REMOVING SHADOWS FROM IMAGES USING COLOR AND NEAR-INFRARED

N. Salamati, A. Germain, S. Süsstrunk

School of Computer and Communication Sciences,
École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
{neda.salamati,arthur.germain,sabine.susstrunk}@epfl.ch

ABSTRACT

Shadows often introduce errors in the performance of computer vision algorithms, such as object detection and tracking. This paper proposes a method to remove shadows from real images based on a probability shadow map. The probability shadow map identifies how much light is impinging on a surface. The lightness of shadowed regions in an image is increased and then the color of that part of the surface is corrected so that it matches the lit part of the surface. The result is compared with two other shadow removal frameworks. The advantage of our method is that after removal, the texture and all the details in the shadowed regions remain intact.

1. INTRODUCTION

Removing shadows from images can significantly improve and facilitate the performance of certain computer vision tasks, such as tracking, segmentation, and object detection, where shadow boundaries are often confused with those of different surfaces or objects. It is therefore of great importance to discover ways of properly detecting shadows and removing them while keeping other details of the original image intact.

A lot of research has been performed to detect shadows. Finlayson *et al.* [1] proposed the invariant image method, which requires the knowledge of the capture devices' characteristics to color calibrate the camera. Color calibration of the camera leads to 1-D co-ordinates, a function of image chromaticities that is invariant to illuminant color and intensity. Projecting a color image into the illumination invariant direction forms a gray-scale image, which is independent of the illumination condition. Afterwards, shadow edges are detected by finding edges in the intensity image that are not in the illumination invariant image.

Levine and Bhattacharyya [2] proposed to manually train a support vector machine to segment an image into shadow and non-shadow regions. After validating the classifier, the shadowed regions are found. In this paper, we use "LB" to refer to Levine and Bhattacharyya's method.

Assuming that the histogram of the illumination is sparse, Weiss in [3] proposes to acquire a sequence of images in which the shadow edges move, i.e., when just illumination changes. The median of this sequence is calculated for each pixel and amounts to the maximum-likelihood estimation of the reflectance only image. Although his method results in very natural looking images, one of the shortcomings of this method is that it is not always practical to acquire a sequence of images where all the objects and surfaces do not move and just changing the illumination generates moving shadow edges.

(A) (B) (C)

Fig. 1. Column (A) is the color image. Column (B) shows the NIR image of the scene and column (C) is their corresponding shadow maps.

Fredembach and Süsstrunk [4] presented a simple though accurate shadow detection method (FS method) that employs the inherent sensitivity of digital camera sensors to the near-infrared (NIR) part of the spectrum. Incorporating the features offered by the NIR band along with color information has recently drawn the attention of researchers in many imaging applications, such as illumination estimation [5], dark flash photography [6], video conferencing [7], object segmentation [8], and scene recognition [9]. Using the property of the NIR band, in which most of the colorants are transparent or have higher reflectance [10], the authors showed that combining the dark map of both visible and NIR images with ratios of the color channels (red, green and blue) to NIR identifies the pixels that are shadow candidates. Fig.1 shows the shadow map of two images.

As it can be seen in Fig.1, the results are accurate in real and complex scenes, including regions that are partially lit (penumbra), and not lit at all (umbra).

After identifying shadows in an image, several methods have been proposed to remove these shadows. One possible way is to form the gradient image and remove the gradient information at the shadow edges. Therefore, the shadow-free image can be formed by solving a Poisson's equation [1,11]. One can also replace the pixel color within the shadow regions with the average color found at the boundaries [2].

In this paper, we propose a new approach to remove the shadows that are detected when first applying the FS shadow detection method [4]. We form a probability map that gives us the information on how much shadow information each pixel contains. Then, we present a method to lighten up the shadowed regions in order to obtain a shadow free image, which is natural looking and keeps all the details and textures intact. The results are compared to the shadow removal frameworks by Fredembach and Finlayson (FF) [11] and Levine and Bhattacharyya (LB) [2].

