

Interactive Removal and Ground Truth for Difficult Shadow Scenes

HAN GONG^{1,*} AND DARREN COSKER²

¹School of Computing Sciences, University of East Anglia, UK

²Department of Computer Science, University of Bath, UK

*Corresponding author: h.gong@uea.ac.uk

Compiled August 9, 2016

A user-centric method for fast, interactive, robust and high-quality shadow removal is presented. Our algorithm can perform detection and removal in a range of difficult cases: such as highly textured and colored shadows. To perform detection an on-the-fly learning approach is adopted guided by two rough user inputs for the pixels of the shadow and the lit area. After detection, shadow removal is performed by registering the penumbra to a normalized frame which allows us efficient estimation of non-uniform shadow illumination changes, resulting in accurate and robust removal. Another major contribution of this work is the first validated and multi-scene category ground truth for shadow removal algorithms. This data set containing 186 images eliminates inconsistencies between shadow and shadow-free images and provides a range of different shadow types such as soft, textured, colored and broken shadow. Using this data, the most thorough comparison of state-of-the-art shadow removal methods to date is performed, showing our proposed new algorithm to outperform the state-of-the-art across several measures and shadow category. To complement our dataset, an online shadow removal benchmark website is also presented to encourage future open comparisons in this challenging field of research.

1. INTRODUCTION

Shadows are ubiquitous in image and video data, and their removal is of interest in both computer vision and graphics. Although shadows can be useful cues, *e.g.* shape from shading, they can also affect the performance of algorithms (*e.g.* in segmentation and tracking). Their removal and editing is also often the pain-staking task of graphical artists. A successful shadow removal method should seamlessly relight the shadow area while keeping the lit area unchanged. The umbra is the darkest part of the shadow whilst the penumbra is the transitional shadow boundary with a non-linear intensity change between the umbra and lit area. The textures in shadowed surface generally become weaker that contrast artifacts can appear in shadow areas due to image post-processing [1].

One of the difficulties in detecting and removing shadows is the large variability in different shadow types. In particular, the following common attributes (*e.g.* Fig.10) of shadows can significantly increase the difficulty of their removal:

- **Texture of cast surface** Strong texture causes higher intensity variation which makes it difficult to extract illumination change from intensity changes. In addition, dark textures can appear similar to shadows, which can often confuse

shadow detection algorithms.

- **Shadow softness** Softness generally relates to the size of a shadows penumbra. Higher softness brings challenges in preserving penumbra texture when removing shadow. When the illumination change becomes much weaker than the intensity change caused by texture, it can be difficult to extract the component of illumination change.
- **Broken shadow** Broken shadows contain variable illumination attributes such as irregular shape, highly varying penumbra size, and overlapping penumbra. Fixed illumination models can find such irregular illumination changes difficult to process.
- **Shadow color** When shadows are not conventionally black but instead colorful, it is not only difficult for machines to detect this appearance change but even humans. Also, even when shadows of this kind are detected, their removal is still difficult as the color in the umbra could be related to the surfaces' reflection.

Given sufficient training data, automatic approaches can robustly remove common shadows, however there are difficul-

ties in dealing with difficult shadow scenes. User-aided approaches give users more control over difficult shadow removal. In this paper, an interactive, high-quality and robust method for fast shadow removal is proposed using two rough user-defined strokes indicating the shadow and lit image areas. The approach sacrifices full autonomy [2, 3] for very broad and simple user input – contrasting with existing manual approaches that require fine-scale input (accurate shadow contours [4, 5]) or highly simplistic inputs (single pixel [6]) that can result in shadow detection artifacts. This on-the-fly learning approach is robust to large variations in user input, and can remove shadows with difficult attributes such as colored and broken shadow. Given detection, reliable shadow removal is delivered – verified with thorough quantitative tests for different types of shadow (for the first time in this area) comparing to previous state of the art approaches. A large high-quality and multi-scene-category ground-truth data set for the evaluation of shadow removal is also presented – consisting of 186 images with reliable ground truth. This overcomes issues with previous data sets – such as inconsistencies between shadow and shadow-free images – the results of which are also quantitatively verified and compared against previous data sets. The approach presented represents what the authors believe to be a state-of-the-art method for shadow removal, with the most robust evaluation of such methods to date across a range of difficult shadow cases.

