

Appearance Harmonization for Single Image Shadow Removal

Li-Qian Ma
TNList, Tsinghua University
Beijing, China
liqianma@qq.com

Jue Wang
Adobe Systems Inc.
juewang@adobe.com

Eli Shechtman
Adobe Systems Inc.
elish@adobe.com

Kalyan Sunkavalli
Adobe Systems Inc.
sunkaval@adobe.com

Shi-Min Hu
TNList, Tsinghua University
Beijing, China
shiminhu@tsinghua.edu.cn

Abstract

Shadow removal is a challenging problem and previous approaches often produce de-shadowed regions that are visually inconsistent with the rest of the image. We propose an automatic shadow region harmonization approach that makes the appearance of a de-shadowed region (produced using any previous technique) compatible with the rest of the image. We use a shadow-guided patch-based image synthesis approach that reconstructs the shadow region using patches sampled from non-shadowed regions. This result is then refined based on the reconstruction confidence to handle unique textures. Qualitative comparisons over a wide range of images, and a quantitative evaluation on a benchmark dataset show that our technique significantly improves upon the state-of-the-art.

1. Introduction

Cast shadows occur in many photography scenarios, and often lead to distracting artifacts that detract from the visual appeal of a photograph. Removing cast shadows from such photographs is often highly desirable, yet difficult to achieve due to its inherently ill-posed nature: it is difficult for computational techniques, without any prior knowledge, to disambiguate shadows from dark textured regions in the scene.

In the past decade, many approaches have been proposed for removing shadows in photographs. However, many of these techniques suffer from *inconsistency* artifacts, i.e. the de-shadowed region is visually incompatible with the rest of the input image. Most previous methods assume simplified shadow models that boil down to a simple color and intensity correction of the shadowed pixels. This assumption typically does not produce good results in presence of soft

shadows, complex spatially-varying textures, complex reflectance properties of the underlying material (e.g., BRDF) and loss of dynamic range in the shadow region (see Fig. 1 top). Postprocessing, such as tone/color adjustment, gamma correction, lossy compression, can also easily violate common shadow models.

In this work we aim at providing new tools that can help users achieve high quality shadow removal results. We propose a new technique called *Shadow Region Harmonization* (SRH), which can effectively remove inconsistency artifacts from existing shadow removal results. Our method is built on the general idea of building correspondence between shadow and non-shadow regions, and enforcing consistent color and texture properties of corresponding regions. Our key insight is that the shadow region often (but not always) contains the same materials as in the non-shadow regions. Thus, correspondence between these two types of regions can be constructed *locally*, instead of globally. We build such correspondence using a guided patch-based image synthesis framework, where shadow regions are reconstructed using non-shadow ones. For each shadow patch, we use its corresponding non-shadow ones to compute a parametric appearance correction model based on color and texture. To handle unique shadow materials, we compute for each patch a *correction confidence*, and use it in an optimization process to ensure that patches without good correspondences can also be corrected properly. We quantitatively evaluate the SRH method on a recent benchmark dataset [10], and show that it can significantly improve the output of existing methods, including the state-of-the-art ones.

2. Related Work

We first briefly review representative works closely related to ours, and then discuss the inconsistency artifacts in previous methods in more detail.

