

Rainbow Flash Camera:Depth Edge Extraction Using Complementary Colors

Taguchi, Y.

TR2012-078 October 2012

Abstract

We present a novel color multiplexing method for extracting depth edges in a scene. It has been shown that casting shadows from different light positions provides a simple yet robust cue for extracting depth edges. Instead of flashing a single light source at a time as in conventional methods, our method flashes all light sources simultaneously to reduce the number of captured images. We use a ring light source around a camera and arrange colors on the ring such that the colors form a hue circle. Because complementary colors are arranged at any position and its antipole on the ring, shadow regions where a half of the hue circle is occluded are colored according to the orientations of depth edges, while non-shadow regions where all the hues are mixed have a neutral color in the captured image. In an ideal situation, the colored shadows in a single image directly provide depth edges and their orientations. In practice, we present a robust depth edge extraction algorithm using an additional image captured by rotating the hue circle with 180 degrees. We demonstrate the advantages of our approach using a camera prototype consisting of a standard camera and 8 color LEDs.

European Conference on Computer Vision (ECCV)

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Rainbow Flash Camera: Depth Edge Extraction Using Complementary Colors

Yuichi Taguchi

Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA

Abstract. We present a novel color multiplexing method for extracting depth edges in a scene. It has been shown that casting shadows from different light positions provides a simple yet robust cue for extracting depth edges. Instead of flashing a single light source at a time as in conventional methods, our method flashes all light sources simultaneously to reduce the number of captured images. We use a ring light source around a camera and arrange colors on the ring such that the colors form a hue circle. Because complementary colors are arranged at any position and its antipole on the ring, shadow regions where a half of the hue circle is occluded are colorized according to the orientations of depth edges, while non-shadow regions where all the hues are mixed have a neutral color in the captured image. In an ideal situation, the colored shadows in a single image directly provide depth edges and their orientations. In practice, we present a robust depth edge extraction algorithm using an additional image captured by rotating the hue circle with 180° . We demonstrate the advantages of our approach using a camera prototype consisting of a standard camera and 8 color LEDs.

Key words: Multi-flash camera, depth edge extraction, color multiplexing, complementary color, hue circle

1 Introduction

Depth edges (discontinuities) in a scene play a key role in various computer vision applications such as segmentation, object detection, and pose estimation. Using a single 2D image, it is difficult to distinguish depth edges from intensity edges. To robustly detect depth edges, Raskar et al. [1] presented an active illumination camera, called multi-flash camera (MFC), that captures multiple images by casting shadows from different light sources and analyzes the shadow profile in the captured images. In the most basic approach, one light source is turned on at a single capture time.

We present a method that multiplexes this process to reduce the number of captured images. Instead of flashing each light source at a time, we simultaneously flash all light sources. Our method is based on the color theory on a hue circle: We exploit the complementary nature of a hue circle to *colorize* shadow regions with different hues depending on the orientations of depth edges, while letting the other non-shadow regions have a neutral (white) color. Our method

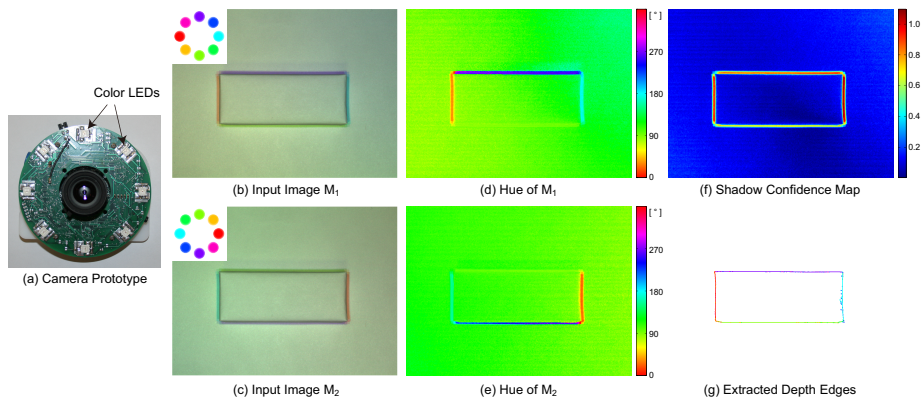


Fig. 1. Overview of our method. (a) Our camera prototype consists of a standard camera and 8 color LEDs. (b, c) Images captured by simultaneously flashing the LEDs with different hues (shown in the insets). (d, e) Hue components of (b) and (c). Note that shadow regions are colorized with the different hues, while non-shadow regions have a neutral color. The colors of the shadow regions directly provide the orientations of the corresponding depth edges. (f) Shadow confidence map, where each pixel value corresponds to the distance between the pixels in the input images on the hue-saturation plane. (g) Extracted depth edges. Each depth edge pixel is depicted with the color of the LED that contributed the most to the depth edge, indicating the edge orientation.

produces colored shadow regions in a single image, which is sufficient to determine depth edges and their orientations in an ideal situation. In practice, to robustly detect depth edges independent of object colors and ambient lights, we propose an approach that uses two images, one captured with a hue circle and the other captured by rotating the hue circle with 180° (i.e., the complementary version of the original hue circle). We present a camera prototype (Figure 1(a)) consisting of a standard camera and 8 color LEDs and compare the performance of our method with that of conventional methods.

