Multiresolution Color Correction

Raja Balasubramanian, Ricardo de Queiroz, Zhigang Fan
Digital Imaging Technology Center
Xerox Corporation, Webster NY 14580
raja,queiroz,fan@wrc.xerox.com

ABSTRACT

In this paper, a color correction system is embedded into a multiresolution representation with the goal of reducing the complexity of 3D look-up table transformations. A framework is assumed wherein the image undergoes a multiresolution decomposition, *e.g.* discrete wavelet transform, for the purpose of image compression or other processing. After the image is reconstructed from its multiresolution representation, color correction is usually required for rendering to a specific device. The color correction process is divided into two phases: a complex multidimensional transform (Phase 1), and a series of essentially 1-D transforms (Phase 2). Phase 1 correction is then moved within the multiresolution reconstruction process in such a way that a small subset of the image samples undergoes the multidimensional correction. Phase 2 correction is then applied to all image samples after the image is reconstructed to its full resolution. The recently proposed spatial CIELAB model is used to evaluate the algorithm. The computational cost incurred by the color correction is considerably reduced, with little loss in image quality.

1. INTRODUCTION

Digital imaging systems typically comprise a sequence of image processing and analysis steps between image acquisition and final output rendering. Multiresolution analysis is one example of a powerful image processing tool used in many applications, including segmentation, filtering, halftone descreening, and image compression. The general idea behind multiresolution processing is that 1) the image is decomposed into a family of representations at different spatial resolutions; 2) application dependent processing is carried out on some subset of these representations; and 3) the output image is synthesized from the processed representations. This work addresses a scenario wherein a color image undergoes multi-resolution processing for a given application (*e.g* compression or descreening) , and is then ultimately rendered on some output color device such as a display or printer. The latter requires that the image is color corrected to compensate for the colorimetric characteristics of the output device. In conventional systems, color correction is usually performed at the final stage, and is applied on a pixel-by-pixel basis to the entire full resolution image.

In [1], a technique was proposed for combining JPEG decompression [2] and color correction to reduce the computational cost. A complex multidimensional color correction is applied to the image corresponding to only the lower order discrete cosine transform (DCT) coefficients, while a much simpler correction is applied to the image corresponding to the higher order terms. The approach was motivated by the fact that the human visual system (HVS) is less sensitive to color errors at high spatial frequencies.

In this paper, a technique for combining multiresolution analysis and color correction is proposed with a similar motivation. To illustrate the idea, a particular example is considered, where a discrete wavelet transform (DWT) [3] is employed for multiresolution analysis, and a 4 colorant (CMYK) printer is used as the output device. Examples of applications utilizing wavelet transforms are compression [4],[5] and descreening or inverse halftoning [6]. The prime advantage afforded by the proposed technique is a reduction in computational cost required by the color correction. Another potential advantage is to reduce the effect of high frequency noise in the color correction. This technique differs from [1] in that it is not restricted to the JPEG framework, but addresses a more general architecture where multiple resolutions of an image (equivalently a frequency band decomposition of the image) are available at some point in the image processing path.

As an example of an application for the proposed method, consider a color copier, where the original to be copied is a halftone pictorial document. In this case, descreening may be an important step, especially if subsequent analysis and processing modules require continuous tone data. Most descreening algorithms generate a lowpass (and hence low resolution) image as an intermediate step. In fact, [6] explicitly uses a DWT-based approach to solve this problem. The proposal in this paper is to utilize this multi-resolution image representation to perform accurate color correction on the reduced image, and then perform a partial correction to the full resolution image prior to printing.

This paper is organized as follows. In Sec. 2, an overview is presented on multi-resolution image analysis using wavelet decomposition. In Sec. 3, a similar overview is given for printer color correction. In Section 4, a technique is described for combining the wavelet transform and device color correction. Experimental results are presented in Section 5, and concluding remarks are given in Section 6.

