

Identifying Image Composites Through Shadow Matte Consistency

Qiguang Liu, Xiaochun Cao, Chao Deng, and Xiaojie Guo

Abstract—In this paper, we propose a framework for detecting tampered digital images based on photometric consistency of illumination in shadows. In particular, we formulate color characteristics of shadows measured by the shadow matte value. The shadow boundaries and the penumbra shadow region in an image are first extracted. Then a simple and efficient method is used to estimate shadow matte values of shadows. Our approach efficiently extracts these constraints from a single view of a target scene and makes use of them for the digital forgery detection. Experimental results on both simulated photos and visually plausible real images demonstrate the effectiveness of the proposed method.

Index Terms—Image authentication, photo composites, shadow matte.

I. INTRODUCTION

PEOPLE are fond of tampering with photos and videos. Operations, such as matting and compositing, are introduced to produce special effects [1]. Since then, image region copy and paste has become one of the most common video editing and manipulating techniques due to its simplicity. With the development of powerful image and video editing tools, it is becoming incredibly easy to generate visually plausible tampered images and videos. At the same time, as internet and multimedia technologies are becoming increasingly mature and popular, there are growing tampered photos and videos flooding the televisions, magazines, and networks, which blocks our eyes towards the truth. Therefore, evaluating the authentication of digital photography has turned out to be a crucial task nowadays.

Over the past few years, both blind [2]–[4] and non-blind approaches for image authentication have been developed. Among the non-blind approaches, watermarking [5], [6] is a popular technique, which inserts a watermark at the recording time and extracts hidden message later to verify the image authenticity. Unfortunately, the non-blind methods need preprocessing to the original image, which is not feasible for images independent of any pre-inserted information.

Manuscript received October 01, 2010; revised March 16, 2011; accepted March 26, 2011. Date of publication April 07, 2011; date of current version August 17, 2011. This work was supported by NSFC (No. 60905019, 61003200), Program for New Century Excellent Talents in University, Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology, SKL of PMTI, SKL of CG&CAD, and Open Projects Program of NLPR. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jiwu Huang.

The authors are with the School of Computer Science and Technology, Tianjin University, Tianjin 300072, China (e-mail: qliu@tju.edu.cn; xcao@tju.edu.cn; dengchao@tju.edu.cn; xguo@tju.edu.cn).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIFS.2011.2139209

Farid [7] roughly groups blind approaches for image authentication into five categories: 1) pixel-based techniques; 2) format-based techniques; 3) camera-based techniques; 4) physically based techniques; and 5) geometric-based techniques. Several approaches in the first group detect statistical anomalies introduced at the pixel level. In [8]–[10], copied regions are detected by grouping and sorting duplicated blocks. Unnatural correlations resulting from the resampling process when tampering an image are identified and utilized as clues of forgery [11]. Anomalies in statistics of Fourier components can be used to detect splicing of images [12]–[14]. However, those methods are all limited to analyze correlations between pixels arising from a specific form of tampering. As in the second categories, the authors in [15]–[17] exploit properties of specific lossy compression schemes for digital forgery detection. The authors in [18] introduce JPEG error analysis to the study of image forensics. They believe that lossy compression of the JPEG format leaves evidences of forgery. Methods belonging to the third approach are also introduced for forgery detection by exploiting artifacts introduced by the camera. Chromatic aberration [19], color filter array [20], camera response [21], and sensor noise [22]–[24] are modeled to infer the source digital cameras and reveal inconsistency in digitally altered images. Other works such as [25] identify doctored images by tracing the entire in-camera and post-camera processing operations. Several approaches in the fifth group are also introduced into photo forgery detection. The authors in [26] discover composites of people by estimating principle points, which requires relative high resolution for the eye region. Wang and Farid [27] argue that the skew parameter can also be compared to detect image re-projections. In [28], two-view geometrical constraints are enforced to find fake regions on pictures targeting the same scene. More recently, segmentation techniques are introduced to improve the rough targeted fake regions [29].

Our work falls into the fourth group, since we apply photometric constraints on shadows for detecting the forgery in an image as shown in Fig. 1. Methods using physical rules detect anomalies of interaction between objects in the target scene. Johnson and Farid [30] approximate two components of light source direction by considering the two-dimensional surface normals at the occluding object boundary and use them to identify the inconsistency in a picture. Because of the missing component of the light direction, this method might suffer from an ambiguity. The authors in [31] estimate a three-dimensional light source direction via analyzing a model of the human eye. However, their method needs the existence of specular highlights on human eyes. In [32], complex lighting environments are modeled by the linear sum of several spherical harmonic

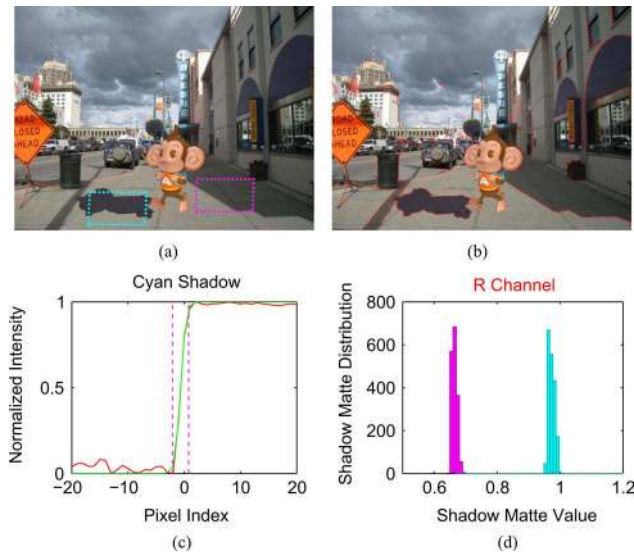


Fig. 1. Shadows are the necessary part of an object when copying and pasting the object into a target image to maintain the integrity. We use shadow matte values extracted from the shadows of scene objects to detect image forgery. (a) Example image to be detected. (b) Detected shadow boundaries. (c) Penumbra shadow width. (d) Distribution of a set of shadow matte obtained by subtraction of two constructed intensity surfaces for the sampled shadows marked by colored rectangles in (a). The mean value of the distribution gives the actual shadow matte value used in our work.

functions. The linear weights of these functions are compared to reveal inconsistencies of lighting within a doctored image.

In our perception of the world, shadows provide important visual cues for depth, shape, contact, and lighting [33]. When tampering with photos, they are integral parts of an image and following consistency of such properties in shadows can be significantly difficult [34], [35]. Based on this observation, we propose a new framework for detecting tampered digital images using shadows. We show how to extract photometric constraints from a single view of a target scene and to make use of them for the digital forgery detection. The width of the penumbra shadow region [Fig. 1(c)] in an image are first extracted. Then we evaluate a scale factor called shadow matte value by analyzing illumination intensity in shadows. We effectively estimate the shadow matte value of full shadows without being affected by shape or texture of the shadow receiving surface. Inconsistency of the matte value between shadows implies a forgery in an image.

Zhang and Cao [36] extract the shadow matte value by calculating the color intensity medians of two regions laid in and out sides of the shadows and then evaluate the ratio between them. This method works pretty well when the shadow receiving surface is flat and not textured. However, shadows in real-life pictures are usually projected onto sophisticatedly textured surface with complex geometric features. Similarly, Fredembach and Finlayson [37] find the shadow matte value by minimizing differences between pixels inside and outside the shadow boundary. This method is much faster than traditional integration and works very well when the shadow is projected to a flat non-textured surface, but still omits the geometric change of the surface. Arbel and Hel-Or [38] improve the method by reconstructing an intensity surface to acquire the shadow matte. In the step of estimating the shadow matte values, we adopt

their mathematical model but with several improvements. We develop an approximation to the thin plate model. As a result, we propose a direct linear solution while [38], to achieve optimization, has to use gradient descent algorithm and to perform iteratively. Instead of reconstructing a single intensity surface without shadow, we recover two intensity surfaces, shadow intensity surface that reflects the intensities of the sampled pixels with full shadow on them and non-shadow that reflects the intensities of the sampled pixels without shadow on them. In addition, our work is different from [38] in three aspects. First, our work focuses on the application of image authentication while they are interested in shadow removal. Second, we estimate the penumbra region width and do not need to provide a penumbra mask as input in their paper. Third, we adopt the shadow detection to automate the image forensic process. Methods in [39] and [40] also remove shadows by estimating the shadow matte value. Instead of using a uniform shadow matte value, they assume the shadow matte values in the full shadow region varies. Utilizing nonuniform shadow matte value is more accurate for shadow removing. However, we argue that a uniform shadow matte value assumption is sufficient and valid to reflect the illumination condition in our image forensic application. The main reason is the observation [39], [41] that the shadow matte values are smooth and only change mostly according to the distance to the shadow boundary. However, our anchor points used to compute the shadow matte values are close to shadow boundary. In addition, the tiny variance in shadow matte values for our anchor points are trivial when compared to the shadow matte differences between the authenticate and composite shadow regions. Many approaches [42]–[45] have managed to derive an image into intrinsic images, from which the shadow matte value can be obtained. We evaluate the shadow matte value without directly recovering the intrinsic images, which is much simpler for implementation and sufficient for image forgery detection.

