* to *RGB* transformation filter ($K = L = 2$). The quantization noise that manifests itself at the high frequency details such as the edges is seen to be reduced by the optimal adaptive linear filter. The *Y*, *U* and *V* coordinate planes of the coded picture were allocated rates of 0.16992, 0.02539 and 0.03906 bpp, respectively. The *R*, *G* and *B* component planes were reconstructed at PSNRs of 32.072, 30.721 and 29.306 dB, respectively, with the standard fixed filter, and at PSNRs of 32.459, 30.966 and 29.756 dB, respectively, with the optimal adaptive linear filter. The overhead rate for the compressed coefficients of the optimal adaptive linear filter was 0.0097 bpp.

Table 4
Number of bitplanes coded and rate allocated to each *KLT* component

	Component	$KL-1$	$KL-2$	$KL-3$
Lenna $\lambda = 100$	No. levels	6	8	9
	Rate	0.09436	0.25110	0.55264
Peppers $\lambda = 100$	No. levels	8	9	9
	Rate	0.48312	0.28900	0.446900
Sailboat $\lambda = 100$	No. levels	8	8	10
	Rate	0.46674	0.47385	1.19843

planes cannot exploit the interplane statistical dependencies, and therefore, does not yield a marked reconstruction quality advantage over the standard fixed filter.

In Tables 4–6, results similar to those reported in Tables 1–3, are tabulated for the case where the coordinate system used to decorrelate the color planes is *KLT* instead of the *YUV*. The standard way to use the *KLT* coordinate system is to design the filter for RGB-KLT transform at the encoder and compress and transmit the filter coefficients to the decoder at a nonzero overhead rate. The symmetric *KLT–RGB* transform is then performed using the same coefficients at the decoder. Even though this method of reconstructing *RGB* components is image adaptive it is suboptimal.

Our proposed optimal adaptive linear filter design method improves the performance for *KLT* at the expense of some computational complexity and negligible overhead rate increase. In this case, the encoder determines the *KLT–RGB* filter coefficients for a support that extends spatially as well as across the *KLT* component planes in order to minimize the mean squared reconstruction error. These coefficients are compressed and transmitted instead of the coefficients of the symmetric filter.

5. Conclusions

In this paper, we have first introduced an optimal adaptive linear filter design method based on the statistics estimated from lossy compressed luminance and chrominance coordinate planes and the original color coordinates planes for converting the reconstructed luminance and chrominance coordinates to reconstructed *RGB* coordinates. Given the filter support, the adaptive linear filter is optimal in the minimum mean squared error sense. The method can be used in a postprocessor at the back end of color image compression systems for reducing the distortion of the lossy compressed image. The image adaptive design of the filter necessitates the transmission of the coefficients to the decoder as side information. We have also