2. MULTIRESOLUTION ANALYSIS USING THE WAVELET TRANSFORM

Numerous methods exist to decompose images into multiple resolutions. We use the wavelet transform [3], [4] as an example of a multiresolution technique, as this is a powerful and general framework that has shown considerable promise in several applications such as compression [3]-[5]. The wavelet transform, assumed to operate independently on each of the color separations, is applied with the use of a cascade of 2-channel filter banks. These filter banks are composed of decimators (*i.e*. downsampling operators) and low- and high-pass filters. At a given level, the image is split into 4 subbands (low- and high-pass in each direction). The wavelet transform is obtained by repeating, at each level, the downsampling and filtering process over the low-pass subband of the previous level. This is shown in a pictorial example in Fig. 1, for 2 levels. The image is then processed or analyzed in the wavelet domain for the given application, and is finally submitted to an inverse wavelet transform in order to reconstruct the spatial domain representation. Note that since each downsampling or upsampling operation is accompanied by a filtering operation, the multiresolution representation can also be conceptualized as a decomposition of the image into a series of frequency bands, *i.e*. a subband decomposition. Such a decomposition trades frequency resolution for spatial resolution. For our purposes, the DWT is assumed to already be already a part of the system for a given application such as compression (see Fig. 2); we simply take advantage of its availability for efficient color correction.

3. COLOR CORRECTION OF OUTPUT DEVICES

The technique proposed in this paper will, in principle, apply for any output device. The current discussion will however be restricted to the case of printers, since this involves the greatest complexity in color correction, and hence the greatest motivation for reducing the computational cost. In particular, the discussion will be confined to the frequently encountered case of 4 colorant (CMYK) printing, although it can be extended to other colorant sets.

Printer color correction for a CMYK printer is essentially a transformation from a 3-D device independent color space to the CMYK amounts required to reproduce the input color. The input space is often based on the CIE colorimetric standard [7]. For our purposes, the color correction may be divided into two phases. The first phase is a transform from the 3-D colorimetric space to CMY space, and takes into account unwanted absorptions and other complex color interactions in the printer colorants [8]. This transform usually requires the use of a 3-D lookup table and interpolation to calculate the ink amounts [9]. The second phase consists of simpler color transformations that are typically one-dimensional. A set of 1-D tone reproduction curves (TRC's) linearize the tone response of the individual printer colorants to a metric such as neutral luminance [9]. In addition, for CMYK printers, some type of undercolor removal (UCR) and gray component replacement strategy needs to be applied prior to the TRC's, in order to convert the three ink amounts to CMYK [8],[9]. A common approach is to use the minimum of C, M, Y to index into two 1-D functions, one for K addition, and one for CMY subtraction.

Qualitatively, one can think of Phase 2 as a first order 1-D luminance tone correction, while Phase 1 is a higher order multidimensional color correction. In terms of computational complexity, Phase 1 is more costly than the Phase 2, as it involves 3-D interpolation (refer to [1] for a detailed complexity analysis). In standard color management solutions, both phases are folded into a single 3-to-4 lookup table, requiring the same order of complexity as that of Phase 1. For the approach we propose, it is important that the two phases are kept separate.

4. COLOR CORRECTION IN THE WAVELET DOMAIN

4.1 Method overview

In an application that requires image processing in the wavelet domain, the standard approach is to first perform the inverse wavelet transform, and then apply the color correction in the spatial domain at the original image resolution, as shown in Fig. 3. It was noted earlier that the wavelet transform is a form of spatially localized frequency decomposition of the image. The proposed technique recognizes the opportunity offered by such a frequency decomposition; namely that the human visual system exhibits relatively low sensitivity to color errors at high spatial frequencies [10]. The idea therefore is to apply Phase 1 color correction to the low resolution images that correspond to low frequency bands, and apply only the simpler Phase 2 correction to the high resolution images that also include information from the high frequency bands. Since the low resolution images are subsampled versions of the original, the savings in computation arise from the fact that the costly multidimensional color transformation in Phase 1 correction is applied only to a subset of the image pixels.

The method just described is implemented by moving Phase 1 color correction inside the inverse wavelet (*i.e.* image reconstruction) transform as shown in Fig. 4. Part of the reconstruction is applied in the device independent color space. Phase 1 color correction is then applied to convert the subsampled image to a device color space, and the remaining wavelet reconstruction is applied in this device space. Finally, Phase 2 correction is applied to the completely reconstructed (*i.e*. full resolution) image.