The contributions of this paper include: 1) We make use of the shadow matte values for the digital forgery detection. 2) We identify the width of the penumbra region of a shadow utilizing the similarity of the shape between the averaged penumbra curve and the sigmoid function. 3) We explore a new simple and efficient method to extract shadow matte values of shadows from a single view of a target scene.

The rest of this paper is organized as follows: After introducing preliminaries of shadow characteristics used in this work in Section II, we describe our method in Section III. The performance of our method is demonstrated by experimental results in Section IV. Finally, Section V concludes this paper.

II. PRELIMINARIES ON SHADOW PROPERTIES

As a necessary part of most real-world photos and an important clue for photographic composite detection, shadows share characteristics on such as generation, composition, and visual appearance. We will discuss them in this section.

A. Geometric Properties of Shadows

Although shadows differ in geometric appearance due to different shapes of the objects from which shadows are casted and various surfaces to which shadows are projected, the composition of shadows is similar.




Fig. 2. Shadow generation. Part of the light from a light source is obstructed by a sphere and cast a shadow on a plane.

As can be noticed in Fig. 2, shadows are caused by the shadow casting object obstructing the direct light from the light source to the shadow receiving object. The darker region on the surface of the shadow casting object facing opposite direction to the light source is called self shadow. The darker area on the surface of the shadow receiving object is referred to as cast shadow. The contour of the shadow casting object forming the boundary of the cast shadow is shadow casting edge. As can be seen, the cast shadow is not solely colored and is divided into two regions. We call the darker area full shadow and the less dark area penumbra shadow.

In a nature scene, illumination is much more complex. Rather than a simplistic distance (small-area) point light source, light reaching to a surface may come from more than one source, such as other light sources, atmosphere, and inter-reflections in a scene. However, for outdoor images, we assume that all the light sources can be divided into two types: the sun and ambient light. The ambient light in a scene is assumed to be the same everywhere and the sun can be seen as a (small-area) point light source. With this assumption of two types of light sources: the sun and the ambient light, shadows are formed by obstructing direct light from the sun, leaving the shadow areas partially illuminated and the idea of full and penumbra shadow can be similarly defined as in Fig. 2.

B. Photometric Properties of Shadows

Shadowed image areas are made up of discrete color or intensity values. These values relate to a various of factors including lighting in the environment, surface properties and geometry, camera optics, scene geometry layout, mutual occlusion along the light rays, and sensor properties. A more detailed treatment of these topics can be found in [40], [46], and other textbooks on computer graphics. In this paper, we adopt the commonly used model [43], [44] on intrinsic images that the intensity at a pixel \mathcal{I}_k for channel $k \in \{R, G, B\}$ is due to reflectance \mathcal{R}_k and illumination \mathcal{L}_k :

$$\mathcal{I}_k = \mathcal{R}_k \cdot \mathcal{L}_k. \quad (1)$$

Shadows are formed by reduction in the illumination field resulting from obstructing part of the direct light received at a surface point. This reduction of illumination is assumed by multiplying a scale factor \mathcal{C}_k called shadow matte value to \mathcal{L}_k field. \mathcal{C}_k is ranging from 0 to 1, denoting between irradiance and radiance regions. Therefore, the actual color intensity sponsored can be expressed as

$$\mathcal{I}_k = \mathcal{R}_k \cdot \mathcal{L}_k \cdot \mathcal{C}_k. \quad (2)$$

If we denote $\mathcal{I}'_k = \mathcal{R}_k \cdot \mathcal{L}_k$ indicating the intensity without shadow, then

$$\mathcal{I}_k = \mathcal{I}'_k \cdot \mathcal{C}_k. \quad (3)$$

This image model simply means that the observed image \mathcal{I} is the shadowless image \mathcal{I}' scaled by the fractional shadow matte value \mathcal{C} . For a region in the full shadow area, the value of \mathcal{C} is assumed to be uniform, whereas \mathcal{C} changes smoothly across the penumbra shadow. Taking the logarithm of both sides of (3) gives a more intuitive understanding that the color intensity is additively obtained:

$$I_k = I'_k + C_k \quad (4)$$

where I , I'_k , and C are the logarithms of \mathcal{I} , \mathcal{I}' , and \mathcal{C} , respectively. Our method is operated over the log domain of the images for each channel. We omit k in all the above notations for simpleness. When we refer to shadow matte value in the rest of the paper, we use the log domain C .

Note that the specular component in the typical bidirectional reflectance distribution function (BRDF) model is omitted. In the shadow region, we assume only the diffuse component (also known as Lambertian or matte reflection) scatters light uniformly in all directions. This assumption is mostly valid as the shadow receiving surfaces in outdoor scenes are typically asphalt, brick, stone, mud, grass, and concrete [47]. In addition to those lighting factors mentioned above, the acquisition and digitization pipeline in the camera also affects the pixel intensity in an image. Sensors will bring noises into the image, information loss also occurs due to clipping of pixel intensities caused by the limited range of camera sensors, and postprocessing is commonly performed in the pipeline. Many of these factors are not consistent with our model. Unfortunately, a further discussion of them is beyond the scope of this paper. We assume these factors do not affect the direct proportion relationship between the pixel intensity in an image and the sponsored intensity of the sensor.

III. METHOD

The illumination of real-world scenes is complex and the shadows formed in the same photo may vary and overlap with each other if there are multiple light sources. Fortunately, in most of the natural scenes, especially outdoor views, there is only one distant direct light source, e.g., the sun. In this work, we focus on the outdoor scenes where the single distant light source assumption is valid. As mentioned above, the shadow receiving surface has to be nearly Lambertian. Although there are specular shadow receiving areas such as marble in the outdoors, the typical ground materials are asphalt, brick, stone, mud, grass, and concrete, which are mostly Lambertian. The attenuation of the lighting varies when the shadow casting object becomes transparent or the distance from the shadow casting object and the shadow receiving surface is too small. It is assumed in this work that the shadow casting object is opaque and that the distances between the shadow casting and shadow receiving objects are nontrivial. This assumption is valid considering the fact that most of the image composites involve the targets such as humans, cars, and animals, which are thick, sizable, and opaque. Note that the shadow receiving surface does




Fig. 3. Shadow sampling. (a) Shadow is sampled bidirectionally along the normal direction of the shadow edge. (b) The sampled pixels generate a rectangle region. (c) 3D surface constructed using intensities of the sampled pixels.

not have to be flat as our method is able to handle curved surfaces. We also assume that there exists another authentic shadow which has similar shadow receiving surface to that of the composite shadow. In addition, the two shadows are assumed to be close so that the nonlinear radiometric responses and perspective distortions of the common camera will trivially affect the values of the shadow matte in the two shadow areas. Therefore, the shadow matte values of these shadows should be close. The inconsistency of shadow matte values in the same image will indicate a forgery.

A. Calculating Shadow Matte Value

The intensity of pixels on a curved surface without the shadow produces an intensity surface. By reconstructing this intensity surface on the shadowed region, we can find how much the shadow affects the intensity of the shadowed pixels and hence acquire the shadow matte values. To achieve this goal, pixels are sampled bidirectionally along the normal directions of the shadow edge as shown in Fig. 3(a). The width of the sampled data for each side of the shadow boundary is the width of the penumbra region. If the penumbra width is very narrow, the sample width is set to a default value which is 20 pixels in our implementation. Then the sampling generates a rectangle containing all the sampled pixels shown in Fig. 3(b). We define the direction along the shadow edge as y axis and the normal direction as x axis. The surface constructed using the sampled pixels is shown in Fig. 3(c). As can be seen, the color intensity drops from the lit area through the penumbra area to the full shadow area.

To estimate the shadow matte value, thin plate model is used by the method in [40]. But they only use the intensity of the non-shadowed area to recover the entire intensity surface without shadow, leaving the shadowed area governed by a smooth term. We believe that only use the non-shadowed area cannot wholly anticipate the shadowed intensity. Also thin plate model requires the solution of a linear system consisting of a linear equation for each sampled pixel, which is both memory consuming and computationally intensive. Instead, we utilize both the shadowed and non-shadowed regions and develop a approximation to the thin plate model that require much less variables. Instead of reconstructing a single intensity surface without shadow, we recover two intensity surfaces, shadowed surface $f_s(x, y)$ that reflects the intensities of the sampled pixels with full shadow on them and non-shadowed surface $f_n(x, y)$ that reflects the intensities of the sampled pixels without shadow on them. Then, the subtraction of the

intensities of the corresponding pixels between the non-shadowed and shadowed surface gives the shadow matte value.