2. USING NIR INFORMATION IN DETECTING SHADOWS

Prior to removing the shadows, we apply the FS shadow detection framework [4], where it has been shown that NIR wavelengths (700-1100 nm) can be imaged and used along with a color image to accurately detect shadows in the image. To this end, we capture NIR and color images from a scene (I_{VIS} and I_{NIR}) [12] and formulate a "dark map" and "ratio map" to find shadow-candidate pixels.

First, a joint dark map of color and NIR images D is formed that is the multiplication of the visible and NIR dark maps (D_{VIS} and D_{NIR}):

$$D_{VIS} = 1 - B; \;\; D_{NIR} = 1 - I_{NIR} \tag{1}$$

$$D = f(D_{VIS})f(D_{NIR}) \tag{2}$$

where B is the brightness of the visible image, which is calculated as the pixels' norm in an RGB cube, and I_{NIR} is the intensity of the NIR image. $f(\cdot)$ is an ascending function that compresses the shadows (larger values) [4]. The argument is that shadow pixels are to be found in the darker regions of the image. Thus, an "AND" operator on the visible and NIR dark maps identifies the pixels in the image that are dark in the two representations of the scene.

In the second step, the ratio image F is formed to be combined with the dark map and generates a final shadow map M.

$$M = DF \tag{3}$$

The physical property of the NIR band (namely, higher reflectance of different materials as well as distinct behavior of most illuminants in the NIR part of the spectrum) set the dark pixels that correspond to dark objects apart from the shadowed pixels in M.

The larger the pixel values in M, the more probable they are to be under the shadow. In [4], it is proposed that the location of the first valley t_1 in the histogram of M specifies the threshold value to generate the final shadow mask S.

$$S(x,y) = \begin{cases} M(x,y) & \text{if } M(x,y) \geq t_1 \\ 0 & \text{if } elsewhere \end{cases} \tag{4}$$

We refer the reader to [4] for more details about the method.

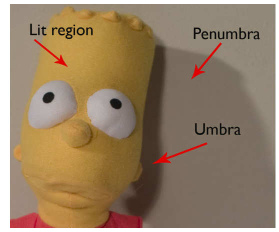

Fig. 2. Umbra and penumbra. A non-point light source will produce three distinct lighting areas; lit regions, partially lit (penumbra), and not lit at all (umbra).

3. SHADOW REMOVAL FRAMEWORK

A non-point-light source creates three distinct illuminated areas in a scene, namely non-lit areas (umbra), partially lit areas (penumbra), and completely lit areas. In detecting the shadowed regions, S is formulated in a way that regions that are not lit (shadowing object blocks all the light from the source) have larger values. Pixels in penumbra regions, as the intensity of light is increasing, have smaller values in S, and the S value in the lit regions is 0. Consequently, S values can be seen as a shadow probability map that gives the probability of each pixel to be under the shadow ($S(x \in \text{umbra}) > S(x \in \text{penumbra}) > S(x \in \text{lit regions})$). Fig. 2 shows the different shadow regions in the visible image of a scene. We divide the shadow regions in S ($S \neq 0$) into the penumbra and umbra according to their pixel values. Analyzing the histogram of S, we can observe two peaks, which correspond to the umbra and penumbra regions. The position of the valley between these peaks (t_2) gives the threshold to segment S into umbra, penumbra, and non-shadow regions.

$$S(x,y) \in \begin{cases} \text{Umbra} & \text{if } S(x,y) \geq t_2 \\ \text{Penumbra} & \text{if } t_1 \leq S(x,y) < t_2 \\ \text{non-shadow} & \text{if } S(x,y) = 0 \end{cases} \tag{5}$$

Removing shadows can be performed by lightening umbra and penumbra regions. Since shadowed regions do not have a constant intensity (the intensity gradually increases from shadow to light), we propose to increase the lightness of the pixels according to the shadow probability map S.

$$L(x,y) = \begin{cases} S(x,y)g_1(S(x,y)) & \text{if } S(x,y) \in \text{umbra} \\ S(x,y)g_2(S(x,y)) & \text{if } S(x,y) \in \text{puenumbra} \\ S(x,y) & \text{if } S(x,y) \in \text{non-shadow} \end{cases} \tag{6}$$

where L is the lightness value of I_{VIS} in CIELab color space and $g(\cdot)$ is a function used to increase the intensity of each pixels under the shadow given the corresponding shadow probability value in S.

$$\forall s \in S \qquad \begin{array}{l} g_1(s) = se^{s^m} \\ g_2(s) = se^{s^n} \end{array} \tag{7}$$

m and n values are parameters that are empirically set at 2 and 4, respectively, for our image dataset.

