Automatic Color-to-Gray Conversion for Digital Images in Gradient Domain

Lu Hao, Jie Feng and Bingfeng Zhou
Institute of Computer Science and Technology, Peking University, Beijing, P.R.China
{*pkuhaolu, feng_jie, zhoubingfeng*}*@pku.edu.cn*

Keywords: Color Removal, Gradient Domain, Image Processing, Color Difference.

Abstract: Color-to-grayscale conversion for digital color images is widely used in many applications. In this paper an automatic gradient domain color-to-gray conversion method is described. By enhancing the luminance gradient with a modulated chromatic difference enhancement in CIELAB space, a gradient field is created to construct the resulting grayscale image using a Poisson equation solver. A sign function for the gradient is defined for isoluminance color images to keep correct color ordering. By introducing a structural similarity index measurement (SSIM), the main parameters of the method are automatically optimized in the sense of human vision. Therefore, this method can automatically produce artifact-free and salience-preserving grayscale images that coincide with human perception for the color difference.

1 INTRODUCTION

Grayscale images are necessary in many application areas such as black-and-white printing, computational photography, video and animation, etc. Hence, converting digital color images into grayscale images without losing of details and distortions in human vision is an important issue in computer graphics. Although some color-to-grayscale conversion algorithms have been successfully used in industry, there are still many problems remain unsolved, e.g. maintaining the color discriminability for isoluminant colors during the conversion.

In the existing literatures, there are mainly three categories of color-to-gray conversion algorithms:

(1) *Linear combination* of the original color channels, typically, the Y component of CIEXYZ system (Ohta and Robertson, 2005). This kind of algorithms are widely used in industry, but they lack the discriminability of isoluminance colors.

(2) *Global optimization* algorithms (Gooch et al., 2005; Kim et al., 2009) try to avoid the problem of category 1, by solving a global optimization problem to modulate the final grayscale representation. However, many of this kind of methods are very time-consuming. Some new researches introduce simplified methods to improve the computational performance (Lu et al., 2012b).

(3) *Local feature enhancement* algorithms (Neumann et al., 2007; Smith et al., 2008; Grundland and Dodgson, 2007; Ancuti et al., 2011) locally enhance the grayscale to preserve the original color and luminance contrasts, but still suffer from the low execution efficiency and the grayscale distortion.

For an ideal color-to-gray conversion algorithm, several requirements should be satisfied. First, the resulting grayscale image have to coincide with the luminance vision of human eyes, which is typically defined by the L component of CIELAB [1] color model (Wyszecki and Stiles, 1982). Second, for an isoluminance color image, all the colors in the image must be discriminable in the resulting grayscale image. Third, no artifacts should be introduced into the resulting grayscale image. For an image generated by a *Poisson Equation Solver* (PES) working in gradient domain, these artifacts usually appears in the form of "halo effect", which must be reduced to be invisible by human eyes (Fattal et al., 2002).

In this paper, we present a new category 3 algorithm based on our previous work (Zhou and Feng, 2012). The new method automatically generates grayscale images from color images under a structural

This work is partially supported by NSFC grants #61170206, #61370112, and Specialized Research Fund for the Doctoral Program of Higher Education #20110001110077.

[1]For simplicity, in this paper, CIELAB refers to CIE 1976 ($L^*a^*b^*$)-Space, and variables a and b are used to stand for a^* and b^* respectively.

Figure 1: Gradient-domain salience-preserving color-to-gray conversion. Left: Original color image. Middle: Grayscale image converted by our method ($\beta = 0.2, \gamma = \infty, \alpha = 0.1$). Right: L in CIELAB.

similarity constrain. It performs color-to-gray conversion in the CIELAB color space, and taking advantage of the gradient domain of the image.

In gradient domain image processing, an image can be presented by the gradients at each pixel, and the gradient can be treated as a partial derivative of the original image. When this partial derivative is solved by a PDE solver such as PES, the original image can be reconstructed. By modifying data in the gradient domain, a different image can be obtained for certain purpose. In Fattal'method (Fattal et al., 2002), this strategy is used to convert a HDR image into a LDR one, so that it can be correctly displayed in a LDR device. Similar applications can also be found in the literature (Pérez et al., 2003; McCann and Pollard, 2008).

Therefore, in our method, we generate such a gradient field and employs a PDE solver to reconstruct the grayscale image from the color image. The gradient at each pixel measures the color differences between its neighbors. By enhancing the luminance difference with the chromatic difference component in CIELAB space, the salience of the original color image caused by color vision can be well preserved in the resulting grayscale image (Fig. 1). By attenuating the amount of the chromatic difference component, grayscale distortions can be minimized and become imperceptible to human eyes. Additionally, a sign function for the color difference is also defined to keep correct color ordering for isoluminance color images.