The non-shadowed surface should fit the intensities of the sampled pixels in the lit area, while the shadowed surface should fit the intensities in the fully shadowed region. So we define an energy function that fits the sampled pixels as

$$E_d(f_n, f_s) = \int \int_{\Omega_n} (f_n(x, y) - I(x, y))^2 dx dy + \int \int_{\Omega_s} (f_s(x, y) - I(x, y))^2 dx dy \quad (5)$$

where Ω_n and Ω_s denote the non-shadowed and fully shadowed sampled regions. Also, the two surfaces should be smooth. So an energy function measuring the smoothness of the surfaces is defined as

$$E_s(f_n, f_s) = \int \int_{\Omega} \left(\frac{\partial^2 f_n}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 f_n}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 f_n}{\partial y^2} \right)^2 + \left(\frac{\partial^2 f_s}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 f_s}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 f_s}{\partial y^2} \right)^2 dx dy \quad (6)$$

where Ω is the entire sampled region. In addition, since we assume a uniform shadow matte value, the variation of the subtraction between the two surfaces should be limited. So the third energy function is defined as

$$E_c(f_n, f_s) = \int \int_{\Omega} \left(\frac{\partial(f_n - f_s)}{\partial x} \right)^2 + \left(\frac{\partial(f_n - f_s)}{\partial y} \right)^2 dx dy. \quad (7)$$

The total energy can then be defined as

$$E = E_d + E_s + E_c. \quad (8)$$

Minimizing E over f_n and f_s needs first to discretize f_n and f_s on a regular grid, resulting in two variables for each pixel, which is both time and space consuming to solve. So we develop a light weighted approximation by using a set of parameterized quadratic splines along the x axis:

$$f_n(x, y) = a_n(y)x^2 + b_n(y)x + c_n(y) \quad (9)$$

$$f_s(x, y) = a_s(y)x^2 + b_s(y)x + c_s(y) \quad (10)$$

where $a_n(y), b_n(y), c_n(y), a_s(y), b_s(y), c_s(y)$ are parameters governing the splines. Since there are only three parameters for each spline, the number of the variables is largely decreased. Another benefit of using parameterized spline is that: the energy function defined by L2 norm is sensitive to outliers. Fortunately, since a spline needs only three parameters, we only use the median half of the sorted intensities along a spline of the sampled pixels, which are sufficient to fit the spline and have low probability to be outliers. The first and second order derivatives along the y axis can be easily obtained by finite element analysis. The first and second order derivatives along the x axis are the first and second order derivative of the quadratic function representing the spline:

$$\frac{\partial(f_*(x, y))}{\partial x} = 2a_n(y)x + b_n(y) \quad (11)$$

$$\frac{\partial(f_*^2(x, y))}{\partial x^2} = 2a_n(y) \quad (12)$$

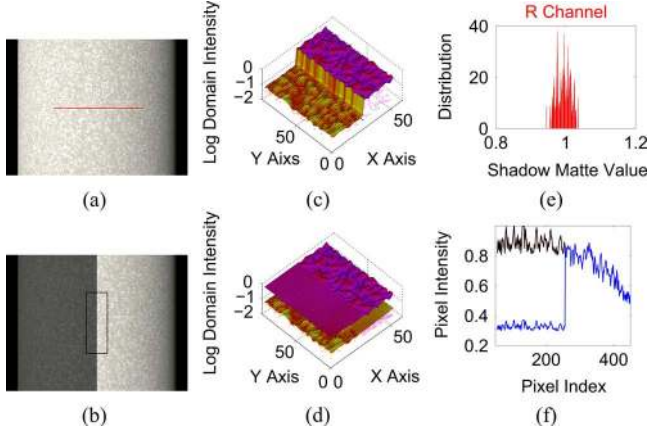


Fig. 4. Calculate C in full shadow region on a simulated cylinder surface. (a) A fully illuminated cylinder surface. (b) The left half of (a) is shaded by a full shadow generated using $C = 1.0$ in (4). (c) Constructed surface using intensity of the sampled pixels. (d) Recovered non-shadowed surface and shadowed surface. The estimated C is 0.9944, which is 0.56% off the ground truth. (e) Recovered one row of the shadowed image using the estimated C .

where $* \in \{n, s\}$. If we define $\mathbf{f} = [a_n(1), b_n(1), c_n(1), \dots, a_n(m), b_n(m), c_n(m), \dots, a_s(1), b_s(1), c_s(1), \dots, a_s(m), b_s(m), c_s(m)]$, where m is the number of the splines, then the total energy can be rewritten as

$$E = \mathbf{f}^T \mathbf{A} \mathbf{f} - 2\mathbf{f}^T \mathbf{b} + c \quad (13)$$

where \mathbf{A} , \mathbf{b} , c are some constant coefficients. Equation (14) can be easily minimized by solving the following linear system of equations:

$$\mathbf{A} \mathbf{f} = \mathbf{b}. \quad (14)$$

Accordingly, f_n and f_s are obtained. Then the shadow matte value can be estimated by

$$C = \text{mean}\{f_n(x, y) - f_s(x, y)\}, \quad (x, y) \in \Omega. \quad (15)$$

To better illustrate this idea, we generate a scene of a fully illuminated cylinder [Fig. 4(a)] and cast an artificial shadow with $C = 1.0$ on it [Fig. 4(b)]. Fig. 4(c) shows the surface constructed using intensities of the sampled pixels in the black rectangle in Fig. 4(b). By performing the proposed method, we obtain the non-shadowed surface [higher surface in Fig. 4(d)] and the shadowed surface [lower surface in Fig. 4(d)]. Fig. 4(e) gives the distribution of $f_n(x, y) - f_s(x, y)$ in R channel. By averaging the distribution, we obtain the shadow matte C as 0.9944, which is 0.56% off the ground truth. The red curve in Fig. 4(f) is the sampled intensities of a row [red line in (a)] in the original image and the blue curve is the sampled intensities of the row in the shadowed image accordingly. With the C calculated and added to the shadowed half of the row [left part of the blue curve in Fig. 4(f)], the recovered intensity curve showed with black color approximately coincides with the red one. This demonstrates the accuracy of our method.

Zhang *et al.* [36] extract the shadow matte value on non-textured surface by calculating the color intensity medians of two regions laid in and out sides of the shadow, respectively, and then calculating the ratio between them. Such simple estimation




Fig. 5. Effectiveness comparison in estimating the shadow matte value C . The first column shows the original shadowed images. The middle column row gives recovered images using C estimated by the method in [36]. The last column shows the recovered images after removing the shadows using C estimated by our method.

cannot work effectively on textured shadow receiving surface. As shown in Fig. 5, we obtain the shadow matte values of the shadows of the images in the first column applying our method and the method proposed by [36] and then remove the shadows using the obtained shadow matte values, respectively. As can be seen, the resulting images (last column in Fig. 5) obtained by our C are visually shadow free, while the shadows in the images (middle column in Fig. 5) generated using the shadow matte values estimated by [36] are still visible, which demonstrates the better efficiency of our approach.

The method in [40] calculates nonuniform shadow matte values for pixels in the full shadow region, while we estimate a uniform shadow matte value. Images in the first and third columns of Fig. 6 give two shadowed images and the shadow removed images using method in [40] (second row) and our method (last row), respectively. The second and last columns shows the distribution of the shadow matte value estimated using their method in three channels. The x coordinate of the magenta line with a circle on the top of the line indicates the shadow matte value estimated by our method. As can be seen, the distribution of their shadow matte value is centered around our value. Therefore, though giving nonuniform shadow matte value may be more accurate for shadow removing, for image forgery detection, the uniform shadow matte value is sufficient to reflect the illumination condition in the shadow.