Note that in order to combine Phase 1 color correction with the wavelet reconstruction in the manner described, the input color space to the Phase 1 transform must be of the same sense and orientation as the output color space. That is, if the input is colorimetric RGB, then the output must be printer RGB (which can be given by R=1-C, G=1-M, B=1-Y). Likewise, if the input is a colorimetric luminance-chrominance space, then the output printer coordinates must also be in a luminance-chrominance form [1]. The reason for this is that the frequency bands that have gone through Phase 1 color correction now have to be combined with the remaining frequency bands that have not been color corrected, and such combinations only make sense if the image data in all the bands are of the same color sense.

4.2 Computational cost vs image quality

Note that when performing N stages on the forward transform over an image of P x P pixels and only N-M stages in the inverse transform, the result is a subsampled image with dimensions P' x P', where $P=2^M P'$. Hence, Phase 1 color correction is applied to only a small fraction ($1/2^{2M}$) of the samples, thus drastically reducing the computational cost. If C_1 and C_2 represent the computational cost per pixel to perform Phase 1 and Phase 2, color correction respectively, the overall cost for the proposed system is

$$C = C_1\, 2^{-2M} + C_2. \qquad (1)$$

This is to be compared with the complexity of the standard color correction, which is on the order of $4C_1/3$ (since 3D interpolation has to be carried for each of C, M, Y, K as opposed to 3 color separations in the proposed method). The wavelet transform is not included in the computation because it is already embedded in the proposed framework (see Fig. 3).

Table 1 shows the computational cost per sample necessary to color correct the image using the new and conventional approaches for M = 1, 2, 3 (*i.e*. Phase 1 applied to 1/4, 1/16, and 1/64 of the samples, respectively). The savings provided by the proposed method is significant, particularly for multiplications, additions, and table lookups.

The parameter M offers the trade-off between computational cost and image quality. Higher values of M result in more computational savings, as well as greater image degradation. Since the degradation is introduced at high spatial frequencies, where the HVS has reduced sensitivity, it is hoped that an acceptable quality-cost trade-off can be reached with this technique.

Furthermore, since the color correction process itself is derived from measurements of targets of uniform patches, it is based on a low frequency characterization of the device. Hence, there seems to be some consistency in applying the color correction to a low frequency (*i.e*. locally averaged) version of the image. The high frequency information must somehow be added back to the image in order to preserve texture and edges. However, since there is no explicit color characterization of the device at these frequencies, a simple approximation can be used to add the information in these bands, exploiting again the frequency response of the HVS.

5. EXPERIMENTAL RESULTS

Tests were carried out to examine the cost benefit trade-off for the proposed scheme. For this experiment, no application dependent processing (*e.g*. compression, segmentation, *etc*.) in the wavelet domain was included. That is, the wavelet transform was computed up to a given stage M (where M=0 represents the original image at full resolution); Phase 1 color correction was performed on the low-pass channel at that stage; the inverse wavelet transform was performed; and finally, Phase 2 color correction was applied on the entire image. A 16-tap QMF (Johnston) filter bank [3], [4] was used for the wavelet analysis. The images were processed through the standard and proposed techniques, and reproduced on a Xerox 5760 CMYK xerographic printer at 400 dpi.

For M=1 or M=2, pictorial images corrected with the proposed algorithm result in quality that is very close to, if not indistinguishable from, that of the standard approach. For M=3, artifacts start to appear in the form of "beating" or "oscillations" in reasonably flat areas of the image. These artifacts are mainly noticeable in the luminance component of the image, so that similar effects may also appear if one would apply the same concept for TRC adjustments of monochrome images. Those artifacts are due to the highly non-linear nature of the printing process and result from the mixing of corrected and uncorrected frequency bands.