This color difference is calculated base on the CIELAB model (Ohta and Robertson, 2005), which is a reflection of the color vision of human eyes, hence the converted grayscale image will be an excellent approximation of the original color image. Furthermore, there are four parameters used for attenuating the chromatic difference and keeping the color ordering. They are automatically optimized utilizing a structural similarity index measurement (SSIM) (Wang et al., 2004) between the converted grayscale image and the original color image, which leads to better preserving of the structural details of the image. Experiments show that our method can produce outstanding results in contrast to the prior works.

2 COLOR-TO-GRAY CONVERSION IN CIELAB-BASED GRADIENT DOMAIN

2.1 Gradient Domain Image Processing

In gradient domain, a grayscale image I is a discretization of a continuous 2D function $I(x,y)$ defined in R^2. It can be represented by the gradient ∇I of the original $I(x,y)$

$$\nabla I = (I_x, I_y) = (\tfrac{\partial I}{\partial x}, \tfrac{\partial I}{\partial y}). \tag{1}$$

The discrete form of Eq.(1) is formulated as

$$\begin{aligned} \nabla I &= (I_x, I_y) \\ &= (I(x+\Delta x, y) - I(x,y), I(x, y+\Delta y) - I(x,y)). \end{aligned} \tag{2}$$

Therefore, this equation can be solved by a PDE solver such as Poisson equation solver (PES) (Fattal

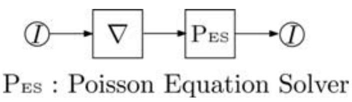

Figure 2: Gradient domain image processing.

et al., 2002; Press et al., 1992) to reconstruct the original image I, as illustrated in Fig. 2. For the problem of color-to-gray conversion, if I refers to the luminance component L of a color image C, then from the gradient field ∇I we can reconstruct the luminance of C (Fig. 3(d)).

(a) (b) (c) (d)

Figure 3: Color-to-gray conversion by enhancing the luminance difference with the chromatic color difference. (a) Original color image. (b) Using chromatic difference enhancement without attenuation ($\beta = 1$, $\gamma = \infty$, $\alpha = 0$). (c) Artifacts removed by using attenuated chromatic difference ($\beta = 1$, $\gamma = \frac{1}{21}$, $\alpha = 0$). (d) No chromatic difference added ($\beta = 0$, $\alpha = 0$).

2.2 The Measurement of Color Difference

CIELAB is a uniform color space, which means: given two points in CIELAB space, their Euclidean distance exactly measures the perceptive feeling of color difference in human eyes for the two colors they represent (Ohta and Robertson, 2005; Shevell, 2003). Therefore, the color difference ΔE of the two colors is defined by:

$$\Delta E = \sqrt{(\Delta L)^2 + (\Delta a)^2 + (\Delta b)^2}. \tag{3}$$

where ΔL, Δa, Δb are the differences along the coordinate axis. Hence, if we use only the luminance component L to reconstruct the grayscale image, the resulting image will not coincide with the color difference that human eyes perceive (Gooch et al., 2005).

It is straightforward to use the color difference in Eq.(3) which includes chromatic difference for instead in constructing the gradient field. Experiments show that more color difference can be successfully preserved in this way (Fig. 1). However, perceptible grayscale distortion may occur at the same time, especially where strong or noisy color differences exists (Fig. 3(b)). In order to remove these artifacts, we add an attenuation function $A(\cdot)$ to Eq.(3), whose details will be given in Section 3. Then, the modulated color difference is formulated as

$$\Delta E = \sqrt{(\Delta L)^2 + \left(A\left(\sqrt{(\Delta a)^2 + (\Delta b)^2}\right)\right)^2}, \tag{4}$$

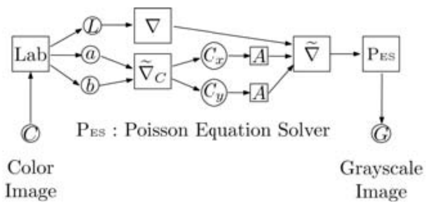

Figure 4: Color-to-gray conversion based on color difference.