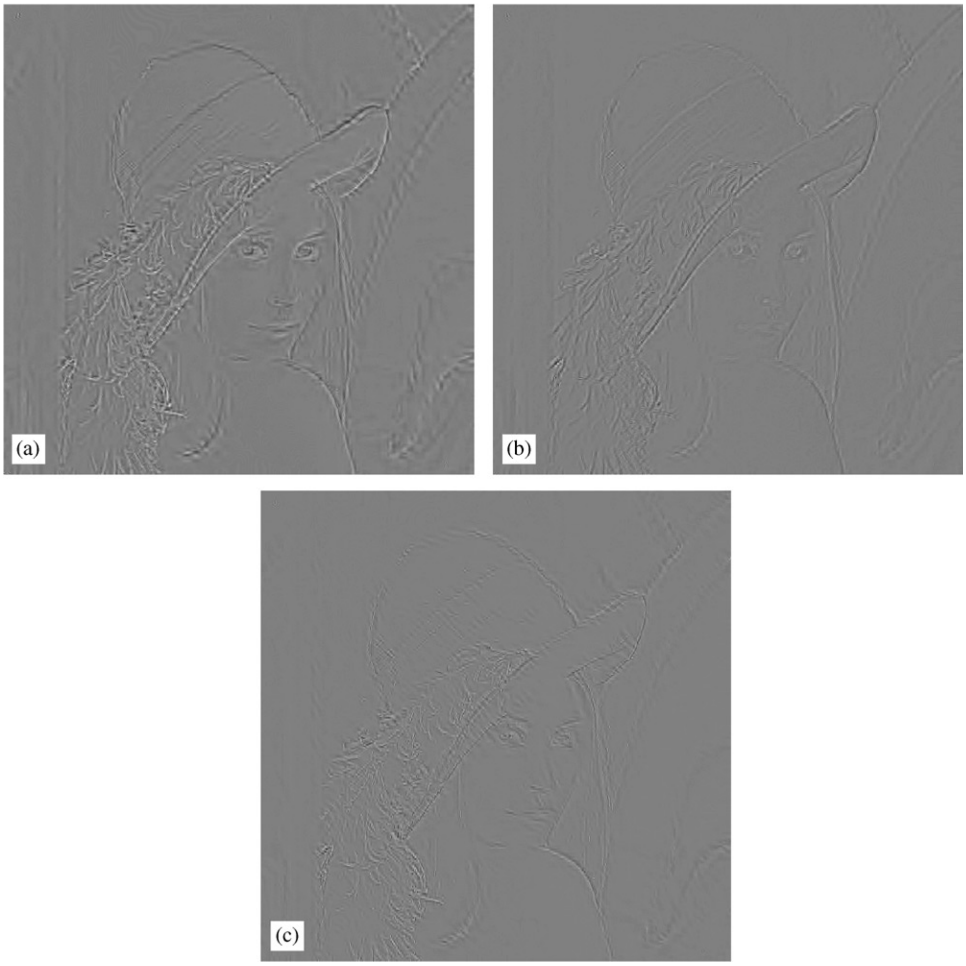

Fig. 3. Grayscale difference planes between the reconstruction error amplitude planes obtained with the optimal adaptive linear filter and the corresponding reconstruction error amplitude planes obtained with the standard, fixed filter ((a) blue, (b) green, (c) red). Lighter areas indicate smaller and darker areas indicate larger reconstruction error amplitude obtained with optimal adaptive linear filter than with standard fixed filter. (All difference values have been scaled by multiplication with a factor of 3.)

outlined a simple yet effective strategy for compressing the filter coefficients by arithmetic coding in order to keep the overhead rate low.

Experimentally, the ratio of the objective reconstruction quality gain obtained by the application of the optimal adaptive linear filter in place of the standard, fixed filter in YUV to RGB transformation to the overhead rate expended turned out to be significantly larger than the marginal return of further (independently) coding the YUV planes. This makes the optimal adaptive linear filtering method a viable alternative to any coding approach that exploits the statistical dependencies of spatial detail among color coordinate planes.

Appendix A

The optimal adaptive linear filter has a filter support that is spatially centered at the location of the target pixel whose luminance and chrominance values are estimated and spans all coordinate planes. The filter support covers the decompressed

Table 5
PSNR of color planes reconstructed by the application of the standard, fixed, optimal adaptive interplane, and adaptive intraplane filters to decompressed KLT planes

	Lenna			Peppers			Sailboat		
	R	G	B	R	G	B	R	G	B
Standard (symmetric)	36.7165	35.1167	33.2964	32.8115	33.7747	33.5703	32.3889	32.5100	32.4323
Adaptive interplane $K=L=0$	36.7210	35.1190	33.3042	32.8122	33.7826	33.5790	32.3936	32.5232	32.4426
Adaptive interplane $K=L=1$ (compressed coef.)	37.0434	35.2869	33.5475	32.8278	33.8688	33.6691	32.5181	32.6282	32.5493
Adaptive interplane $K=L=1$ (floating precision)	37.0674	35.2873	33.5472	32.8285	33.8726	33.6690	32.5242	32.6290	32.5494
Adaptive interplane $K=L=2$	37.1672	35.3831	33.6125	32.9268	33.9366	33.7415	32.6092	32.7849	32.6246
Adaptive intraplane $K=L=1$	36.8279	35.1395	33.2705	32.8067	33.7834	33.5993	32.4359	32.5440	32.4316