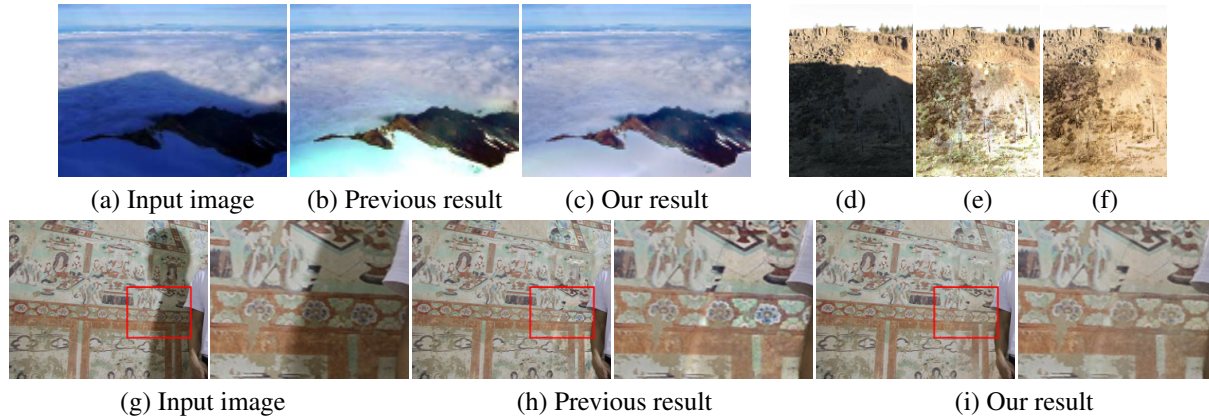


Figure 1. Top row: input image (a)(d); the state-of-the-art shadow removal method of Liu and Gliercher. [17] produces results with color inconsistencies (b)(e); our shadow region harmonization (SRH) method automatically corrects these issues (c)(f). Bottom row: input image (g); the state-of-the-art shadow removal method of Liu and Gliercher. [25] produces results with texture inconsistencies (h); our shadow region harmonization (SRH) method automatically corrects these issues (i).

Shadow removal is an extensively studied problem and modern approaches are well summarized in recent surveys [26, 20]. Shadow analysis is also closely related to intrinsic image decomposition [4, 12] – the problem of separating an image into reflectance and illumination components – though shadow removal focuses on the illumination variation caused by occluded light sources.

A common way to detect shadows is to use *illuminant invariant features* [6, 8], which help to detect shadow boundaries. Shadow removal can then be achieved by reconstructing the image with the shadow gradients edited [8, 9]. Baba et al. [1] estimate shadow density directly using patch lightness. Gryka et al. [13] extract soft shadows by learning a regression function from image patches to shadow mattes. However, both these techniques use simple shadow models that are often inadequate for the non-linear images that constitute the vast majority of photographs. Laplacian pyramid [21] and gradient domain processing [17] have been used to improve the consistency of textures between well-lit and shadowed regions. These techniques have limitations when there are multiple textures in the same shadow region. Guo et al. [14] build a classifier to detect shadowed and non-shadowed region pair of the same texture, however it is based on image segmentation, which itself could be fragile on complex scenes.

Because of the inherently ambiguous nature of shadow detection and removal, many previous approaches require manual specification of the shadow region [21, 17, 19]. Given this input, shadow removal can be posed as a matting [24] or labeling [18] problem.

Patch-based synthesis has shown great success in image and video completion [22] and other editing tasks since the introduction of *PatchMatch* [2] – a fast approximate method for computing patch-based dense correspondences. PatchMatch has been generalized to support scaling and ro-

tation of patches [3] as well as gain/bias of each individual color channels [15]. This family of techniques have been widely used for finding patch correspondence in a rotation and scale-invariant manner, and can also handle differences in illumination conditions.

Our SRH algorithm uses guided patch-based synthesis [22] to reconstruct the shadow region using non-shadow patches, using the initial shadow removal result as guidance. This is similar to the Image Analogies framework [16] that was used for guided texture synthesis and image enhancement. In our method we address a unique challenge: the shadow region may contain unique structures/materials, and thus patch synthesis can only be partially successful. We use an optimization approach that gracefully combines traditional shadow color correction with patch-based synthesis to generate consistent removal results in the entire shadow region. Interestingly, Gryka et al. [13] also used patch-based image completion [22] to compute one of the features for its learning based framework. However the synthesis was not guided so the completed content might end up looking very different than the real one.

Inconsistency artifacts of shadow removal Existing shadow removal approaches often produce inconsistency artifacts in recovered shadow regions, due to the violation of the simplified shadow models they use; these include both color and texture inconsistencies. Specifically, most approaches cannot model the loss of dynamic range in shadow regions [7, 19], which leads to inconsistent noise properties and texture characteristics between recovered shadow regions and non-shadow regions, such as the examples in Fig. 4. Pixel-based approaches [8, 7, 1] suffer from inaccuracies in the estimation of the shadow parameters, leading to color shifts or residual shadows in the recovered shadow regions (Fig. 1(top)). To correct such artifacts some methods leverage shadow/non-shadow region corre-

spondence [21, 17, 10], or region-based color transfer [11]. However, they are still not robust to complex spatially-varying textures, complex reflection and shading properties. Fig. 3(e) shows an example with a colorful translucent occluder, where the green color can not be eliminated using any existing shadow removal methods.

Our SRH approach builds dense correspondence between shadow and non-shadow regions to enforce both color and texture consistency, as shown in Fig. 1. It is an automatic post-processing method that is independent of specific shadow removal approach being applied first, and thus can be applied widely.

3. Our Approach

Our SRH approach takes the original image, I , and an initial shadow removal result, I^N , as input, and produces a higher quality result, I^H .