2.3 Color-difference-Based Color-to-Gray Conversion

Based on the idea of Section 2.2, we propose a new framework that introduces chromatic color difference into color-to-gray conversion. As shown in Fig. 4, the input color image C is represented in $L(x,y)$, $a(x,y)$, $b(x,y)$ channels of CIELAB model. Its gradient field $\widetilde{\nabla}C$ is composed of a luminance gradient component ∇L and a chromatic gradient component $\widetilde{\nabla}_C(a,b)$. The former is calculated as in Eq.(2), and the latter is obtained by:

$$\widetilde{\nabla}_C(a,b) = (C_x, C_y), \tag{5}$$

where,

$$\begin{aligned} C_x &= \sqrt{(a(x+\Delta x,y) - a(x,y))^2 + (b(x+\Delta x,y) - b(x,y))^2}, \\ C_y &= \sqrt{(a(x,y+\Delta y) - a(x,y))^2 + (b(x,y+\Delta y) - b(x,y))^2}. \end{aligned} \tag{6}$$

Then, $\widetilde{\nabla}C$ can be calculated from ∇L and $\widetilde{\nabla}_C(a,b)$ using Eq.(7), before it is fed into the PES to reconstruct the grayscale image G:

$$\begin{aligned} \widetilde{\nabla}C = \\ \Big(& sign(L_x, a(x+\Delta x,y), a(x,y), b(x+\Delta x,y), b(x,y)) \cdot \sqrt{L_x^2 + A^2(C_x)}, \\ & sign(L_y, a(x,y+\Delta y), a(x,y), b(x,y+\Delta y), b(x,y)) \cdot \sqrt{L_y^2 + A^2(C_y)}\Big). \end{aligned} \tag{7}$$

Here, $A(\cdot)$ is the attenuation function for the chromatic differences C_x and C_y, used to remove grayscale distortions caused by the PES. Function $sign(\cdot)$ defines the sign of the gradient. It is used to determine the color ordering for isoluminance color images (Section 4).

3 ARTIFACT REMOVAL

When creating new images with PES, a common problem is the existence of artifacts. In color-to-grayscale conversion, the artifacts lead to the grayscale distortion as shown in Fig. 3(b). There are many works aim to solve this problem, e.g. Fattal (Fattal et al., 2002) employs a multi-scale schema and Neumann (Neumann et al., 2007) removes the inconsistency of the gradient field. In our method, we employ a single-scale method and selectively attenuate the gradient enhancement to the chromatic differences to remove the artifacts. Experiments show that this scheme is fast and efficient (Fig. 3(c)).

The attenuation of gradient enhancement takes the form of a attenuation function $A(\cdot)$ as mentioned in Section 2.2, which is defined as:

$$A(x) = x \cdot \left(\beta \left(1 - \left(\frac{x}{cx_{max}} \right)^{\gamma} \right) \right) = x \cdot A_0(x), \quad (8)$$

where, $x \in [0, x_{max}], c \in [1, \infty), \beta \in [0, \infty)$ and $\gamma \in (0, \infty)$. The function works only on chromatic difference C_x and C_y, therefore the enhancement to the luminance difference is always valid. Function $A(\cdot)$ scales down the input signal x by a scaling function $A_0(x)$. As illustrated in Fig. 5, larger value of γ will preserve more high chromatic differences, while smaller γ will attenuate the high chromatic difference and preserve low chromatic differences. The constant c is used to ensure that the largest chromatic difference will not be completely scaled down. In our implementation, we choose $c = 2.0$. Here β and γ are two key parameters to reduce grayscale distortions. Their values are automatically optimized according to a structural similarity function which measures the degree of distortion (Section 5).

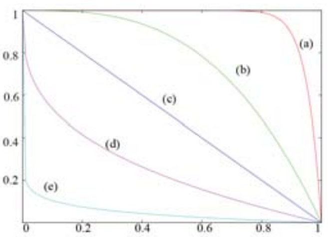

Figure 5: Image of the scaling function $A_0(x)$ in Eq.(8). Here: $c = x_{max} = 1$ and (a):$\gamma = \frac{1}{21}$, (b):$\gamma = \frac{1}{3}$, (c):$\gamma = 1$, (d):$\gamma = 3$, (e):$\gamma = 21$.

4 COLOR ORDERING FOR ISOLUMINANCE IMAGE

Keeping correct ordering for isoluminance colors is a challenge for color-to-gray conversion. In a converted grayscale image, the difference between the colors with different luminance are easier to preserve. However, it is difficult for isoluminance colors, since they are not discriminable in luminance. In our method, we determine the color orders by defining a sign function for the gradient field $\widetilde{\nabla}C$.

