

Real-Time Video Decolorization Using Bilateral Filtering*

Yibing Song Linchao Bao Qingxiong Yang
Department of Computer Science
Multimedia software Engineering Research Centre (MERC)
City University of Hong Kong, Hong Kong, China
MERC-Shenzhen, Guangdong, China

yb.song@my.cityu.edu.hk linchaobao@gmail.com qiyang@cityu.edu.hk

Abstract

This paper presents a real-time decolorization method. Given the human visual systems preference for luminance information, the luminance should be preserved as much as possible during decolorization. As a result, the proposed decolorization method measures the amount of color contrast/detail lost when converting color to luminance. The detail loss is estimated by computing the difference between two intermediate images: one obtained by applying bilateral filter to the original color image, and the other obtained by applying joint bilateral filter to the original color image with its luminance as the guidance image. The estimated detail loss is then mapped to a grayscale image named residual image by minimizing the difference between the image gradients of the input color image and the objective grayscale image that is the sum of the residual image and the luminance. Apparently, the residual image will contain pixels with all zero values (that is the two intermediate images will be the same) only when no visual detail is missing in the luminance. Unlike most previous methods, the proposed decolorization method preserves both contrast in the color image and the luminance. Quantitative evaluation shows that it is the top performer on the standard test suite. Meanwhile it is very robust and can be directly used to convert videos while maintaining the temporal coherence. Specifically it can convert a high-resolution video (1280×720) in real time (about 28 Hz) on a 3.4 GHz i7 CPU.

1. Introduction

Color-to-gray conversion is widely used in single-channel image processing applications. The conversion is

*This work was supported in part by a GRF grant from the Research Grants Council of Hong Kong under Grant U 122212 and Adobe Gift Fund #9220066

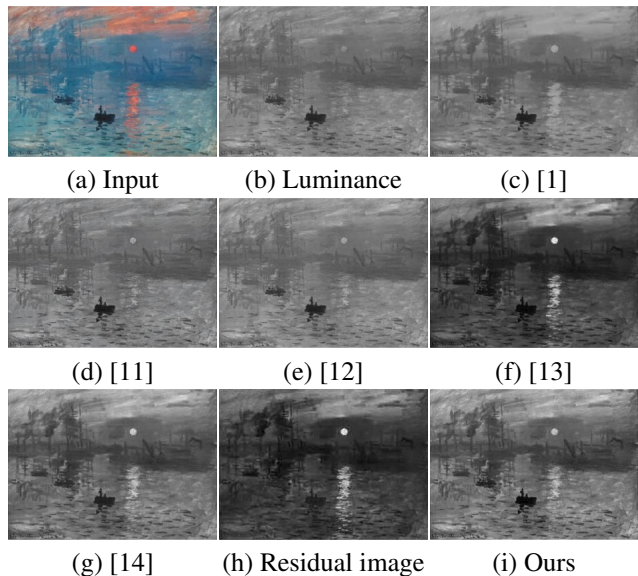


Figure 1. A color image (a) often reveals important visual details missing from its luminance (b). Recently, a number of color-to-gray conversion methods have been proposed to preserve contrast with respect to the original color image. However, the contrast of the converted grayscale image will be either lower than the original color image (see (c)-(e)) or the contrast in the luminance (b) will be lost (see (f)). This paper proposes to combine the luminance (b) and a residual image (h) derived from the color image to a grayscale image (i) that preserves contrast in both the color image (a) and the luminance (b). Note: the residual image (h) is scaled for visualization purpose.

a dimensionality reduction process that three dimensional data has to be reflected in one dimension in the same range. This reduction inevitably suffers from information loss. For instance, the contrast around the sun and its reflection in Figure 1 (a) is not clearly reflected in its luminance (b) due to the removal of chromaticity. It deserves to analysis how we should organize data in the limited space to preserve these visual details.

