

A Reflectance Based Method For Shadow Detection and Removal

Sri Kalyan Yarlagadda, Fengqing Zhu

School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana, USA.

Abstract—Shadows are common aspect of images and when left undetected can hinder scene understanding and visual processing. We propose a simple yet effective approach based on reflectance to detect shadows from single image. An image is first segmented and based on the reflectance, illumination and texture characteristics, segments pairs are identified as shadow and non-shadow pairs. The proposed method is tested on two publicly available and widely used datasets. Our method achieves higher accuracy in detecting shadows compared to previous reported methods despite requiring fewer parameters. We also show results of shadow-free images by relighting the pixels in the detected shadow regions.

Keywords—shadow detection, reflectance classifier, shadow removal, image enhancement.

I. INTRODUCTION

Shadows are ubiquitous. They are formed when light is partially or fully occluded by objects. Shadows provide information about lighting direction [1], scene geometry and scene understanding [2] in images and are crucial for tracking objects [3] in videos. They also form an integral part of aerial images [4]. However, shadows can also complicate tasks such as object detection, feature extraction and scene parsing [5].

There have been many methods proposed to detect shadows from images and videos [3], [6], [5], [7], [8], [9], [10]. In this paper we focus on detecting shadows from color images. With the recent boom in data driven approaches, machine learning based methods have been applied to detect shadows [5], [7], [8]. In [5] Conditional Random Fields consisting of 2490 parameters are used to detect shadows in gray scale images using features such as intensity, skewness, texture, gradient similarity etc. In [7] Convolutional Neural Networks consisting of 1000's of parameters are used to detect shadows. In [6] intensity information around edges is used to detect shadow boundaries. In [8] image is first segmented and various classifiers are used to detect regions similar in color and texture by comparing different segments with each other.




Figure 1: Note that both the surfaces are dark but one is due to shadow and another is due to shading. Such examples complicate shadow detection but can be solved using neighborhood information

In this paper we propose a non-training based shadow detection method which requires fewer parameters, yet

achieves high accuracy compared to previous methods [8], [5], [7]. We differ from [8] in the features and classifiers used for comparing regions and also in the approach of using these comparisons to obtain the shadow mask. Every surface is characterized by two features: its reflectance and its texture. When a shadow is cast on a surface its illuminance reduces, but its reflectance remains the same. Due to the reduction in the illuminance, there will also be some loss in texture information. By examining a surface, it is difficult to tell whether it is dark due to the effects of shadow or shading. An example of this is given in Figure 1. By comparing surfaces with each other we can detect shadows with greater confidence. Hence, by pairing different regions of an image based on their reflectance, texture and illumination characteristics we can detect shadows efficiently.

II. SHADOW DETECTION

Our goal is to group different regions of an image based on their reflectance, texture and illumination characteristics. To group pixels with similar properties into different regions, we first segment an image using the Quickshift method [11] with a Gaussian kernel size of 9. Our assumption is that a single segment should contain pixels with similar reflectance and illumination. An example of segmentation result is shown in Figure 2. In the subsections below we explain how we design the reflectance, texture and illumination classifiers to label each segment as shadow or non-shadow.



(a) Original Image (b) Segmented Image

Figure 2: An example of a test image segmented using Quickshift with a kernel size of 9. The segmentation correctly separates the boundaries between the shadow and non-shadow regions.

A. Reflectance classifier

Consider the illumination model used in [8],

$$I_i = (t_i \cos(\theta) L_d + L_e) R_i \quad (1)$$

where I_i is the vector representing the i^{th} pixel in RGB space, L_d and L_e are vectors representing the direct light and reflected light from the environment, respectively. θ is the angle between direct light and surface normal and R_i is the




Figure 3: An illustration of color vectors is given in 3a. In 3b visually region A (the shadow region) on the road appears blue, not gray. This is due to the large chromaticity difference in direct light L_d and the reflected light L_e leading to a large θ_1 . However, the angle between I_D and I_{NS} (θ_2) will be small so that we can classify region A and region B as shadow non-shadow pairs.

reflectance vector. The value of t_i indicates whether the pixel belongs to shadow or non-shadow segment. When $t_i = 0$ the pixel is in the shadow segment and vice versa. Two segments belonging to the same surface but under different illumination can be modeled as $t_i = 0$ for all the pixels in the shadow segment and $t_i = 1$ for all the pixels in the non-shadow segment. Assuming that direct light and environment light are constant in magnitude and direction over the two segments, we can see that the reflectance property of the surface remains constant in both cases. Taking the respective median of all pixels in RGB color space in each of the segments, we have the following,

$$I_{NS} - \text{Median color of non-shadow segment}$$

$$I_S - \text{Median color of shadow segment}$$

$$I_D = I_{NS} - I_S = (\cos(\theta)L_d)R^{median}$$

In the case when L_e is similar to L_d in terms of chromaticity, the angle between the vectors I_D and I_{NS} should be zero. However, in practice L_e differs from L_d , hence these two vectors will have a small angle provided they are of the same material and a large angle if they are of different material with different reflectance properties. By thresholding the angle between the color vectors I_D and I_{NS} , we can decide whether two segments with different illumination conditions belong to the same material. We call this the “angle criterion.” Notice that we don’t look at the angle between I_{NS} and I_S because in the case where L_e is significantly different from L_d in terms of chromaticity, the angle between these two vectors will be very large even if they represent the same surface. We set the angle threshold to be 10° . An example of this case is illustrated in Figure 3.