$\widetilde{\nabla}C$ is constructed from the modulated color difference (Eq.(4)), hence it is not a signed value by itself. If there is luminance difference between a pixel and its neighbor, the sign of the gradient at that pixel can be reasonably defined as the sign of the luminance difference. But that do not work for a pixel that has equal luminance with its neighbors. Instead, we employ a similar schema as Gooch's method (Gooch et al., 2005). By competing the luminance difference ΔL with the chromatic difference $\vec{\Delta_C}$, our sign function is defined as:

$$sign(\Delta L, a_2, a_1, b_2, b_1) = sign(\Delta L + \alpha \cdot (\vec{v_\theta} \cdot \vec{\Delta_C})), \quad (9)$$

where, $(L_1, a_1, b_1), (L_2, a_2, b_2)$ are CIELAB coordinates of two colors, $\Delta L = L_2 - L_1$, $\vec{v_\theta} = (\cos\theta, \sin\theta)$, and $\vec{\Delta_C} = (a_2 - a_1, b_2 - b_1)$. $\alpha \in [0, 1]$ defines the strength of the chromatic difference affecting the sign of the gradient , and $\theta \in [0, 2\pi)$ defines a direction in *a-b* plane of CIELAB space.

Fig. 6 demostrates the effect of our sign function. In the original image (Fig. 6(a)), the chromatic color differences between neighboring pixels are larger then their luminance difference, hence the sign function helps to reveal the color-blindness testing patterns in the converted grayscale images.

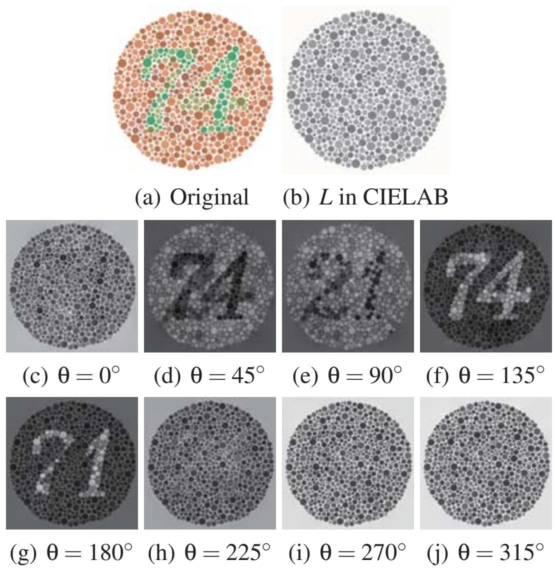

(a) Original (b) L in CIELAB

(c) $\theta = 0°$ (d) $\theta = 45°$ (e) $\theta = 90°$ (f) $\theta = 135°$

(g) $\theta = 180°$ (h) $\theta = 225°$ (i) $\theta = 270°$ (j) $\theta = 315°$

Figure 6: Color discriminability testing on a color blindness testing chart (Ishihara, 1917). Different θ reveals different patterns ($\alpha = 1, \beta = 1, \gamma = \infty$).

5 AUTOMATIC PARAMETER OPTIMIZATION

Our color-to-gray method are able to produce perfect salience-preserving grayscale images. Preliminary results prove the validity of our method (Fig.7). However, inappropriate parameters will result in grayscale distortions in the converted image. In order to better coincide with the human vision of color difference and reduce artifacts, the four key parameters in our algorithm, α, β, γ and θ, are automatically decided according to a *Structural Similarity Index Measurement* (SSIM) between the converted image and the input image.

SSIM is a useful method for measuring perceptual difference between two grayscale images (Wang et al., 2004). It considers the structural information of an image as a feature independent of the average luminance and contrast. To measure the similarity of image X and Y, SSIM is defined as

$$SSIM(X,Y) = [l(X,Y)]^{k_1} \cdot [c(X,Y)]^{k_2} \cdot [s(X,Y)]^{k_3}, \quad (10)$$

in which $l(X,Y)$, $c(X,Y)$ and $s(X,Y)$ are the luminance comparison, the contrast comparison and the structure comparison, respectively. *SSIM* works on an 11×11 window (X_j, Y_j), which moves over the two images X and Y. Finally, a *Mean Structural Similarity Index Measurement* (MSSIM) (Wang et al., 2004) is calculated as the similarity of the two images:

$$MSSIM(X,Y) = \frac{1}{M} \sum_{j=1}^{M} SSIM(X_j, Y_j). \quad (11)$$

Here we extend the MSSIM to measure the similarity of two RGB color images, by computing the average MSSIM value of the three channels. Since a grayscale image is also a degenerated RGB image with equal values in the three channels, we can employ Eq.(12) to measure the similarity between the input color image and the converted grayscale image.

$$MSSIM = \frac{1}{3}(MSSIM_r + MSSIM_g + MSSIM_b) \quad (12)$$

Then, by solving an optimizing problem, we can find a group of optimal parameters $\hat{\alpha}$, $\hat{\beta}$, $\hat{\gamma}$ and $\hat{\theta}$ that produce the maximum measurement value of MSSIM:

$$(\hat{\alpha}, \hat{\beta}, \hat{\gamma}, \hat{\theta}) = \arg\max MSSIM(C, G(\alpha, \beta, \gamma, \theta)). \quad (13)$$

This optimizing problem is solved by a commonly used multi-dimension downhill simplex method (Press et al., 1992).