This paper proposes a color-to-gray conversion method that is simple, effective, efficient and robust. Because the human visual system is more sensitive to luminance than the chromaticity values, the luminance information ought to be kept as much as possible. This paper thus aims at recovering the color contrast/detail lost in the luminance. The loss is estimated as a color image using the bilateral filter and next linearly mapped to a grayscale image named *residual image*. The mapping function is the same at each pixel location and is obtained by minimizing the difference between the image gradients of the input color image and the objective grayscale image that is the sum of the residual image and the luminance. As a result, the converted grayscale image preserves the visual details in both the luminance and the original color image. This is proved numerically using the color contrast preserving ratio (CCPR) proposed in [13]: quantitative evaluation shows that the proposed method is the top performer on the 24 tested images provided by [4].

Besides being effective, a decolorization method needs to be efficient in order to process a video in limited time frame. The main computation involved in the proposed method is the bilateral filter used to estimate detail loss in the luminance. The bilateral filtering method proposed in [23] is adapted in this paper. Due to the lack of image structure, the size of the bilateral filter kernel is set to be as large as the input image to cover every pixel. Large filter kernel allows high compression on the spatial domain in a bilateral filter [20] and enables the proposed color-to-gray conversion method to run in real time on a 3.4 GHz i7 CPU.

Temporal consistency is ubiquitous in video data, and need to be taken into account in video decolorization. The grayscale image converted using the proposed method is actually the sum of the residual image and the luminance of the original color image, thus we just need to make sure that temporally coherent residual image can be obtained. The residual image is linearly mapped from the estimated detail loss and the linear mapping function is computed from all the image pixels and thus robust to temporal variations. As a result, the uncertainty can only come of the detail loss estimated from bilateral filtering. Nevertheless, bilateral filter is a very robust filter, and temporal consistency can be guaranteed as demonstrated in Section 4.3.

The proposed decolorization method has the following advantages over the state-of-the-art methods:

1. Real-time performance: it can convert a 1280×720 resolution image in real-time on a 3.4GHz CPU, which is comparable to the currently fastest method [14].
2. Higher quality: quantitative evaluation [13] on the standard dataset [4] demonstrates it is the top performer.
3. Robust: it can be directly applied to convert every

frame during video decolorization to achieve temporal consistency.

2. Related Work

An overview of current state-of-the-art color-to-gray conversion methods is given in Section 2.1. A brief overview of the bilateral filter is given in Section 2.2.

2.1. Color-to-gray Image Conversions

Traditional color to grayscale conversion such as utilizing only the luminance information fails for images with isoluminant changes. The state-of-the-art decolorization methods can be categorized as local and global mapping. In local mapping pixels are processed spatially. Contrast can be enhanced in local region. In [17] high frequency components of chrominance is added to luminance in order to enhance color edges. In [1] an optimization approach is introduced that iteratively searched the gray levels that best represented the local contrast between all color pairs. In [11] chrominance edges is enhanced by using adaptively-weighted multi-scale unsharp masking. These local approaches may not maintain the constant color regions and visual artifacts may occur. In [19] bilateral filtering is conducted by quantitative measuring the lost color contrast in the luminance and identify proper coefficients of the color transformation model.

In global methods [15] analyzes color differences by predominant component analysis. The lightness and color order could be better preserved by restraining the added chrominance. They did not take into account spatially distant chromatic differences, causing different colors into similar grayscale values. In [12] a nonlinear global mapping method is proposed. The parameters were estimated by minimizing cost function preserving color difference in CIELab color space. In [18] a global energy function is proposed and variationally optimized. In [6] a energy function is defined on a clustered color image, which enabled different color space transformations. In [13] a global optimization approach is established aiming at maximally preserving the original color contrast. An approximation solution to [13] is proposed in [14]. In [16] luminance and chrominance is merged to obtain color difference while chromatic contrast is enhanced. However they had to select offset angle for images.