B. Performance Under Different Lighting Conditions

To verify the applicability of the proposed approach, we apply the method in different cases where shadows exist: under the street lamp at night with a simple textured shadow receiving surface (Fig. 7 Upper), in the sunlight in the day with simple (Fig. 7 Middle), and complex (Fig. 7 Bottom) textured shadow receiving surfaces. In all the cases, we first capture an image (Fig. 7, first column) of a scene with a shadow in it using a camera shown in the first column of Fig. 7. With the camera standing still, we move the shadow casting object out of the view. Then we capture an image of the same scene as shown in the second column of Fig. 7. The interval between the time when the two images are taken is very short. Therefore, they are assumed to share the same lighting condition. Then, the proposed method is applied to the shadowed image to estimate the

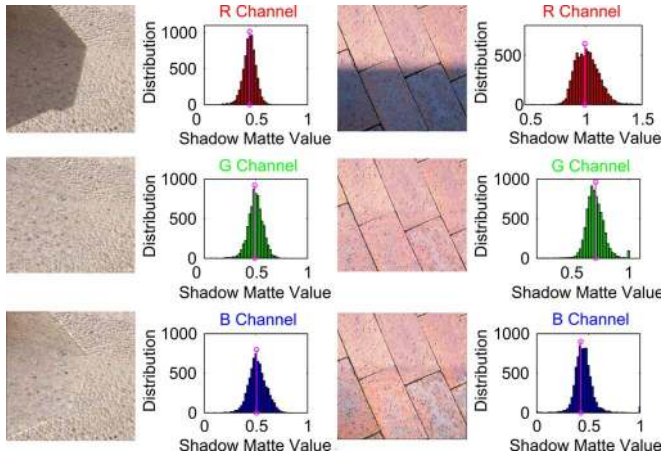


Fig. 6. Compare of the nonuniform [40] and uniform shadow matte value. The distribution of nonuniform shadow matte value is centered around our uniform shadow matte value.




Fig. 7. Effectiveness evaluation of the proposed method under different conditions. The images in the first column capture scenes with shadows in them, and the images in the second column capture the same scenes with the shadows moving out of the scenes. The interval of the time between the two images are taken with the camera standing stationary is very short. The third column shows images after we remove the shadows in the first column using the proposed method.

shadow matte value based on which the shadow is removed, resulting in the third column in Fig. 7.

The similarity in the shadow region between recovered image after removing the shadow using the proposed method and the originally non-shadowed image indicates the effectiveness of the method. The similarity can be identified by three statistical indicators: the mean and the standard deviation of the pixel wise difference rate (MDR and SDR) between the intensities of the two images in the shadow region, the normalized cross correlation (NCC) of the intensity between the two images in the shadow region. MDR and SDR are defined by the following equations:

$$MDR = \sum \frac{\mathcal{I}_r(i) - \mathcal{I}_n(i)}{N \cdot \mathcal{I}_n(i)} \quad (16)$$

$$SDR = \sqrt{\frac{1}{N} \sum \left(\frac{\mathcal{I}_r(i) - \mathcal{I}_n(i)}{\mathcal{I}_n(i)} - MDR \right)^2} \quad (17)$$

where \mathcal{I}_r and \mathcal{I}_n represent the intensity of the recovered shadow free image and the originally non-shadowed image, respectively, i denotes the index of the pixels in the shadow

TABLE I
PERFORMANCE EVALUATION UNDER DIFFERENT CONDITIONS

Condition	Channel	MDR	SDR	NCC
Night	Red	0.0271	0.0605	0.9077
	Green	0.0227	0.0438	0.9101
	Blue	0.0375	0.0721	0.9178
Day	Red	-0.0398	0.0899	0.9435
	Green	-0.0329	0.0831	0.9476
	Blue	-0.0307	0.0789	0.9461
Texture	Red	-0.1038	0.0869	0.9245
	Green	-0.1029	0.0791	0.9241
	Blue	-0.0901	0.0787	0.9156



Fig. 8. Examples of shadow detection. The detected shadow boundaries are marked by red pixels in the images.

region, and N is the number of the pixels in the shadow region. MDR and SDR are calculated in each channel separately.

Table I shows the statistical results of MDR, SDR, and NCC for the three channels of each case shown in the images of Fig. 7. Low value of MDR and SDR and high value of NCC indicate low difference between the recovered shadow free image and the originally non-textured image, which demonstrates the effectiveness of the proposed method in different conditions.

C. Shadow Boundary Detection and Penumbra Shadow Identification

In the above subsections, we obtain the shadow matte value utilizing pixels sampled near the shadow boundaries and demonstrate the effectiveness of the proposed method using both synthesized and real images. To implement the proposed method, however, three problems have to be solved: 1) How do we get the boundary of the shadows? 2) How can the normals of the pixels on the boundaries be estimated? 3) How can we identify the penumbra shadow region?

The shadow boundaries can be marked manually. However, users are not always well trained and they should not be bothered to mark them themselves. So, automation in detecting shadow boundaries is desired. Any shadow detection algorithm that suits the scene of the image can be used, such as approaches based on invariant color models [48] and methods utilizing prior knowledge of the scene [49]. In our implementation, we apply a probabilistic approach suggested in [47]. Fig. 8 shows some results of the shadow detection. The detected shadow boundaries are marked by red pixels in the images.

Detecting shadows automatically is a really tough task in current research state [50], not only because the complex layout of the nature images, but also due to many near black objects that are similar to shadows. So, the whole boundary of a shadow may not be successively extracted and some of the boundaries of the near black objects may be falsely recognized




Fig. 9. Calculating normals on the shadow boundary. (a) Detected non-smooth shadow boundary. (b) Shadow boundary represented in binary image. (c) Use many ten-pixel line segments to fit the shadow boundary in (b). (d) Calculated normals of every line segment in (c).

as the boundaries of shadows. Automatically sampling shadows of two boundaries and comparing them has a non-trivial possibility of comparing between shadows and dark objects. So an interaction is introduced. After automatically obtaining the shadow boundaries in an image, we select one point in the suspected shadow and another outside the shadow. Then m points along the shadow boundary around the intersection point between the line connecting the selected two points and the shadow boundary are chosen for sampling.

The extracted shadow boundaries are not always smooth [Fig. 9(a) and (b)]. To find the normal of the points on a shadow boundary, simply using a polynomial to fit the curve of the boundary cannot work well. A low degree polynomial cannot capture the necessary detail of the curve, while a high degree polynomial might lead to serious over-fitting problem. To find the normals on the boundary with satisfying accuracy as well as computation simplicity, we divide the shadow boundary into many sub-curves with ten successive pixels, then fit each of them with a line segment [Fig. 9(c)]. Because of the aliasing, some of line segments do not look very straight. The normal of every point on the line segments are set to be the normal of its corresponding line segment [Fig. 9(d)].

The identification of the penumbra shadow region is crucial in our method. In our implementation, the width of the penumbra shadow w along the boundary of the shadow is assumed to be constant. So the penumbra shadow region is obtained by extending w pixels bidirectionally in the normal directions of pixels on the shadow boundary. The blue curve in Fig. 10(a) shows a typical log domain intensity distribution near the shadow boundary along a single spline. The pixel index indicates the offset between each pixel on the spline and the shadow boundary pixel. By averaging the intensity of all the splines, we get a smoother curve shown by the red curve in Fig. 10(a). It is observed that the average log domain intensity in both the full shadow and lit regions do not oscillate largely, while it changes gradually from lower values in full shadow region to higher values in lit region. Noticing that the shape of the red curve is similar to that of rescaled and shifted logistic sigmoid function defined in (18), we fit the red curve using this function:

$$S(x) = \frac{1}{1 + \exp(-\alpha(i - \beta))} \quad (18)$$

where α and β are coefficients governing the expansion and shift of the standard logistic sigmoid function on the horizontal axis. Since the domain of the logistic sigmoid function is $[0, 1]$,




Fig. 10. Penumbra shadow identification. (a) A typical curve of the log domain intensity of a single spline (blue curve) and the curve of the average log domain intensity over all the splines (red curve). (b) The normalized average log domain intensity curve (red curve) and the fitted curve of logistic sigmoid function (green curve). The interval between the two magenta dashed lines represents the position of the penumbra shadow region.

to fit the red curve in Fig. 10(a), the average log domain intensity should be normalized first shown using the following equation:

$$ns(i) = \frac{\bar{s}(i)}{\max(\bar{s}) - \min(\bar{s})} \quad (19)$$

where s denotes the average log domain intensity over all the splines on a shadow, ns indicates the normalized intensity, and i is the intensity index. The curves of \bar{s} and ns are shown by the red curve in Fig. 10(a) and the red curve in Fig. 10(b), respectively. α and β can be evaluated after fitting the normalized average log domain intensity curve ns using the logistic sigmoid function.

Our goal is to find the width of the penumbra shadow region, which can be obtained by estimating the width of the gradually changing region of the logistic sigmoid function approximating the curve of normalized average log domain intensity. By analysis, the width of the penumbra shadow region can be identified as $4(\alpha\sqrt{\pi/8})^{-1}$. Taking the offset into consideration, the interval of the index of the penumbra shadow region should be $[-2(\alpha\sqrt{\pi/8})^{-1} + \beta, 2(\alpha\sqrt{\pi/8})^{-1} + \beta]$, which is shown by the interval between the two red dashed lines in the case shown in Fig. 10(b). More results can be found in Figs. 1, 11–13, and 17.