3.1. The Harmonization Model

In contrast to previous shadow removal methods that use pixelwise shadow models, the SRH method is based on image patches that capture better local color and texture characteristics. It computes a parametric *Appearance Harmonization Model*, which describes for each shadow patch, how to change its color and texture to make it more consistent with its corresponding non-shadow patches. It contains three components: (1) color correction parameters, (2) texture correction parameters, and (3) a correction confidence.

Color Correction We model the effect of a shadow as a per-channel affine transform on the pixel values. Shadows tend to be spatially smooth, and we incorporate this prior by assuming that this affine transform is constant within a small patch (e.g. 5×5). This gives us the following color correction model:

$$I_c^H(P) = g^c(P) \cdot I_c^N(P) + b^c(P), \quad (1)$$

where c is the color channel index, $I^H(P)$ and $I^N(P)$ correspond to patches P in I^H and I^N , and $g^c(P)$ and $b^c(P)$ are per-patch, per-channel gain and bias values that together constitute our affine color correction term.

Texture Correction The standard deviation of color channels within an image patch is often used to describe local texture and image details [21, 17]. We denote the standard deviation of patch P and color channel c in I^N and I^H as $\sigma_c^N(P)$ and $\sigma_c^H(P)$, respectively. We scale $\sigma^N(P)$ as:

$$\sigma_c^H(P) = s_c(P) \cdot \sigma_c^N(P), \quad (2)$$

so that shadow patches have the same level of local color contrast as the corresponding non-shadow ones. Note that the scale s_c is different from the gain in Eqn. 1 as it controls the deviation around the mean color (vs. 0).

Correction Confidence To estimate the above correction parameters for each patch, we use a patch synthesis approach to match shadow and non-shadow patches of the same material. Due to limitations of patch synthesis, not all matches are reliable, and the estimated parameters at these patches are incorrect. To describe the reliability of the correction parameters of a specific shadow patch P , we compute an extra confidence value $C(P)$ in $[0, 1]$ and add it to the parameter list. We will describe how to compute this value in Sec. 3.2.2.

Final Model Allowing gain, bias and texture scaling in every color channel will result in a 10 dimensional parameter vector for each pixel. To maintain a good balance between model complexity and flexibility, in practice we choose to apply corrections in CIE Lab space, and only enable gain in the L channel, bias in the a and b channels for color correction, and only the L channel for texture correction. This gives us a 5-channel parameter map $S = (g^L, b^A, b^B, s^L, C)$, which we refer to as the *shadow correction map*. Other parameter combinations will be discussed and compared in Sec. 4.3.

3.2. Shadow Correction Map Generation

We now describe how to estimate the shadow correction map S , given the source image I and an initial shadow removal result I^N . Our key idea is to build dense correspondences between shadow and non-shadow patches, and derive the correction parameters from them in a way that unique shadow structures and materials can be properly handled.

A binary mask is needed to indicate which pixels are inside the shadow region that need to be corrected. In our work we assume the only difference between I^N and I is the color of shadow pixels. We apply a small threshold on the difference image $I^N - I$ to generate the correction region \mathcal{R}_c , and further apply a small dilation operation using a 3×3 kernel to remove occasional small holes inside it. This is shown in Fig. 2(c).

3.2.1 Patch-based Synthesis

We use patch-based synthesis to synthesize a new shadow-free image, I^S , where the correction region \mathcal{R}_c has been filled using patches from the source region $\mathcal{R}_s = I \setminus \mathcal{R}_c$. We use a guided variant of the *Image Melding* algorithm [5] as the basis for this synthesis task, given its support for patch scaling, rotation, reflection, and color gain and bias. It is applied in a coarse-to-fine manner on an image pyramid.

Despite its color and texture inconsistencies, I^N usually gives us a good initial estimation of the final result that can be used as guidance for the synthesis process. We thus use I^N in two ways. First, we use it to initialize the synthesis result, I^S , at the coarsest level. Second, in the spirit of



(a) Input image (b) Initial result (c) Shadow mask (d) Patch synthesis, I^S (e) Confidence, C (f) Final result, I^C

Figure 2. Given an input image I (a), we compute an initial shadow removal result, I^N (b) using current shadow removal methods (top, [17], bottom, [18]). We use this result to estimate a shadow mask (c), and a patch-based synthesized image, I^S (d). This result is inaccurate in regions that are not observed in the non-shadow region and this is captured by the confidence map, C (e). Combining I^C and I^S using C gives us the final correction, I^C , (f) where the inconsistency artifacts in I^N are removed.

