Shadow Removal Using Intensity Surfaces and Texture Anchor Points

Eli Arbel and Hagit Hel-Or, *Member*, *IEEE*

Abstract—Removal of shadows from a single image is a challenging problem. Producing a high-quality shadow-free image which is indistinguishable from a reproduction of a true shadow-free scene is even more difficult. Shadows in images are typically affected by several phenomena in the scene, including physical phenomena such as lighting conditions, type and behavior of shadowed surfaces, occluding objects, etc. Additionally, shadow regions may undergo postacquisition image processing transformations, e.g., contrast enhancement, which may introduce noticeable artifacts in the shadow-free images. We argue that the assumptions introduced in most studies arise from the complexity of the problem of shadow removal from a single image and limit the class of shadow images which can be handled by these methods. The purpose of this paper is twofold: First, it provides a comprehensive survey of the problems and challenges which may occur when removing shadows from a single image. In the second part of the paper, we present our framework for shadow removal, in which we attempt to overcome some of the fundamental problems described in the first part of the paper. Experimental results demonstrating the capabilities of our algorithm are presented.

Index Terms—Shadow removal, shading, color, texture, shadow detection, region growing, enhancement.

1 Introduction

The problem of shadow removal has been studied by researchers in many contexts: derivation of illumination intrinsic images, reconstruction of photographic-quality shadow-free images, improving performance of various algorithms such as recognition algorithms, shadow matting, etc. In this paper, we consider the following goal: Given a single shadow image, produce a high-quality shadow-free image which is perceived as having been acquired in the same scene but without shadows.

Following the formulation in [1], an image $\mathcal{I}(x,y)$ is considered to be composed of the albedo $\mathcal{R}(x,y)$ and illumination $\mathcal{L}(x,y)$ fields as follows:

$$\mathcal{I}_k(x,y) = \mathcal{R}_k(x,y) \cdot \mathcal{L}_k(x,y), \tag{1}$$

where $k \in \{R,G,B\}$ and $\cdot$ denotes pixelwise multiplication.

Denoting $\hat{\mathcal{L}}_k(x,y)$ as the illumination field without shadows, $\mathcal{L}_k(x,y)$ can be expressed as

$$\mathcal{L}_k(x,y) = \hat{\mathcal{L}}_k(x,y) \cdot \mathcal{C}_k(x,y), \tag{2}$$

where $\mathcal{C}_k(x,y)$ represents the shadow intensities or shadow scale factors of channel k. This gives rise to the common shadow image formulation:

$$\mathcal{I}_k(x,y) = \mathcal{R}_k(x,y) \cdot \hat{\mathcal{L}}_k(x,y) \cdot \mathcal{C}_k(x,y). \tag{3}$$

Shadow removal is often performed in the log domain, thus (3) is reformulated:

$$I_k(x,y) = R_k(x,y) + \hat{L}_k(x,y) + C_k(x,y), \tag{4}$$

where I, R, $\hat{L}$, and C are the logarithms of $\mathcal{I}$, $\mathcal{R}$, $\hat{\mathcal{L}}$, and $\mathcal{C}$, respectively.

The classic approach to shadow removal in a single image is to estimate the shadow scale factors, $\mathcal{C}_k(x,y)$, whether explicitly in the image domain or implicitly in the gradient domain, and then to remove the shadows by canceling the effect of the shadow scale factors in the image (by division using (3) or by subtraction using (4)).

While former studies have demonstrated impressive results on specific examples, virtually all studies embody assumptions under which the methods can be applied successfully. These might include lighting conditions, acquisition device properties, type and behavior of shadowed surfaces, statistics of shadow pixels, and many more. Consequently, most methods are capable of producing good results only on a subset of possible images, namely, those in which the assumptions hold.

The goal of the first part of this paper is to provide a comprehensive survey of the problems and challenges related to shadow removal from a single image. As demonstrated in this part, some of the issues are fundamental to shadow removal and have a significant influence on the design of shadow removal algorithms, as well as on the quality of the shadow-free images obtained by the algorithms.

In the second part of the paper, we present a novel framework for shadow removal which was developed in light of the issues presented in the first part. We show that if each image channel is considered as an intensity surface, approximating the shape of the intensity surface in shadow regions can aid in obtaining shadow-free images of high quality, regardless of whether shadows are uniform or nonuniform, and regardless of intensity surface geometry. Furthermore, we show that this approach enables the preservation of the original texture in shadow-free regions

• *The authors are with the Department of Computer Science, University of Haifa, Haifa 31905, Israel.*
E-mail: eliarbel@gmail.com, hagit@cs.haifa.ac.il.

Manuscript received 18 July 2009; revised 7 Jan. 2010; accepted 29 May 2010; published online 18 Aug. 2010.
Recommended for acceptance by R. Ramamoorthi.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2009-07-0461.
Digital Object Identifier no. 10.1109/TPAMI.2010.157.

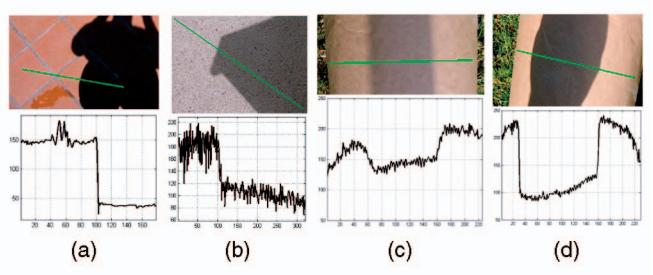

Fig. 1. Examples of uniform and nonuniform shadows. Plots show intensity along the line marked in the image above. (a) Uniform shadow cast on flat surface. (b) Nonuniform shadow cast on flat surface. Note that the shadow is darker in the right part of the image. (c) Uniform shadow cast on curved surface. The geometry of the surface is preserved in the shadow region. (d) Nonuniform shadow cast on curved surface. Intensity change is inconsistent with the curvature of the surface.

and wide penumbra areas. As a final and independent step in our shadow removal algorithm, we perform shadow-free region enhancement, which aims at further improving the quality of the final shadow-free image, reducing the effects of noise and image processing transformations which may introduce artifacts in shadow-free regions. Finally, although we concentrate, in this paper, on the problem of shadow removal, we also describe a simple and effective method for shadow mask derivation. The process relies on minimal user input that provides cues on shadow and nonshadow surfaces, which are later used to classify shadow and nonshadow pixels.

2 Problems and Challenges in Shadow Removal

In this section, we enumerate the various problems and challenges related to the task of shadow removal. It is worth noting that a given shadow image does not necessarily include all of the phenomena mentioned below, and indeed, in many of the images we explored only a subset of phenomena occurs. However, in order to develop a robust shadow removal algorithm which can effectively handle shadow images acquired under different conditions and of different scene types, any shadow removal algorithms should account for the various types of possible phenomena which may affect the final result.

2.1 Physical Phenomena

Physical phenomena occur in the physical world and obviously affect the digital representation of the scene.

2.1.1 Shadow Intensity

A shadowed surface is a part of the surface which is occluded from at least one direct light source in the scene. As a result, a reduction in light intensity is observed in shadow regions. Many methods attempt to remove shadows by first estimating (either explicitly or implicitly) the amount of intensity reduction in the shadow region (the *shadow intensity*) and deducing the corresponding *shadow scale factors* (i.e., $C(x, y)$ in (3)). The shadows are then removed by applying the inverse transformation on the shadow regions according to the shadow scale factors.

Two possible cases may be considered with respect to shadow intensity: The first is where shadow intensity is uniform in the shadow region, resulting in a *uniform shadow*. The second case is where shadow intensities vary across a shadow region, yielding a *nonuniform shadow*. See Fig. 1. The phenomenon of varying shadow intensities usually occurs due to ambient light and is most common in scenes where the occluding object is close to the shadowed surface; thus, less ambient light reaches the inner regions of the shadow than the outer parts. Interreflections are another source of nonuniformity of shadows and can be caused by the occluding object itself or by other objects in the scene.

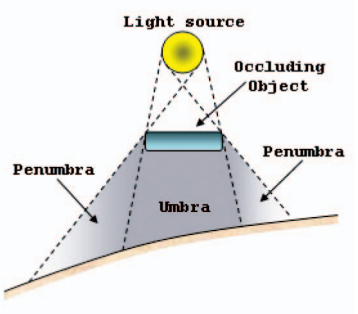

Fig. 2. Umbra and penumbra of a shadow.

Determining shadow intensity usually involves estimation of the shadow scale factor. In the case of a uniform shadow, the scale factor is a single unknown; however, in the case of a nonuniform shadow, the scale factor is spatially varying and a per-pixel estimate must be determined. Figs. 8a and 20b show the effects of incorrectly assuming a uniform shadow when attempting to remove a nonuniform shadow.

2.1.2 Umbra and Penumbra

A shadow region can be partitioned into umbra and penumbra regions. Fig. 2 illustrates the formation of shadow umbra and penumbra. The umbra of a shadow is the part of the shadowed surface in which the direct light source is completely obscured by the occluding object. The penumbra of a shadow is the part of the surface where the light source is only partially occluded. Shadow intensities typically change smoothly in penumbra regions when transitioning from the umbra to the nonshadowed region of the surface. Penumbra occurs when the light source is not a point source or due to diffraction of light rays caused by the occluding object [2], [3].

Regardless of whether the shadow intensity is uniform or not in the umbra, by definition, shadow intensities vary in penumbra regions. The width of the penumbra, as well as the rate of illumination change across the penumbra, varies in a given shadow region and among different shadow regions (consider Figs. 3 and 4). In some cases, penumbra width is very small and difficult to detect in digital images, in which case the penumbra is referred to as a *hard shadow edge*. However, in many natural images, the penumbra is noticeably wide; thus, special handling in the shadow removal process might be required.

2.1.3 The Light Source

The type and shape of the light source is another factor which may influence shadow removal algorithms. Algorithms that

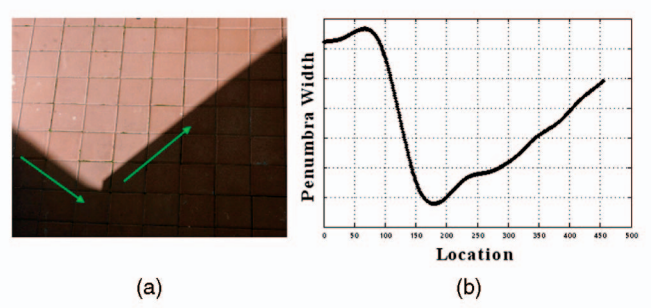

Fig. 3. Varying width of shadow penumbra. (a) Shadow image. (b) Plot of penumbra width along the shadow boundary in the direction depicted by the arrows in the left image.

assume a certain spectral power distribution, for instance, a Planckian light source [4], may fail in handling indoor shadow images acquired under artificial illumination. In cases where the light source is not a point source or more than one light source exists, as can occur in indoor images, for example, complex soft shadows may appear (for example, Fig. 22) that affect shadow removal algorithms considerably.