D. Forgery Detection Using the Shadow Matte Value

In our work, we apply the shadow matte value C as the cue for forgery detection. We sample shadows in an image and extract their shadow matte values, respectively. By comparing the difference between the C values of two shadows, we can find if there exists an inconsistency and suspect whether the image is tampered.

Estimating the shadow matte value of a shadow is performed on all the three channels of an image separately, hence forming a three-dimensional vector $\mathbf{C} = \{C_r, C_g, C_b\}$ referred to as shadow matte vector. We define the inconsistency vector $\mathbf{D} = \{D_r, D_g, D_b\}$ between according element of two shadow matte vectors $\mathbf{C}(x)$ and $\mathbf{C}(y)$ as follows:

$$D_k = \exp(\sigma) \cdot (\exp(-C_k(x)) - \exp(-C_k(y))) \quad (20)$$

where $k \in \{r, g, b\}$. We omit k and denote D_k as D in the rest of this paper. The exponential functions in (20) are utilized to draw back the nonlinear distortion of the difference caused by

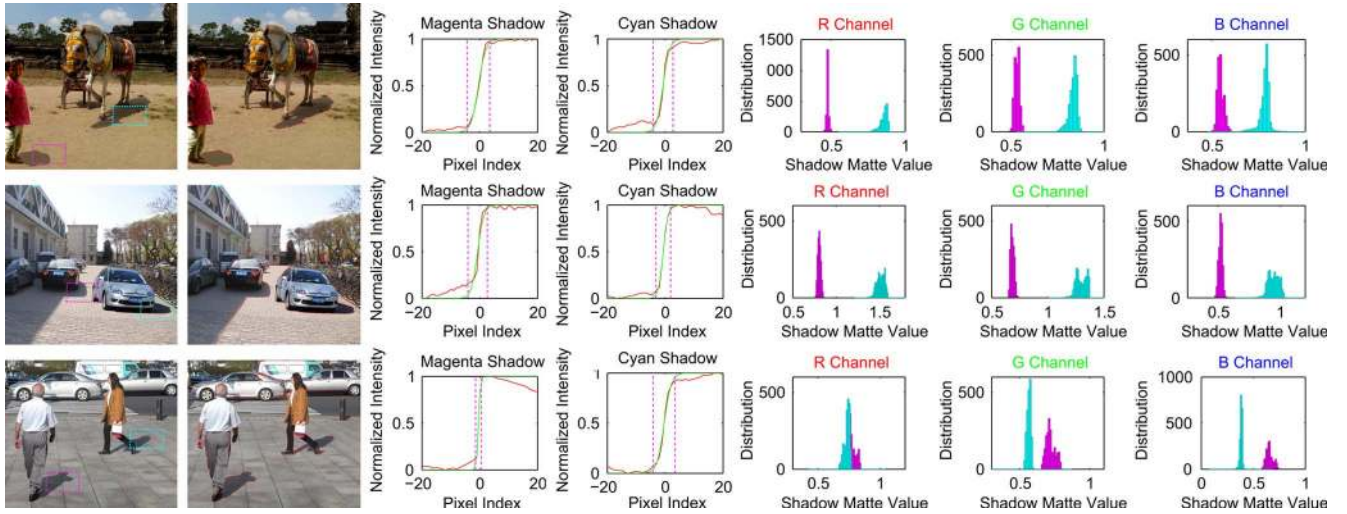


Fig. 11. Experimental results on forged images with simply textured shadow receiving surfaces. The first column shows the image to be detected. The child in the first row, the front car in the second row, and the woman in the third row are forged. The shadows sampled in each image are marked by magenta and cyan boxes. The second column shows the detected shadows. The following two columns show the penumbra width. The last three columns are the distributions of $f_n - f_s$ for the marked shadows of red, green, and blue channels, respectively.

the log domain. σ is a scale factor, which is 2.7 obtained by experiments. We identify the inconsistency due to measurement error to satisfy a Gaussian distribution with the density function $D \sim N(0, 1)$:

$$\varphi(D) = \frac{1}{\sqrt{2\pi}} e^{-D^2/2}. \quad (21)$$

The inconsistency due to measurement error that fall out of the interval $[-1.96, 1.96]$ has a very low probability to happen. So we define the tolerance interval as $[-1.96, 1.96]$ and any image having inconsistency fall out of the tolerance interval is considered to be forged.

IV. EXPERIMENTAL RESULTS

In this section, we utilize our method to image forgery detection and verify the effectiveness of our proposed method using real photos. Fig. 11 shows examples of forgery detection on forged images where shadows are casted onto shadow receiving surfaces with simple texture. The first column in Fig. 11 shows the images to be detected. Two shadows in each image are sampled and marked by magenta and cyan boxes, respectively. The second column shows the detected shadow boundaries in each image. The following two columns give the fitted sigmoid functions and the penumbra width for the two shadows. From the results of Figs. 11–13, as well as Fig. 17, it is observed that the logistic sigmoid function approximates the normalized log domain intensity pretty well. Then, our method is applied to the sampled shadows to find their C values. The last three columns are plots of the distribution of $f_n - f_s$ in red, green, and blue channels generated when performing our method. The two distributions of the sampled shadows of the same image in each channel are drawn in one plot and colored with the same colors marking the shadows. The difference between the two distributions intuitively describes the difference of the two shadows in that channel.

TABLE II
FORGERY DETECTION ON IMAGES IN FIG. 11

Image	Channel	Magenta C	Cyan C	Inconsistency	IN/OUT
Top	Red	0.4785	0.8574	2.9083	OUT
	Green	0.5355	0.8354	2.2569	OUT
	Blue	0.5474	0.7845	1.8172	IN
Middle	Red	0.8064	1.5157	3.3743	OUT
	Green	0.6851	1.2977	3.4352	OUT
	Blue	0.5213	0.9456	3.0551	OUT
Bottom	Red	0.7687	0.7339	0.2442	IN
	Green	0.7196	0.5653	-1.2086	IN
	Blue	0.6561	0.3829	-2.4253	OUT

Table II shows the possibility of the inconsistency for the above experiments. The fifth column in the table shows the inconsistency calculated by (20) using the third and fourth columns. If the inconsistency falls into the tolerance interval, the last column will be “IN”. Otherwise, it will be “OUT”. If the last column of either of the three channels of a image is “OUT”, the image is considered to be fake.

Fig. 12 and Table III show examples of forgery detection on forged images where shadows are casted onto shadow receiving surfaces with complex texture. Fig. 12 and Table III are organized the same as Fig. 11 and Table II.

If an object from another image is copied and pasted onto a target image together with an artificial shadow generated for it, the shadow matte values of the artificial shadow and another sampled shadow should indicate an inconsistency. However, if the two sampled shadows are all real and originally exist in the target image, our method should not detect an inconsistency. To verify this, we sampled two shadows marked by cyan and magenta boxes in each of the scenes in Fig. 13 which are known as authentic. The results are shown in Fig. 13 and Table IV.

To obtain quantitative evaluation, we perform our method on a database. The database is generated using ten source images and ten target images. Both the source images and the target images should have cast shadows in them. One object with its shadow in each of the source image is extracted manually. Some

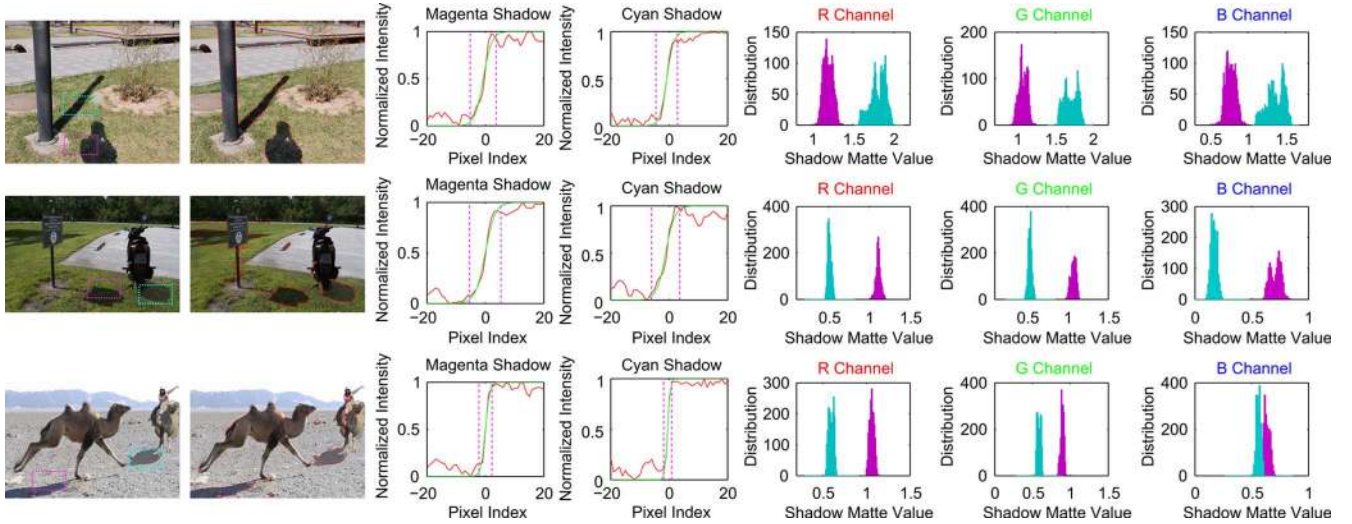


Fig. 12. Experimental results on forged images with complexly textured shadow receiving surfaces. The human shadow, the motorcycle, and the camel riding by a girl are all forged. The figure is organized the same to Fig. 11.