1.1 Contributions

This paper makes the following contributions:

- We present a novel multiplexing method for multi-flash-based depth edge extraction that exploits the complementary nature of a hue circle.
- We describe a robust depth edge extraction algorithm using two images captured with a hue circle and its complementary version.
- We demonstrate a camera prototype using 8 color LEDs and validate our method under various conditions.

1.2 Related Work

Depth Edge Extraction: Raskar et al. [1] introduced a multi-flash camera (MFC) for extracting depth edges by casting shadows from different light posi-

tions and applied it to non-photorealistic rendering. Depth edges obtained using an MFC have been also used for silhouette-based surface reconstruction [2], depth-edge-preserving stereo [3], and object detection and pose estimation in robotic applications [4–6]. The above works captured multiple images by using one flash at a time.

Feris et al. [7] presented a color multiplexed MFC using three red, green, and blue light sources. They proposed a single-shot method based on learning shadow color transitions for specific scenes, and a two-shot method using an additional reference image captured with white light sources for general scenes. Because their method encodes shadows from the three light sources separately into each RGB channel of the camera, it is only applicable to three light source positions. We present a novel color multiplexing method built on the color theory on a continuous hue circle, thus applicable to any number of light sources that approximate the hue circle. In [8, 9], a single-shot method was presented based on frequency multiplexing. Their method projects multiple sinusoidal patterns such that their frequencies are maintained independent of the scene geometry, and then performs frequency demultiplexing to detect shadow regions [8] or to recover individually illuminated images [9]. Their method requires multiple projectors as light sources, while our method uses multiple color LEDs, resulting in a simpler and more inexpensive system. Their method also sacrifices spatial resolution due to frequency computation in local neighborhoods, while our method enables pixel-wise depth edge extraction as in conventional MFC approaches.

Illumination multiplexing has been used for various active illumination applications, including photometric stereo [10–13], structured light [12], image-based relighting [14, 15, 12], object material segmentation [15], as well as depth edge extraction [7]. The multiplexing was done using standard three (RGB) color channels [7, 10, 11] or more channels with multispectral illumination [15, 13] to reduce the number of captured images. Schechner et al. [14] used multiplexing to improve the signal-to-noise ratio using the same number of captured images. De Decker et al. [12] used both color and time multiplexing and showed that at least $\frac{n+2}{3}$ images are required to demultiplex n light sources. The goal of these methods is to demultiplex captured images to obtain multiple images as if the scene were lit by individual light sources. In contrast, our method directly uses the captured images to obtain depth edges *without* demultiplexing. Our method does not have the limitation on the number of light sources; the more discrete light sources we have, the better we can approximate a continuous hue circle.

Complementary colors were used to obtain object material colors in active illumination systems [16, 12]. They captured two consecutive frames by using complementary colors for each projector pixel [16] or each light source [12], and then added the two frames to simulate the case as if the scene were lit by white light sources. Minomo et al. [17] presented a system that projects two images from two projectors such that they have pixel-wise complementary colors on a plane. The system was used for artistic visualization to colorize the shadows of people interacting with the system. We present a novel application of complementary colors for depth edge extraction.




Fig. 2. Depth edge extraction in 2D using two light sources. (a) A conventional MFC flashes a single light source at a time and extracts shadow regions that appear black from each image. (b) Our method simultaneously flashes the light sources having complementary colors. Shadow regions have colors corresponding to the light sources, while non-shadow regions have a neutral (white) color due to the mixture of complementary colors.




Fig. 3. (Left) Depth edge extraction in 3D using a ring light source having a continuous hue circle. At a depth edge, a half of the hue circle is occluded and the hues in the other half are integrated in the shadow region. This produces colored shadow, whose hue corresponds to the average of the integrated hues. The shadow color directly provides the orientation of the depth edge. On the other hand, all the hues are mixed in non-shadow regions, producing a natural (white) color. (Right) Robust depth edge extraction using two images. We capture the two images by using a hue circle H and its complementary version \bar{H} and compute the distance between the shadow colors on the hue-saturation plane to obtain a pixel-wise shadow confidence map.

2 Depth Edge Extraction Using Complementary Colors

This section describes the principle of our depth edge extraction method using complementary colors. In this paper, we refer to image regions that are occluded from at least one of discrete light sources or a part of a continuous light source as *shadow* regions, while those lit by all discrete/continuous light sources as *non-shadow* regions.