2.2. Bilateral Filter

The bilateral filter is a robust edge-preserving filter proposed in [5]. It has been used in many computer vision and computer graphics tasks, and a general overview of the applications can be found in [20]. A bilateral filter has two filter kernels: a spatial filter kernel and a range kernel for measuring the spatial and range distance between the center pixel and its neighbors, respectively. The two filter ker-

nels are traditionally based on a Gaussian distribution [9]. Specifically, let $\mathbf{I}_{\mathbf{p}}$ be the color at pixel \mathbf{p} and $\mathbf{I}_{\mathbf{p}}^{\mathbf{I}}$ be the filtered value, we want $\mathbf{I}_{\mathbf{p}}^{\mathbf{I}}$ to be

$$\mathbf{I}^{\mathbf{I}}(\mathbf{p}) = \frac{\sum_{\mathbf{q} \in \Omega_{\mathbf{p}}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|\mathbf{I}(\mathbf{p}) - \mathbf{I}(\mathbf{q})\|) \mathbf{I}(\mathbf{q})}{\sum_{\mathbf{q} \in \Omega_{\mathbf{p}}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|\mathbf{I}(\mathbf{p}) - \mathbf{I}(\mathbf{q})\|)}, \quad (1)$$

where \mathbf{q} is a pixel in the neighborhood $\Omega_{\mathbf{p}}$ of pixel \mathbf{p} , and G_{σ_s} and G_{σ_r} are the spatial and range filter kernels measuring the spatial and range/color similarities. The parameter σ_s defines the size of the spatial neighborhood used to filter a pixel, and σ_r controls how much an adjacent pixel is down-weighted because of the color difference. A joint (or cross) bilateral filter [10, 7] is the same as the bilateral filter except that its range filter kernel G_{σ_r} is computed from another image named guidance image. Let \mathbf{J} denote the guidance image, the joint bilateral filtered value at pixel \mathbf{p} is

$$\mathbf{I}^{\mathbf{J}}(\mathbf{p}) = \frac{\sum_{\mathbf{q} \in \Omega_{\mathbf{p}}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|\mathbf{J}(\mathbf{p}) - \mathbf{J}(\mathbf{q})\|) \mathbf{I}(\mathbf{q})}{\sum_{\mathbf{q} \in \Omega_{\mathbf{p}}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|\mathbf{J}(\mathbf{p}) - \mathbf{J}(\mathbf{q})\|)}. \quad (2)$$

Note that the joint bilateral filter ensures the texture of the filtered image $\mathbf{I}^{\mathbf{J}}$ to follow the texture of the guidance image \mathbf{J} . The efficient implementation of bilateral filtering can be found in [8, 3, 2, 23, 21, 22]

3. Decolorization Using Bilateral Filtering

An overview of the proposed method is presented in Figure 2. We first estimate the detail loss in the luminance due to the removal of chromaticity using bilateral filtering. Figure 2 (a) and (b) are the input color image \mathbf{I} and its luminance L , respectively, and (c) is the bilateral filtered image $\mathbf{I}^{\mathbf{I}}$. As an edge-preserving filter, the filtered image well preserves the contrast between the green dots and the two blue numbers. σ_r is set to 0.02 in this experiment. *The normalized image coordinate is used in this paper such that it resides in [0, 1], and $\sigma_s = 2$ for all the experiments conducted. Image color/intensity is also normalized such that it ranges from 0 to 1.*

Figure 2 (d) presents the joint bilateral filtered image \mathbf{I}^L with its luminance L as the guidance image. A joint bilateral filter ensures that the texture of the filtered image to follow the texture of the guidance image. Hence, because the contrast between the green dots and the two blue numbers is lost in the luminance image L , it also disappears in the joint bilateral filtered image \mathbf{I}^L . The estimate of the detail loss \mathbf{D} is then presented in Figure 2 (e) as a color image by subtracting \mathbf{I}^L in (d) from $\mathbf{I}^{\mathbf{I}}$ in (c):

$$\mathbf{D}_c = \mathbf{I}_c^{\mathbf{I}} - \mathbf{I}_c^L, c \in \{r, g, b\}. \quad (3)$$

The estimate \mathbf{D} is next linearly mapped to a grayscale image to form the residual image R in Figure 2 (f) by minimizing