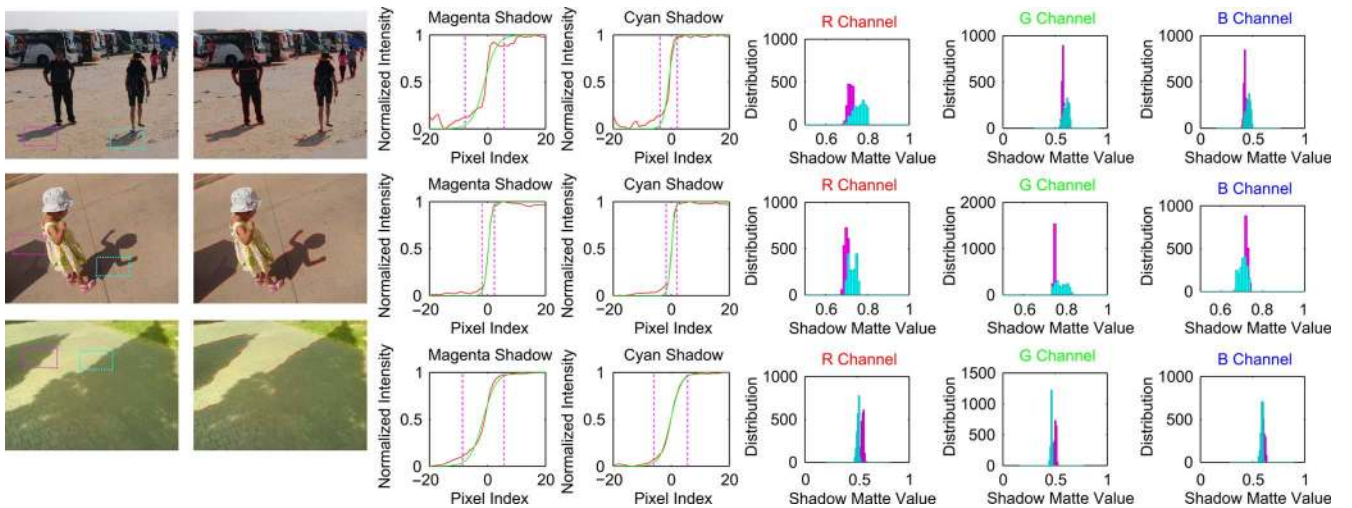


Fig. 13. Experimental results on images with two authentic shadows. The figure is organized the same as Fig. 11.

TABLE III
FORGERY DETECTION ON IMAGES IN FIG. 12

Image	Channel	Magenta C	Cyan C	Inconsistency	IN/OUT
Top	Red	1.1894	1.8066	2.0858	OUT
	Green	1.0737	1.7045	2.3789	OUT
	Blue	0.7707	1.3649	3.0839	OUT
Middle	Red	1.1080	0.5071	-4.0473	OUT
	Green	1.0711	0.5336	-3.6285	OUT
	Blue	0.7095	0.1666	-5.2768	OUT
Bottom	Red	1.0620	0.5964	-3.0500	OUT
	Green	0.8914	0.5834	-2.2007	OUT
	Blue	0.6348	0.5599	-0.6128	IN

TABLE IV
FORGERY DETECTION ON IMAGES IN FIG. 13

Image	Channel	Magenta C	Cyan C	Inconsistency	IN/OUT
Top	Red	0.7209	0.7600	0.2771	IN
	Green	0.5767	0.6067	0.2470	IN
	Blue	0.4260	0.4571	0.2979	IN
Middle	Red	0.6981	0.7282	0.2195	IN
	Green	0.7486	0.7802	0.2188	IN
	Blue	0.7175	0.7040	-0.0982	IN
Bottom	Red	0.5546	0.5094	-0.3956	IN
	Green	0.5032	0.4659	-0.3423	IN
	Blue	0.6025	0.5887	-0.1121	IN

examples of the extracted objects with their shadows are shown in the first column in Fig. 14. The target images are chosen with different lighting conditions, textures on the ground, and scene layout. Examples of them are shown in the second column of Fig. 14. Then the objects with their shadows are copied from the source images and pasted onto the target images, resulting in 100 source and target image pairs. From each image pair, a fake image is generated. Operations such as resizing, stretching,

and turnover are used on the source objects to make the generated images more visually pleasing. The third column in Fig. 14 gives some examples of the generated images. Also, another 40 shadowed nature images under different illumination conditions without forgery are added into the database, examples of which are shown in the last column in Fig. 14.

After obtaining the image database, the proposed method is performed on it. True Positives are measured as the number of



Fig. 14. Fake photo database generated using ten source images and ten target images. The first column shows examples of the source images. The second column shows selected images in the target images. The third column is examples of images forged by copying and pasting the objects in the source images into the target images. The last column shows examples of images without forgery added into the database.




Fig. 15. ROC curve generated given different values of σ in (20).

genuine images recognized as images without forgery and the False Positives are measured as the number of forged images recognized as genuine. The ROC curve obtain using our method is shown by the blue curve in Fig. 15. The green curve is obtained using the method in [38]. The adjusting parameter used for generating the ROC curve is σ in (20) ranging from 2.0 to 5.0. We define the accuracy as

$$Accuracy = \frac{TP + \frac{100}{40} * TN}{200} \quad (22)$$

where TP is the number of True Positives and TN is the number of True Negatives. Though the images in the database are complex in illumination conditions, textures of the ground, and scene layout, our method still obtains the very pleasing accuracy as 87.0% when $\sigma = 2.7$.

Utilizing the constructed database, we also perform statistics on the shadow matte value of real shadows and forged shadows. The distributions of the three channels for real shadows and forged shadows in the database are shown in the first two columns of Fig. 16, respectively. As can be seen, the shadow matte values for both real shadows and forged shadows are irregularly distributed. Therefore, one can hardly anticipate the shadow matte value in an image by experience. As a result, when forging images, it is hard to maintain the consistency of shadows. Statistics is also performed on the shadow matte value

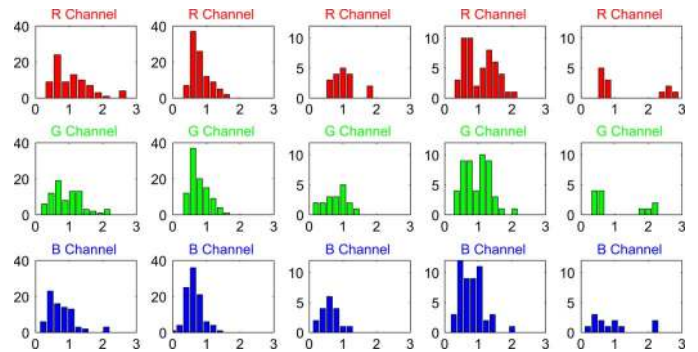


Fig. 16. Shadow matte value distribution for real and forged shadows and for shadows on different grounds. The horizontal axis is the shadow matte value and the vertical axis gives the distribution. The first two columns are the distributions for real shadows and forged shadows. The last three columns are the distributions for shadows on earth, concrete ground, and grass, respectively.

of different ground plane. The last three columns of Fig. 16 give the shadow matte distributions on the ground of earth, concrete, and grass. Shown in the figure, the shadow matte values in each kind of the ground is irregularly distributed.