2.1 Principle in 2D

Let us first explain the principle of depth edge extraction in 2D using two light sources. For now we suppose that there is no ambient light and the scene has a neutral (white) color. We also assume that the scene is Lambertian, and the

light sources are located sufficiently far from the scene, illuminating the scene uniformly. As shown in Figure 2(a), a conventional method [1] flashes a single white light source at a time, captures two images, and extracts depth edges as boundaries between shadow (black) and non-shadow (white) regions. There are two such boundaries in this scene, corresponding to a depth edge (depth discontinuity) and a shadow edge (boundary between shadow and non-shadow regions on the same plane). These edges can be distinguished by finding white-to-black transitions in the captured image along the epipolar line defined by the light source position [1].

Our method uses complementary colors (red and cyan in Figure 2(b)) for the two light sources and captures a single image by flashing the two light sources at the same time. Because the mixture of complementary colors in an additive color system results in white, non-shadow regions, lit by the two light sources, appear white in the captured image. In contrast, shadow regions, lit by only a single light source due to occlusion, are colorized with the colors of the light sources. Note that the shadow color indicates the orientation of a depth edge.

2.2 Principle in 3D

Here we extend our discussion in 3D. A conventional MFC places N (typically 2, 4, or 8) white light sources around a camera, captures N images by flashing a single light source at a time, and extracts depth edges corresponding to N orientations from each of the captured image. As we use more light source positions, we can distinguish more depth edge orientations in the scene, but we need to take more images.

We build our theory on a continuous light source, instead of discrete light sources. Consider a ring light source around a camera that realizes a continuous, maximally saturated hue circle with the same brightness, as shown in Figure 3 (left). We capture a single image by flashing the entire light source including different hues. In the shadow region corresponding to a straight depth edge of an object¹, a half of the hue circle is occluded by the object, while the other half is integrated in the shadow region.

The additive mixture of any number of colors is determined as the weighted average of the positions of the original colors on the hue-saturation plane, referred to as Newton’s geometrical weighting [18]. The brightness of each original color corresponds to the weight. Thus, in our case, the mixture of colors in the half of the hue circle results in a color that has the hue at the center of the half circle and is slightly unsaturated (called saturation cost [18]). As a consequence, shadow regions in the single captured image are colorized according to the orientations of the corresponding depth edges. In non-shadow regions, all colors in the hue circle are mixed, producing a white color.

Note that in an ideal situation, our method can detect shadow regions and corresponding depth edges, as well as distinguish an infinite number of depth edge orientations, from a single captured image by simply looking up the hues

¹ For curved depth edges, we consider the tangent line at each depth edge point.




Fig. 4. Colors captured by the camera against the duty cycle for each RGB sub-LED.

in the image. In contrast, conventional methods require more images to detect and distinguish more depth edge orientations².

In practice, however, object colors and ambient lights are typically unknown, which makes the single-shot method unstable. In the next section, we present a practical camera implementation and a robust depth edge extraction algorithm that uses two images, one captured using a hue circle and the other captured by rotating the hue circle with 180° .

3 Camera Prototype

Figure 1(a) shows our camera prototype, consisting of a standard camera (Point Grey Dragonfly) surrounded by 8 color LEDs. Each color LED is composed of three sub-LEDs of RGB components. We control the brightness of the individual RGB components by software pulse-width modulation (PWM) using a PIC microcontroller. All LEDs are synchronized with the camera via a trigger signal sent from the camera to the PIC.

Color Calibration: We measured response functions for the RGB components by changing the duty cycle of the PWM signal. We assumed that the 8 RGB components have the same response function for each color channel. We captured multiple images by flashing a single R, G, or B component of all the 8 color LEDs simultaneously with different duty cycles. Figure 4 shows the mean observed color in the captured image at each duty cycle for each RGB component. Our prototype has approximately linear response functions with some color cross talks. To obtain a desired color (r_0, g_0, b_0) for each color LED, we fit linear functions to the responses and computed the duty cycles required for each RGB component by solving a linear least-squares problem.

3.1 Robust Depth Edge Extraction Using Two Images

As described earlier, only using hues in a single image for depth edge extraction would be unstable due to object colors, ambient lights, and noise. We propose a robust approach that captures two images M_1 and M_2 , one by using a hue circle

² Vaquero et al. [19] showed that three light positions are sufficient to cast shadows for all depth edges in general. However, it only provides three depth edge orientations.

H as described in Section 2.2, and the other by rotating the hue circle with 180° , i.e., using the complementary version \bar{H} of the original hue circle H . As shown in Figure 3 (right), the color of a shadow region obtained with H and the color of the same shadow region obtained with \bar{H} become complementary; i.e., the distance between the two colors on the hue-saturation plane is maximized. We use this distance as a confidence value to detect shadow regions.

Note that instead of capturing the second image with \bar{H} , one could capture an ambient image (no light source on) or an image with making the colors of all light sources white [7]. However, they produce neutral color illuminations, corresponding to the center of the hue-saturation plane. Thus the distance measure becomes less reliable than our method. Any other rotation of the original hue circle H also produces a less reliable distance measure, because it makes the angle between the two shadow colors less than 180° , while the saturations of the shadow colors are always the same due to the integration of a half circle.