2.2 Scene Characteristics

In addition to the physical phenomena described above, the nature of the acquired scene and the objects in it can also play a crucial role in shadow removal. While it is fairly reasonable to assume certain behavior of illumination, and subsequently, of shadows, e.g., that shadow intensities inside the umbra are locally constant [5], [6], it is unrealistic to assume global behavior of all possible scenes. This implies that in the context of shadow removal, the complexity of a shadow image is derived virtually from the complexity of the acquired scene. The type of surfaces in the scene, object geometry and configuration, and their interaction with the light source all have influence on shadow removal algorithms.

2.2.1 Self-Shadows and Shading

Shadows and shading take prominent roles in our visual understanding of the world. They supply numerous cues which assist in depth perception, object positioning [7], etc. Our perception of object geometry is also greatly affected by shading and illumination. In particular, it is the *self-shading* which gives us the strongest cues about object geometry.

Since self-shadows and shading usually arise from a direct light source and rarely by ambient light, they are

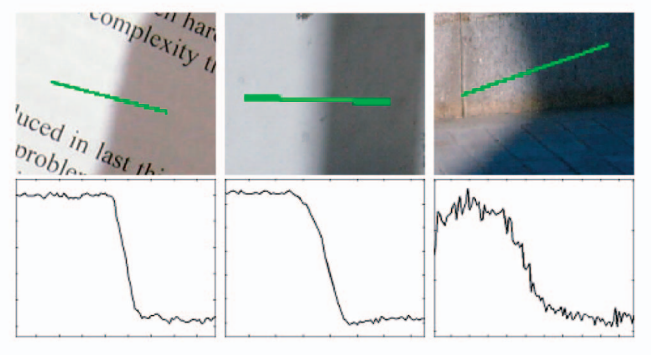

Fig. 4. Examples of different penumbra cross sections.

Fig. 5. Shadow penumbra reconstruction—mean value is reconstructed correctly but high-order textural information is lost.

absent in shadowed regions where the light source is occluded, causing information loss in those areas. Removing the shadow does not restore the shading cues since the information is inherently missing in the original shadow image. In these cases, an unnatural shadow-free image is often produced. An example is shown in Fig. 8b.

2.2.2 Complexity of Shadowed Surface

Several factors related to the surfaces upon which the shadow is cast contribute to the complexity of the shadow removal problem. First, algorithms that rely on mean pixel values may fail in handling shadows that span different surfaces since this implies differences of pixel value statistics.

An additional factor to consider is the textural content of the shadowed surface. Since textured surfaces usually incorporate high-order statistical information, removing shadows from textured surfaces using a linear shadow removal process might not yield satisfactory results. For example, in addition to using scale factors for removing a shadow from a given image, the variance and higher order statistics must be reconstructed in the shadow-free region. To further illustrate this, consider Figs. 8c and 8d. As can be seen in these images, the mean values are correctly reconstructed (using the scale factors approach), however, the shadow-free regions still contain artifacts that require a high-order reconstruction. Another example is shown in Fig. 5, where the mean value is reconstructed in the penumbra but high-order textural information is lost.

In addition, shadow surfaces with texture may undergo various image processing transformations that would require special handling in the shadow removal process. This issue is further discussed in Section 2.3.

2.2.3 Geometry of the Shadowed Surface

Shadows cast on curved surfaces (e.g., the common case of highlights and shadows on faces—see Fig. 28) might also pose problems for shadow removal algorithms. Linear methods that rely on first-order statistics can fail in removing shadows on curved surfaces since first-order statistics such as mean values vary across curved surfaces in shadow regions as well as in nonshadow regions (see Figs. 1c and 1d). An example of removing shadow cast on curved surface using pixel statistics is given in Fig. 7.

2.2.4 Intersection of Shadow and Reflectance Boundaries

The process of shadow removal invariably involves dealing with the shadow boundary. This necessitates distinguishing

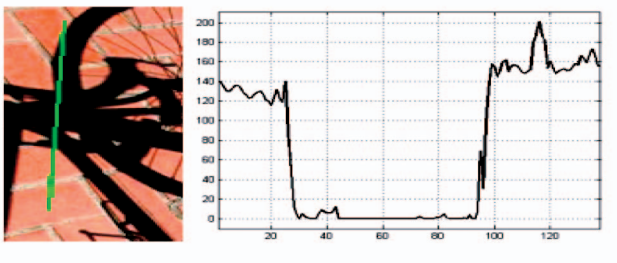

Fig. 6. Example of intensity clipping in shadow pixels. The intensity variation due to the brick texture in the shadow region along the marked line is lost due to clipping and binning of dark values.

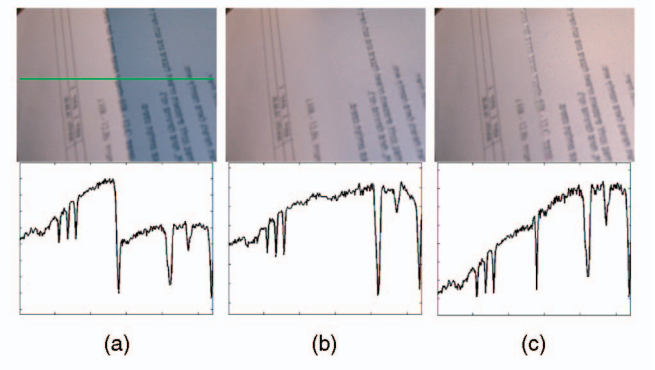

Fig. 7. Example of problems with gradient domain shadow removal from curved surface. (a) Shadow image. (b) Shadow-free image obtained by nullifying shadow gradients. Note that the text at the shadow boundary is missing since reflectance gradients are also nullified. Also note that the geometry of the surface is not preserved (consider the corresponding cross section). This is due to the fact that the scale factor is estimated using only first-order pixel statistics. (c) Shadow-free image obtained using the method described in [9]. Note that the geometry and texture of the surface is preserved since a high-order model is used without nullifying image gradients.

between shadow edges and reflectance edges. Reflectance edges that cross or coincide with shadow edges (whether sharp shadow edges or wide penumbra) must be restored consistently with the same type of reflectance edges outside the shadow. This difficulty is especially significant for algorithms that work in the gradient domain since the gradients in such edges are composed of both reflectance and shadow changes [6], [8], thus requiring the algorithm to modify only the shadow term of the shadow edge gradient. Fig. 7 demonstrates the effect of information loss when nullifying shadow gradients that intersect with reflectance gradients. Note that shadows on textured surfaces fall into this category as well.

2.3 Image Acquisition and Processing

The phenomena described above occur in the physical world independent of the capturing device. In this section, we describe issues related to the acquisition and digitization pipeline which may influence shadow removal algorithms.

2.3.1 *Capturing Device*

Perhaps the most prominent influence of the capturing device in shadow images is the presence of sensor noise introduced in dark shadow regions, yielding low signal-to-noise ratio. While this noise may be scarcely visible in the original shadow image due to low intensities, it may be enhanced by the shadow removal algorithm and strongly affect the quality of the final shadow-free image (see Fig. 8c).

Information loss also occurs due to clipping of pixel intensities caused by the limited range of camera sensors (see Fig. 6), as well as by quantization of similar valued dark points in shadow regions into the same quantization bin. Shadow removal algorithms that scale the values in shadow regions can not overcome the quantization effects and may produce clipping artifacts and possibly false contours.

2.3.2 *Postprocessing*

All cameras, whether high or low end, perform some form of image processing within the acquisition pipeline in an attempt to produce high-quality and pleasing images. This includes producing high-quality shadow images. The acquisition pipeline typically involves processing such as color balancing, tone mapping, and highlight and shadow toning. While improving the quality of shadow images, such transformations may pose challenging problems to shadow removal algorithms. These transformations are often inconsistent with the shadow model and the processing used by shadow removal algorithms. Fig. 9 shows an example where image contrast enhancement of the original image produces unpleasing contrast effects in the shadow-free image.

In addition to the processing within the acquisition pipeline, images are commonly compressed by the camera into some standard output format—the most popular by far being the JPEG compression standard. JPEG compression introduces noticeable artifacts in images. While these artifacts may be unnoticed in dark shadow regions, they form a significant problem in shadow removal as their appearance may be enhanced as a result of the shadow removal process, as can be seen in Fig. 8d.

To conclude this section, we note that a rigorous statistical analysis on the frequency and significance of the affects described in this section is in order. Unfortunately, it is beyond the scope of this paper. We note, however, that the observations described in this section are based on our experience with several hundred images. We have observed that over 50 percent of the shadows are nonuniform shadows, and most of them require postprocessing enhancement. Nonuniform shadows were found to be common in scenes acquired at mid to short range with more dominant affects in scenes with man-made structures. Shadows rarely straddled more than one surface; however, textured shadowed surfaces were found to be very common. Penumbra width, in our images, varied between 2 pixels and 15 pixels (as in the soft shadows in Fig. 22). Finally, we find that effects of image acquisition appear in various combinations and affect shadow removal to various degrees. The most difficult affect in our images was due to image contrast enhancement, which required postprocessing enhancement.