Figure 4: The segmented image in 4a is grouped using the luminance classifier and the result is shown in 4b.

B. Luminance Classifier

Shadows are formed when direct light is partially or fully occluded and hence have lower illumination. The decrease in illumination depends on the relative intensities of L_d and L_e . A large decrease in illumination intensity darkens the shadow. To build an effective luminance classifier, we need to be able to detect the decrease in illumination and be able to attribute that decrease to obstruction of light and not due to some noise. In order to model this, we look at the luminance values of all pixels in the LAB color space. We compute the median luminance of all segments in the LAB space and compute the histogram of the median luminance values. The peaks of the histogram give us an estimate of the number of different illumination regions in the image.

We then split the image into regions by grouping segments based on their proximity to the peaks. Segments within the same region are not compared because they have similar illumination intensity while segments from different groups are allowed for comparison to detect shadows. This step is useful because it adaptively groups segments into regions with similar illumination. An example of grouping segments into regions based on their luminance is shown in Figure 4. In addition to the grouping criteria, for two segments to be shadow non-shadow pairs, the ratio of their median luminance T in LAB space has to be above the threshold of 1.2 in order to avoid comparing segments with similar illumination. T can be anywhere between 1 and ∞ and the closer it is to 1 the closer the illumination intensities of the two segments are. Shadow non-shadow pairs will have a high values of T compared to segments with similar illumination intensities.

C. Texture classifier

Since shadow and the corresponding non-shadow segments are of the same material their texture characteristics will be similar. However, due to the reduction in illumination intensity of shadow segments, some texture information is lost. To capture this phenomenon, we look for texture similarity between the segments under comparison provided that their T is not very high, because if its high a lot of texture information would have been lost. We compute the Earth Mover Distance between the histograms of the

texture maps [12] of both segments and threshold it to find whether the two segments have similar texture. However, if T is greater than 2.4 we do not compare them for texture similarity as a lot of texture information is lost in the shadow segment due to the decrease in illumination.

D. Implementation

In this subsection, we describe how we use the above three classifiers to detect shadow non-shadow segment pairs. Each segment is compared to its neighboring segments using the reflectance, texture and luminance classifiers discussed above. If all the classifiers label the pair as a shadow non-shadow pair, we store that connection. We use these connections to connect more segments. For every shadow non-shadow pair, we take all the non-classified neighbors of the shadow segment and compare them to non-shadow segment using the above classifiers. We repeat this process 3 times. The reason is that some shadow segments may have neighbors which are also shadow segments themselves. Such segments will not be detected in the first iteration. In order to connect them to the already labeled shadow segments, we repeat the process by using the information obtained from the initial connections. The process is illustrated in Figure 5.




Figure 5: The connections obtained by first iteration are marked with the white lines (5a) and the connections obtained by second iteration are marked by the blue lines (5b).

E. Refinement

The above implementation detects shadow non-shadow pairs with similar reflectance, texture and different luminance but does not put any constraint on how bright the shadow segment should be. Without such constraint, two very bright segments can be misclassified as shadow non-shadow pairs. In order to avoid this, we limit the shadow region to have a gray scale value lower than the Otsu threshold of the image. We segment the image again with a Gaussian kernel of size 3 (smaller than the Gaussian kernel used in the initial segmentation) and look for segments which contain shadow pixels using the initial shadow mask. A finer segmentation mask leads to better modeling of the shadow non-shadow boundaries. Given a segment contains shadow pixels, if more than 70% of pixels in that segment have a gray scale value less than the Otsu threshold, we

label the entire segment as a shadow segment and if not we label the entire segment as non-shadow.

III. EXPERIMENTAL RESULTS

The proposed method is evaluated on two publicly available datasets, the UIUC dataset [8] and the UCF dataset [5].

A. UIUC Dataset

The UIUC dataset consists of 108 images with shadows, out of which 32 images have been used for training and 76 for testing by [8]. We have evaluated our method on the 76 test images. In addition to computing the per class accuracy, we also show the Balanced Error Rate (BER) for our method which is computed as the following,

$$\text{BER} = 1 - \frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \quad (2)$$

where FP is False Positives, FN is False Negatives, TP is True Positives and TN is True Negatives. The lower the BER the better the method. BER is used because there are fewer shadow pixels than non-shadow pixels in the images. The results of our methods and others are shown in table I. We have achieved the highest accuracy detecting shadows and a very close BER compared to [7] which has the smallest BER of all three methods.