Detailed Steps: For the prototype, we used a discrete set of hues $H_d = \{45^\circ \times i \mid i = 0, \dots, 7\}$ and its complementary set $\bar{H}_d = \{180^\circ + 45^\circ \times i \pmod{360^\circ} \mid i = 0, \dots, 7\}$ with the maximum saturation. The same brightness (value) was set for the 8 HSV colors such that the maximum RGB values stay within the maximum duty cycle obtained in the color calibration process. Our depth edge extraction algorithm is summarized as follows:

1. Capture two images M_1 and M_2 with the hue sets H_d and \bar{H}_d , respectively.
2. Convert M_1 and M_2 into the HSV color space, resulting in M_1^{HSV} and M_2^{HSV} .
3. For each pixel, compute the distance between M_1^{HSV} and M_2^{HSV} on the hue-saturation plane to produce a *shadow confidence map*.
4. In the shadow confidence map, for each of the 8 light sources,
 - (a) Traverse pixels along the epipolar line corresponding to the light position.
 - (b) Find step edges with negative transitions.
 - (c) For each edge pixel \mathbf{x} , if the hues of $M_1^{\text{HSV}}(\mathbf{x})$ and $M_2^{\text{HSV}}(\mathbf{x})$ are within a threshold of the hues of the target light source, mark the pixel \mathbf{x} as a depth edge candidate.
5. For the depth edge candidates, perform connected-component analysis and hysteresis thresholding, similar to the Canny edge detector [20].

For the above Step 4, we used a 3×3 Sobel filter whose direction is aligned with the illumination direction. We also used non-maximum suppression for each direction to obtain thin edges that better localize the depth discontinuities. Step 4c is required because the negative transitions could be a shadow edge of an antipodal light source. See Figure 1(f)—If we traverse the center row of the image from left to right to find depth edges for the red light source, there are two negative transitions; one of them is a depth edge for the red light source, while the other is a shadow edge for the cyan light source.

For visualization purpose, we color-coded our depth edge results with the color of the light source that provided the maximum Sobel filter response. Thus the color indicates the depth edge orientation. Note that the conventional method also extracts depth edges with N orientations; however, we depicted them with black color without showing the orientations in the results shown in this paper. We also dilated the extracted depth edges by one pixel for better visibility.

4 Experiments

We evaluated our method under various conditions by comparing it with the conventional MFC method [1] and Feris et al.’s RGB multiplexing method [7].

4.1 Comparisons with the Conventional Method

To obtain conventional MFC images, we turned all RGB components of a color LED on to produce a white light source and captured N flash images, I_1, \dots, I_N . The conventional method also requires one ambient image, I_{N+1} , to compensate for object colors and ambient lights. To obtain depth edges, we performed the following steps [1]:

1. Subtract the ambient image from flash images: $I_i^s = I_i - I_{N+1}$ ($i = 1, \dots, N$).
2. Compute a pixel-wise maximum image, $I_{\max}(\mathbf{x}) = \max_i(I_i^s(\mathbf{x}))$, to generate a *shadow-free* image.
3. Create a ratio image for each flash position: $R_i(\mathbf{x}) = I_i^s(\mathbf{x})/I_{\max}(\mathbf{x})$.

The ratio images R_i act as shadow confidence maps for each light position in the conventional method. We generated conventional depth edges by applying the same edge detection algorithm used in our method (Steps 4 and 5), except the color check process (Step 4c), to the ratio images.

Number of Flashes: Figure 5 shows the effect of using different numbers of flashes for depth edge extraction. For our method, we used a reduced set of color LEDs to capture the input images, as shown in the insets of Figure 5. For the conventional method, we used the same reduced set of the LED positions as our method. To extract depth edges using a reduced set of flashes, we made the number of directions used for the epipolar line traversal (Step 4) the same as the number of flashes used. Both methods extract a larger number of depth edges when we use more flashes. Note that the conventional method requires $N + 1$ images to use N light sources, while our method always uses 2 images independent of the number of light sources. Using a larger number of color LEDs, our method can better approximate a continuous hue circle, which helps obtain more depth edge orientations using the same number of captured images.

Scene Textures: Figure 6 compares results for highly textured scenes. Both our method and the conventional method extract depth edges in the presence of complex object colors and intensity edges. Standard intensity-based edge detectors [20] would extract both depth edges and intensity edges from such scenes. Note that our method requires shadow regions to reflect all the colors of the light sources; otherwise it will miss depth edges corresponding to orientations of non-reflected colors. In contrast, the conventional method requires shadow regions to reflect some colors (spectrums) of white light sources; thus it is more robust to object colors.