Note that in the object cut-and-paste operations, the target texture under shadow might be different from the one under shadow in the source image. Therefore, cut-and-paste may introduce dissimilar textures or photometric profiles across the shadow boundary. It is possible to resort to other shadow invariant features as cues for composite authentication. In the proposed framework, the estimated C may not represent shadow matte value specifically. However, the C value consistency remains hard to be enforced for the forger. Instead of cut-and-paste, the more recent solution to compose a shadow onto a host image is through matting [51], [41]. Although it is getting mature in the academic community, it might not be easy for an average user to stick to the shadow consistency before an easy-to-use commercial tool is available.

Failures of our method exist for both the doctored images and authentic images. For the forged images, it is hard for our method to work when the composite shadows are consistent with the real target shadows. For example, see the upper row in Fig. 17 and Table V. In the real world, our method would also fail when a shadow is copied and pasted to another position in the same image as shown in Fig. 18. Fortunately, this forgery can be detected by other complementary detection, such as identical objects recognition [52], [53]. For the authentic images, our method fails to work when the shadow matte consistency assumption is not satisfied. Bottom rows in Fig. 17 and Table V give a very challenging example. The shadows of trees in this scene are very complex. The shadow is very blurry, and many holes exist inside of the shadow boundaries. Our method fails in this case.

V. CONCLUSION

In this paper, we have proposed a new framework for detecting tampered digital images based on photometric consistency of illumination in shadows. We detected the shadow boundaries in an image and also the penumbra region width of the shadow. Then we estimated the shadow matte values for each of the sampled shadows in an image to be tested and used the consistency of them to inform if the image is doctored. Both

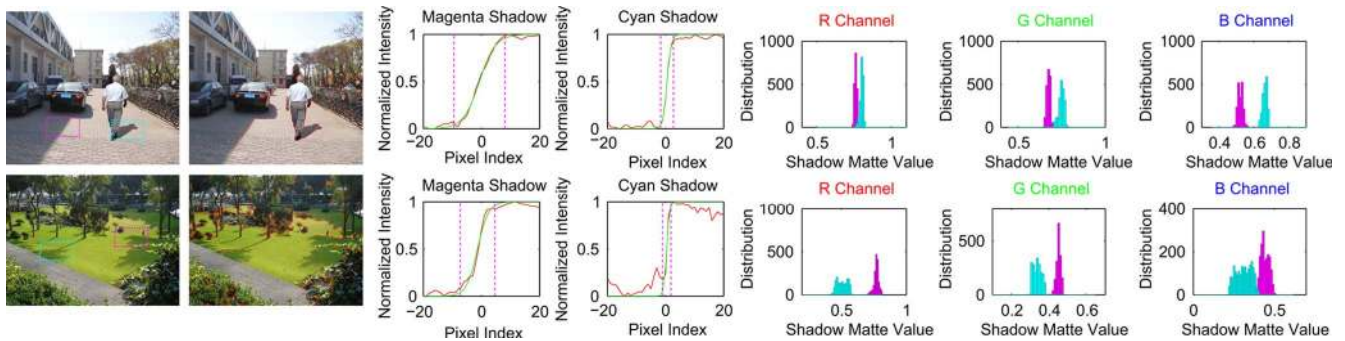


Fig. 17. Failed examples using our method. The figure is organized the same as Fig. 11.




Fig. 18. Our method cannot work in the situation that the shadow is copied and pasted from the same image as the target image.

TABLE V
FORGERY DETECTION ON IMAGES IN FIG. 17

Image	Channel	Magenta C	Cyan C	Inconsistency	IN/OUT
Top	Red	0.7611	0.8005	0.2684	IN
	Green	0.6740	0.7484	0.5437	IN
	Blue	0.5221	0.6630	1.1604	IN
Bottom	Red	0.7696	0.5095	-2.0471	OUT
	Green	0.4518	0.3383	-1.1383	IN
	Blue	0.4441	0.3129	-1.3379	IN

simulated images and visually plausible real images are used to demonstrate our approach. Our method is computational simple and effective. As has been discussed, many approaches have been proposed for digital forgery detection. However, no method can be perfect and solve all kinds of forgery. Though our method can identify whether an image is tampered, one limitation of our method is that it can not determine which part of the image is doctored. What we need is to integrate many different algorithms for this complicated problem and draw a more cumulative and convincing conclusion. We argue that shadows in the real world carry much more clues for forgery identification that we have not noticed, especially in the physical aspect. Our future work is to extract these features and utilize them in digital image forgery detection.

REFERENCES

- [1] A. R. Smith and J. F. Blinn, "Blue screen matting," *Proc. ACM SIGGRAPH*, vol. 5, pp. 259–268, 1996.
- [2] A. Swaminathan, M. Wu, and K. J. R. Liu, "Non-intrusive forensic analysis of visual sensors using output images," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2006, pp. 401–404.
- [3] A. Swaminathan, M. Wu, and K. J. R. Liu, "Image tampering identification using blind deconvolution," in *Proc. Int. Conf. Image Processing*, 2006, vol. 1.
- [4] A. Swaminathan, M. Wu, and K. J. R. Liu, "Component forensics for digital camera: A non-intrusive approach," in *Proc. Conf. Info. Sciences and Systems*, 2006, pp. 1194–1199.
- [5] I. Cox, M. Miller, and J. Bloom, *Digital Watermarking*. San Mateo, CA: Morgan Kaufmann, 2002.
- [6] W. Luo, F. Huang, and J. Huang, "Edge adaptive image steganography based on LSB matching revisited," *IEEE Trans. Inf. Forensics Security*, vol. 5, no. 3, pp. 201–214, 2010.
- [7] H. Farid, "A survey of image forgery detection," *IEEE Signal Process. Mag.*, vol. 2, no. 26, pp. 16–25, 2009.
- [8] B. Mahdian and S. Saic, "Detection of copy-move forgery using a method based on blur moment invariants," *Forensic Sci. Int.*, vol. 171, no. 2, pp. 180–189, 2007.
- [9] A. Popescu and H. Farid, Exposing Digital Forgeries by Detecting Duplicated Image Regions, Dept. Comput. Sci., Dartmouth College, 2004, Tech. Rep. TR2004-515.
- [10] J. Fridrich, D. Soukal, and J. Lukas, "Detection of copy-move forgery in digital images," in *Proc. Digital Forensic Research Workshop*, Cleveland, OH, 2003.
- [11] A. Popescu and H. Farid, "Exposing digital forgeries by detecting traces of re-sampling," *IEEE Trans. Signal Process.*, vol. 53, no. 2, pp. 758–767, Feb. 2005.
- [12] T. T. Ng and S. F. Chang, "A model for image splicing," in *Proc. Int. Conf. Image Processing*, 2004, pp. 24–27.
- [13] T. T. Ng and S. F. Chang, "Blind detection of photomontage using higher order statistics," in *Proc. IEEE Int. Symp. Circuits and Systems*, 2004, pp. 688–691.
- [14] W. Chen, Y. Q. Shi, and W. Su, "Image splicing detection using 2-D phase congruency and statistical moments of characteristic function," in *Proc. Society of Photo-Optical Instrumentation Engineers Conf.*, San Jose, CA, Feb. 2007.
- [15] H. Farid, Digital Image Ballistics From JPEG Quantization: A Followup Study, Dept. Comput. Sci., Dartmouth College, 2008, Tech. Rep. TR2008-638.
- [16] J. Lukás and J. Fridrich, "Estimation of primary quantization matrix in double compressed JPEG images," in *Proc. Digital Forensics Research Workshop*, Cleveland, OH, Aug. 2003, pp. 5–8.
- [17] A. C. Popescu and H. Farid, "Statistical tools for digital forensics," in *Proc. 6th Int. Workshop Information Hiding*, Berlin-Heidelberg, Germany, 2004, pp. 128–147.
- [18] W. Luo, J. Huang, and G. Qiu, "JPEG error analysis and its applications to digital image forensics," *IEEE Trans. Inf. Forensics Security*, vol. 5, no. 3, pp. 480–491, 2010.
- [19] M. K. Johnson and H. Farid, "Exposing digital forgeries through chromatic aberration," in *Proc. Workshop Multimedia and Security*, New York, 2006, pp. 48–55.
- [20] S. Bayram, H. T. Sencar, and N. Memon, "Source camera identification based on CFA interpolation," in *Proc. Int. Conf. Image Processing*, 2005, vol. 3, pp. 69–72.
- [21] Z. Lin, R. Wang, X. Tang, and H. Y. Shum, "Detecting doctored images using camera response normality and consistency," in *Proc. IEEE Int. Conf. Computer Vision Pattern Recognition*, 2005, pp. 1087–1092.
- [22] H. Gou, A. Swaminathan, and M. Wu, "Noise features for image tampering detection and steganalysis," in *Proc. Int. Conf. Image Processing*, 2007, vol. 6, pp. 97–100.
- [23] M. Chen, J. Fridrich, M. Goljan, and J. Lukás, "Determining image origin and integrity using sensor noise," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 74–90, 2008.
- [24] J. Lukás, J. Fridrich, and M. Goljan, "Determining digital image origin using sensor imperfections," in *Proc. Society of Photo-Optical Instrumentation Engineers Conf.*, Mar. 2005, vol. 5685, pp. 249–260.