Image Analogies [16], we use it as a guidance layer during the synthesis process. Specifically, the distance between a patch P in \mathcal{R}_c and a patch Q in \mathcal{R}_s is defined as:

$$d(P, Q) = \|I^S(P), I^S(Q)\|_2 + \beta \|I^N(P) \cdot \mathbf{g}(P) + \mathbf{b}(P), I^S(P)\|_2 + \gamma \mathcal{E}(\mathbf{g}(P), \mathbf{b}(P)), \quad (3)$$

where $\|I(P), I(Q)\|_2$ is the average $L2$ color distance of two patches, and β and γ are balancing weights. The second term is our guidance term; it constrains the synthesized patch $I^S(P)$ to be similar to the initial shadow removal result $I^N(P)$. The additional gain $\mathbf{g}(P)$ and bias $\mathbf{b}(P)$ are introduced in this term to compensate for the possible color and intensity inconsistencies in I^N . Since we expect I^N to be reasonably close to the final result, we also add a third term $\mathcal{E}(\mathbf{g}(P_i), \mathbf{b}(P_i))$ to punish unrealistically large gain and bias, defined as (P is omitted):

$$\mathcal{E}(\mathbf{g}, \mathbf{b}) = \sum_c (|\mathbf{g}^c - 1| + |\mathbf{b}^c - 0|), c \in \{L, A, B\}. \quad (4)$$

We use this patch synthesis process produces to reconstruct I^S , which will be used in the next step to derive the correction map. Fig. 2(d) shows two synthesis results.

Parameter settings We decompose the input image into a pyramid with a coarsest scale of size 30 pixels (smaller dimension), and increase the scale by a factor of 1.4 for each pyramid level. β is set to 30. Patch distance is computed in CIELab space, and gain (L channel), bias (a & b channels) ranges are set to $[0.9, 1.11]$, $[-0.05, 0.05]$, respectively. γ is set to 4.

3.2.2 Computing Correction Parameters

The synthesis result I^S can not be directly used as the final output for several reasons. Firstly, patch synthesis does not

perform well for regions with unique structures, as shown in Fig. 2(d)(top). Secondly, patch synthesis may not converge well especially for highly textured regions, resulting in blurry results, as shown in the example in Fig. 2(d)(top and bottom). Nevertheless, I^S contains good color and texture information in a large portion of the shadow region, that can be used to estimate the corresponding parameters in the correction map.

Color Parameters For each patch P in the correction region, the color correction parameters can be derived from Eqn. 1 as:

$$\begin{aligned} \mathbf{g}^L(P) &= \sum_{p \in P} I_L^S(p) / \sum_{p \in P} I_L^N(p), \\ \mathbf{b}^A(P) &= \frac{1}{\|P\|} \left(\sum_{p \in P} I_A^S(p) - \sum_{p \in P} I_A^N(p) \right), \\ \mathbf{b}^B(P) &= \frac{1}{\|P\|} \left(\sum_{p \in P} I_B^S(p) - \sum_{p \in P} I_B^N(p) \right), \end{aligned} \quad (5)$$

where L, A, B denote the luminance and chrominance channels.

Texture Parameters Directly computing the texture correction parameter from the synthesis result I^S is sub-optimal given that it may be blurry and lack image details (see Fig. 2(d)(bottom)). We instead directly use the patch correspondence, instead of final synthesis result, to estimate this parameter. That is, for patch P in \mathcal{R}_c , we get a source region patch Q that is used to vote for the final synthesis result. We assume P should have the same texture characteristics as Q , So the texture consistency parameter can be computed as:

$$s^L(P) = \sigma_Q / \sigma_P. \quad (6)$$

Correction Confidence Given the original shadow removal result I^N and the patch synthesis result I^S , the syn-

thesis confidence $C(P)$ for each patch P is defined as:

$$C(P) = 1 - \frac{\min_{\mathbf{g}, \mathbf{b}} \|I^N(P) * \mathbf{g}(P) + \mathbf{b}(P) - I^S(P)\|_2}{\|I^N(P)\|_2 + \varepsilon}, \quad (7)$$

where the distance between $I^N(P)$ and $I^S(P)$ is minimized by searching the best gain for the L channel, and the best bias for the a and b channels per patch. The confidence is normalized by the average pixel luminance $\|I^N(P)\|_2$ to avoid a bias in dark regions. ε is a small constant to avoid division by zero. Intuitively, if $I^N(P_i)$ and $I^S(P_i)$ contain the same structure and only differ by a global color transform, we have high confidence that the patch synthesis result is correct, and the correction parameters is reliable. Otherwise if $I^N(P_i)$ and $I^S(P_i)$ contain structural differences, the confidence value will be low, indicating incorrect patch synthesis. Example confidence maps are shown in Fig. 2(e); note that unique structures in the shadow region such as the dark brown mountain in Fig. 2(top) cannot be synthesized well, and consequently pixels inside these structures have very small confidence values.