Strong specular highlights (Figure 7) violate our assumption that all the hues are mixed in non-shadow regions, resulting in a neutral color. They also violate the assumption in the conventional method that non-shadow regions have

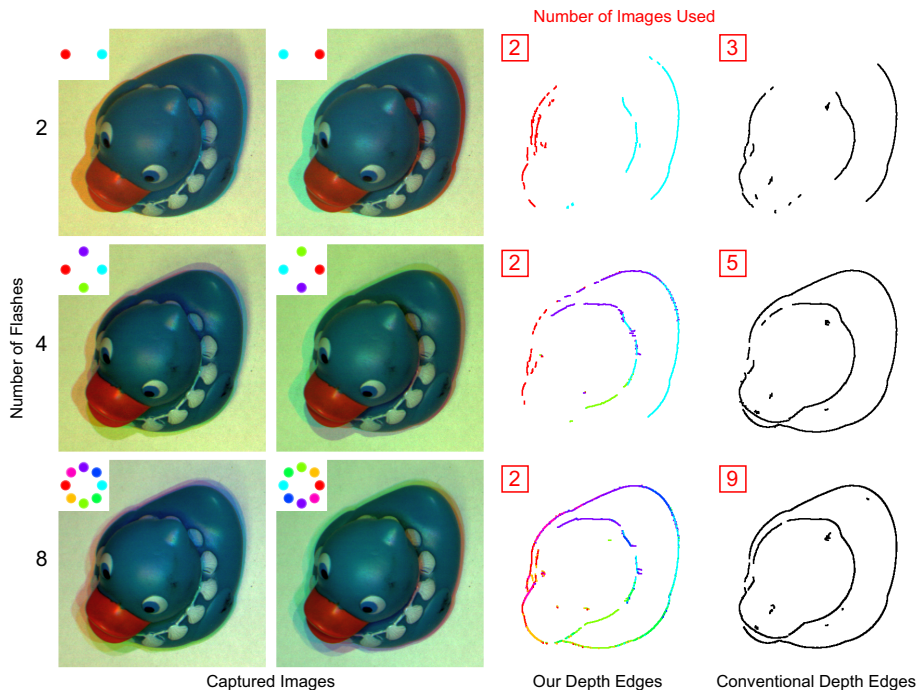


Fig. 5. Depth edge extraction using different numbers (2, 4, and 8) of flashes. For our method, we used a reduced set of color LEDs as shown in the insets to capture the two images. Note that the shadow regions in the captured images include more colors as we increase the number of color LEDs. For the conventional method, we used the same reduced set of the LED positions as our method. As we increase the number of flashes, both methods extract more depth edges at different orientations. Note that the conventional method requires $N + 1$ images, while our method always uses 2 images independent of the number of flashes.

the same appearance in all the flash images. The conventional method, however, can exploit the fact that specular highlights change their positions in the flash images, since each image is captured with a single light source at a different position. In [1, 21], instead of using a pixel-wise maximum value to compute the shadow-free image, the median intensity or median of gradients at each pixel was used to remove the effect of specular highlights. Here we simply used the median intensity to compute the shadow-free image and obtained the depth edges shown in Figure 7 for the conventional method, which do not include artifacts from specular highlights. On the other hand, our method uses all the flashes simultaneously. This produces specular highlights at fixed positions in the two captured images, causing false depth edges.

Semi-Specular Surfaces: Our method as well as the conventional method can tolerate semi-specular surfaces. Figure 8 shows an example. Note that the hues in non-shadow regions (object surfaces) change in the two captured images

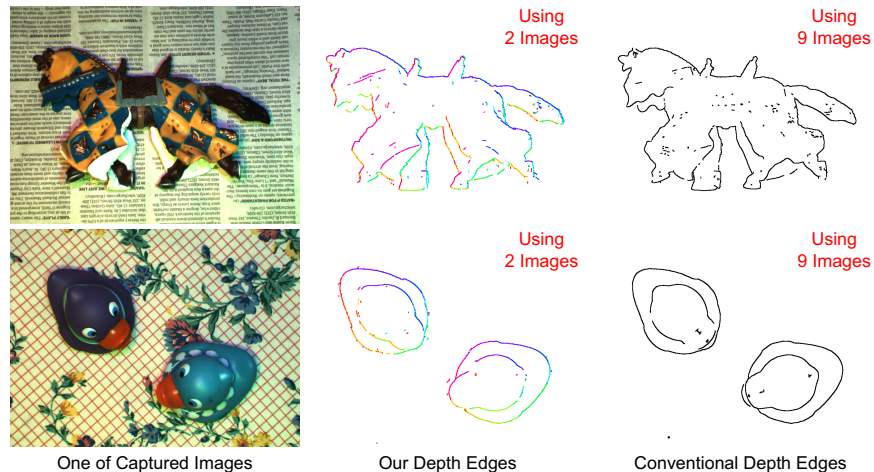


Fig. 6. Results for highly textured scenes. Both our method and the conventional method can robustly extract depth edges for scenes including many intensity edges.

because of specular reflections. Such object surfaces also increase the dynamic range of the scene: Depending on the normals of the surfaces, some regions reflect a large amount of light and become bright, while other regions remain dark. Both of these too bright or dark regions lead to low-saturation colors that are close to the center of the hue-saturation plane, where the estimation of the hue becomes unstable. As shown in Figure 8(c), only using the absolute difference of the hues as a shadow confidence map would be unstable for such regions. Our shadow confidence map (Figure 8(d)) defined by the distance on the hue-saturation plane provides a more robust measure. Our depth edges (Figure 8(f)) are comparable to the conventional depth edges (Figure 8(e)). We applied an object detection and pose estimation algorithm using the fast directional chamfer matching (FDCM) [5, 6] to our depth edges. Figure 8(g) demonstrates that we can correctly estimate the poses of objects using our depth edges.