1.A.Related work

A shadow is generally defined as having an umbra and penumbra area – denoted by the central shadow region and its border (penumbra) transitioning illumination between the fully dark and lit area. A shadow image I_c^+ can be considered to be a Hadamard product [7] of a shadow scale layer S_c and a shadow-free image I_c^* as follows:

$$I_c^+ = I_c^* \circ S_c \quad (1)$$

where c is a RGB channel. For a lit pixel, the illumination is constant in both shadow and shadow-free images. For a shadow pixel, its intensity in a shadow image is lower than its intensity in the shadow-free image. Therefore, the scales S_c of the lit area are 1 and other areas’ scales are between 0 and 1.

Approaches to shadow removal can be categorized as either automated or user-aided. The differentiation between fully automated or user-aided relates to initial detection of the shadow – with removal itself (after detection) being a largely automatic task. In any case, both removal and detection are ill-posed problems and difficult to reliably achieve. We summarize the features and requirements of the recent shadow removal methods in Table 1.

1.A.1.Shadow detection

As for automated shadow detection, intrinsic image [10] based methods and illumination-invariant image [8, 9, 17] based methods are one such popular approach to the problem. The decomposition of intrinsic images provides shading and reflectance information but can be unreliable leading to over-processed results. Intrinsic image based methods generally assume that the illumination change leads to smooth intensity change and the neighboring pixels have similar chromaticities. Illumination-invariant image is fast to compute but it only provides reflectance information. The derivation of illumination-invariant image assumes that the image is linear (not rendered). However, most of the images found on the Internet are rendered, *e.g.*, compressed and

Table 1. Feature comparison of shadow removal methods: “Illumination Preserving” refers to the ability to preserve the original illumination in the lit area. “Texture Preserving” refers to the preservation of the correct surface texture under the penumbra after removal. “Color Correction” refers to the ability to correct color artifacts caused by image post-processing after removal.

a) Automatic Shadow Removal					
	Preservation		Color Correction	Special Setup	
	Illum.	Text.			
Finlayson [8]	No	No	No	None	
Fredembach [9]	No	No	No	None	
Yang [10]	No	Yes	No	None	
Guo [2]	Yes	Yes	No	None	
Drew [11]	No	Yes	No	Flash	
Salamaty [12]	Yes	Yes	No	NIR/RGB	
Khan [3]	Yes	Yes	No	None	
b) User-Aided Shadow Removal					
	User Input	Preservation		Color Correction	
		Illum.	Text.		
Wu [13]	multiple strokes for each shadow	Yes	Yes	No	
Liu [5]	precise boundary mark for each shadow	Yes	Yes	Yes	
Shor [6]	single click for each shadow	Yes	No	No	
Su [4]	rough boundary mark for each shadow	No	Yes	No	
Arbel [1]	multiple clicks for each shadow	Yes	Yes	Yes	
Gong [14]	single stroke for each shadow	Yes	Yes	Yes	
Xiao [15]	two types of scribbles for sampling shadow and lit intensities	Yes	Yes	Yes	
Zhang [16]	two types scribbles for sampling shadow and lit intensities	Yes	Yes	Yes	

gamma corrected. The non-linearity caused by image rendering can break the algorithm. Besides, their shadow detection relies on comparing the difference between the edge detection results of a shadow image and its illumination-invariant image. This property makes these methods incapable of removing soft or light shadows. Shadow detection can also be achieved by shadow features learning [2, 3, 18–23]. However, shadow detection is constrained by the range of training data and quality of classifier, image edge detection and segmentation there-in. [18–22] detect shadows by classifying edges in an image according to shadow edge features such as changes in intensity, texture, and color ratio. Guo *et al.* [2] adopt similar features but detect shadows by classifying segments in an image and pairing shadow and lit segments globally. This method is more robust because segment pairing correlates both neighboring and non-neighboring surfaces. Some recent methods [3, 23] adopted Con-

volution Neural Network (CNN) for detecting shadows from single images. Based on training from massive data, CNN-based shadow detection provides fast speed and high accuracy. However, the science behind a CNN remains unexplainable. Some methods require additional controllable light sources to capture shadow-less objects, *e.g.*, by comparing flash and no-flash image pairs [11]. However, active lighting restricts the applicable type of scenes - as moving lights around and using special lighting setups outdoors is often not practical. Some methods adopt optical filters to obtain multi-spectral images for illumination detection, *e.g.* by comparing NIR and RGB images [12] and by comparing RGB and single-color-filtered image [24], but these methods assume some special scenarios, *e.g.* sunlight and non-black surfaces. They are thus not applicable to the removal of normal single RGB images.