3 Related Work

Shadow detection and removal has been approached from numerous aspects including shadow detection and removal

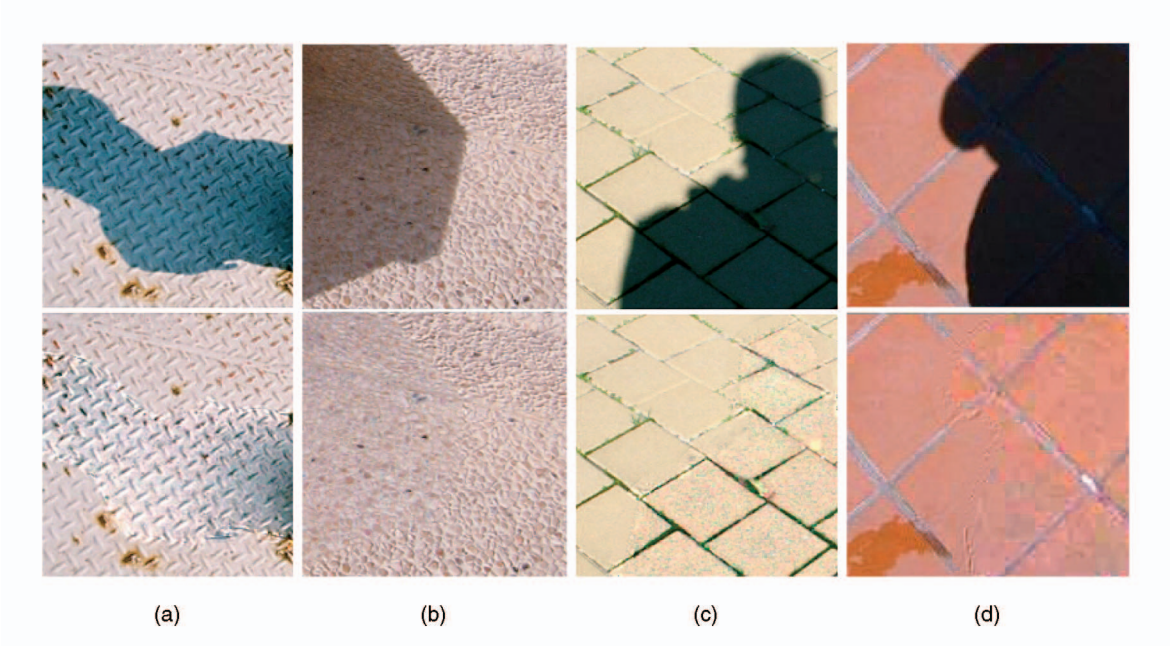

Fig. 8. Examples of problems in shadow removal. (a) Removing nonuniform shadow assuming uniform shadow intensity. (b) Absence of self-shading in the shadow-free region. (c) Enhancement of noise in the shadow-free region. (d) JPEG artifacts.

from multiple images and video sequences [10], [11], [12], based on special camera filters [13], based on object models—in the context of tracking, e.g., cars and pedestrians [14], [15], [16], [17], and shadow removal from projected environments [18]. In this paper, we focus on the most common use, namely, shadow removal from a single image.

Shadow removal involves two basic stages: detection of shadow regions, typically expressed in the form of detecting shadow edges, and the removal of shadows from an image. The shadow problems and issues described in Section 2 affect both the detection and the removal of shadows; however, in this paper, we focus on the removal of shadows, and the method presented in Section 4 assumes that shadow boundaries are given. Methods for automatic shadow detection in a single image can be found in [4], [19], [20], [21]. A user-guided method for extracting shadow regions is described in Section 6.

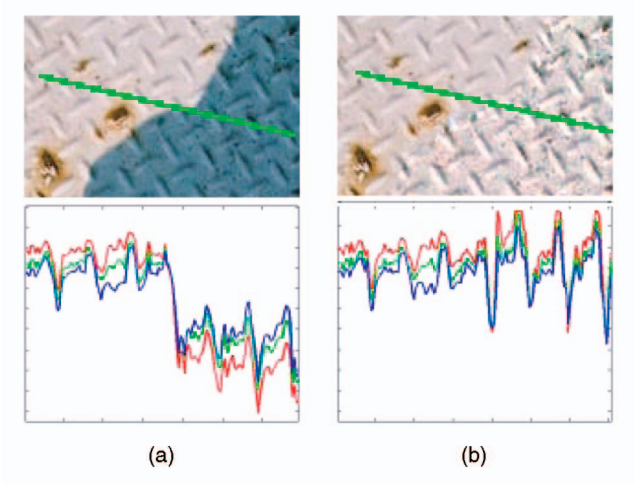

Fig. 9. The effect of image processing on shadow-free regions. (a) Details are enhanced in the shadow region, yet the image appears natural and pleasing. (b) Removing the shadow yields an unnatural image in which the contrast in the shadow-free region is inconsistent with that of the nonshadow region, as depicted in the corresponding cross sections.

Shadow removal methods for a single image can be classified into two categories: methods operating in the gradient domain [4], [6], [8], [22], [23], [24], [25] and methods operating in the image intensity domain [5], [9], [26], [27], [28], [29], [30], [31].

Shadow removal based on the gradient domain was suggested by Finlayson et al. [4], [23]. The core idea in these studies is to nullify gradients of shadow edges and then reconstruct the shadow-free image by integration, assuming a certain type of light source and special properties of camera sensors. In [24] and [25], studies that relax the camera properties assumption are presented. While making a big leap in automatically removing shadows from a single image, this approach suffers from an inherent problem of the gradient-based shadow removal algorithms, which is related to the global integration step [10]. The integration usually results in artifacts such as changes in color balance and global smoothness of the reconstructed image (see Fig. 26). Being aware of the problems due to global integration, Fredembach and Finlayson [22] suggested a shadow removal algorithm in which 1D integration is performed instead of global integration.

Although impressive results are presented, the nullification of shadow edge gradients causes textural information loss in penumbra regions that must be restored artificially, whether by nullified gradients [4] or in-painting [22], [23]. Removal of shadows cast on curved surfaces with wide penumbra regions is also strongly affected by nullification of shadow edges. Fig. 7 shows an example of shadow removal from a curved surface using nullification of shadow boundary gradients.

In an attempt to develop shadow removal algorithms based on the gradient domain that preserve textural information in penumbra regions, Xu et al. [6] and Mohan et al. [8] suggested shadow removal methods which naturally extend that of Finlayson et al.'s in [4]. The method described in [6] reconstructs the penumbra regions by clipping large gradients, assuming that they are due to object boundaries, whereas small gradients are due to changes in illumination. In [8], the authors suggest a user-guided approach which tries to model soft shadow edges in the intensity domain by assuming symmetric, sigmoidal-like behavior of shadows across penumbra regions (consider Fig. 4 for a counter example), and then using the derivatives of the shadow model to remove gradients that are due to shadow boundaries.

While the examples given in [6] and [8] (see Fig. 26) show that the methods are effective in handling shadows with soft boundaries, being based on the gradient domain introduces inherent problems with the approach. The gradient domain methods for shadow removal modify only the gradients in penumbra regions; thus, these methods cannot handle nonuniform shadows as this implies changes in illumination inside the shadow region, and not only at the shadow boundaries, as assumed by these methods. Since postacquisition image processing transformations may introduce artifacts in umbra regions, these gradient-based algorithms do not handle such artifacts as well.

Another approach to shadow removal from a single image is based on the intensity domain. A simple intensity domain shadow removal method was proposed by Baba et al. [27] and [28]. The method is based on color and variance adjustment of shadow pixels in RGB space, assuming a single flat texture shadow surface.

The authors of [32] describe a method in which light occlusion factors are used for shadow removal. Occlusion factors are estimated in the intensity domain and further smoothed in the gradient domain to obtain a smooth shadow mask. The initial estimation of the occlusion factors is obtained by assuming planar and roughly constant-value surfaces on which shadows are cast.

Two intensity domain methods have been proposed by Finlayson et al. [29], [26]. In the study described in [29], the Retinex theory is employed for shadow removal where large changes in intensities which are due to shadow boundaries are ignored in the Retinex computation. The method in [26] is based on the estimation of shadow scale factors assuming uniform shadow intensities and hard shadows. Both methods use in-painting for completion of missing information in shadow boundary regions.

In [31], a method for shadow removal that uses a Pulse Coupled Neural Network is presented. As with the studies of [26], [27], [28], this method relies on first-order statistics for determining a single scale factor of shadow regions, and thus cannot properly handle nonuniform shadows and shadows cast on curved surfaces.

Wu et al. [5], [30] described a method for shadow removal in the context of shadow matting. The authors suggest a method that estimates shadow intensities based on shadow and nonshadow intensity ratios in the umbra, and use a Bayesian framework for regularization of shadow scale factors in the umbra and penumbra regions. The method is capable of removing soft shadows while preserving texture at shadow boundaries, assuming roughly uniform shadow in umbra regions.

Another intensity-based method for shadow removal which is capable of handling shadows with wide penumbra cast on textured and curved surfaces was proposed in our earlier study [9]. We suggested a model in which pixel intensities in the image form a so-called intensity surface and, by assuming uniform shadows, the proposed method finds shadow scale factors using a high-order model, namely, cubic smoothing splines across penumbra regions.

Although high-quality results have been demonstrated by some of the studies described above, they do not seem to handle some of the fundamental problems described in Section 2, namely, nonuniform shadows, curved and textured surfaces, etc. Furthermore, in our experiments, we observed that many shadow images undergo post-acquisition transformations which severely affect the final shadow-free results. Previously suggested shadow removal methods do not address this problem at all.

In this paper, we suggest a novel shadow removal algorithm which is capable of handling nonuniform shadows, as opposed to our previous work in [9] that can only handle uniform shadow. Our suggested algorithm can handle shadows cast on flat or curved surfaces which may also be textured. Note that by curved surfaces we refer to surfaces where pixel intensities vary considerably across the surface, regardless of the actual form of the surface in the physical scene (for example, consider Figs. 1c and 1d). In addition, we outline an approach for enhancing shadow-free regions containing artifacts of postacquisition transformations and noise.

4 Shadow Removal

Shadow removal is typically performed in two stages: 1) the detection stage in which shadow regions are detected, specifically by determining the shadow boundaries, and 2) the reconstruction stage in which the shadow is actually removed and a shadow-free image is produced. In this section, we suggest a novel reconstruction stage for shadow removal, i.e., removing the shadows in an image once they have been detected. Any shadow detection algorithm can be used, but since our method is not confined to images with certain illumination conditions such as outdoor scene images, one could use a shadow detection algorithm that best suits the illumination conditions in a given image.

The main theme in our approach is the notion that image data should not be nullified at any stage of the process; rather, image content should be preserved and, if necessary, modified. In addition to the unpleasing global effects in the image, the gradient-based methods often produce irrecoverable artifacts along the shadow edge (see Fig. 5), and they only modify the gradients of penumbra regions. We require that our method be capable of dealing with nonuniform shadows, i.e., modify the gradients in umbra regions as well. Thus, we use an intensity rather than a gradient approach. We also require that our method be able to handle varying penumbra width as well as profile and shadows on curved and textured surfaces.