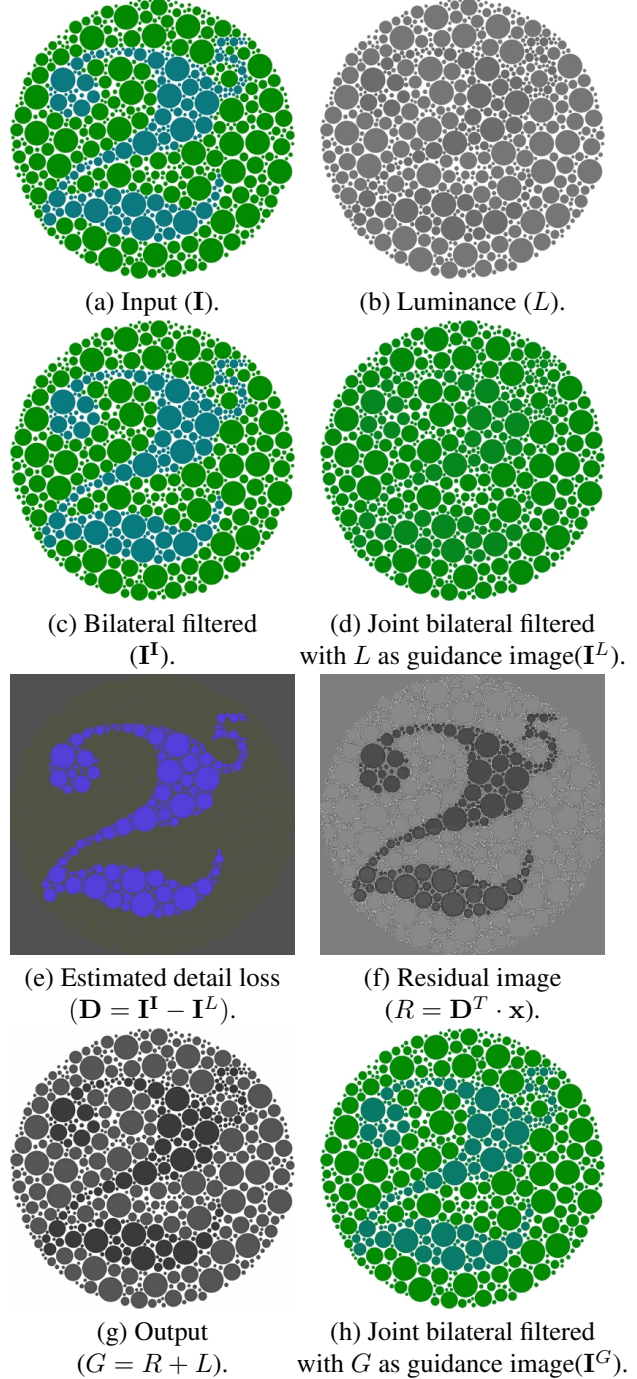


Figure 2. The proposed method. (a) is the input color image and (b) is its luminance. (c) is the bilateral filtered image of (a) and (d) is the joint bilateral filtered image using (b) as the guidance image. The difference between (c) and (d) is presented in (e) and is used as the estimate of the detail lost in the luminance (b). This estimate (e) is then linearly mapped to a grayscale image named residual image (f). The output of the proposed method is the sum of this residual image and the luminance and is presented in (g). (h) is the joint bilateral filtered image obtained using the converted grayscale image (g) as the guidance image. (h) demonstrates the improvement in preserving the contrast between the green dots and the two blue numbers in the original color image (a).




Figure 3. Converted grayscale images of Figure 2 (a). The first row is computed using the detail loss estimate \mathbf{D} in Equation 3 and the second row is computed using \mathbf{D}^{App} in Equation 9. Note that the two converted grayscale images are very similar when σ_r is small, and the converted grayscale image obtained from the approximation method (using \mathbf{D}^{App}) preserves the color contrast lost in the luminance even when σ_r is relatively large.

the difference between image gradients of the input color image \mathbf{I} and the objective grayscale image G that is the sum of the residual image R and the luminance L . Specifically, let $\mathbf{x} = [x_r, x_g, x_b]^T$ denote the mapping function, then at each pixel location \mathbf{p}

$$R(\mathbf{p}) = \mathbf{D}(\mathbf{p})^T \cdot \mathbf{x}. \quad (4)$$

Let ∇ denote the image gradient operator, and

$$\mathbf{A} = \begin{pmatrix} \vdots \\ [\nabla \mathbf{D}_r(\mathbf{p}), \nabla \mathbf{D}_g(\mathbf{p}), \nabla \mathbf{D}_b(\mathbf{p})] \\ \vdots \end{pmatrix} \quad (5)$$