Table 6
Overhead rate for the transmission of filter coefficients for the KLT to RGB transformation methods

	Lenna	Peppers	Sailboat
Standard	0.0004959	0.0005035	0.0004959
Adaptive interplane $K=L=0$	0.0004959	0.0005035	0.0004959
Adaptive interplane $K=L=1$ (compressed coef.)	0.002846	0.002621	0.002583
Adaptive interplane $K=L=1$ (floating precision)	0.009888	0.009888	0.009888
Adaptive interplane $K=L=2$	0.006649	0.006523	0.006424
Adaptive intraplane $K=L=1$	0.001022	0.001011	0.001019

luminance and chrominance coordinates of the target pixel as well as its neighbors. In this appendix, we show that the optimal adaptive linear filter yields PSNR gain with respect to the standard, fixed filter. This gain is due to the correlations between the luminance and chrominance coordinates of the target pixel and the luminance and chrominance coordinates of its neighbors. Since the correlation coefficient between the luminance (or corresponding chrominance) coordinates of two contiguous pixels is rather high ($\rho \approx 0.9$) for most natural images, excluding the neighbors' luminance and chrominance coordinates from the filter's support is not desirable.

To facilitate our analysis here, we decompose the optimal adaptive linear filter into two parts (i.e. the overall transformation matrix is the product of two matrices). The first part reconstructs improved luminance and chrominance coordinates from the decompressed values of the luminance and chrominance coordinates. The second part is the standard, fixed filter of Eq. (1) that converts the improved luminance and chrominance coordinates to the final reconstruction values for the RGB coordinates. Hence, the first part may be obtained by post-multiplying the transformation matrix of the optimal adaptive linear filter by the inverse of the transformation matrix of the standard, fixed filter. We analyze the mean squared error in the luminance/chrominance coordinate space yielded by the first part which is identical to the mean squared error in the *RGB* coordinate space yielded by the optimal adaptive linear filter due to the orthogonality of the standard, fixed filter.

Let X_γ and $\hat{X}_\gamma$ be the original luminance or chrominance coordinate of the target pixel, and its best linear estimate, respectively. Without loss of generality, we consider a single neighbor of the target pixel. Observation vector U is written as

$$U = (U_1^{\mathrm{T}} \vdots U_2^{\mathrm{T}})^{\mathrm{T}},$$

which consists of U_1, the vector of luminance and chrominance coordinates of the target pixel contaminated with additive quantization noise and U_2, the vector of luminance and chrominance coordinates of the target pixel's neighbor contaminated with additive quantization noise. The noise is assumed to be white and independent of the luminance and chrominance coordinates,[2] and the data and noise are assumed to have zero mean. The estimation error can then be expressed as

$$\varepsilon = E[(X_\gamma - \hat{X}_\gamma)X_\gamma] = r_{X_\gamma X_\gamma} - r_{X_\gamma U} r_{UU}^{-1} r_{UX_\gamma},$$

where $r_{X_\gamma X_\gamma}$ is the variance of X_γ, $r_{X_\gamma U}$ is the vector of correlation values between X_γ and each component of U, and r_{UU}^{-1} is the inverse of the correlation matrix for U. The vector $r_{X_\gamma U}$ can be partitioned as

$$r_{X_\gamma U} = r_{UX_\gamma}^{\mathrm{T}} = (r_{X_\gamma U_1}^{\mathrm{T}} \vdots r_{X_\gamma U_2}^{\mathrm{T}})^{\mathrm{T}}.$$

Similarly, r_{UU}^{-1} can be partitioned as

$$r_{UU}^{-1} = \begin{bmatrix} R_1 \vdots R_2 \\ \cdots\cdots \\ R_3 \vdots R_4 \end{bmatrix}.$$