Processing time of our method was almost the same as that of the conventional method, since both methods use similar building blocks for depth edge extraction. To extract depth edges using 8 light sources from 1024×768 pixel images, both methods took about 200 milliseconds using a C++ implementation on a standard PC with an Intel Core i7-950 processor.

4.2 Comparisons with the RGB Multiplexing Method

Feris et al. [7] captured two images for general scenes: one with three red, green, and blue light sources on and the other by replacing them with white light sources. We realized it by assigning the colors to three LEDs at the positions of $\{0^\circ, 135^\circ, 225^\circ\}$ as shown in the insets of Figure 9. We computed a shadow confidence map by taking the ratio between each channel of the two intensity-normalized images as described in [7], and generated depth edges by applying

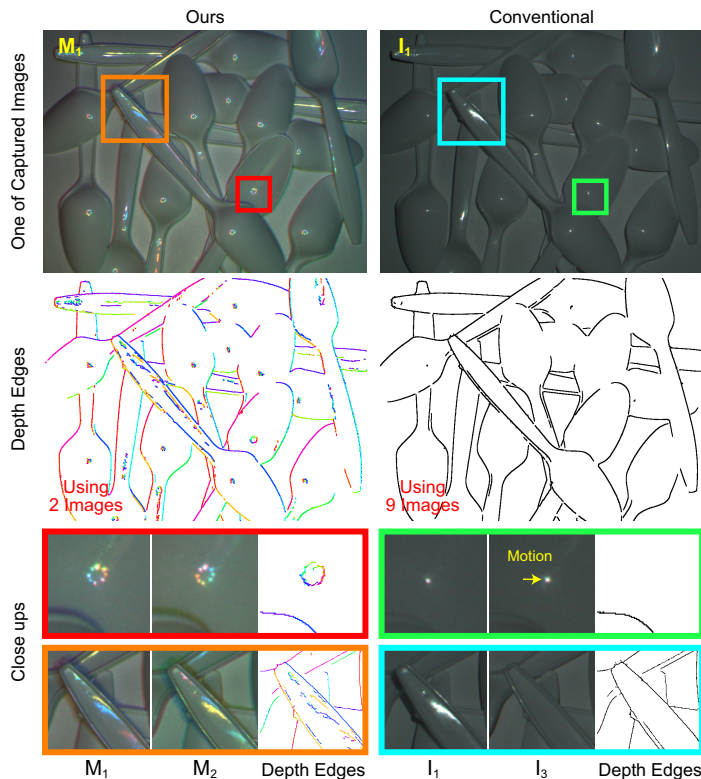


Fig. 7. Effects of specular highlights for (left) our method and (right) the conventional method. (Top) One of the captured images. (Middle) Extracted depth edges. (Bottom) Close ups of two captured images (M_1 and M_2 in our method, I_1 and I_3 in the conventional method) and depth edges. In conventional MFC images captured by using a single light source at a time, the positions of specular highlights typically change depending on the positions of the light sources. Conventional approaches [1, 21] exploit this fact to remove the effect of specular highlights. On the other hand, our method uses all flashes simultaneously, thus producing specular highlights at fixed positions in the two captured images. This confuses our algorithm and generates false depth edges.

Steps 4 and 5 in our method. We used only the first image in the color check process (Step 4c), since the second image does not convey color information.

If the three RGB light sources were positioned at $\{0^\circ, 120^\circ, 240^\circ\}$, Feris et al.’s method becomes a special case of our method, capturing the first image with a hue circle approximated by three light sources and the second image with white light sources. However, as discussed in Section 3.1, it provides a less reliable distance measure than ours, resulting in noisy depth edges as shown in Figure 9 (top). Moreover, because of the limited color information available only in the first image, orientations of some depth edges are misclassified. If we used three light sources, we would capture the second image with complementary col-