Fig. 3 illustrates an image where the lightness is corrected applying our method. Increasing the lightness of the shadowed

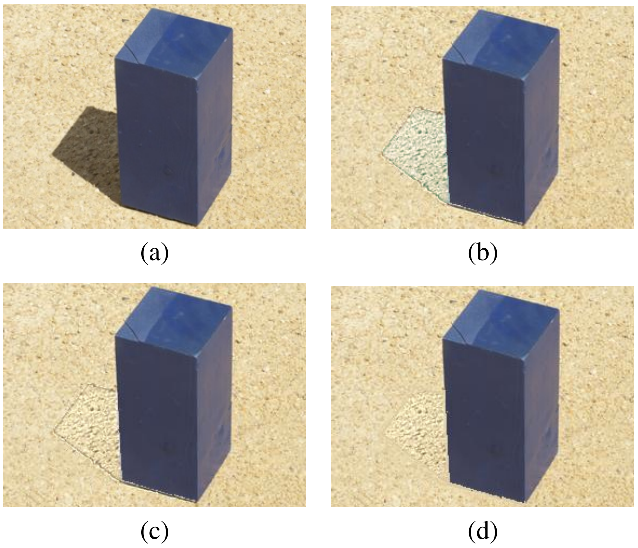

Fig. 3. (a): Original image. (b): Lightness corrected. (c): Color corrected. (d): Borders corrected.

parts using the proposed method keeps the details of the shadowed surfaces intact. As we can see in Fig. 3, although the lightness of the shadowed parts are corrected, some shadowed surfaces in the image still do not look similar to the lit parts. The reason is that given two pixels located on the sides of a shadow boundary with the same surface reflectance, the ratios of the two pixels are not the same in all three color channels because of the ambient light. These two pixels will be different not only in intensity, but also in hue and saturation. Thus, correcting just the intensity of the shadowed pixels does not remove the shadow and we need to correct the chromaticity values as well. In this paper, we refer to a^* and b^* values in CIELAB color space as the chromaticity attributes.

Applying a mean shift algorithm we segment the entire image according to its color values. We start with the penumbra pixels in the color image, having the hypothesis that segments in penumbra are certainly adjacent to a non-shadow region. For each segment in penumbra P we consider all its neighbor segments LIT_j in the lit part of the image ($LIT_j \in$ lit part & $Di(P) \cap LIT_j \neq \emptyset$, where $Di(\cdot)$ is a dilation function). Among all the neighbor segments, we choose the one that is closest in chromaticity to our segment of interest. This segment is further denoted as LIT.

Afterwards, we rescale the shadowed segment's chromaticity values so that the average of the chromaticity in that segment P matches the average of the chromaticity in the aforesaid lit segment LIT.

$$\begin{aligned} a^*_P &:= a^*_P \frac{<a^*_L>}{<a^*_P>} \\ b^*_P &:= b^*_P \frac{<b^*_L>}{<b^*_P>} \end{aligned} \tag{8}$$

where a^*_P and b^*_P are the a^* and b^* attributes of the corresponding segment and $< \cdot >$ is the average operator. The chromaticity correction is valid for the surfaces that are partly lit and partly under shadow. For such regions, the chromaticity can be corrected so that the shadow part of the surface will have the same chromaticity as the lit part of the surface. Nevertheless, there might exist some surfaces that are completely under shadow. Changing the chromaticity values of such surfaces to the closest adjacent lit segment will introduce false colors. To prevent this effect, if the chromaticity difference between P and LIT is not small enough, the chromaticity value of segment P will not be changed.