If

$$r_{UU} = \begin{bmatrix} r_{U_1U_1} & \vdots & r_{U_1U_2} \\ \cdots & & \cdots \\ r_{U_2U_1} & \vdots & r_{U_2U_2} \end{bmatrix},$$

with $r_{U_iU_j}$ defined as the cross-correlation matrix of vectors U_i and U_j, then the blocks of r_{UU}^{-1} may be related to the blocks of r_{UU} as

$$R_1 = r_{U_1U_1}^{-1} + r_{U_1U_1}^{-1} r_{U_1U_2} (r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1} \times r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1},$$

$$R_2 = -r_{U_1U_1}^{-1} r_{U_1U_2} (r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1},$$

$$R_3 = -(r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1} r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1},$$

$$R_4 = (r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1},$$

where $r_{U_2U_1} = r_{U_1U_2}^{\mathrm{T}}$ for r_{UU} to be a real symmetric correlation matrix. The expected estimation error can then be expressed as

$$\begin{aligned} \varepsilon &= r_{X_\gamma X_\gamma} - r_{X_\gamma U} r_{UU}^{-1} r_{UX_\gamma} \\ &= r_{X_\gamma X_\gamma} - r_{X_\gamma U_1}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{X_\gamma U_1} \\ &\quad - (r_{X_\gamma U_2}^{\mathrm{T}} - r_{X_\gamma U_1}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2}) \\ &\quad \times (r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1} \\ &\quad \times (r_{X_\gamma U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{X_\gamma U_1}). \end{aligned}$$

In the above equation, the term $(r_{X_\gamma U_2}^{\mathrm{T}} - r_{X_\gamma U_1}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})(r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1}(r_{X_\gamma U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{X_\gamma U_1})$ is due to the inclusion of the contaminated coordinates of the neighbors into the observation vector U and therefore corresponds to the gain of the optimal adaptive linear filter with $K = L > 0$ over the gain with $K = L = 0$. Since $r_{U_1U_1}^{-1}$, $r_{U_2U_2}$ are real symmetric matrices, $(r_{U_2U_2} - r_{U_1U_2}^{\mathrm{T}} r_{U_1U_1}^{-1} r_{U_1U_2})^{-1}$ is also real symmetric. Since the above term is a quadratic form, the gain must be greater than or equal to zero. Therefore, we conclude that incorporating the neighbor's contaminated coordinates into the observation vector U can help reduce the estimation error.

To gain further insight, let us assume that the data is wide sense stationary and that the ratio of the correlation between a coordinate of the target pixel and a coordinate of its neighbor to the correlation between corresponding coordinates of the target pixel is given by ρ. One can then express $r_{X_\gamma U}$ and r_{UU} as

$$r_{X_\gamma U} = r_{UX_\gamma}^{\mathrm{T}} = (r_{X_\gamma U_1}^{\mathrm{T}} \vdots r_{X_\gamma U_2}^{\mathrm{T}})^{\mathrm{T}} = (r_{X_\gamma U_1}^{\mathrm{T}} \vdots \rho r_{X_\gamma U_1}^{\mathrm{T}})^{\mathrm{T}},$$

$$r_{UU} = \begin{bmatrix} r_{U_1U_1} & \vdots & r_{U_1U_2} \\ \cdots & & \cdots \\ r_{U_2U_1} & \vdots & r_{U_2U_2} \end{bmatrix} = \begin{bmatrix} r_{ZZ} & \vdots & \rho r_{ZZ} \\ \cdots & & \cdots \\ \rho r_{ZZ} & \vdots & r_{ZZ} \end{bmatrix} + \mathrm{Diag}[\sigma_{N_Y}^2, \sigma_{N_U}^2, \sigma_{N_V}^2, \sigma_{N_Y}^2, \sigma_{N_U}^2, \sigma_{N_V}^2],$$

[2] In reality, the quantization noise is signal dependent.