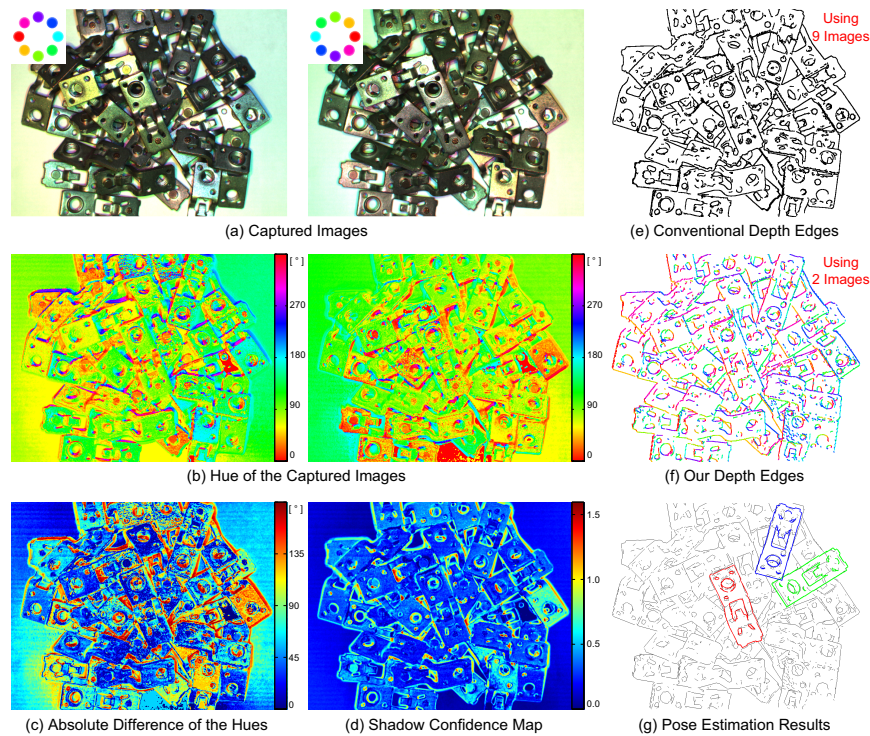


Fig. 8. Results for semi-specular objects. (a) Captured images. (b) Hue of the captured images. Note that non-Lambertian reflections on object surfaces violate the assumption that non-shadow regions have a fixed hue in the two images. (c) Only using the absolute difference of the hues does not provide a robust shadow confidence map, because colors on semi-specular object surfaces have small saturations, at which the estimation of the hue becomes unstable. (d) The distance on the hue-saturation plane provides a better shadow confidence map. (e) Depth edges extracted with the conventional method, requiring 9 images. (f) Depth edges extracted with our method using 2 images. (g) Object detection and pose estimation results using the fast directional chamfer matching (FDCM) algorithm [5, 6] for our depth edges. Best three estimated poses (red, green, and blue) are superimposed on the line representation of the depth edges.

ors of RGB (i.e., CMY) to make the distance measure more robust, producing better depth edges as shown in Figure 9 (middle). While the standard RGB multiplexing is applicable only to three light sources, our method allows any number of light sources multiplexed to approximate a hue circle. Our method using 8 light sources further improves the depth edges as shown in Figure 9 (bottom).

5 Conclusions

We presented a novel color multiplexing method for depth edge extraction that exploits the complementary nature of a hue circle. Our method colorizes shadow

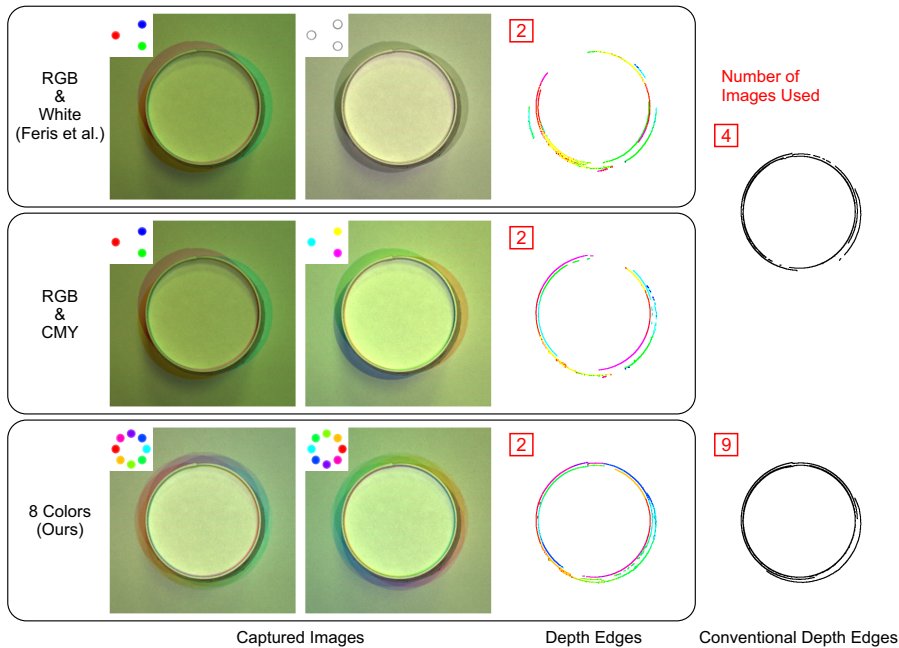


Fig. 9. Comparison with Feris et al. [7] for a scene including the bottom of a paper cup. (Top) The depth edges computed by Feris et al.’s method are noisy and include incorrect orientations due to the limited color information available only in the first image. (Middle) Instead of using the white light sources, we can capture the second image using the complementary colors (CMY) to make the distance measure more robust, producing better depth edges. (Bottom) Our method allows any number of light sources multiplexed. Using 8 light sources further improves the depth edges. Conventional depth edges extracted using 3 and 8 light sources are also shown as references.