In addition to a qualitative comparison, a quantitative technique for evaluating the algorithms was desired. A simple metric would be to compute the root mean square (RMS) difference between the CMYK images resulting from the standard and proposed schemes. However, CMYK differences are not representative of visually perceived error. To alleviate this problem, the CMYK images were transformed into CIELAB space using a printer model [10]. (To maintain consistency throughout the system, the same printer model was used to derive the color correction transforms.) However, while the colors are now in a perceptually meaningful representation, the CIELAB coordinate system suffers from the drawback that it is intended only for comparison of large uniform patches, and does not account for the sensitivity of the HVS to different spatial frequencies. This limits the utility of the CIELAB space for evaluating differences between reproductions of natural scenes. Since the proposed algorithm exploits frequency characteristics of the HVS, an error metric was desired that was also spatial frequency dependent. To this end, the recently proposed spatial CIELAB (sCIELAB) model [11] was used. This model utilizes the luminance and chrominance spatial contrast sensitivity functions, as well as the traditional CIELAB color space, to derive a color difference metric that accounts for both spatial frequency characteristics and color difference sensitivity of the HVS. For the color difference term, the original sCIELAB model uses CIELAB 1976 ΔE [11]. In this work, we used the CIE 1994 color difference equations [12] as a superior metric for perceived color difference. The final sCIELAB error metric is in the form of an error image that represents the perceived difference between two images. An aggregate error criterion was selected as

$$\Delta E_{rms} = \sqrt{\frac{1}{S}\sum_{ij} \Delta E_{ij}^2}\ , \qquad (2)$$

where S is the number of image pixels, and ΔE_{ij} are the individual sCIELAB ΔE errors for each pixel. Tests were carried out for M=1, M=2, and M=3, for 3 images with varying pictorial content. A viewing distance of 14 inches and a 400 dpi resolution were assumed. The average ΔE across the 3 images is shown as a function of M in Fig. 5. In general, $\Delta E < 1$ is considered to be below the visual detectability threshold. Noise, printer instability, and other imperfections in the system can easily give rise to reproduction errors between $\Delta E = 1$ and $\Delta E = 3$. Given this margin, and the qualitative observations, we believe that $M = 2$ provides substantial savings with acceptable image quality. We remark finally that to make the evaluation more realistic, the effect of wavelet processing for the given application (*e.g*. compression, descreening, etc) must also be considered.

6. CONCLUSIONS

A method has been proposed for combining multiresolution image analysis and device color correction. It is based on the observation that the human visual system is most sensitive to color errors in low frequency channels. The color correction is divided into a computationally complex phase of multidimensional transforms, followed by a simple phase of 1-D transforms. Both phases are carried out on a downsampled version of the image, in order to achieve accurate color correction in the low frequency bands. Only the simpler second phase is carried out on the full resolution image. This results in computational savings with very little loss in visual quality. The spatial CIELAB model has been found to be an effective tool for evaluating image quality for this application. Future work includes the investigation of different wavelet basis functions, and different color spaces to be used in the Phase 1 correction.

7. REFERENCES

1. R. Balasubramanian, R. V. Klassen, and R. deQueiroz, "Fast color correction of JPEG Compressed Images," *Proc. IS&T/SID's 5th Color Imaging Conf.*, Scottsdale, AZ, Nov. 1997.
2. W. P. Pennebaker and J. L. Mitchell, *JPEG: Still Image Compression Standard*, Van Nostrand-Reinhold, New York, NY, 1993.
3. G. Strang and T. Q. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge Press, Wellesley, MA, 1996.
4. M. Vetterli and J. Kovacevic, *Wavelets and Subband Coding*, Prentice-Hall, Englewood Cliffs, NJ, 1995.
5. A. Said and W. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. on Circuits and System for Video Technology*, vol. 6, pp. 243-250, June 1996. .
6. J. Luo, R. L. de Queiroz and Z. Fan ``Universal descreening technique via wavelet analysis," *Proc. IS&T/SPIE Symp. on Electronic Imaging: Science and Technology*, San Jose, CA, SPIE Vol. 3018, pp. 18-29, Feb. 1997.
7. R. W. G. Hunt, *Measuring Color*, Ellis Horwood, Chichester, 1995.
8. J. A. C. Yule, *Principles of Color Reproduction*, John Wiley & Sons, New York, 1967.
9. H. R. Kang, *Color Technology for Electronic Imaging Devices*, SPIE, 1997.
10. R. Balasubramanian, "A printer model for dot-on-dot halftone screens", *SPIE Proc.*, vol. 2413, pp. 356-364, 1995.
11. X. M. Zhang, and B. A. Wandell., "A spatial extension to CIELAB for digital color image reproduction", *Proceedings of the SID (Society for Information Display) Symposium*, 1996.
12. *Industrial colour-difference evaluation*, CIE TC 1-29 technical report, 1995.

Table 1: Comparison of computational cost per image sample, using standard and proposed approaches.