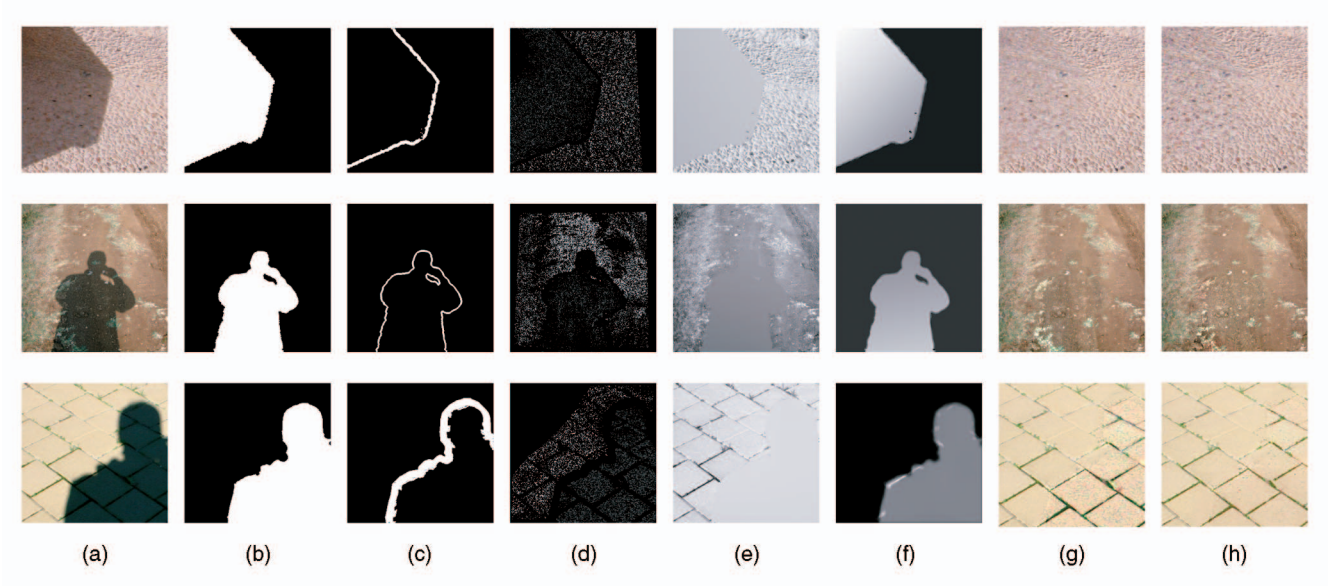

Fig. 10. Stages of the proposed method. (a) Original. (b) Shadow mask. (c) Penumbra mask. (d) Anchor points. (e) Reconstructed intensity surface. (f) Shadow scale factor image. (g) Shadow-free image before enhancement. (h) Final shadow-free image following region enhancement.

The proposed method for shadow removal is based on our previous method described in [9] in which Cubic smoothing splines are used across the penumbra regions for finding the correct scale factor. The method in [9] assumes similar pixel intensities on both sides of the shadow boundary; thus, it is capable of removing uniform shadows only. This presents a significant restriction on shadow removal as uniformity cannot be assumed in many scenes. Fig. 20 shows the failure of the method in [9] when shadows are nonuniform. In this paper, we follow the same model as in [9], where pixel intensities form an intensity surface. However, we adopt a global encompassing approach that computes the entire intensity surface concurrently. Using a single model for approximation of shadow regions allows an estimate of a per-pixel scale factor, i.e., *handling nonuniform shadows*. This approach accommodates penumbra of varying width and profile and is able to deal with shadows on curved surfaces as well as preserving texture across shadow boundaries.

4.1 Algorithm Summary

Algorithm 1 outlines the steps and the general data flow of the shadow removal algorithm described in this paper. The various steps of the process are detailed in the following sections. Fig. 10 shows the stages of the algorithm implemented on a shadowed image.

Algorithm 1. Shadow removal algorithm flow
Input: I - RGB image
Output: $I_{ShadowFree}$ - Shadow-free RGB image
1: $\{\text{ShdwPix}, \text{NonShdwPix}\} \Leftarrow$ getUserShadowPix(I)
2: $\{M_s, M_r\} \Leftarrow$ getMasksBySVM(ShdwPix,NonShdwPix,I)
3: $M_p \Leftarrow$ MRFpenumbraLabel(M_s, I_k) // $M_p \subseteq M_s$
4: $\{\hat{M}_s, \hat{M}_r\} \Leftarrow$ getAnchorPointsMasks(M_s, M_r, I)
5: $S_{edge} \Leftarrow$ EdgeDetector(M_s)
6: **for** $k \in R, G, B$ **do**
7: $\hat{C}_k(x,y) \Leftarrow$ calcScaleFactors($\hat{M}_s, \hat{M}_r, I_k$)
8: $C_k(x,y) \Leftarrow$ fillNonAnchorPixels($\hat{C}_k(x,y)$)
9: $C_k(x,y) \Leftarrow$ directionalSmooth($M_p, S_{edge}, C_k(x,y)$)
10: $I_{Fixed} \Leftarrow$ applyScaleFactors($I_k, C_k(x,y)$)
11: $I_k \Leftarrow M_s \cdot I_{Fixed} \cup (1 - M_s) \cdot I_k$
12: **end for**
13: $I \Leftarrow I_R \cup I_G \cup I_B$
14: $I_{ShadowFree} \Leftarrow$ shadowFreeRegionEnhancement(I)

The algorithm proceeds as follows: The user marks shadow and nonshadow coordinates interactively on the input image (line 1). The selected RGB values are then used for shadow and nonshadow mask extraction, using SVM and region growing as described in Section 6 (line 2). Using MRF labeling [9] (see the Appendix), the penumbra mask is calculated from the shadow mask and image (line 3). Given the shadow and surround masks, anchor point selection is performed as described in Section 4.4 (line 4). Finally, shadow edges are extracted from the shadow mask (line 5) to be used later for directional smoothing of penumbra scale factors.

The core of the process is the shadow removal phase, which consists of several steps (lines 7-11) applied to each channel independently (following the description in Sections 4.2 and 4.3). First, the shadow scale factors in the umbra M_s are calculated, as well as an initial estimate of the scale factors in the penumbra M_p using the intensity surface approximation method and (7)-(10) (lines 7 and 8). Directional smoothing of the shadow scale factors within the penumbra is then performed (line 9) using the shadow edge information extracted in line 5 (see Section 4.5). The shadow in the image is removed by adding the channel image with the scale factors image as in (4) (line 10) and combining with the original nonshadow pixels (line 11).

The three processed RGB channels are combined to produce a shadow-free image (line 13). The final shadow-free image is produced by applying the shadow-free region enhancement algorithm described in Section 5 (line 14).

4.2 Intensity Surface Approximation

The main part of the proposed shadow removal process involves the estimation of the shadow scale factors $C(x,y)$ in (4) (lines 7-11 in Algorithm 1). Solving for $C(x,y)$ in (4) is

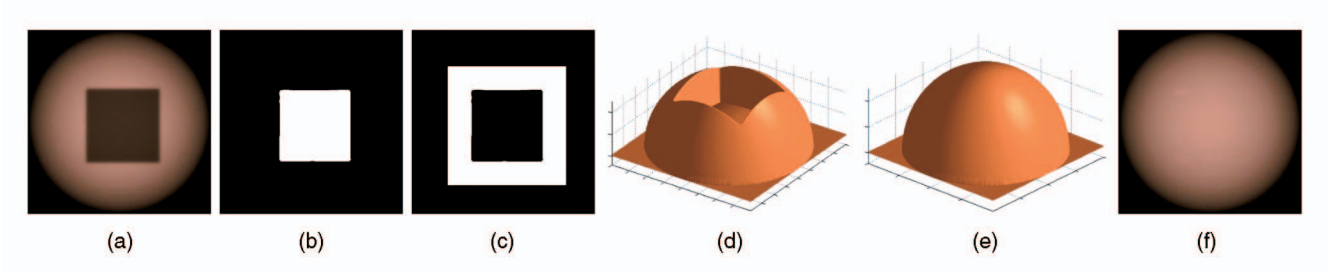

Fig. 11. An example of an approximated shadow-free intensity surface. (a) Synthetic image of sphere with shadow. (b) Shadow mask. (c) Surround mask. (d) Intensity surface of the shadow image. (e) Intensity surface of the approximated shadow-free surface. Note that the approximated surface is smooth in the shadow region and correctly exhibits the sphere geometry. (f) Reconstructed shadow-free image.

an ill-posed problem, and a smoothing constraint must be introduced. This is performed by considering the pixel intensities of each channel as forming an intensity surface[1] (see Figs. 11d and 12b) and deriving a smooth approximation of the shadow-free intensity surface within the umbra and penumbra regions from which smooth shadow scale factors can be determined. The approximation is obtained using a thin-plate surface model (e.g., [33], [34], [35], [36]).

Let M_s and M_r denote shadow and surround masks, respectively (Figs. 11b and 11c). Note that M_s should include both umbra and penumbra regions of the shadow. We assume that a surround mask exists and is wide enough to capture scene statistics, and that not all shadow boundaries coincide with reflectance boundaries.

Denote by $z = f(x, y)$ the shadow-free intensity surface represented by a smooth thin-plate surface in the shadow region. We define an energy function on f that measures the smoothness of the surface:

$$E_s(f) = \int\int_{\Omega} \left(\frac{\partial^2 f}{\partial x^2}\right)^2 + 2\left(\frac{\partial^2 f}{\partial x \partial y}\right)^2 + \left(\frac{\partial^2 f}{\partial y^2}\right)^2 dxdy, \qquad (5)$$

where Ω is the spatial domain over which f is defined (in our case, $M_s \cup M_r$). Thus, a good approximation of a shadow-free intensity surface is one that minimizes E_s, i.e., finds a smooth surface with minimal curvature, in the shadow region.

The approximated surface should coincide with the intensity surface in the nonshadowed regions surrounding the shadow (M_r). Thus, the following data term is minimized over the surround pixels defined by mask M_r:

$$E_d(f) = \int\int_{\Omega} \omega(x, y) \cdot [f(x, y) - I(x, y)]^2 dxdy,$$
$$\omega(x, y) = \begin{cases} 1, & (x, y) \in M_r, \\ 0, & \text{otherwise.} \end{cases} \qquad (6)$$

Approximation of the shadow-free intensity surface is found by solving the following functional:

$$f = \underset{\hat{f}}{\operatorname{argmin}} \; E_s(\hat{f}) + E_d(\hat{f}). \qquad (7)$$

A scale factors field can then be calculated over the mask M_s:

$$C(x, y) = I(x, y) - f(x, y) \qquad (8)$$

from which a shadow-free image can be derived.