From Fig. 7 we can see that the converted results with excessive enhancement of details (Column 2) turn out to have lower MSSIM values than our optimized results (Column 3). That is because, the grayscale distortions will affect the structural features of the converted images, and our extended MSSIM has the ability of detecting these distortions and maintaining structural similarity. Hence, it helps to find better parameters for color-to-gray conversion.

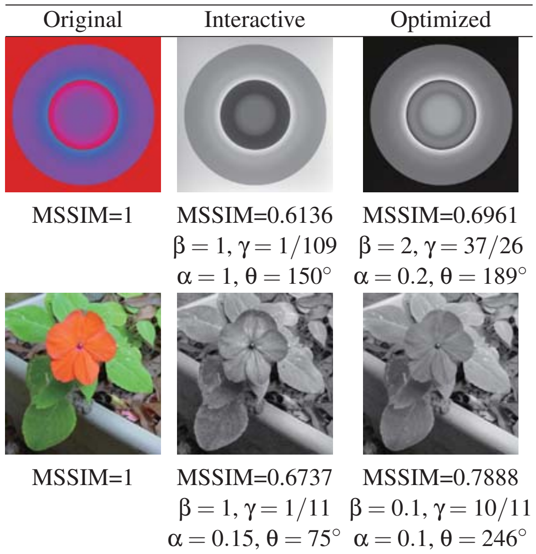

Figure 7: Gray-scale images converted by our method. Column 2: Interactively chosen parameters excessively enhance the details and result in lower MSSIM; Column 3: Automatically optimized parameters obtain higher MSSIM and visually better results.

Table 1: Corresponding parameters of our method for Fig. 8, where columns labeled (1) through (6) correspond to the images from left to right.

Image Number	(1)	(2)	(3)	(4)	(5)	(6)
β	0.11	0.2	0.1	0.2	0.3	0.1
γ	$\frac{2}{7}$	$\frac{3}{26}$	$\frac{12}{33}$	1	$\frac{19}{36}$	$\frac{4}{9}$
α	0.12	0.2	0.1	0.2	0.3	0.2
θ	318°	314°	230°	189°	308°	221°

6 EXPERIMENTAL RESULTS

We implemented our gradient domain color-to-gray conversion method and perform it on a number of RGB images. Input images are first converted to CIEXYZ and then to CIELAB. The RGB color is in PAL-RGB standard and reference white is D_{65} (Ohta and Robertson, 2005; Pascale, 2003). After the L channel of the grayscale image is reconstructed, it is converted back into RGB color, with the dynamic range scaled to $[0, 255]$, and the chrominance value of all the pixels set as that of D_{65}.

In order to evaluate the quality of our method, we compared it with 7 prior works (Lu et al., 2012a; Kim