denote the gradient of image \mathbf{D} . Also let

$$m(\mathbf{p}) = \mathbf{argmax}_{c \in \{r, g, b\}} |\mathbf{D}_c(\mathbf{p})|, \quad (6)$$

denote the channel that has the largest amount of contrast loss at pixel \mathbf{p} and

$$\mathbf{B} = \begin{pmatrix} \vdots \\ \nabla \mathbf{I}_{m(\mathbf{p})}(\mathbf{p}) \\ \vdots \end{pmatrix} \quad (7)$$

denote the gradient of image \mathbf{I} at the corresponding channel, the mapping function \mathbf{x} is computed by solving the following function

$$\mathbf{A}\mathbf{x} + \nabla L = \mathbf{B} \Rightarrow \mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T (\mathbf{B} - \nabla L), \quad (8)$$

where $\mathbf{A}^T \mathbf{A}$ is a 3×3 matrix and $\mathbf{A}^T (\mathbf{B} - \nabla L)$ is a 3×1 vector. Equation 8 ensures the similarity of image gradients

between the converted grayscale image $G = R + L$ and the input color image \mathbf{I} .

The joint bilateral filtered image \mathbf{I}^G obtained using the converted grayscale image G as the guidance image is presented in Figure 2 (h). Apparently, \mathbf{I}^G in Figure 2 (h) can be exactly the same as the bilateral filtered image \mathbf{I}^I in (c) only when all the contrast loss are successfully recovered in the converted grayscale image G . However, this is normally impossible for natural images due to dimension reduction. Nevertheless, visual comparison of Figure 2 (d) and (h) demonstrates great improvement in preserving the contrast between the green dots and the two blue numbers in the original color image in (a).

3.1. Fast Approximations

In practice, we use the original color image \mathbf{I} to approximate its bilateral filtered image \mathbf{I}^I in Equation 3 to reduce the computational complexity, which means that

$$\mathbf{D}_c^{App} = \mathbf{I}_c - \mathbf{I}_c^L, c \in \{r, g, b\} \quad (9)$$

is used as the estimate of detail lost \mathbf{D} .

As an edge-preserving filter, the bilateral filtered image \mathbf{I}^I is very close to \mathbf{I} when σ_r is small. In this paper, we use the peak signal-to-noise ratio (PSNR) to measure the similarity between image \mathbf{I} and \mathbf{I}^I :

$$\text{PSNR} = 10 \log_{10} \left(\frac{h \cdot w}{\sum_{\mathbf{p}} \|\mathbf{I}_{\mathbf{p}} - \mathbf{I}_{\mathbf{p}}^I\|^2} \right), \quad (10)$$

where h and w are the height and width of the images. Figure 4 presents the PSNR values computed from 24 tested images provided by [4] and the corresponding bilateral filtered image \mathbf{I}^I obtained using different σ_r parameters. The




Figure 4. PSNR values computed from the 24 tested images and the bilateral filtered images obtained with $\sigma_r \in [0.01, 0.20]$.

pink solid curve in Figure 4 is the mean/average PSNR values computed from all the tested images, and the green and blue curves are the maximum and minimum PSNR values, respectively. As can be seen, the mean PSNR value is larger than 40 dB when $\sigma_r \leq 0.03$, thus there is almost no visible difference between the two images according to [20] and it will be safe to use the original color image \mathbf{I} to approximate the bilateral filtered image $\mathbf{I}^{\mathbf{I}}$ in Equation 3. The converted grayscale images obtained with $\sigma_r = 0.03$ to 0.25 are presented in Figure 3. From top to bottom are the exact and the approximated results obtained from \mathbf{D}^{App} , respectively. Note that when $\sigma_r = 0.03$ (see the 1st column in Figure 3), the approximated grayscale image is very close to the exact one. Also, when σ_r is relatively large (e.g., $\sigma_r = 0.15$, the 4th column in Figure 3), the assumption that the input image I is very similar to its bilateral filtered image $\mathbf{I}^{\mathbf{I}}$ is violated. However the converted grayscale image obtained from the approximation method still correctly preserves the color contrast lost in the luminance. If we change \mathbf{I}^L to $\mathbf{I}^{\mathbf{I}}$ in Equation 9, the original image I is separated into two layers according to [9], $\mathbf{I}^{\mathbf{I}}$ will be the base layer encoding large-scale variations and \mathbf{D}_c^{App} is the detail layer. Now change $\mathbf{I}^{\mathbf{I}}$ back to \mathbf{I}^L , then \mathbf{D}^{App} computed from Equation 9 contains the high-contrast details of the color image lost in the luminance, thus \mathbf{D}^{App} is a robust estimate of the detail lost. As a result, the grayscale image converted using \mathbf{D}^{App} will also effectively reflect the detail loss in the luminance.