1. Note that the intensity surface describes the change in pixel values and not necessarily the surface geometry, although the two are correlated.

An example is shown in Fig. 11. The reconstructed approximated surface f follows the global geometry of the intensity surface, emphasizing the advantage of using the thin-plate model in reconstructing shadow-free images in cases of shadows on curved surfaces as well as nonuniform shadows, as can be seen in Figs. 24c and 21.

4.3 Textured Surfaces

Consider the textured surface in Fig. 12a. Solving (7), a smooth intensity surface f is obtained at pixels in M_s (Fig. 12b). Although the global geometry of the shadow-free surface is obtained, f poorly approximates the desired shadow-free intensity surface as the textural information of the shadow region is lost. Accordingly, applying (8) produces a nonsmooth shadow factors field $C(x, y)$, as shown in Fig. 12c, and an incorrect shadow-free reconstruction. This case often arises in shadow images, namely, when shadows are cast on textured or highly structured surfaces.

The approach taken in this study is to carefully select *anchor points* in M_s and M_r on which the thin-plate minimization will be applied. Specifically, anchor points in M_s and M_r should originate from the same intensity distribution, thus supplying the same information source within the shadow and outside the shadow. A schematic description is shown in Fig. 13. A cross section of the intensity surface of Fig. 12a (blue) is overlaid with the corresponding cross section of the smooth intensity surface f obtained naively based on all surround pixel points M_s (black curve)—Fig. 13a (top). Due to the textured content in the shadow region, the resulting scale factors $C(x, y)$ are not smooth—Fig. 13a (bottom). Carefully choosing corresponding anchor points in M_s and M_r produces a smooth intensity surface (red) and a smooth

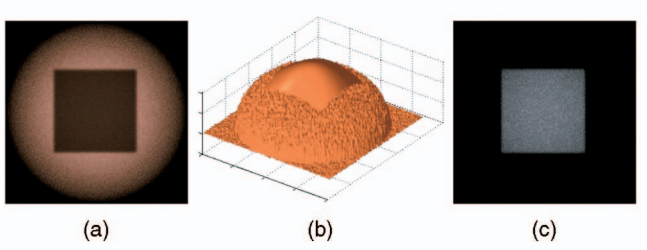

Fig. 12. Intensity surface reconstruction with texture. (a) Synthetic image of noisy sphere with shadow. (b) Intensity surface of the approximated shadow-free surface. (c) The shadow scale factors. Note that the approximated surface is smooth, creating a nonsmooth shadow scale factor field.

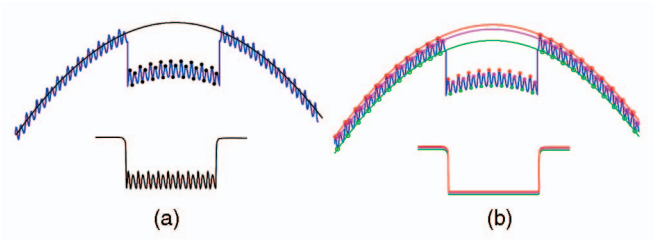

Fig. 13. Approximating a textured surface using several sets of anchor points. (a) Cross section of textured curved surface with shadow region (blue) overlaid with the corresponding cross section of the recovered smooth intensity surface f (black). The resulting scale factors field $C(x,y)$ is shown below. (b) Selection of appropriate anchor points in M_s and M_r results in a smooth f (red) and a smooth scale factors field shown below. Three intensity surface reconstructions based on three sets of anchor points are shown. The recovered intensity surfaces differ, but the resulting scale factor fields shown below are consistent.

scale factors field—Fig. 13b. An extension to this approach would involve using several disjunct sets of anchor points, each of which produces an approximate intensity surface and corresponding partial scale factor field as shown in Fig. 13c (green and purple). Note that the intensity surfaces differ between sets of anchor points; however, the scale factor field should be consistent. A more robust scale factor field $C(x,y)$ can thus be obtained.

Denote by $\hat{M}_s$ and $\hat{M}_r$ the masks representing the collection of anchor pixels in the shadow and surround area, respectively. Replacing M_s and M_r in the approximation process described in Section 4.2, (7) is applied using the new masks $\hat{M}_s$ and $\hat{M}_r$ and using

$$\omega(x,y) = \begin{cases} 1, & (x,y) \in \hat{M}_r, \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

The scale factors field is then obtained using (8) applied only on pixels in $\hat{M}_s$ and $\hat{M}_r$. This produces a smooth yet incomplete scale factors field $\hat{C}(x,y)$. To compute the complete scale factors field $C(x,y)$, the missing values in $\hat{C}(x,y)$ (associated with pixels in $M_s - \hat{M}_s$) must be reconstructed. The thin-plate approach is used again and a functional, similar to (7), is computed with

$$E_d(f) = \int\int_{\Omega} \omega(x,y) \cdot \left[C(x,y) - \hat{C}(x,y)\right]^2 dxdy,$$
$$\omega(x,y) = \begin{cases} 1, & (x,y) \in \hat{M}_s, \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

This step guarantees a smooth $C(x,y)$ field in the umbra due to the smoothness term E_s in (7).

An example is shown in Fig. 14. Considering pixels of varying intensities (bright brick pixels and dark pixels between the bricks) as data points for the thin-plate minimization produces a nonsmooth scale factors field $C(x,y)$ (Fig. 14d (top)). Naively smoothing $C(x,y)$ and then applying to $I(x,y)$ yields unsatisfactory results and introduces Mach-Band-type artifacts in the shadow-free image at the shadow boundaries and in textured regions, as can be seen in Fig. 14e. Selecting appropriate anchor points, originating from the bricks alone and not from the dark pixels between the bricks, produces the new masks $\hat{M}_s$ and $\hat{M}_r$ (Figs. 14b and 14c (bottom)) and results in a smooth scale factors field $C(x,y)$ (Fig. 14d (bottom)) and a pleasing shadow-free image (Fig. 14f).

4.4 Anchor Point Selection

Selecting the appropriate anchor points is based on the assumption that shadows preserve *monotonicity* of pixel intensities in the umbra, i.e., the order of two nonshadowed pixels with respect to their intensities does not change when the pixels are shadowed. Exploiting this property, we derived a simple heuristic that allows us to select anchor pixels that correctly maintain the smoothness of $C(x,y)$. The histograms of the original M_s and M_r pixel values are calculated. Based on the monotonicity property, these two intensity distributions should display strong correlation, as shown in Fig. 15, displaying the histograms associated with the shadow and surround of Fig. 14a. A collection of pixels in shadow (M_s) and surround (M_r) that are likely to have originated from the same source is chosen by selecting correlating pixels in these histograms. A histogram shaping approach can be adopted; however, we found that it is sufficient to simply select pixels in both M_s and M_r with occurrence probability above a specific threshold. In Fig. 14a, pixels with probability above 0.5 were considered as anchor points (marked by dashed lines in Fig. 15). These collections of pixels define new shadow and surround masks $\hat{M}_s$ and

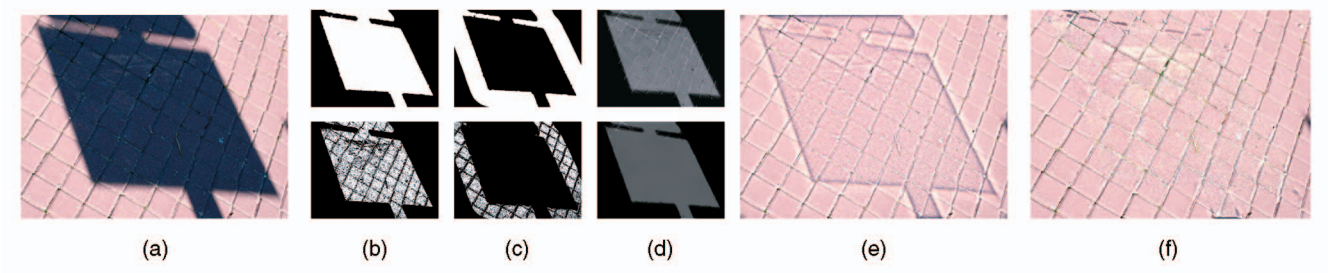

Fig. 14. Shadow removal using anchor points. (a) Shadow image. (b) and (c) Top: Shadow mask M_s and surround mask M_r. (d) Top: Resulting scale factors image, $C(x,y)$. Note the undesirable high-frequency content due to the dark pixels between the bricks. (e) Naively smoothing the scale factors image results in a shadow-free image containing Mach bands and slightly smoothed texture in the shadow-free region. (b) and (c) Bottom: Carefully selecting anchor points in shadow and surround produces new shadow and surround masks $\hat{M}_s$ and $\hat{M}_r$. Note that the dark pixels between the bricks are not included. (d) Bottom: Resulting smooth scale factors image $C(x,y)$ based on anchor points. (f) Resulting shadow-free image.

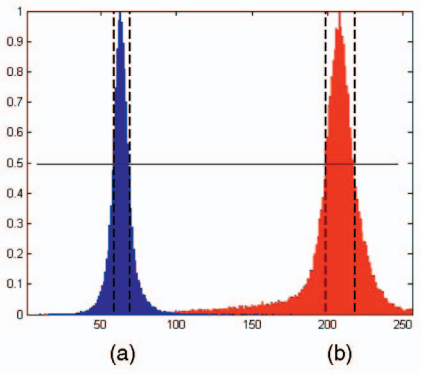

Fig. 15. Histograms associated with the shadow (blue) and surround (red) of Fig. 14a. Correlating pixels in these histograms are chosen as anchor points for the shadow and surround masks. Specifically, pixels in M_s and M_r with occurrence probability above 0.5 (marked by dashed lines) were chosen.

$\hat{M}_r$ and are used as anchor pixels in the intensity surface estimation (see Figs. 14f and 14g). It can be seen that the darker pixels between the bricks have been discarded and are not used as anchor points.