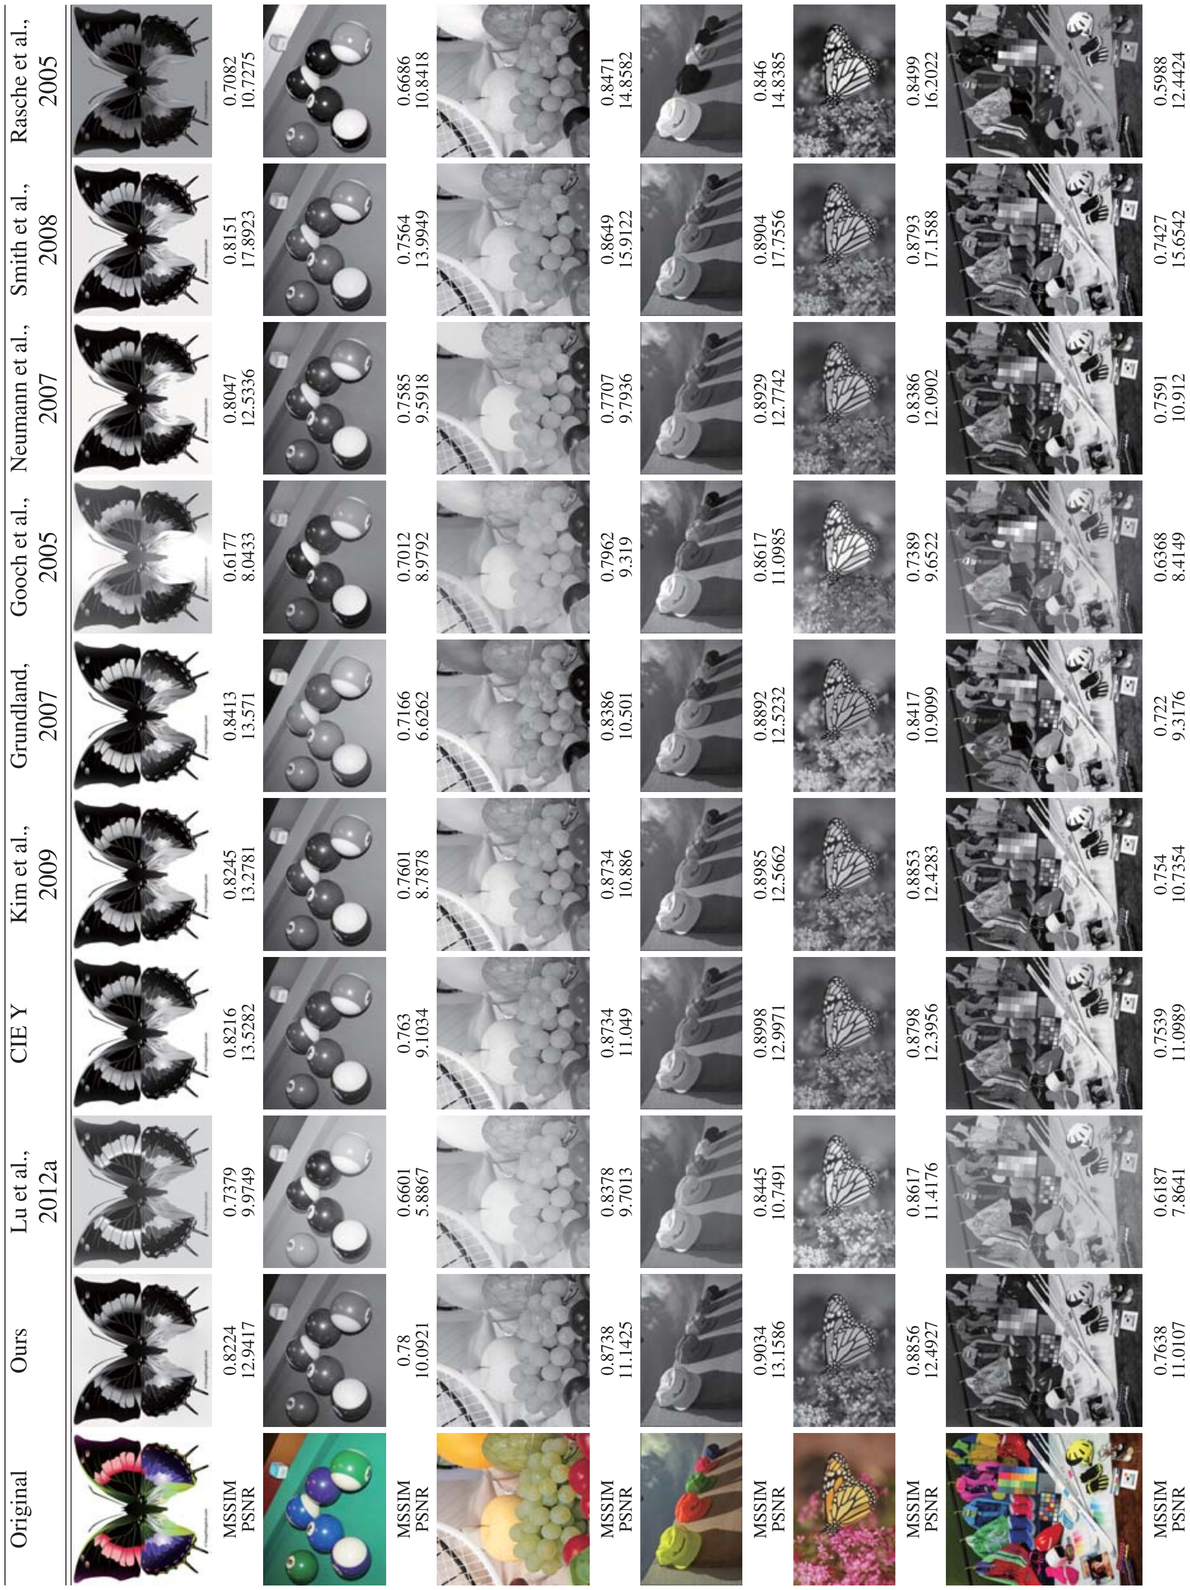

Figure 8: Comparison of our method with some of the prior works. MSSIM and PSNR values are given below each converted image. The corresponding parameters of our method are given in Table 1.