User-aided methods generally achieve higher accuracy in shadow detection at the practical expense of varying degrees of manual assistance. Wu *et al.* [13] require extensive user input where the user needs to define multiple regions of shadow, lit area, uncertainty and exclusion. They estimate the probability that a pixel is part of a shadow according to a 3D Gaussian Mixture Model (GMM) [25] generated from the supplied pixel samples. Liu and Gleicher [5] require fine input defining the accurate shadow boundary. Su and Chen [4] require a rougher shadow boundary as input and align the shadow boundary according to intensity gradient. However, it is cumbersome to define the boundaries of broken shadows in [5] and [4]. Arbel and Hel-Or [1] require users to specify multiple texture anchor points to detect a shadow mask but the number of user input can significantly increase when shadow regions are multiple and scattered. Xiao *et al.* [15] and Zhang *et al.* [16] require two types of user scribbles for sampling shadow and lit intensities. Given enough scribbles, their detection methods can robustly produce shadow masks. However, their methods also require a shadow matte (guided by the scribbles) to identify shadows, which is sensitive to user-scribbles because their image matting is affected by pixel location. Shorl and Lischinski [6] only require one shadow pixel as input. The algorithm detects shadow using image matting from a grown shadow seed. But it has limitations in cases where the other shadowed surfaces are not surrounded by the initially detected surface or when the penumbra is too wide.

1.A.2.Shadow relighting

Shadow relighting is another difficult problem, especially for the recovery of penumbra. Some methods apply zero-penumbra-gradient-filling [26] or native in-painting [6, 9, 17] for penumbra recovery which result in penumbra texture loss. Finlayson *et al.* [8,27] apply an iterative diffusion process to smoothly fill in the derivatives in penumbra which but can still cause texture loss in penumbra. Liu and Gleicher [5] apply a curve fitting method and a global alignment of gradients to acquire shadow scales but has issues when relighting the umbra and can introduce artifacts at uneven boundaries. Arbel and Hel-Or [1] apply a thin-plate model to fit the intensity surface and the algorithm is designed for removing shadows from curved surfaces. Guo *et al.* [2] remove shadows by image matting which treats shadows as the foreground. Their approach can usually introduce penumbra artifacts. Gong *et al.* [14] also apply a curve fitting model and they adopt an intelligent sampling scheme to improve the quality of intensity samples for illumination estimation. Su and Chen [4] estimate shadow scales by using dynamic programming. These data fitting based methods [1, 5, 14] assume

highly-constrained curve or surface functions for illumination change which limit their range of removable shadows. Xiao *et al.* [15] apply a multi-scale adaptive illumination transfer which performs well for removing shadows cast on strong texture surfaces. Zhang *et al.* [16] remove shadows by aligning the texture and illumination details between corresponding shadow and lit patches. However, both [15] and [16] are very sensitive to variable user inputs. Khan *et al.* [3] apply a Bayesian formulation to robustly remove common shadows. However, this method is unable to process difficult shadows such as non-uniform (or broken) shadows. It is also computationally expensive due to a large number of unknown parameters.

1.A.3.Shadow removal ground truth

To date, most shadow removal methods (*e.g.* [1, 4, 13-16]) have been evaluated by visual inspection on some selected images - with only a few exceptions performing quantitative evaluation. This is in part due to a lack of high-quality, varied, and public ground truth data. Shorl and Lischinski [6] perform a quantitative test but comparison is difficult due to their data not being publicly available. Guo *et al.* [2] provide the first public ground truth data set for shadow removal and perform quantitative testing. However, the difficulty of collecting such a data set is highlighted in their work, with the appearance of some global illumination changes and mis-registration between the shadow and shadow-free images being a difficult factor to control. This can make quantitative testing on such data somewhat difficult, as these errors can influence shadow removal results.