4. Experiments

In this section we conduct comparison experiments for numerical and perceptual evaluation on public and real-world dataset. Also temporal evaluation is conducted by representing video performance compared with other methods.

4.1. Numerical Evaluation

We evaluate the proposed method using the 24 tested images provided by [4]. The color contrast preserving ratio (CCPR) proposed in [13] is adopted in this paper for numerical evaluation. The average CCPR values obtained from




Figure 5. Quantitative evaluation using color contrast preserving ratio (CCPR). As can be seen, our method (red solid curve) outperforms all the others on standard test suite.

Table 1. Computational cost evaluation of different methods. We use a 1280×720 color image as input and obtain the time cost of each state-of-the-art method.

Methods	[12]	[16]	[13]	[14]	Ours
Runtime (Sec)	1.224	3.000	2.048	0.035	0.036

different conversion methods are presented in Figure 5. The same parameter setting ($\sigma_s=2$ and $\sigma_r=0.15$) is used in all the experiments presented in this section, and the CCPR values were computed using the source code provided by the authors of [13]. As can be seen, the proposed method is the top performer (the red solid curve).

The computational cost performance is shown in Table 1 where we use a 1280×720 color image as input. Table 1 indicates the proposed method is comparable to the currently fastest method [14]. The computational time of the proposed method is linearly proportional to the number of pixels in the image. This real-time performance enables online high resolution decolorization.

4.2. Perceptual Evaluation

Visual evaluation is presented in Figure 6 and 7. Figure 6 visually compares the recent color-to-grayscale conversion methods with the proposed method using images containing large amount of isoluminant changes. Figure 6 (b) shows that most of the details in the color images in (a) are lost in the luminance. The loss is correctly estimated in the proposed residual images in (h) and is successfully recovered in the converted grayscale images in (i). Figure 7 presents visual comparison using natural images where most of the color contrast is preserved in the luminance in Figure 7 (b). In this case, the residual images in Figure 7 (h) may be relatively flat (and close to zero) and the converted grayscale images in Figure 7 (i) are close to the luminance images in Figure 7 (b). More conversion results of natural images are presented in Figure 8. We compare our method with recent techniques [13, 14]. For natural images, the results obtained using the proposed method perform favorably against [13, 14] on average.




Figure 6. Color-to-grayscale conversion of images with isoluminant changes. Note that the details lost in the luminance in (b) is successfully estimated in the residual image in (i) and recovered in the converted grayscale image in (j). The results are best viewed on high-resolution displays.




Figure 7. Color-to-grayscale conversion of natural images. Note that most of the color contrast is preserved in the luminance in (b), thus the residual images in (i) may be relatively flat (e.g., 1st row of (i)) and the converted grayscale images in (j) are close to the luminance images in (b). The results are best viewed on high-resolution displays.

4.3. Temporal evaluation

Besides images, another view of robust evaluation is temporal coherence in video. If we convert sequential frames to grayscale frames, the perception of altering in adjacent frames should be in accordance with that in the original frames. [12] proposed temporally coherent video conversion, which already demonstrates weakly temporal robustness of their original image decolorization method. Unlike [12], the proposed method automatically preserves

the temporal coherence in a converted grayscale video as shown in Figure 9 (d).