4.5 Determining Penumbra Scale Factors

The assumption of scale factor smoothness and the assumption of monotonicity of pixel intensities (used for extracting anchor points) do not hold in the penumbra. As a consequence, the method described in Sections 4.2-4.4 produces artifacts in the penumbra regions. An approach to overcome these problems is by attempting to smoothly model the penumbra and then calculate a smooth scale factors field using the substraction scheme of (8). Smooth modeling of penumbra profiles was suggested in [8] in which the authors assume a symmetric, sigmoidal-like shape of penumbra cross sections. However, as can be seen in Fig. 4, profiles of penumbra regions do not follow a particular model such as linear or sigmoidal. In the proposed method, instead of assuming or modeling a specific penumbra profile, we only assume that scale factors in penumbra regions are locally smooth in the direction tangent to the shadow edge.

Given the penumbra mask M_p (we use MRF labeling for determining the penumbra mask; see [9] for details and a summary in the Appendix) and given the reconstructed intensity surface f, penumbra scale factors are determined as follows: An initial estimate of the penumbra scale factors is obtained by applying (8) restricted to pixels in M_p. To enforce local smoothness of the penumbra scale factors field in the direction tangent to the shadow edge, *directional smoothing* is applied to $C(x, y)$ over the pixels of M_p. The shadow edge information is used to compute the direction of the smoothing vector. Finally, the penumbra scale factors are combined with the umbra scale factors calculated in Sections 4.2-4.4, to form a complete scale factors field, which is smooth in the umbra and locally smooth in the penumbra. The smooth $C(x, y)$ is then used in (4) to reconstruct the shadow-free image. As can be seen in all the results in Section 7, the proposed method for handling penumbra regions works well in preserving the textural information, without assuming a penumbra model.

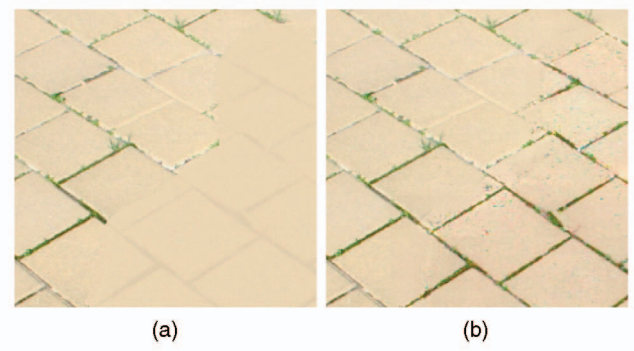

Fig. 16. Shadow removal from a noisy image (refer to Fig. 8c). (a) Enhancement using local variance adjustment. (b) Enhancement using local histogram specification.

5 Shadow-Free Region Enhancement

As discussed in Section 2.3, images undergo various transformations in the acquisition pipeline and postprocessing by imaging software. These transformations often affect shadow images in a manner that is inconsistent with shadow removal algorithm assumptions so that artifacts are introduced in the shadow-free image. An example is shown in Fig. 9 in which the shadow-free region displays high contrast compared to the nonshadow region. Furthermore, noise (e.g., sensor noise) in shadow regions is often enhanced and emphasized in the shadow-free image (see Fig. 8c). Thus, there is a need for an enhancement algorithm that attenuates the noise effects in the shadow-free regions and attempts to equate their appearance with that of their nonshadow counterparts.

A successful algorithm for shadow-free region enhancement should fulfill two basic requirements: It must preserve the original texture in the shadow region and it should be general enough to handle various types of transformations a shadow region might undergo, as well as handling shadow-free regions with noise. In [28], a scheme for shadow removal is suggested in which the mean and variance of pixels in shadow regions are adjusted based on pixels of the corresponding nonshadow surface. While this approach improves the similarity in appearance of shadow-free regions and their nonshadow counterparts in many images, it fails in noisy images. An example is given in Fig. 16a. It can be seen that although noise is significantly attenuated, the contrast in the shadow-free region is globally reduced, producing unpleasing low contrast in the dark regions between the bricks. To achieve better results, an adaptive process should be used. In Fig. 16, for example, the noisy brick regions should be enhanced but the dark regions between the bricks should be left untouched.

Our proposed method for enhancing shadow-free regions is based on the assumption that two matching patches, one inside the shadow-free region and one outside the region, should have similar statistical behavior. Thus, we perform *histogram specification* [37] independently on each patch in the shadow-free region. Histogram specification allows control of the pixel statistics of each patch, and more importantly, it is consistent with the shadow monotonicity property. Denote by u and v two discrete random variables that take values x_i and y_j, with probabilities $P_u(x_i)$ and $P_v(y_j)$ $(i, j = 0 \ldots L-1)$, respectively. We define the following:

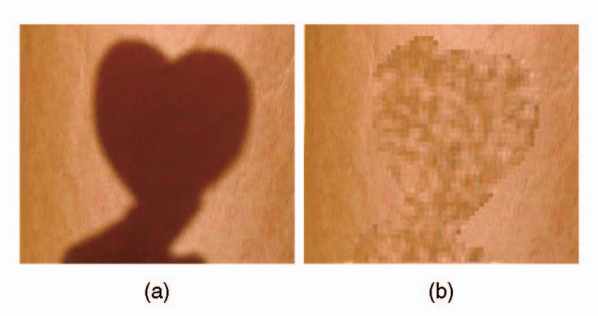

Fig. 17. Shadow-free image enhancement. (a) Shadow on curved surface. (b) Shadow-free region enhancement without preserving local mean values.

$$w_l \triangleq \sum_{i=0}^{l} P_u(x_i) \tag{11}$$

and

$$\tilde{w}_k \triangleq \sum_{j=0}^{k} P_v(y_j). \tag{12}$$

Histogram Specification is performed by mapping each value x_i of u to y_{n_i} of v such that n_i is the smallest for which $\tilde{w}_{n_i} \geq w_i$. Note that if the values x_i and y_j are sorted in ascending order, the procedure performs a nondecreasing linear mapping between u and v since, for each pair $x_i \leq x_{i+1}$, it holds that $n_i \leq n_{i+1}$, which implies that $y_{n_i} \leq y_{n_{i+1}}$. Such a mapping on the intensities preserves the original texture. In our implementation, matching of shadow-free patches to corresponding nonshadow patches is performed using fast normalized cross correlation [38]. Fig. 16b shows an example of enhancing a shadow-free region using the patch-based histogram specification approach. Note that since this enhancement method is performed on shadow-free images as a postprocessing step, it can be used following any shadow removal algorithm.

A subtle point in performing the histogram specification is whether to preserve the original mean value of the patches in the shadow-free region. It is assumed that the correct mean values were reconstructed by the shadow removal algorithm. Thus, it is typically desired to maintain the mean values of the patches. This is specifically true in shadow regions cast on curved surfaces where pixel intensities vary across the shadow region. Fig. 17 displays an example of enhancing a shadow-free region on a curved surface when the mean values of shadow-free patches are allowed to change. In other cases, however, specifically when removing shadows cast on flat surfaces, performing histogram specification without preserving the patch mean values may yield equal or better results. Two examples of shadow-free region enhancement without preserving mean values are given in Fig. 26.

6 User-Guided Extraction of Shadow Masks

Automatic detection of shadow regions in a single image is a challenging problem and several approaches have been proposed in recent years, such as [4], [20], [39]. Although in this work we concentrate in acquiring shadow-free images of high quality, we propose a simple yet effective method for extracting the corresponding masks of shadow regions based on user input. Since automation is not our primary concern but rather the quality of the final output is, requiring the user to supply initial cues for the system seems reasonable and may fit well into photo editing software.

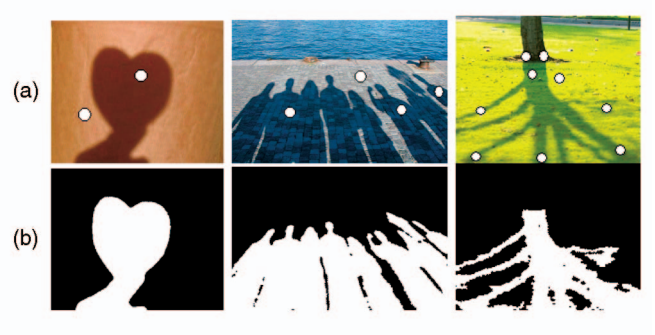

Fig. 18. Examples of user-guided extraction of shadow masks based on region growing and SVM. (a) Shadow images. The circles signify the supplied shadow and nonshadow observations. (b) Shadow masks extracted using region growing and SVM.

Shadow mask derivation is performed by *region growing*, using SVM [40] for pixel classification. The RGB color space is used as the feature space for the SVM, and a supervised learning problem is constructed by asking the user to supply the coordinates of pixels (via mouse clicks) in different shadow and nonshadow regions. The RGB values within a neighborhood around the supplied coordinates are considered as shadow and nonshadow observations on which SVM training is performed. Following training, all image pixels are classified as shadow or nonshadow. A region growing phase is initiated with the coordinates of the shadow observations supplied by the user as initial seeds. New pixels labeled as shadow are added to the shadow mask in each region growing iteration. Fig. 18 shows several examples of shadow masks derived using this method. The observation vectors were taken as 3×3 pixel neighborhoods around the user-selected coordinates. A polynomial kernel of degree 3 was used for the SVM.

The shadow mask is then used to derive the surround mask (M_r in Section 4.2) by expanding a wide band along the shadow mask boundaries. In cases where shadow boundaries coincide with object boundaries (e.g., as occurs in Fig. 21), an additional object mask is derived in a similar manner and used for refining the derived shadow and surround masks.

7 Results

In this section, we give example results of our shadow removal approach.[2] We start by examining the ability of the proposed method to remove nonuniform shadows. Fig. 19 contains a nonuniform shadow cast on a flat surface, exhibiting little texture. A high-quality shadow-free result is obtained using our proposed method. Another example of nonuniform shadow cast on flat surface is given in Fig. 20. The surface in this example exhibits high-frequency texture. Fig. 20b contains a shadow-free result obtained assuming a uniform shadow (i.e., using a global scale factor in umbra [9]). It can be seen that the shadow is not removed

2. High-resolution color images appearing in this section can be viewed online at http://cs.haifa.ac.il/~hagit/papers/ShadowRemoval.

Fig. 19. Shadow removal using our proposed method.

completely. The result of our proposed method is given in Fig. 20c.

Two more examples of shadows cast on flat surfaces are given in Figs. 21 and 22. Fig. 21 contains nonuniform shadow with soft regions (the basketball net shadow). Fig. 22 contains complex nonuniform shadow with soft regions, especially in the right part of the image. The results of these examples demonstrate the ability of our proposed method to handle complex nonuniform shadows with soft shadow regions.