Next, we describe how to use the confidence map to refine the parameters in the correction map.

3.2.3 Correction Map Refinement

The initial correction map is not reliable for all shadow patches, and thus cannot be directly applied to I^N . In this section, we show how to refine it based on the computed correction confidence values (Eqn. 7). We treat color and texture parameters differently at this stage due to their inherently different nature.

Color Parameter Refinement Experiments show that the color correction parameters (gain/bias of CIELAB channel) are generally quite smooth in the shadow regions. To refine color parameters, we propagate them from high confidence patches to low confidence ones, by optimizing a quadratic objective function. Specifically, for channel k ($k = L, A, B$) of the original color correction parameters, denoted by S_0^k , we find new parameters S^k that minimize the following energy function:

$$E^k = \sum_P C(P) \cdot (S^k(P) - S_0^k(P))^2 + \lambda_s \sum_{Q \in \mathcal{N}(P)} (1 - C(P)) \cdot \mathcal{A}(P, Q) \cdot (S^k(P) - S^k(Q))^2, \quad (8)$$

where $\mathcal{N}(P)$ is the set of the neighboring overlapping patches of P . The first term is the data term that constrains S^k to be close to the original estimation S_0^k if the synthesis confidence $C(P)$ is high. The second term is the smoothness term, weighted by the color difference of neighboring

patches, along with $1 - C$. The affinity weight $\mathcal{A}(P, Q)$ is defined as:

$$\mathcal{A}(P, Q) = \mathcal{G}(\|I_n(P) - I_n(Q)\|_2, \sigma_m), \quad (9)$$

where $\mathcal{G}(x, \sigma) = \exp(-x^2/\sigma^2)$ is the Gaussian function. The smooth term allows low confidence patches to receive parameter values from neighboring patches that have similar colors to them. In our implementation we set $\lambda_s = 10$ and $\sigma_m = 0.2$, and solve the linear system using Conjugate Gradient.

Texture Parameter Refinement Unlike the color parameters, texture parameters are very noisy over the whole image. We thus cannot refine them using a similar optimization. Instead for patches with low correction confidence, we resort to the initial shadow removal result for computing the texture parameters. In other words, for regions where patch synthesis is not reliable, we maintain their original texture characteristics in the initial shadow removal result, to avoid introducing additional artifacts. Specifically, we use the correction confidence as the interpolation coefficient to compute the refined texture parameter $s^L(P)$ as:

$$s^L(P) = C(P) \cdot s_0^L(P) + (1 - C(P)) \cdot 1, \quad (10)$$

Where $s_0^L(P)$ is the computed texture parameter using Eqn. 6 on the synthesized image.

3.3. Applying the Correction Map

The final recovery result is obtained by applying the shadow correction map S on the initial shadow removal result I^N . Specifically, to remove color inconsistency, we apply color correction parameters ($\mathbf{g}^L, \mathbf{b}^A, \mathbf{b}^B$) of S to each patch and vote for the output image I^D as:

$$I_L^D(p) = \frac{1}{\|P\|} \sum_{q \in P} I_L^N(q) * \mathbf{g}^L(Q),$$


$$I_A^D(p) = \frac{1}{\|P\|} \sum_{q \in P} (I_A^N(q) + \mathbf{b}^A(Q)),$$

$$I_B^D(p) = \frac{1}{\|P\|} \sum_{q \in P} (I_B^N(q) + \mathbf{b}^B(Q)), \quad (11)$$

where P and Q are patches centered at pixel p and q , respectively. Furthermore, to remove texture inconsistency, we apply texture correction parameters \mathbf{s}^L of each patch and vote for the final recovery result I^C :

$$I_L^C(p) = \frac{1}{\|P\|} \sum_{p \in Q} \left(\frac{\mathbf{s}^L(Q) \sigma_L^N(Q)}{\sigma_L^D(Q)} (I_L^D(p) - \mu_L^D(Q)) + \mu_L^D(Q) \right), \quad (12)$$

where $\mu^D(Q)$, $\sigma^D(Q)$ are the average color and standard derivation of patch Q in I^D , and $\sigma_L^N(Q)$ is the standard derivation of patch Q in I^N .



Source image Previous result I^N Our result I^C

Figure 3. Improving previous shadow removal results with color inconsistency using the proposed SRH algorithm. (a) [7], (b) [21], (c) [11], (d) [10], (e) [13].