5. Conclusion

This paper presents a real-time decolorization method. Color contrast lost in the luminance is estimated using bilateral filtering and then linearly mapped to a grayscale image named residual image. The sum of the residual image and the luminance is the objective grayscale image

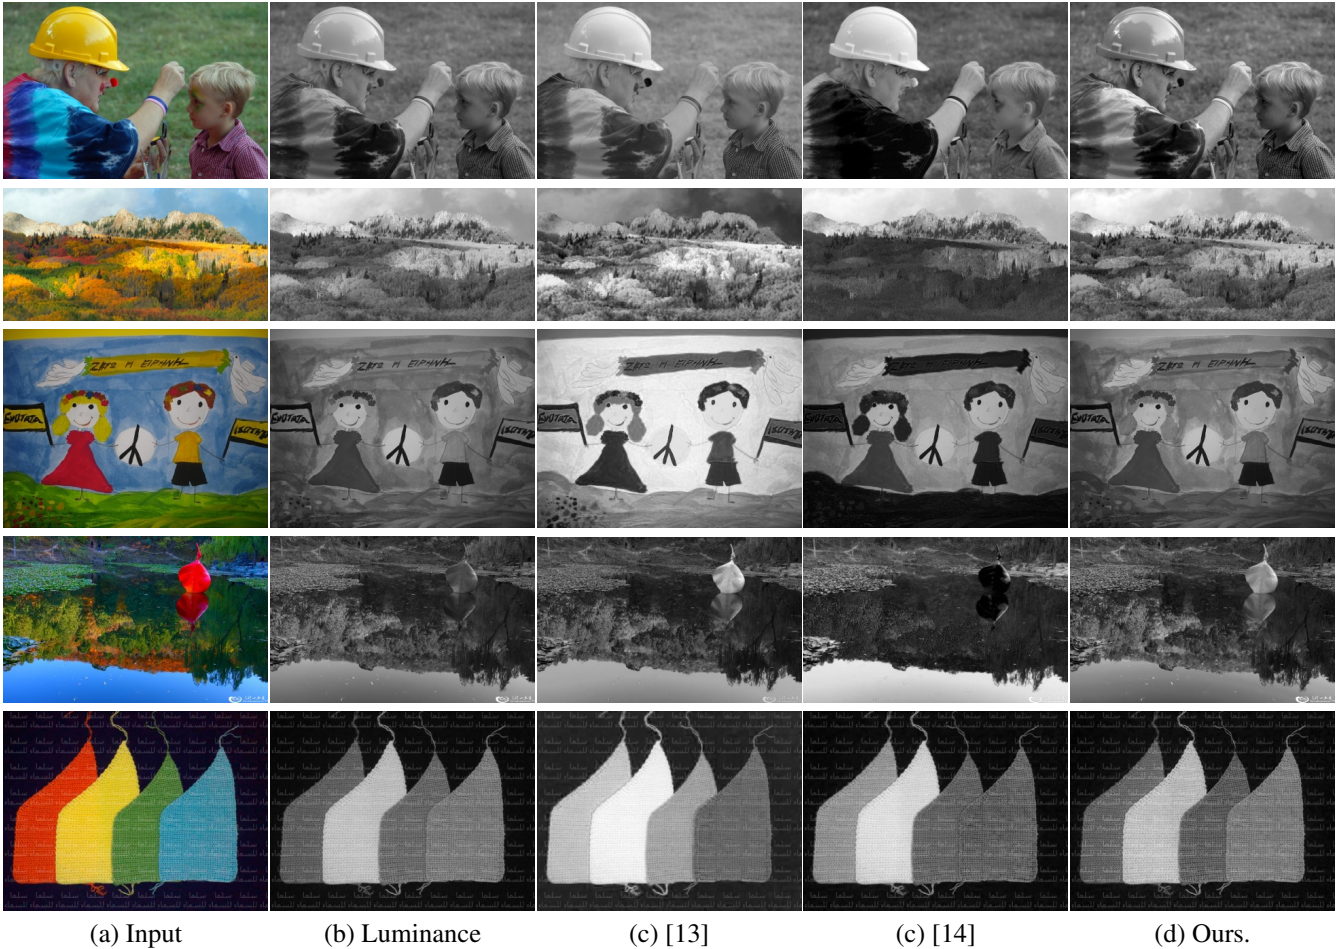


Figure 8. More results on natural images. The proposed method performs favorably against the recent techniques [13, 14]. The results are best viewed on high-resolution displays.

that preserves both the luminance information and the color contrast. Quantitative evaluation on the standard dataset demonstrates that the proposed method outperforms the existing decolorization methods in quality. Meanwhile it is very robust in that it can be directly applied to convert every frame during video decolorization. In addition, it can deal with high resolution image in real time thus can better meet current demand.