where r_{ZZ} is the correlation matrix of luminance and chrominance coordinates before contamination with quantization noise and $\sigma^2_{N_\gamma}$ is the variance of the quantization noise added to luminance/chrominance coordinate γ. If there is no quantization noise, then $r_{U_1U_2} = \rho r_{ZZ}, r_{U_1U_1} = r_{U_2U_2} = r_{ZZ}$, hence $r_{X_\gamma U_2} - r^{\mathrm{T}}_{U_1U_2} r^{-1}_{U_1U_1} r_{X_\gamma U_1} = 0$, but $r_{U_2U_2} - r^{\mathrm{T}}_{U_1U_2} r^{-1}_{U_1U_1} r_{U_1U_2} = (1-\rho^2) r_{ZZ}$ so that the estimation error

$$\varepsilon = r_{X_\gamma X_\gamma} - r^{\mathrm{T}}_{X_\gamma U_1} r^{-1}_{U_1U_1} r_{X_\gamma U_1}$$

is that achieved by the optimal adaptive linear filter with $K = L = 0$ (and also the standard, fixed filter of Eq. (1)). Likewise, when $\rho = 0$, but noise is present, we have $r_{U_1U_2} = 0$, $r_{X_\gamma U_2} = 0$, hence $r_{X_\gamma U_2} - r^{\mathrm{T}}_{U_1U_2} r^{-1}_{U_1U_1} r_{X_\gamma U_1} = 0$, $r_{U_2U_2} - r^{\mathrm{T}}_{U_1U_2} \times r^{-1}_{U_1U_1} r_{U_1U_2} = r_{ZZ}$ so that the estimation error is again that achieved by the standard, fixed filter. On the other hand, when there is luminance/chrominance coordinate quantization and the correlation between contiguous pixels' coordinates is nonzero, optimal adaptive linear filter can reduce estimation error.

References

[1] G. Braudaway, A procedure for optimum choice of a small number of colors from a large color palette for color imaging, in: Electronic Imaging '87, San Francisco, CA, 1987.

[2] P. Heckbert, Color image quantization for frame buffer display, Comput. Graphics 16 (3) (July 1982) 297–307.

[3] S. Mitra, S. Pemmaraju, S. Kompella, S. Meadows, Efficient color image compression using integrated fuzzy neural networks for vector quantization, in: Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 1997. Comput. Cybernet. Simulation 1 (October 1997) 184–188.

[4] D. Mukherjee, S.K. Mitra, Arithmetic coded vector SPIHT with classified tree-multistage VQ for color image coding, in: Proceedings of the IEEE Second Workshop on Multimedia Signal Processing, 1998, pp. 444–449.

[5] A.N. Netravali, C.B. Rubinstein, Luminance adaptive coding of chrominance signals, IEEE Trans. Commun. 27 (4) (April 1979) 703–710.

[6] A. Nosratinia, Optimal subband synthesis, Proceedings of the IEEE International Conference on Image Processing, Santa Barbara, CA, October 1997, pp. 641–644.

[7] M.T. Orchard, C.A. Bouman, Color quantization of images, IEEE Trans. Signal Process. 39 (12) (December 1991) 2677–2690.

[8] E.A. Riskin, Optimum bit allocation via the generalized BFOS algorithm, IEEE Trans. Inform. Theory 37 (March 1991) 400–402.

[9] A. Said, W.A. Pearlman, A new, fast, and efficient image codec based on set partitioning in hierarchical trees, IEEE Trans. Circuits Systems Video Technol. 6 (3) (June 1996) 243–250.

[10] J.M. Shapiro, Embedded image coding using zerotrees of wavelet coefficients, IEEE Trans. Signal Process. 41 (December 1993) 3445–3462.

[11] K. Shen, E.J. Delp, Color image compression using an embedded rate scalable approach, in: Proceedings of the IEEE International Conference on Image Processing, Santa Barbara, CA, October 1997, pp. 26–29.

[12] K. Shen, E.J. Delp, Wavelet based rate scalable video compression, IEEE Trans. Circuits Systems Video Technol. 9 (1) (February 1999) 109–122.

[13] William K. Pratt, Digital Image Processing, Wiley-Interscience, New York, 1991, pp. 66–67.

[14] G. Zamora, S.K. Mitra, Lossless coding of color images using color space transformations, Proceedings 11th IEEE Symposium on Computer-Based Medical Systems, June 1998, pp. 13–18.