Another desirable property as yet not explored by existing data sets or fully explored in work on detection and removal is the categorization of shadows. Such attributes are important to consider as these different shadow types can present their own unique challenges, *e.g.* removal of colored shadows (*i.e.* through a glass bottle) are more difficult than consistent unbroken shadows (*i.e.* a human silhouette). Universal shadow removal approaches should therefore be capable of handling these multiple cases. In addition, having such categories in a ground truth data set is also important - if only to allow us to evaluate different algorithm performance in a range of scenarios and scene types.

1.B.Contributions

Given our review of previous work, 4 main contributions are proposed:

- 1) A rigorous, highly-varied and categorized shadow removal ground truth data set:** Our quantitatively verified high quality data set contains 186 ground truth images organized into common shadow categories. Based on this data, our method is quantitatively evaluated against other state-of-the-art algorithms on different shadow category types.
- 2) Simple and robust user input based shadow detection:** Our shadow detection component requires only two rough user scribbles marking samples of lit and shadow pixels. Our approach differs from previous work requiring more complex user-inputs [1, 4, 5, 13] or simpler inputs [6] that compromise robustness and quality.
- 3) High quality and fast shadow removal:** Unlike existing methods requiring slow pixel-wise optimization [2-4, 6, 13] or an inflexible fitting model [1-5, 15, 16], penumbra unwrapping is introduced upon which multi-scale smoothing is performed to derive sparse shadow scales across the penumbra. This allows robust and efficient estimation of illumination changes without requiring prior training and any assumed illumination change

models. This method is simple and fast yet offers state-of-the-art shadow removal quality.

4) Robust color correction: Post-processing effects may cause inconsistency in shadow corrected areas compared with the lit areas both in tone and contrast. A robust multi-scale color correction is proposed to amend these artifacts.

To summarize, the authors believe these contributions to be important to this area of research due to their significant improvements over the state-of-the-art in shadow removal in a wide range of repeatable tests.

2. SHADOW REMOVAL GROUND TRUTH

A thorough quantitative evaluation of shadow removal performance requires a high-quality, diverse shadow-free ground truth. The first public data set was supplied by [2]. However, while this is a valuable resource for evaluating shadow removal and the first of its kind, there are many opportunities for expansion and several improvements are presented in our new data set. Firstly, the concept of shadow categories is introduced for the first time in our proposed data set, and a wide range of new types proposed. Secondly, ground truth is constructed and verified in a careful manner so as to remove irregularities between test and validation images. In terms of the latter, we note environmental illumination and registration errors between some shadow and ground truth images in existing data sets. An example of comparison is shown in Fig.1. Our new data set avoids these issues and represents, we believe, the most stable and thorough data set for shadow removal evaluation available today. In order to highlight the benefits of our rigorous data protocol, in §4 the quality of our ground truth data is quantitatively compared to another state-of-the-art dataset [2].

Shadow images and their ground truth are captured using a camera with a tripod and a remote trigger. This rig minimizes misalignment due to camera shake. To minimize illumination variance, images are captured within a very short interval of time using a manual capture mode with fixed ISO and exposure settings. When collecting data, environmental effects are often unavoidable, e.g., wind can cause camera shake or the sun might move behind the clouds. Such failed acquisitions are rejected from our data set using a quantitative assessment outlined in §4. For evaluation, our shadow data is also categorized according to 4 different attributes: degree of texture, shadow softness, brokenness of shadow, and color variation. All the validated shadow images are manually categorized according to 4 shadow categories and 3 intensity degrees. The labeling was performed by 5 users and their average responses are rounded to the nearest intensity degree numbers (e.g. 1 for “weak” and 3 for “strong”). In total, our final data set after rejection consists of 186 test cases. For comparison, the previous state of the art from [2] consists only of 28 test cases after applying our strict rejection measure. Examples of images in each category are shown in §4.E.

3. INTERACTIVE SHADOW REMOVAL

In this section, our algorithm is first explained in brief. Technical details for each of its components are then expanded upon in following sections. Our algorithm consists of 4 steps (see Fig.2):

1) Pre-processing (§3.A) An initial shadow mask (Fig.2(b)) is detected using a KNN classifier trained from data from two rough user inputs (e.g. Fig.2(a)). A *fusion image*, which magnifies illumination discontinuities around shadow boundaries, is generated by fusing channels of YCrCb color space and suppressing

texture (Fig.2(c)).