- [25] A. Swaminathan, M. Wu, and K. J. R. Liu, "Digital image forensics via intrinsic fingerprints," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 101–117, 2008.
- [26] M. K. Johnson and H. Farid, "Detecting photographic composites of people," in *Digital Watermarking*. New York: Springer, 2008, pp. 19–33.
- [27] W. Wang and H. Farid, "Detecting re-projected video," in *Proc. Int. Workshop Information Hiding*, Santa Barbara, CA, 2008.
- [28] W. Zhang, X. Cao, Z. Feng, J. Zhang, and P. Wang, "Detecting photographic composites using two-view geometrical constraints," in *Proc. IEEE Int. Conf. Multimedia and Expo*, 2009, pp. 1078–1081.
- [29] W. Zhang, X. Cao, Y. Qu, Y. Hou, H. Zhao, and C. Zhang, "Detecting and extracting the photo composites using planar homography and graph cut," *IEEE Trans. Inf. Forensics Security*, vol. 5, pp. 544–555, Sep. 2010.
- [30] M. Johnson and H. Farid, "Exposing digital forgeries by detecting inconsistencies in lighting," in *Proc. ACM Multimedia and Security Workshop*, New York, 2005.
- [31] M. K. Johnson and H. Farid, T. Furon, F. Cayre, G. J. Doërr, and P. Bas, Eds., "Exposing digital forgeries through specular highlights on the eye," in *Proc. Int. Workshop Information Hiding, ser. Lecture Notes in Computer Science*, 2008, vol. 4567, pp. 311–325.
- [32] M. K. Johnson and H. Farid, "Exposing digital forgeries in complex lighting environments," *IEEE Trans. Inf. Forensics Security*, vol. 2, no. 3, pp. 450–461, 2007.
- [33] D. Kersten, P. Mamassian, and D. C. Knill, "Moving cast shadows induce apparent motion in depth," *Perception*, vol. 26, pp. 171–192, 1997.
- [34] Y.-Y. Chuang, D. B. Goldman, B. Curless, D. H. Salesin, and R. Szeliski, "Shadow matting and compositing," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 494–500, 2003.
- [35] X. Cao, Y. Shen, M. Shah, and H. Foroosh, "Single view compositing with shadows," *Visual Comput.*, vol. 21, pp. 639–648, 2005.
- [36] W. Zhang, X. Cao, J. Zhang, J. Zhu, and P. Wang, "Detecting photographic composites using shadows," in *Proc. IEEE Int. Conf. Multimedia and Expo*, 2009, pp. 1042–1045.
- [37] C. Fredembach and G. Finlayson, "Simple shadow removal," in *Proc. Int. Conf. Pattern Recognition*, 2006, pp. 823–836.
- [38] E. Arbel and H. Hel, "Texture-preserving shadow removal in color images containing curved surfaces," in *Proc. IEEE Int. Conf. Computer Vision Pattern Recognition*, 2007, pp. 1–8.
- [39] Y.-F. Su and H. H. Chen, "A three-stage approach to shadow field estimation from partial boundary information," *IEEE Trans. Image Process.*, vol. 19, no. 10, pp. 2749–2760, Oct. 2010.
- [40] E. A. Hel-Or, "Shadow removal using intensity surfaces and texture anchor points," *IEEE Trans. Pattern Anal. Mach. Intell.*, to be published.
- [41] T.-P. Wu, C.-K. Tang, M. S. Brown, and H.-Y. Shum, "Natural shadow matting," *ACM Trans. Graph.*, vol. 26, no. 2, pp. 1–21, Jun. 2007.
- [42] G. D. Finlayson, M. S. Drew, and C. Lu, "Intrinsic images by entropy minimization," in *Proc. Eur. Conf. Computer Vision*, 2004, pp. 582–595.
- [43] M. F. Tappen, W. T. Freeman, and E. H. Adelson, "Recovering intrinsic images from a single image," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 9, pp. 1459–1472, Sep. 2005.
- [44] Y. Weiss, "Deriving intrinsic images from image sequences," in *Proc. IEEE Int. Conf. Computer Vision*, 2001, pp. 68–75.
- [45] Y. Yacoob and L. S. Davis, "Segmentation using appearance of mesostructure roughness," *Int. J. Comput. Vision*, vol. 83, no. 3, pp. 248–273, 2009.
- [46] T. Weyrich, J. Lawrence, H. Lensch, S. Rusinkiewicz, and T. Zickler, "Principles of appearance acquisition and representation," *Found. Trends Comput. Graph. Comput. Vision*, vol. 4, pp. 75–191, 2008.
- [47] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan, "Detecting ground shadows in outdoor consumer photographs," in *Proc. Eur. Conf. Computer Vision*, Crete, Greece, 2010.
- [48] E. Salvador, A. Cavallaro, and T. Ebrahimi, "Cast shadow segmentation using invariant color features," *Comput. Vision Image Understand.*, vol. 95, no. 2, pp. 238–259, 2004.
- [49] M. D. Levine and J. Bhattacharyya, "Removing shadows," *Pattern Recognit.*, vol. 26, no. 3, pp. 251–265, 2005.
- [50] J. Zhu, K. G. G. Samuel, S. Z. Masood, and M. F. Tappen, "Learning to recognize shadows in monochromatic natural images," in *Proc. IEEE Int. Conf. Computer Vision Pattern Recognition*, 2010, pp. 223–230.
- [51] Y.-Y. Chuang, D. B. Goldman, B. Curless, D. H. Salesin, R. Szeliski, and B. C. D. H. Salesin, "Shadow matting and compositing," *ACM Trans. Graph.*, vol. 22, pp. 494–500, 2003.
- [52] M. Cho, Y. M. Shin, and K. M. Lee, "Unsupervised detection and segmentation of identical objects," in *Proc. IEEE Int. Conf. Computer Vision Pattern Recognition*, 2010, pp. 1617–1624.
- [53] C. Zhang, X. Guo, and X. Cao, "Duplication localization and segmentation," in *Proc. Pacific-Rim Conf. Multimedia*, 2010, pp. 578–589.



Qiguang Liu was born in Gansu, China, in 1988. He received the B.E. degree in software engineering in the School of Computer Software, Tianjin University, Tianjin, China. He is currently pursuing the M.E. degree with School of Computer Science and Technology, Tianjin University.

He has been working as a research assistant in the Computer Vision Lab since April 2010. His research interests include computer vision, computer graphics, and multimedia forensics.



Xiaochun Cao received the B.E. and M.E. degrees, both in computer science, from Beihang University (BUAA), Beijing, China. He received the Ph.D. degree in computer science from the University of Central Florida, Orlando, with his dissertation nominated for the university-level award for Outstanding Dissertation.

After graduation, he spent about two and half years at ObjectVideo Inc. as a Research Scientist. Since August 2008, he has been with Tianjin University, Tianjin, China, where he is currently a Professor of computer science. He has authored and coauthored over 40 peer-reviewed journal and conference papers and has been in the organizing and the technical committees of several international colloquia.

Dr. Cao was the recipient of the Piero Zamperoni best student paper award at the International Conference on Pattern Recognition in 2004 and 2010.



Chao Deng was born in China in 1988. He received the B.E. degree in software engineering from the School of Computer Software, Tianjin University (TJU), Tianjin, China, in 2009. He is currently pursuing the M.E. degree with the School of Computer Science and Technology, TJU.

He has been working as a research assistant in the Computer Vision Laboratory since 2009. His research interests lie in the computer vision field, covering mainly image processing, multimedia forensic, and pattern recognition.



Xiaojie Guo was born in China in 1986. He received the B.E. degree in software engineering from the School of Computer Science and Technology, Wuhan University of Technology, Wuhan, China, in 2008. He is currently pursuing the Ph.D. degree with the School of Computer Science and Technology, Tianjin University (TJU), Tianjin, China.

He has been working as a research assistant in the Computer Vision Lab since 2008. His research interest lies in computer vision, covering mainly image understanding, pattern recognition, and multimedia

forensics.

Mr. Guo was the recipient of the Piero Zamperoni best student paper award in the International Conference on Pattern Recognition 2010.