After "correcting" the color of penumbra regions, we continue rescaling the chromaticity values of umbra regions by applying the same framework.

Finally, all boundaries between shadowed regions and neighboring lit regions are smoothed by convolving them with a Gaussian mask. Thus, we introduce a uniform transition between shadowed regions that were lightened and neighboring non-shadowed regions.

4. RESULT AND DISCUSSION

Fig. 4 shows the result of applying our framework on some images. The results are compared to the state of the art shadow removal frameworks in [2,11]. The two different methods are applied on the images given the FS shadow map [4]. The first algorithm is LB [2] that assigns the average color of the lit region to the adjacent shadowed region. The second framework (FF) [11] removes the shadow edges from the gradient image and re-integrates the gradient image in all three color channels to derive a shadow free image [11]. As it can be seen in Fig. 4, with the LB framework [2] all high frequency details from the shadow regions are removed. The FF results look acceptable but in some cases the overall color is different. This is because of the Neumann boundary conditions in solving the Poisson's equation. Additionally, the portion of the information erased at the shadow edges in the gradient image can result in a smudge effect in the re-integrated image. To remove the gradient information on shadow edges, the assumption is that the reflectance on the shadow edges does not change, which is not correct in textured surfaces.

One of the main advantages of our proposed method though is that the texture of the surface that was under the shadow is preserved to a good extend and no harsh transition between the shadowed parts and non-shadowed parts can be seen. The shortcoming of our method is that the result is dependent on the values of the g_1 and g_2 parameters. The proposed values for these parameters do not always give the best result (see last image in Fig. 4 for illustration).

5. CONCLUSION

This paper describes a shadow removal method for real images based on the shadow map proposed by Fredembach and Süsstrunk [4]. We increased the lightness of shadowed regions in an image knowing the probability shadow map. The color of that part of the surface is then corrected so that it matches the lit part of the surface. Our algorithm worked successfully in removing both umbra and penumbra shadows.