regions in captured images because of the occlusion of a half of the colors in the hue circle, while letting non-shadow regions have a neutral color because of the mixture of all the colors in the hue circle. Although our theory provides a simple way to extract depth edges with orientations by looking up the hues in a single image, in practice, we presented a robust depth edge extraction algorithm using two images captured with a hue circle and its complementary version. We demonstrated a camera prototype using 8 color LEDs and discussed the advantages and limitations of our approach compared to conventional methods under various scenes.

Acknowledgments: The author thanks Jay Thornton, Ramesh Raskar, Amit Agrawal, Oncel Tuzel, Srikumar Ramalingam, Tim K. Marks, Ming-Yu Liu, Makito Seki, and Yukiyasu Domae for valuable discussions and their support, and the anonymous reviewers for their helpful comments. Special thanks to John Barnwell, William Yerazunis, and Abraham Goldsmith for their help in developing the camera prototype.

References

1. Raskar, R., Tan, K.H., Feris, R., Yu, J., Turk, M.: Non-photorealistic camera: Depth edge detection and stylized rendering using multi-flash imaging. *ACM Trans. Graphics* **23** (2004) 679–688
2. Crispell, D., Lanman, D., Sibley, P.G., Zhao, Y., Taubin, G.: Beyond silhouettes: Surface reconstruction using multi-flash photography. In: 3DPVT. (2006) 405–412
3. Feris, R., Raskar, R., Chen, L., Tan, K.H., Turk, M.: Multiflash stereopsis: Depth-edge-preserving stereo with small baseline illumination. *IEEE Trans. Pattern Anal. Mach. Intell.* **30** (2008) 147–159
4. Agrawal, A., Sun, Y., Barnwell, J., Raskar, R.: Vision-guided robot system for picking objects by casting shadows. *Int'l J. Robotics Research* **29** (2010) 155–173
5. Liu, M.Y., Tuzel, O., Veeraraghavan, A., Chellappa, R., Agrawal, A., Okuda, H.: Pose estimation in heavy clutter using a multi-flash camera. In: ICRA. (2010) 2028–2035
6. Liu, M.Y., Tuzel, O., Veeraraghavan, A., Taguchi, Y., Marks, T.K., Chellappa, R.: Fast object localization and pose estimation in heavy clutter for robotic bin picking. *Int'l J. Robotics Research* **31** (2012) 951–973
7. Feris, R., Turk, M., Raskar, R.: Dealing with multi-scale depth changes and motion in depth edge detection. In: SIBGRAPI. (2006) 3–10
8. Vaquero, D.A., Raskar, R., Feris, R.S., Turk, M.: A projector-camera setup for geometry-invariant frequency demultiplexing. In: CVPR. (2009) 2082–2089
9. Chen, C., Vaquero, D., Turk, M.: Illumination demultiplexing from a single image. In: ICCV. (2011)
10. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Opt. Eng.* **19** (1980) 139–144
11. Hernandez, C., Vogiatzis, G., Brostow, G.J., Stenger, B., Cipolla, R.: Non-rigid photometric stereo with colored lights. In: ICCV. (2007) 1–8
12. De Decker, B., Kautz, J., Mertens, T., Bekaert, P.: Capturing multiple illumination conditions using time and color multiplexing. In: CVPR. (2009) 2536–2543
13. Fyffe, G., Yu, X., Debevec, P.: Single-shot photometric stereo by spectral multiplexing. In: ICCP. (2011) 1–6
14. Schechner, Y., Nayar, S., Belhumeur, P.: A theory of multiplexed illumination. In: ICCV. Volume 2. (2003) 808–815
15. Park, J.I., Lee, M.H., Grossberg, M.D., Nayar, S.K.: Multispectral imaging using multiplexed illumination. In: ICCV. (2007)
16. Sá, A.M., Carvalho, P.C.P., Velho, L.: (b, s) -BCSL: Structured light color boundary coding for 3D photography. In: VMV. (2002)
17. Minomo, Y., Kakehi, Y., Iida, M., Naemura, T.: Transforming your shadow into colorful visual media: Multiprojection of complementary colors. *Comput. Entertain.* **4** (2006)
18. MacEvoy, B.: Color vision. <http://www.handprint.com/LS/CVS/color.html> (2008)
19. Vaquero, D.A., Feris, R.S., Turk, M., Raskar, R.: Characterizing the shadow space of camera-light pairs. In: CVPR. (2008)
20. Canny, J.: A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **8** (1986) 679–698
21. Feris, R., Raskar, R., Tan, K.H., Turk, M.: Specular reflection reduction with multi-flash imaging. In: SIBGRAPI. (2004) 316–321