References

- [1] G. Amy, O. Sven, T. Jack, and G. Bruce. Color2gray: salience-preserving color removal. *ACM Transactions on Graphics (SIGGRAPH 2005)*.
- [2] A. Andrew, G. Natasha, D. Jennifer, and L. Marc. Gaussian kd-trees for fast high-dimensional filtering. *ACM Transactions on Graphics (TOG)*, 2009.
- [3] W. Ben. Fast median and bilateral filtering. In *ACM Transactions on Graphics (SIGGRAPH 2006)*.
- [4] M. Čadík. Perceptual evaluation of color-to-grayscale image conversions. *Computer Graphics Forum*, 2008.
- [5] T. Carlo and M. Roberto. Bilateral filtering for gray and color images. In *CVPR 1998*.
- [6] L. Cheryl, H. Wolfgang, and M. Rafal. Cluster-based color space optimizations. In *ICCV 2011*.
- [7] E. Elmar and D. Frédo. Flash photography enhancement via intrinsic relighting. *ACM transactions on graphics (SIGGRAPH 2004)*.
- [8] P. Fatih. Constant time $O(1)$ bilateral filtering. In *CVPR 2008*.
- [9] D. Frédo and D. Julie. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics (SIGGRAPH 2002)*.
- [10] P. Georg, S. Richard, A. Maneesh, C. Michael, H. Hugues, and T. Kentaro. Digital photography with flash and no-flash image pairs. In *ACM transactions on graphics (SIGGRAPH 2004)*.
- [11] S. Kaleigh, L. Pierre-Edouard, T. Joëlle, and M. Karol. Apparent greyscale: A simple and fast conversion to perceptually accurate images and video. In *Computer Graphics Forum*. Wiley Online Library, 2008.




Figure 9. Robustness to temporal coherence. The proposed method performs favorably against the recent techniques [13, 14]. The results are best viewed on high-resolution displays.

- [12] Y. Kim, C. Jang, J. Demouth, and S. Lee. Robust color-to-gray via nonlinear global mapping. *ACM Transactions on Graphics (SIGGRAPH ASIA 2009)*.
- [13] C. Lu, L. Xu, and J. Jia. Contrast preserving decolorization. In *ICCP 2012*.
- [14] C. Lu, L. Xu, and J. Jia. Real-time contrast preserving decolorization. In *SIGGRAPH Asia Technical Briefs*. ACM, 2012.
- [15] G. Mark and D. N. A. Decolorize: Fast, contrast enhancing, color to grayscale conversion. *Pattern Recognition*, 2007.
- [16] A. C. Ormiana, A. Cosmin, and B. Philippe. Enhancing by saliency-guided decolorization. In *CVPR 2011*.
- [17] B. Raja and E. Reiner. Spatial color-to-grayscale transform preserving chrominance edge information. In *Color and Imaging Conference*, 2004.
- [18] M. Song, D. Tao, C. Chen, X. Li, and C. Chen. Color to gray: visual cue preservation. *PAMI*, 2010.
- [19] Y. Song, L. Bao, X. Xu, and Q. Yang. Decolorization: Is `rgb2gray()` out? In *SIGGRAPH Asia Technical Briefs*. ACM, 2013.
- [20] P. Sylvain and D. Frédo. A fast approximation of the bilateral filter using a signal processing approach. *IJCV 2009*.
- [21] Q. Yang. Recursive bilateral filtering. In *ECCV*, 2012.
- [22] Q. Yang. Hardware-efficient bilateral filtering for stereo matching. *PAMI*, 2013.
- [23] Q. Yang, K.-H. Tan, and N. Ahuja. Real-time o (1) bilateral filtering. In *CVPR 2009*.