	Look-ups	Adds	Multiplies	Compares	Shifts
Standard	19.00	14.00	12.00	2.50	2.00
M=1	14.75	9.75	2.25	2.63	1.25
M=2	11.94	7.67	0.56	2.15	1.06
M=3	11.23	7.17	0.14	2.04	1.02

Figure 1: Pictorial depiction of a two-level wavelet transform

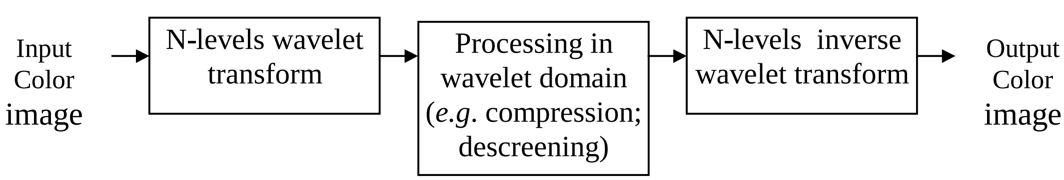

Figure 2: Processing in wavelet domain

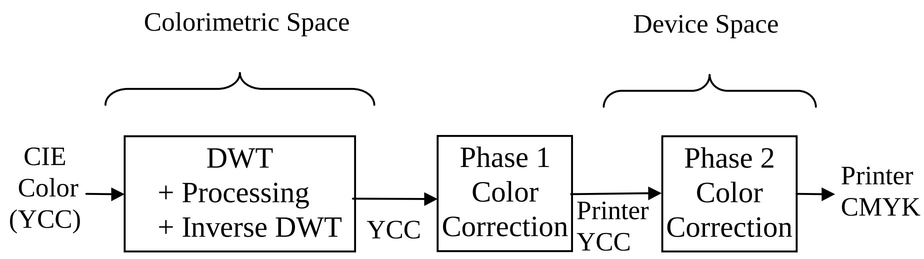

Figure 3: The standard approach of wavelet processing, followed by two phase color correction in the spatial domain at full resolution. An example of a color space (YCC) is used for illustration.

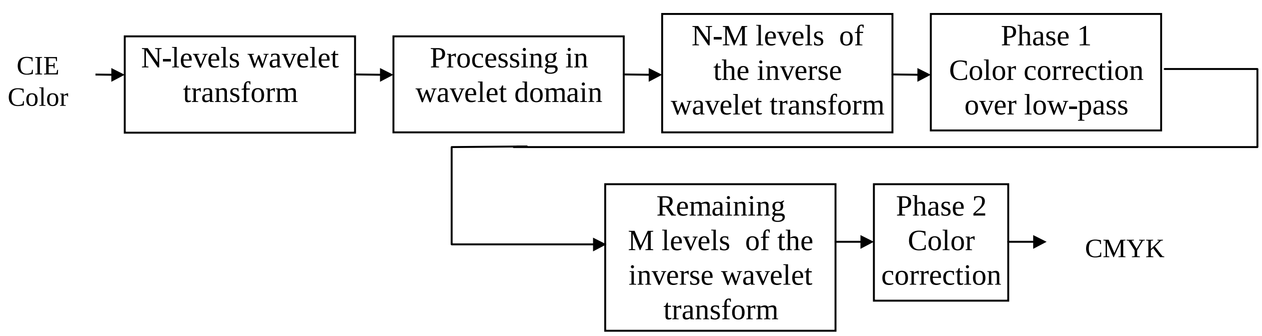

Figure 4: Proposed approach where-in Phase 1 correction is applied on a low-pass version of the image, and Phase 2 correction is applied on the entire image.

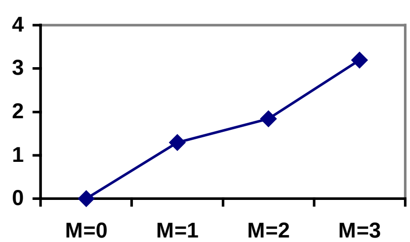

Figure 5: sCIELAB ΔE error between proposed scheme and the standard approach, M=0 (*i.e.* all samples undergo Phase 1 correction). Printer resolution is 400 dpi resolution, and the assumed viewing distance is 14 inches. Results averaged over 3 images.