Table I: Results Our Proposed Method Compared to Other Methods On UIUC Dataset

Methods	Shadows	Non-Shadows	BER
Unary + Pairwise([8])	.716	.952	.166
ConvNet([7])	.847	.955	.099
Our method	.906	.855	.119

B. UCF Dataset

The UCF dataset is also widely used for testing shadow detection methods. It consists of 355 images which are more diverse and complex than the UIUC dataset. In [5] 120 images were used for testing. We have tested our method on 236 images. Out of the 236 images, for 162 of them we have followed the proposed method, but for 74 images from OIRDS [13] dataset we have chosen a threshold of .35 instead of using the Otsu threshold for limiting the gray scale of the shadow pixels. This is because OIRDS dataset contains aerial images with very dark shadow regions. The results are reported in Table II and comparisons to other methods are shown in Table III. In comparison to other methods, our method achieved the highest accuracy in detecting shadows and also has the best BER.

Table II: Detection Confusion Matrices of Our Proposed Method On UCF Dataset

74 images from OIRDS dataset	Shadow	Non Shadow
Shadow	.899	.101
Non - Shadow	.116	.884
162 images from UCF dataset	Shadow	Non Shadow
Shadow	.922	.078
Non - Shadow	.191	.809

Table III: Detection Confusion Matrices of Our Proposed Method Compared to Other Methods On UCF Dataset

Methods	Shadows	Non-Shadows	BER
BDT-BCRF[5]	.639	.934	.2135
Unary + Pairwise([8])	.733	.937	.165
ConvNet([7])	.780	.926	.147
Our method	.920	.827	.1265

IV. SHADOW REMOVAL

To remove shadows we follow the same approach as described in [8]. Some examples of shadow removal are shown in Figure 6.




Figure 6: Sample shadow removal results.

V. CONCLUSION

We proposed a simple yet effective shadow detection method requiring few parameters. Each image was first segmented and segment pairs were identified as shadow non-shadow pairs based on their reflectance, illumination and texture characteristics. Experimental results showed that our method was effective for detecting shadows but had a lower accuracy in identifying non-shadows. The connections between the detected shadow and non-shadow pairs were used to successfully remove shadows in test images.

REFERENCES

- [1] A. A. Efros and S. G. Narasimhan, "Estimating Natural Illumination from a Single Outdoor Image," *2009 IEEE 12th International Conference on Computer Vision*, pp. 183–190, Sept 2009.
- [2] S. Wehrwein, K. Bala, and N. Snavely, "Shadow Detection and Sun Direction in Photo Collections," *2015 International Conference on 3D Vision (3DV)*, pp. 460–468, Oct 2015.
- [3] J. Gao, J. Dai, and P. Zhang, "Region-based Moving Shadow Detection Using Watershed Algorithm," *2016 International Symposium on Computer, Consumer and Control (IS3C)*, pp. 846–849, July 2016.
- [4] Y. I. Shedlovska and V. V. Hnatushenko, "Shadow Detection and Removal Using a Shadow Formation Model," *2016 IEEE First International Conference on Data Stream Mining Processing (DSMP)*, pp. 187–190, August 2016.
- [5] J. Zhu, K. G. G. Samuel, S. Z. Masood, and M. F. Tappen, "Learning to recognize shadows in monochromatic natural images," *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 223–230, June 2010.
- [6] J. F. Lalonde, A. a. Efros, and S. G. Narasimhan, "Detecting ground shadows in outdoor consumer photographs," *European Conference on Computer Vision*, vol. 6312 LNCS, no. PART 2, pp. 322–335, 2010.
- [7] S. Khan, M. Bennamoun, F. Sohel, and R. Togneri, "Automatic Shadow Detection and Removal from a Single Image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 3, pp. 431–446, March 2016.
- [8] R. Guo, Q. Dai, and D. Hoiem, "Paired regions for shadow detection and removal," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 12, pp. 2956–2967, Dec 2013.
- [9] K. L. Chung, Y. R. Lin, and Y. H. Huang, "Efficient shadow detection of color aerial images based on successive thresholding scheme," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 2, pp. 671–682, Feb 2009.
- [10] W. Qi, Z. Wende, and B. V. K. V. Kumar, "Strong shadow removal via patch-based shadow edge detection," *Robot. Autom. (ICRA), 2012 IEEE Int. Conf.*, pp. 2177–2182, May 2012.
- [11] A. Vedaldi and S. Soatto, "Quick shift and kernel methods for mode seeking," *Proceedings of the European Conference on Computer Vision (ECCV)*, 2008.
- [12] D. R. Martin, C. C. Fowlkes, and J. Malik, "Learning to detect natural image boundaries using local brightness, color, and texture cues," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 5, pp. 530–549, May 2004.
- [13] F. Tanner, B. Colder, C. Pullen, D. Heagy, C. Oertel, and P. Sallee, "Overhead imagery research data set (OIRDS) an annotated dat a library and tools to aid in the development of computer vision algorithms, 2009."