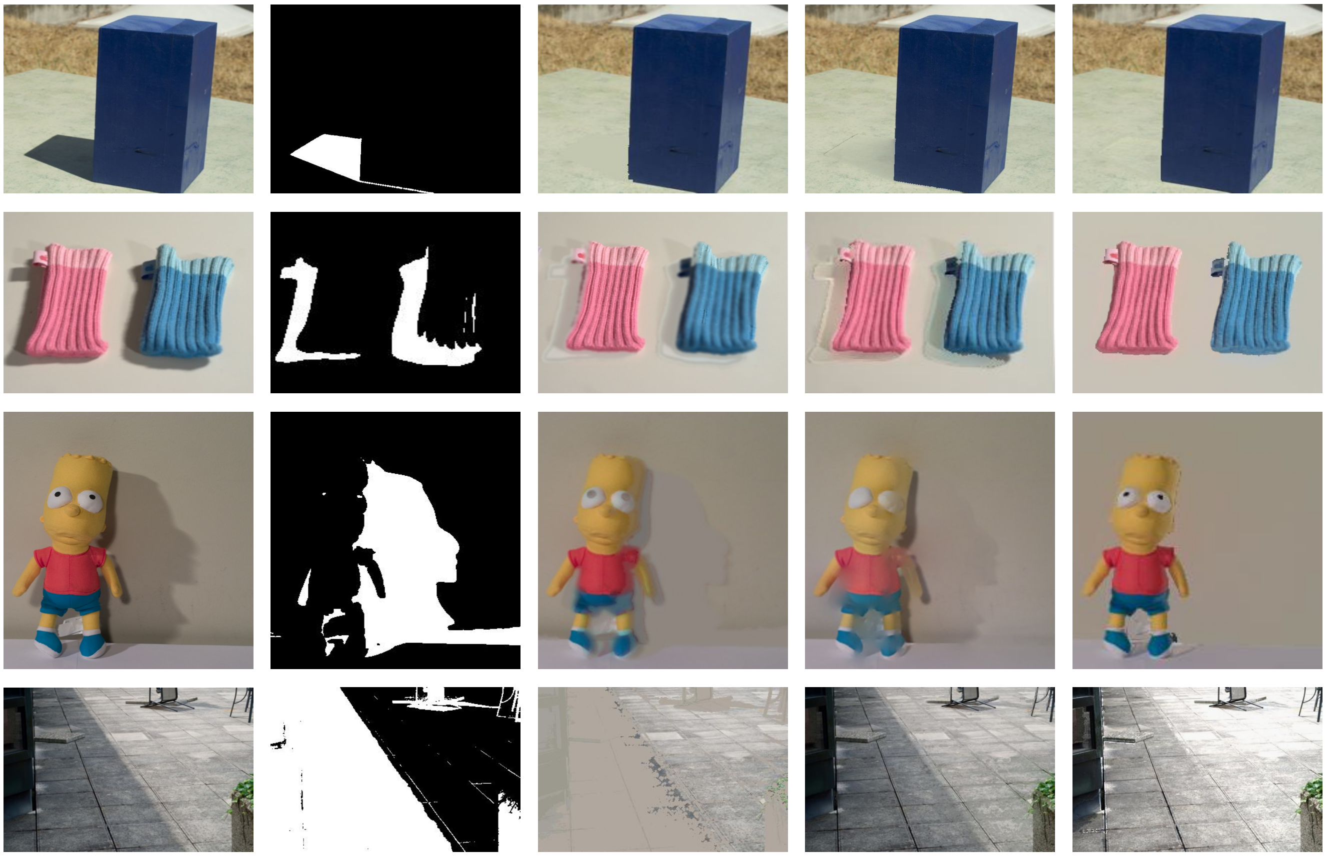

(A) (B) (C) (D) (E)

Fig. 4. Column (A) is the original image, (B) shows the shadow map, (C) shows the results with LB [2], (D) is FF results [11], and (E) shows the result with the algorithm proposed in this paper. We can see that the solution we propose preserves not only the colors, but also the textures of the lightened parts.

6. ACKNOWLEDGMENT

The work presented in this paper was supported by the Swiss National Science Foundation under grant number 200021-124796/1.

7. REFERENCES

[1] G. D. Finlayson, S. D. Hordley, and M. S. Drew, Removing shadows from images, *In Proc. of 7th European Conference on Computer Vision* London, 2002.

[2] M. D. Levine and J. Bhattacharyya, Removing shadows, *Pattern Recognition Letters* vol. 26, no. 3, pp. 307-312, 2005.

[3] Y. Weiss, Deriving intrinsic images from image sequences, *In Proc. of International Conference on Computer Vision* 2001.

[4] C. Fredembach and S. Süsstrunk, Automatic and accurate shadow detection from (potentially) a single image using near-infrared information, *EPFL Technical Report 165527* 2010.

[5] C. Fredembach and S. Süsstrunk, Illuminant estimation and detection using near infrared, *In Proc. of the IS& T/SID Electronic Imaging, Digital Photography V* vol. 7250, 2009.

[6] Krishnan and R. Fergus, Dark flash photography, *ACM Transactions on Graphics, SIGGRAPH Conference Proceedings* 2009.

[7] P. Gunawardane, T. Malzbender, R. Samadani, A. McReynolds, D. Gelb, and J. Davis, Invisible light : Using infrared for video conference relighting *In Proc. of IEEE International Conference on Image Processing* pp. 40054008, 2010.

[8] N. Salamati and S. Süsstrunk, Material-based object segmentation using near-infrared information, *In Proc. of the IS& T/SID Color Imaging Conference* 2010.

[9] M. Brown and S. Süsstrunk, Multi-spectral SIFT for scene category recognition, *In Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR)* 2011.

[10] D.A. Burns and E.W. Ciurczak, *Handbook of near-infrared analysis,* Marceld Ekkerin Inc., 2001.

[11] C. Fredembach and G. D. Finlayson, Simple shadow removal, *In Proc. of 18th International Conference on Pattern Recognition* vol. 1, pp. 137-142, 2007.

[12] C. Fredembach and S. Süsstrunk, Colouring the near infrared, *In Proc. of the IS& T/SID Color Imaging Conference* Portland, 2008.