Fig. 2(f) shows examples of the corrected shadow regions. We can see that the color and texture inconsistency in the initial shadow removal results have been highly suppressed, resulting in more natural shadow removal results.

4. Results and Evaluations

We have implemented our algorithm in C++. On a PC with 3.4GHz CPU and 2G RAM, for a 600×450 image, our single-threaded implementation of the SRH algorithm takes about 2 minutes for the patch-based synthesis step, and 2 seconds for the rest - correction map construction, refinement and obtaining final result. In this section, we evaluate the SRH method by showing both visual examples and quantitative evaluation results on a benchmark dataset.

4.1. Visual Comparisons

Figs. 1(top), and 3, show shadow removal results from previous methods that contain significant color inconsistencies in the recovered shadow regions, including color shifts (Fig. 1(top), Fig. 3((b)(d)) and residual shadows (Fig. 3(a)(c)). We have experimented with a wide range of methods include [7, 21, 17, 11, 10, 13]. The inputs for our SRH algorithm and the results of these methods are all taken

from their original papers ¹. SRH successfully removes the color inconsistencies in the original results, and produces results that are more natural-looking.

In Fig. 1(bottom), Fig. 2(bottom) and Fig. 4 we show shadow removal results on previous methods that contain texture inconsistencies, such as inconsistent noise properties and texture contrast. These methods include [7, 23, 11, 25, 10]. Again, our SRH method successfully corrects these texture artifacts and produces more consistent shadow removal results.

In Fig. 5 we show the results of some algorithms on images from a recently proposed shadow removal benchmark dataset [10], along with improved results using our method. Again, the SRH method successfully suppresses both color and texture inconsistencies.

4.2. Benchmark Evaluation

We comprehensively evaluate our algorithm on the benchmark dataset mentioned above. This dataset consists of 214 test images, and provides quantitative errors of shadow removal results according to four attributes (more details in [10]): texture, brokenness, colorfulness and softness. The authors have also published shadow removal results using their technique as well as two other algorithms [14, 11] for all 214 test examples. We apply SRH on all the test cases using their shadow removal results as input, and report new errors and improvements in Table 1. Due to limited space, we only show the average score of each attribute.

The results show that our SRH method reduces shadow removal errors for all categories and all the three previous methods. Note that our algorithm cannot improve cases with detection errors, where the shadow region is wrongly detected (more in Sec. 4.4). Given that all these approaches introduce detection errors in some test cases, the performance improvement on examples with well-detected shadows are even higher. Some visual comparisons are shown in Fig. 5.

4.3. Robustness and Parameter Settings

Sensitivity to the initial result I^N - To evaluate the robustness of the SRH method, we manually generate synthetic examples of the initial result, I^N , with varying amount of residual shadow, by linearly interpolating shadow images with ground-truth shadow-free images using a fixed alpha matte α . Fig. 6 illustrates the performance of the SRH method on one such example on images generated with different α values and reports the errors of the corrected results against the ground-truth. For

¹Ideally we should have compared all the methods against a common set of examples. However, we did not have access to the code of many of these methods. Instead, we compare against each method on its own successful examples reported in the original paper.




Figure 4. Improving previous results with texture inconsistency using SRH. (a) [7], (b) [23], (c) [11], (d) [10], (e) [13].




Figure 5. Improving state-of-the-art results on examples in the benchmark dataset [10]. (top) Original input images; (middle) Initial shadow removal results I^N using [10]. (bottom) Our harmonized result I^C .

low values of α (0.2), i.e., medium shadow residuals, our algorithm can still generate a visually compelling result. At high values of α (0.5), the initial result is significantly corrupted, and as expected, visual artifacts arise and errors increase. We also evaluate the case of medium residual shadows ($\alpha = 0.2$) and a strong Gaussian noise. Our algorithm successfully corrects the color distribution and texture details of the shadow region (Fig. 6(g)), suggesting it is robust

against noise due to the use of patch statistics.

Color space and gain/bias settings - As described in Sec. 3.1, our shadow harmonization model uses CIELab color space, and enables gain in the L channel and bias in the a and b channels for color correction. We denote this model as Model 0). Here we compare its performance with other commonly-used color spaces and gain/bias settings: Model 1 (CIELAB, enabling L gain and L, a, b bias);




Figure 6. Results generated by our SRH method with synthetic I^N . (a) Full-shadow image & Shadow-free image; (b) and (d) Synthetic I^N generated by linear interpolation of full-shadow image and shadow-free image with $\alpha = 0.2, 0.5$; (c) and (e) SRH result applied on (b) and (d); Our algorithm is robust to an initial results with 20% residual shadow. (f) Close-up of a synthetic I^N with 20% residual shadow and strong Gaussian noise; (g) SRH result applied on (f). Our algorithm is robust to the apparent noise. Average errors w.r.t ground truth are reported in the bottom.