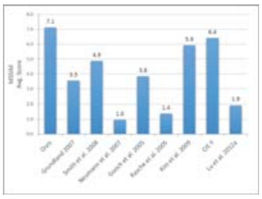

(a)

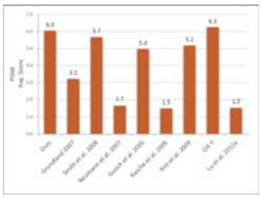

(b)

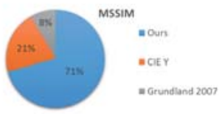

(c)

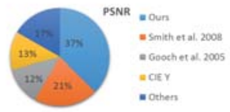

(d)

Figure 9: MSSIM/PSNR score statistic. (a) MSSIM average score of each method; (b) PSNR average score of each method; (c) MSSIM No.1-hit percentage; (d) PSNR No.1-hit percentage.

Original	Our result	L in CIELAB

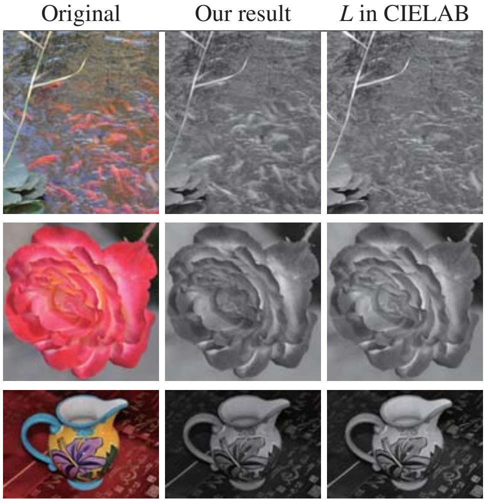

Figure 10: Salience-preserving color-to-gray. Parameters for each image: (1): $\beta = 1$, $\gamma = \frac{1}{41}$, $\alpha = 1$, $\theta = 0°$; (2): $\beta = 1$, $\gamma = \frac{1}{51}$, $\alpha = 1$, $\theta = 270°$; (3): $\beta = 1$, $\gamma = \frac{1}{21}$, $\alpha = 0$.

et al., 2009; Smith et al., 2008; Grundland and Dodgson, 2007; Neumann et al., 2007; Gooch et al., 2005; Rasche et al., 2005) and CIE Y on 24 different input images. A part of the results are shown in Fig. 8, whose corresponding parameters are listed in Table 1.

All of the converted grayscale images are compared with the original color images, and the similarity is evaluated by both MSSIM and PSNR. For each input image, the results of the 9 methods are sorted by their MSSIM/PSNR value and scores are given according to the ranking, i.e. the result with the highest MSSIM/PSNR value has a score of 8, while the lowest corresponds to 0. For each method, how many times it wins the first place is also recorded, and named as the "No.1-hits". The statistical result in Fig.

Original	Our result	L in CIELAB

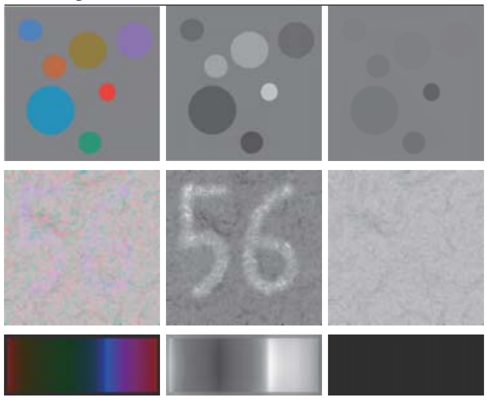

Figure 11: Color discriminability of the algorithm. Parameters: (1) $\beta = 1$, $\gamma = \frac{1}{21}$, $\alpha = 1$, $\theta = 80°$. (2) $\beta = 1$, $\gamma = \infty$, $\alpha = 1$, $\theta = 270°$. (3) $\beta = 1$, $\gamma = \frac{1}{66}$, $\alpha = 1$, $\theta = 315°$.

9 shows that our method has the highest average score and No.1-hits for MSSIM, and the highest No.1 hits for PSNR. CIE Y has an advantage in average PSNR scores because it is a linear combination of R, G and B. Our method is competitive with it and better than all other methods.

From the experimental results, we can see that our method shows a satisfying salience-preserving ability. As demonstrated in Fig. 10, many details, e.g. the fishes in the first row, can be seen more clearly in our results. On another aspect, grayscale distortions are well controlled and no visible artifacts appear in the resulting images.

Fig. 11 shows the color discriminability of our method. Images in the middle column are our results and the right are obtained by using L channel in the CIELAB model. Our method shows perfect color discriminability and ordering for both discrete color (the 1st image) and continuous color (the 2nd and the 3rd). Especially, the 3rd image are computer-designed isolumminance image where $L = 50$. The colors cannot be distinguished in the results of L channel, while our method are able to produce grayscale image with clear and correct color ordering.