Figs. 23 and 24 contain nonuniform shadows cast on curved surfaces. Note that the shadow in Fig. 23 has a wide penumbra and that the shadow in Fig. 24 exhibits high nonuniformity (see Fig. 24b). Nevertheless, due to the high-order model proposed in our method, the correct geometry of the curved surfaces is obtained, yielding high-quality results.

Fig. 25 shows the benefit of using the shadow-free region enhancement suggested in Section 5. Fig. 25a contains a nonuniform shadow cast on a textured surface. The shadow-free region in Fig. 25b has high contrast, as illustrated in Fig. 9. A more pleasing result is obtained using the proposed enhancement process, as can be seen in Fig. 25b. Another example is shown in Fig. 10, which demonstrates how the enhancement process can compensate for lack of self-shading. It can be seen in Fig. 10g that the stone texture in the shadow-free region appears "flat" since self-shading is absent. The shadow-free enhancement process yields a more natural shadow-free image as demonstrated in Fig. 10h.

Fig. 21. Removing nonuniform shadow cast on a flat surface. Note the nonuniform shadows of the player and the basketball net.

Fig. 22. Complex nonuniform shadow with soft shadow regions.

Fig. 26 compares our method to the methods suggested in [6] and [8]. Since the methods of [6] and [8] operate in the gradient domain and a global Poisson equation is solved during the reconstruction phase, the color balance and global smoothness of the reconstructed image is affected, as can be seen in Fig. 26b (taken from [6] and [8]). Fig. 26c displays the results of our proposed algorithm while preserving mean values during the shadow-free enhancement process, and Fig. 26d displays results without preserving mean values.

Fig. 20. Removing nonuniform shadow cast on flat textured surface. (a) Shadow image. (b) Shadow removal assuming uniform shadow. (c) Shadow removal using our proposed method. Texture is preserved in the shadow-free region.

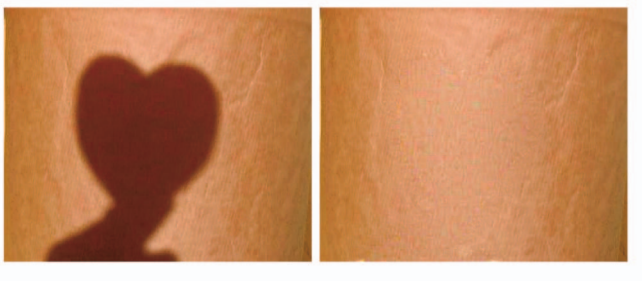

Fig. 23. Removing shadow cast on curved textured surface.

Fig. 27 compares our method with matting approaches: the method of Levine and Weiss [41] in the top row and Wu et al. [30] bottom row. Smoothness assumptions inherent to the approach (as in [41]) strongly affect the texture in the shadow-free image (Fig. 27b). The bottom row example is strongly affected by the bumpy texture's shading, which is lost in the shadow-free image.

Finally, in Fig. 28, we show that the proposed method can be used to remove highlight regions in images. Despite the fact that the highlight region exhibits high nonuniformity, the proposed method can be used without any changes to recover a highlight-free image.

8 Discussion and Future Work

Removal of shadows from a single image is a challenging problem affected by physical phenomena, scene characteristics, and device parameters, as were reviewed in Section 2. In light of these issues, a shadow removal algorithm was introduced in this paper. As demonstrated by a variety of examples, the method is capable of producing high-quality results on many types of shadow images, coping with some of the fundamental problems in shadow removal described in Section 2. In addition, the proposed shadow-free region enhancement process can greatly increase algorithm robustness in handling shadow images that have undergone image processing transformations in shadow regions, as well as shadow images with noise.

Overall runtime of the algorithm is primarily dependent on the size of the shadow region. This affects both the approximation time and the shadow-free region enhancement algorithm. In particular, most of the algorithm runtime is consumed by the sliding window procedure of the enhancement algorithm in which fast normalized cross correlation is performed. For images of size 500×500 pixels processed on a 1 GB Pentium 4, a typical run takes several minutes to complete.

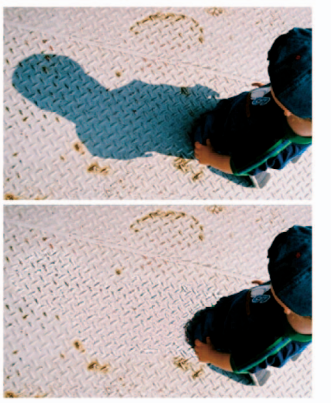

Fig. 25. Shadow removal example: region enhancement (refer to Fig. 9).

Additional concerns should be addressed such as removing shadows cast on different types of surfaces and handling shadow boundaries that coincide with object boundaries. Examples of shadow boundaries that coincide with object boundaries can be seen in Figs. 21, 25, and 28. In these figures, the seam between shadow and objects seems artificial.

Appendix

An MRF labeling approach is used for determining the penumbra mask M_p. Details can be found in [9]. A brief summary is given here.

Given an image I and its gradient magnitude field $\|\nabla I\|$, define the gradient magnitude distribution image P^I as

$$P^I_{xy} = \Pr(\|\nabla I(x,y)\|), \tag{13}$$

where $\Pr(\|\nabla I(x,y)\|)$ is the probability of the gradient magnitude at pixel (x,y) in the image.

Labeling penumbra pixels by direct thresholding of P^I produces many false alarms and misses. The suggested approach exploits pixels with strong evidence of being edge pixels, and propagates this evidence to their neighboring pixels. Thus, low-evidence pixels are supported by

Fig. 24. Shadow removal example. (a) Nonuniform shadow cast on curved surface. (b) Shadow-free image assuming uniform shadow. Note the artifacts and high contrast in the shadow-free region. (c) Shadow-free image using our method and following postprocessing enhancement.

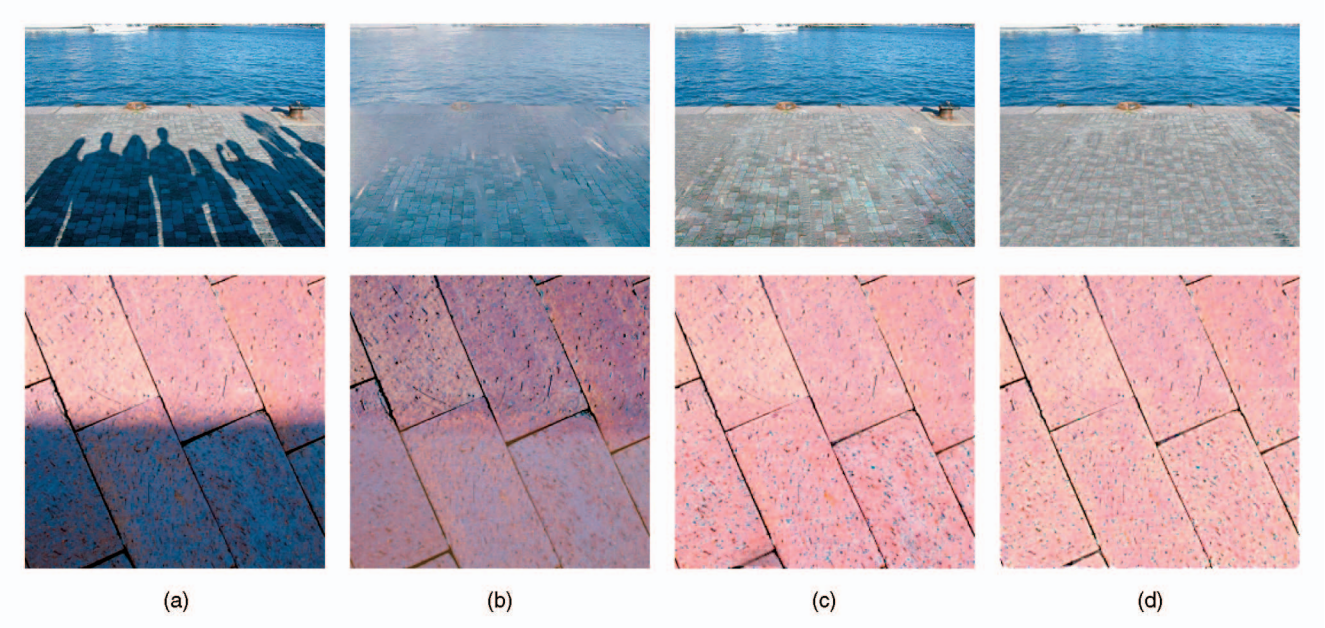

Fig. 26. Comparison of the proposed method with the method of [6] (top row) and [8] (bottom row). (a) and (b) Shadow and shadow-free images taken from [6] and [8]. (c) Shadow-free image produced by the proposed method. (d) Applying the postprocessing enhancement when mean values of shadow-free patches are allowed to change.

neighboring edge pixels. We implement this scheme using a Markov Random Field (MRF) [42] over P^I. A unique random variable is associated with each pixel in P^I defined over $\{1, 0\}$, denoting whether the underlying pixel should be labeled as an edge pixel or not, respectively.

Let g_{xy} be the MRF random variable at location (x, y). We define the *posterior energy* [42] of the field g as follows:

$$\sum_{x,y}(1-g_{xy})\left[(1-P^I_{xy})+\sum_{g_{x'y'}\in\mathcal{N}_{xy}}\psi(g_{x'y'},g_{xy})\right]+\alpha g_{xy}, \quad (14)$$

where $\mathcal{N}_{xy}$ represents the 4-neighborhood of pixel (x, y). The term $(1 - P^I_{xy})$ is the *prior energy* related to the probability of a pixel being an edge pixel. $\psi(g_{x'y'}, g_{xy})$ is the *likelihood energy* of a pixel, which depends on its neighboring pixels:

$$\psi(g_{x'y'},g_{xy})=\begin{cases}1, & g_{x'y'}\neq g_{xy}\wedge|P^I_{x'y'}-P^I_{xy}|\leq t1,\\ & g_{x'y'}=g_{xy}\wedge|P^I_{x'y'}-P^I_{xy}|\geq t2,\\ 0, & \text{otherwise.}\end{cases} \quad (15)$$

The defined MRF is parameterized by α, $t1$, and $t2$. Parameter α bounds the local energy of a pixel when labeled as an edge pixel. Parameters $t1$ and $t2$ are thresholds, $t1 < t2$. The likelihood energy penalizes for neighboring pixels differing in label but with similar edge probabilities, and neighboring pixels of similar label but differing in edge probabilities.