	Gong [10]		Gong [11]		Guo [14]	
	E_r	E_{r*}	E_r	E_{r*}	E_r	E_{r*}
Tex.	0.34 : 0.32 (6%)	0.21 : 0.17 (19%)	0.47 : 0.43 (9%)	0.36 : 0.31 (14%)	0.61 : 0.56 (8%)	0.51 : 0.44 (14%)
Soft.	0.40 : 0.37 (8%)	0.23 : 0.22 (12%)	0.51 : 0.45 (12%)	0.40 : 0.32 (20%)	0.77 : 0.69 (10%)	0.68 : 0.56 (18%)
Bro.	0.42 : 0.40 (5%)	0.25 : 0.22 (12%)	0.59 : 0.53 (10%)	0.52 : 0.41 (21%)	0.81 : 0.77 (5%)	0.76 : 0.70 (8%)
Col.	0.44 : 0.40 (9%)	0.29 : 0.23 (21%)	0.75 : 0.72 (4%)	0.69 : 0.65 (6%)	1.12 : 1.01 (10%)	1.09 : 0.93 (15%)
Other	0.40 : 0.36 (10%)	0.26 : 0.21 (19%)	0.57 : 0.51 (11%)	0.48 : 0.40 (17%)	0.72 : 0.66 (8%)	0.65 : 0.56 (14%)

Table 1. Quantitative results of the SRH algorithm on the benchmark dataset [10]. The quality of the results is evaluated w.r.t four attributes: texture, brokenness, colorfulness and softness. E_r is the error for shadow region only, E_{r*} is the error for the entire image. Results format - Error before harmonization : Error after harmonization with SRH (relative error decrease).

Model 2 (RGB, enabling R, G, B gain); Model 3 (RGB, enabling R, G, B gain and bias); Model 4 (HLS, enabling L gain and H, S bias). For each model, Eqn. 5 and Eqn. 12 are modified depending on whether gain/bias is enabled for each channel. If both gain and bias are enabled for a channel, the two values are computed by matching the mean and standard variation of the patch pairs. We compare these different color models on the benchmark dataset, using [10]’s shadow removal results as I^N . Resulting errors of each model are reported in supplemental material, which suggest that Model 0 achieves slightly better performance than other models.

Parameter ranges Parameter range settings (gain of L channel, bias of a, b channel) are important parameters in the patch-based synthesis process of SRH (Sec. 3.2.1). If the ranges are too narrow, PatchMatch may not have sufficient freedom to correct errors in the initial shadow removal results. On the other hand, if the ranges are too wide, patches of different materials are more likely to be matched. To find the best parameter ranges, we test different range settings on the benchmark dataset, again using [10]’s shadow removal results as I^N . Specifically, we test three parameter range settings: narrow (L gain [0.99, 1.01], a & b bias $[-0.01, 0.01]$), middle (L gain [0.9, 1.11], A & b bias $[-0.05, 0.05]$), and wide (L gain [0.8, 1.25], a & b bias $[-0.1, 0.1]$). The resulting errors are reported in supplemental material in detail. The middle range setting achieves the best performance and we fix the range to this range when generating all results reported in the paper. Also, the performance difference between the three settings are small,

indicating the robustness of the algorithm against these parameter settings.

4.4. Limitations

Our approach has several limitations. Firstly, the SRH method cannot correct errors introduced by shadow detection failure. For the example showed in Fig. 7(top), the shadow detection process of [11] fails and generates removal results with strong artifacts in the non-shadow region. SRH removes some of them, but noticeable artifacts still persist.

We have also observed that previous shadow removal methods sometimes generate significant boundary discontinuities. Fig. 7(bottom) shows one such example. Our method is mainly designed for harmonizing the interior of the shadow region, and is more limited at correcting such boundary discontinuities. As shown in the figure, it successfully removes color and texture inconsistencies inside the shadow region, but leaves some amount of boundary discontinuity in the final result.

5. Conclusion

We propose a fully automatic Shadow Region Harmonization approach for removing color and texture inconsistencies introduced by previous shadow removal methods. This technique is based on a parametric correction model, whose parameters are estimated by reconstructing the shadow region using non-shadow patches through a patch-based, guided image synthesis process. We also introduce synthesis confidence to deal with unique structures



(a) Input image (b) Initial result (c) Our result

Figure 7. Failure cases. (top) Our algorithm does not handle significant shadow detection failure as shown in this example from [11]. (bottom) The SRH algorithm harmonizes the interior of the shadow region, but cannot fix the significant boundary discontinuities on this example from [10].

and materials inside the shadow region. Extensive evaluation shows the effectiveness and robustness of the proposed method.