Hence, the 4-parameter model in our method provides the ability of detail extracting and enhancing from the color images; and optimal values for the parameters, in the sense of human vision, can be automatically calculated with the MSSIM constrain. Our method is able to preserve more details and coincide with the human perception of color difference.

We implemented our color-to-gray conversion method using a PES given by Press(Press et al., 1992). As the PES has a linear time complexity (Fattal et al., 2002), our method has a steady execution speed of around 2 seconds per mega pixel (1024×1024 pixels in RGB) in each optimizing iteration on a computer with Intel Core Dou CPU 2.2GHz and 2GB memory. The total computing time depends on the complexity of the input image and the number of the iterations.

7 CONCLUSION

In this paper we explored the gradient domain color-to-gray conversion. By controlling the strength of chromatic enhancement to the luminance gradient, we are able to obtain a salience-preserving grayscale image with no visible grayscale distortion. It is based on an observation that grayscale distortion is mainly caused by strong chromatic differences, and Eq.(8) aims to attenuate these strong gradient. Experiments have proven the validity of the observation. By defining a sign function for the enhanced gradient, our method is also able to keep correct color ordering for isoluminance images.

Our method support automatic optimization of the main parameters according to a structural similarity measurement between the converted image and the original one. This method is effective and can generate grayscale images that coincide with human vision. However, the computing efficiency of current optimizing process is not high enough for real-time applications. That is what we need to improve in future works.

REFERENCES

Ancuti, C., Ancuti, C., and Bekaert, P. (2011). Enhancing by saliency-guided decolorization. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 257–264.

Fattal, R., Lischinski, D., and Werman, M. (2002). Gradient domain high dynamic range compression. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 249–256, New York, NY, USA. ACM.

Gooch, A. A., Olsen, S. C., Tumblin, J., and Gooch, B. (2005). Color2gray: Salience-preserving color removal. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 634–639, New York, NY, USA. ACM.

Grundland, M. and Dodgson, N. A. (2007). Decolorize: Fast, contrast enhancing, color to grayscale conversion. *Pattern Recogn.*, 40(11):2891–2896.

Ishihara, S. (1917). Test for coiour-blindness. *Tokyo: Hongo Harukicho*.

Kim, Y., Jang, C., Demouth, J., and Lee, S. (2009). Robust color-to-gray via nonlinear global mapping. In *SIGGRAPH Asia '09: ACM SIGGRAPH Asia 2009 papers*, pages 1–4, New York, NY, USA. ACM.

Lu, C., Xu, L., and Jia, J. (2012a). Contrast preserving decolorization. In *Computational Photography (ICCP), 2012 IEEE International Conference on*, pages 1–7. IEEE.

Lu, C., Xu, L., and Jia, J. (2012b). Real-time contrast preserving decolorization. In *SIGGRAPH Asia 2012 Technical Briefs*, SA '12, pages 34:1–34:4, New York, NY, USA. ACM.

McCann, J. and Pollard, N. S. (2008). Real-time gradient-domain painting. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–7, New York, NY, USA. ACM.

Neumann, L., Čadík, M., and Nemcsics, A. (2007). An efficient perception-based adaptive color to gray transformation. In *Proceedings of Computational Aesthetics 2007*, pages 73– 80, Banff, Canada. Eurographics Association.

Ohta, N. and Robertson (2005). *Colorimetry: Fundamentals and Applications*. Wiley& Sons, New York.

Pascale, D. (2003). A review of rgb color spaces ... from xyY to R'G'B'. *Babel Color*.

Pérez, P., Gangnet, M., and Blake, A. (2003). Poisson image editing. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, pages 313–318, New York, NY, USA. ACM.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA.

Rasche, K., Geist, R., and Westall, J. (2005). Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum*, 24(3):423–432.

Shevell, S. K. (2003). *The Science of Color*. Elsevier, Oxford, UK.

Smith, K., Landes, P.-E., Thollot, J., and Myszkowski, K. (2008). Apparent greyscale: A simple and fast conversion to perceptually accurate images and video. *Computer Graphics Forum*, 27(2):193–200.

Wang, Z., Bovik, A., Sheikh, H., and Simoncelli, E. (2004). Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 13(4):600–612.

Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae.* Wiley-Interscience, New York, NY, USA, 2 edition.

Zhou, B. and Feng, J. (2012). Gradient domain salience-preserving color-to-gray conversion. In *SIGGRAPH Asia 2012 Technical Briefs*, SA '12, pages 8:1–8:4, New York, NY, USA. ACM.