Equation (14) is minimized over the field g. Given the minimizing g, penumbra pixels are extracted by finding binary connected components on g originating from pixels that appear both in the shadow edge image and in g. Fig. 10c displays examples of penumbra masks produced using this approach.

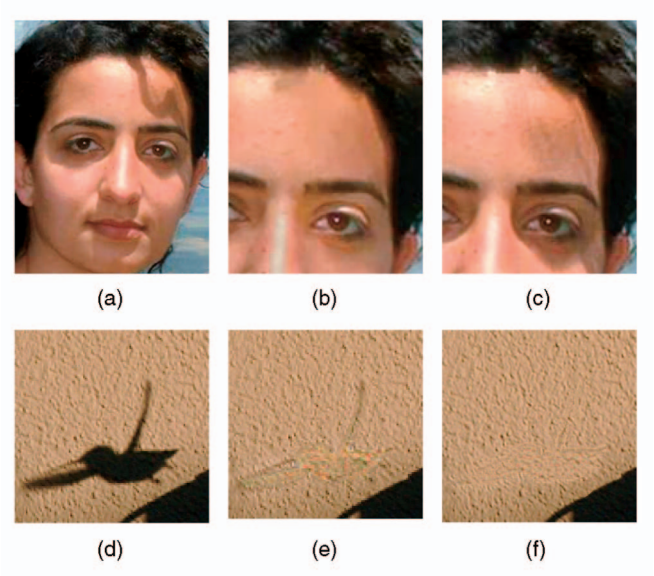

Fig. 27. Comparison with matting approaches. Top row: comparison with [41]. Bottom row: comparison with [30]. (a) and (b) Shadow and shadow-free images taken from sources. (c) Shadow-free image produced by the proposed method. (b) and (c) are scaled to show texture.

Fig. 28. Highlight removed using our proposed method.

REFERENCES

[1] H. Barrow and J. Tenenbaum, "Recovering Intrinsic Scene Characteristics from Images," *Computer Vision Systems,* A. Hanson and E. Riseman, eds., Academic Press, 1978.

[2] J.B. Keller, "Geometrical Theory of Diffraction," *J. Optical Soc. Am.,* vol. 52, no. 2, pp. 116-130, 1962.

[3] H. Primack, H. Schanz, U. Smilansky, and I. Ussishkin, "Penumbra Diffraction in the Quantization of Dispersing Billiards," *Physical Rev. Letters,* vol. 76, pp. 1615-1618, Mar. 1996.

[4] G.D. Finlayson, S.D. Hordley, and M.S. Drew, "Removing Shadows from Images," *Proc. European Conf. Computer Vision,* vol. IV, pp. 823-836, 2002.

[5] T.-P. Wu and C.-K. Tang, "A Bayesian Approach for Shadow Extraction from a Single Image," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 480-487, 2005.

[6] L. Xu, F. Qi, and R. Jiang, "Shadow Removal from a Single Image," *Proc. IEEE Int'l Conf. Intelligent Systems Design and Applications,* pp. 1049-1054, 2006.

[7] P. Mamassian, D.C. Knill, and D. Kersten, "The Perception of Cast Shadows," *Trends in Cognitive Sciences,* vol. 2, no. 8, pp. 288-295, 1998.

[8] A. Mohan, J. Tumblin, and P. Choudhury, "Editing Soft Shadows in a Digital Photograph," *IEEE Computer Graphics and Applications,* vol. 27, no. 2, pp. 23-31, Mar./Apr. 2007.

[9] E. Arbel and H. Hel-Or, "Texture-Preserving Shadow Removal in Color Images Containing Curved Surfaces," *Proc. IEEE Conf. Comp. Vision and Pattern Recognition,* pp. 1-8, 2007.

[10] Y. Weiss, "Deriving Intrinsic Images from Image Sequences," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 68-75, 2001.

[11] Y.-Y. Chuang, D.B. Goldman, B. Curless, D.H. Salesin, and R. Szeliski, "Shadow Matting and Compositing," *Proc. ACM SIGGRAPH,* pp. 494-500, 2003.

[12] A. Prati, I. Mikic, M. Trivedi, and R. Cucchiara, "Detecting Moving Shadows: Algorithms and Evaluation," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 25, no. 7, pp. 918-923, July 2003.

[13] G. Finlayson, C. Fredembach, and M. Drew, "Detecting Illumination in Images," *Proc. IEEE Int'l Conf. Computer Vision,* 2007.

[14] A. Leone and C. Distante, "Shadow Detection for Moving Objects Based on Texture Analysis," *Pattern Recognition,* vol. 40, no. 4, pp. 1222-1233, 2007.

[15] A. Bevilacqua, "Effective Shadow Detection in Traffic Monitoring Applications," *Winter School of Computer Graphics,* vol. 11, pp. 57-64, 2003.

[16] P. Rosin and T. Ellis, "Image Difference Threshold Strategies and Shadow Detection," *Proc. British Machine Vision Conf.,* pp. 347-356, 1995.

[17] A.W. So, K.-Y.K. Wong, R.H.Y. Chung, and F.Y.L. Chin, "Shadow Detection for Vehicles by Locating the Object-Shadow Boundary," *Proc. IASTED Conf. Signal and Image Processing,* pp. 315-319, 2005.

[18] D. Grest, J.-M. Frahm, and R. Koch, "A Color Similarity Measure for Robust Shadow Removal in Real Time," *Proc. Vision, Modeling, and Visualization Conf.,* pp. 253-260, 2003.

[19] E. Salvador, A. Cavallaro, and T. Ebrahimi, "Cast Shadow Segmentation Using Invariant Color Features," *Computer Vision and Image Understanding,* vol. 95, no. 2, pp. 238-259, 2004.

[20] M.D. Levine and J. Bhattacharyya, "Removing Shadows," *Pattern Recognition Letters,* vol. 26, no. 3, pp. 251-265, 2005.

[21] K. Barnard and G.D. Finlayson, "Shadow Identification Using Colour Ratios," *Proc. IS&T/SID Color Imaging Conf.: Color Science, Systems and Applications,* pp. 97-101, 2000.

[22] C. Fredembach and G. Finlayson, "Hamiltonian Path Based Shadow Removal," *Proc. British Machine Vision Conf.,* pp. 970-980, 2005.

[23] G. Finlayson, S. Hordley, C. Lu, and M. Drew, "On the Removal of Shadows from Images," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 28, no. 1, pp. 59-68, Jan. 2006.

[24] G. Finlayson, M. Drew, and C. Lu, "Intrinsic Images by Entropy Minimization," *Proc. European Conf. Computer Vision,* pp. 582-595, 2004.

[25] Z. Figov, Y. Tal, and M. Koppel, "Detecting and Removing Shadows," *Proc. IASTED Int'l Conf. Computer Graphics and Imaging,* 2004.

[26] C. Fredembach and G. Finlayson, "Simple Shadow Removal," *Proc. IEEE Int'l Conf. Pattern Recognition,* pp. 832-835, 2006.

[27] M. Baba, M. Mukunoki, and N. Asada, "Shadow Removal from a Real Image Based on Shadow Density," *Proc. ACM SIGGRAPH,* 2004.

[28] M. Baba and N. Asada, "Shadow Removal from a Real Picture," *Proc. ACM SIGGRAPH,* 2003.

[29] G. Finlayson, S. Hordley, and M. Drew, "Removing Shadows from Images Using Retinex," *Proc. IS&T/SID Color Imaging Conf.: Color Science, Systems and Applications,* pp. 73-79, 2002.

[30] T.-P. Wu, C.-K. Tang, M.S. Brown, and H.-Y. Shum, "Natural Shadow Matting," *ACM Trans. Graphics,* vol. 26, no. 2, p. 8, 2007.

[31] X. Gu, D. Yu, and L. Zhang, "Image Shadow Removal Using Pulse Coupled Neural Network," *IEEE Trans. Neural Networks,* vol. 16, no. 3, pp. 692-698, May 2005.

[32] Z. Du, X. Qin, H. Lin, and H. Bao, "Shadow Removal in Gradient Domain," *Proc. Int'l Conf. Image Analysis and Recognition,* pp. 107-115, 2005.

[33] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically Deformable Models," *Proc. Ann. Conf. Computer Graphics and Interactive Techniques,* pp. 205-214, 1987.

[34] C. Lürig, L. Kobbelt, and T. Ertl, "Hierarchical Solutions for the Deformable Surface Problem in Visualization," *Graphical Models,* vol. 62, no. 1, pp. 2-18, 2000.

[35] L.D. Cohen and I. Cohen, "Finite-Element Methods for Active Contour Models and Balloons for 2-D and 3-D Images," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 15, no. 11, pp. 1131-1147, Nov. 1993.

[36] D. Terzopoulos, A. Witkin, and M. Kass, "Symmetry-Seeking Models for 3D Object Recognition," *Int'l J. Computer Vision,* vol. 1, no. 3, pp. 211-221, Oct. 1987.

[37] A.K. Jain, *Fundamentals of Digital Image Processing.* Prentice-Hall, Inc., 1989.

[38] J. Lewis, "Fast Normalized Cross-Correlation," *Proc. Conf. Vision Interface,* pp. 120-123, 1995.

[39] E. Salvador, A. Cavallaro, and T. Ebrahimi, "Shadow Identification and Classification Using Invariant Color Models," *Proc. IEEE Int'l Conf. Acoustics, Speech, and Signal Processing,* pp. 1545-1548, 2001.

[40] C. Cortes and V. Vapnik, "Support-Vector Networks," *Machine Learning,* vol. 20, no. 3, pp. 273-297, 1995.

[41] A. Levin and Y. Weiss, "User Assisted Separation of Reflections from a Single Image Using a Sparsity Prior," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 29, no. 9, pp. 1647-1654, Sept. 2007.

[42] S.Z. Li, *Markov Random Field Modeling in Image Analysis.* Springer-Verlag New York, Inc., 2001.

Eli Arbel received the BSc and MSc degrees in computer science (both summa cum laude) from the University of Haifa, Israel, in 2005 and 2008, respectively. He is currently a research staff member at the IBM Research Lab in Haifa, Israel. His research interests include computer vision, image processing, and formal methods for synthesis and verification of hardware designs.

Hagit Hel-Or received the PhD degree in computer science from the Hebrew University of Jerusalem, Israel, in 1994. She is a faculty member in the Department of Computer Science at the University of Haifa, Israel. Her research interests in the area of image processing and computer vision include pattern recognition, color vision, imaging technologies, and computational and human vision. She is a member of the IEEE.