References

- [1] M. Baba, M. Mukunoki, and N. Asada. Shadow removal from a real image based on shadow density. In *ACM SIGGRAPH 2004 Posters*, 2004.
- [2] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman. Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3), 2009.
- [3] C. Barnes, E. Shechtman, D. B. Goldman, and A. Finkelstein. The generalized patchmatch correspondence algorithm. In *Proc. ECCV*, 2010.
- [4] H. G. Barrow and J. M. Tenenbaum. Recovering intrinsic scene characteristics from images. In *Int. Conf. on Computer Vision Systems*, 1978.
- [5] S. Darabi, E. Shechtman, C. Barnes, D. B. Goldman, and P. Sen. Image Melding: Combining inconsistent images using patch-based synthesis. *ACM Trans. Graph.*, 31(4), 2012.
- [6] A. Ecsins, C. Fermller, and Y. Aloimonos. *Shadow free segmentation in still images using local density measure*. IEEE Computer Society, 2014.
- [7] G. Finlayson, M. Drew, and C. Lu. Intrinsic images by entropy minimization. In *Proc. ECCV*, 2004.
- [8] G. D. Finlayson, S. D. Hordley, C. Lu, and M. S. Drew. Removing shadows from images. In *Proc. ECCV*, 2002.
- [9] C. Fredembach and G. D. Finlayson. Hamiltonian path-based shadow removal. In *Proc. BMVC*, 2005.
- [10] H. Gong and D. Cosker. Interactive shadow removal and ground truth for variable scene categories. In *Proc. BMVC*. BMVA Press, 2014.
- [11] H. Gong, D. Cosker, C. Li, and M. Brown. User-aided single image shadow removal. In *Proc. ICME*, 2013.
- [12] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. Ground-truth dataset and baseline evaluations for intrinsic image algorithms. In *Proc. ICCV*, 2009.
- [13] M. Gryka, M. Terry, and G. J. Brostow. Learning to remove soft shadows. *ACM Trans. Graph.*, 2015.
- [14] R. Guo, Q. Dai, and D. Hoiem. Single-image shadow detection and removal using paired regions. In *Proc. CVPR*, 2011.
- [15] Y. HaCohen, E. Shechtman, D. B. Goldman, and D. Lischinski. Non-rigid dense correspondence with applications for image enhancement. *ACM Trans. Graph.*, 30(4), 2011.
- [16] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin. Image analogies. In *Proc. ACM SIGGRAPH*, 2001.
- [17] F. Liu and M. Gleicher. Texture-consistent shadow removal. In *Proc. ECCV*, 2008.
- [18] D. Miyazaki, Y. Matsushita, and K. Ikeuchi. Interactive shadow removal from a single image using hierarchical graph cut. In *Proc. ACCV*, volume 5994. 2010.
- [19] A. Mohan, J. Tumblin, and P. Choudhury. Editing soft shadows in a digital photograph. *IEEE Computer Graphics and Applications*, 27(2), 2007.
- [20] A. Sanin, C. Sanderson, and B. C. Lovell. Shadow detection: A survey and comparative evaluation of recent methods. *Pattern Recogn.*, 45(4), 2012.
- [21] Y. Shor and D. Lischinski. The shadow meets the mask: Pyramid-based shadow removal. *Computer Graphics Forum*, 27(2), 2008.
- [22] Y. Wexler, E. Shechtman, and M. Irani. Space-time completion of video. *IEEE Trans. PAMI*, 29(3), 2007.
- [23] T.-P. Wu and C.-K. Tang. A bayesian approach for shadow extraction from a single image. In *Proc. ICCV*, volume 1, 2005.
- [24] T.-P. Wu, C.-K. Tang, M. S. Brown, and H.-Y. Shum. Natural shadow matting. *ACM Trans. Graph.*, 26(2), 2007.
- [25] C. Xiao, D. Xiao, L. Zhang, and L. Chen. Efficient shadow removal using subregion matching illumination transfer. *Computer Graphics Forum*, 32(7), 2013.
- [26] L. Xu, F. Qi, R. Jiang, Y. Hao, G. Wu, L. Xu, F. Qi, R. Jiang, Y. Hao, and G. Wu. Shadow detection and removal in real images: A survey, 2006.