2) Penumbra Unwrapping (§3.B) Based on the detected shadow mask and fusion image, pixel intensities of sampling lines are sampled perpendicular to the shadow boundary (Fig.2(d)).

Noisy samples are removed and remaining columns stored as the initial penumbra strip (Fig.2(e)). The initial columns’ illumination changes are also aligned (Fig.2(f)) by a fine-scale alignment.

3) Relighting (§3.C) From the penumbra strip, a multi-scale shadow scale estimation is applied to quickly and robustly estimate the illumination change along sampling lines and derive the sparse scales for all sampled sites (Fig.2(g)) which are propagated to form a dense scale field (Fig.2(h)). Shadows are removed by inverse scaling using this non-uniform field (Fig.2(i)).

4) Color Correction (§3.D) Post-processing effects may cause inconsistent tone and contrast in shadow removed areas compared with the lit areas’. Without introducing additional artifacts, a multi-scale color correction is proposed to remove these inconsistencies (Fig.2(j)).

Our shadow removal approach includes some standard algorithms which require parameters. These required parameters are denoted throughout the paper and are determined by genetic optimization based parameter learning in §3.E. In this paper, we denote 6 undetermined parameters as h_1, h_2, \dots, h_6 .

3.A. Pre-Processing

Pre-processing provides a detected shadow mask and a fusion image to assist penumbra unwrapping. Although there have been automatic methods for shadow detection (e.g. [2, 3, 18–20, 28]), results are dependent on training data quality and variation. Instead, our method requires no prior training or learning – only two user-supplied rough inputs indicating sample lit and shadow pixels (Fig.2(a)). Highlighted pixels’ RGB intensities in the Log domain are supplied as training features and used to construct a KNN classifier ($K = 3$). Euclidean distance is used as the distance measure and the majority rule with nearest point tie-break as the classification measure. Spatial filtering with a Gaussian kernel (size = h_1 , standard deviation = $h_1/2$) is applied to the obtained image of posterior probability and binarize the filtered image using a threshold of 0.5 (e.g. Fig.2(b)).

Although detection errors along the boundary, as well as post-filtering, can result in intensity samples with unaligned illumination changes at sharp boundaries, our penumbra unwrapping and alignment step (§3.B) can compensate for this.

Thus, our shadow *removal* method is somewhat robust to noise in the initially detected shadow mask, and would also be applicable to alternative (e.g. automatic) detection methods. While our user input format is identical to the two types of scribbles adopted in [15, 16], our method is found more robust to rougher (or fewer) user inputs since our method does require an image matting process that is also guided by sampled pixel location (see §4.A for a test example).

To assist unwrapping of the penumbra, an image is derived that magnifies illumination discontinuities around the shadow boundary – also assisting penumbra location – which is called the *fusion image* (e.g. Fig.2(c)). There are 2 steps in this process:

1) Magnification of Illumination Discontinuity An initial fusion image F is derived that maximizes the contrast between shadow and lit areas by linearly fusing the three channels (C_l) of YCbCr space as follows:

$$F = \sum_{l=1}^3 a_l C_l \text{ subject to } \sum_{l=1}^3 a_l = 1 \quad (2)$$

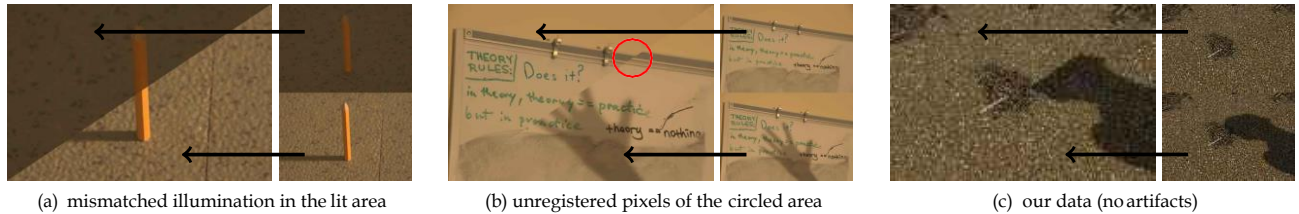
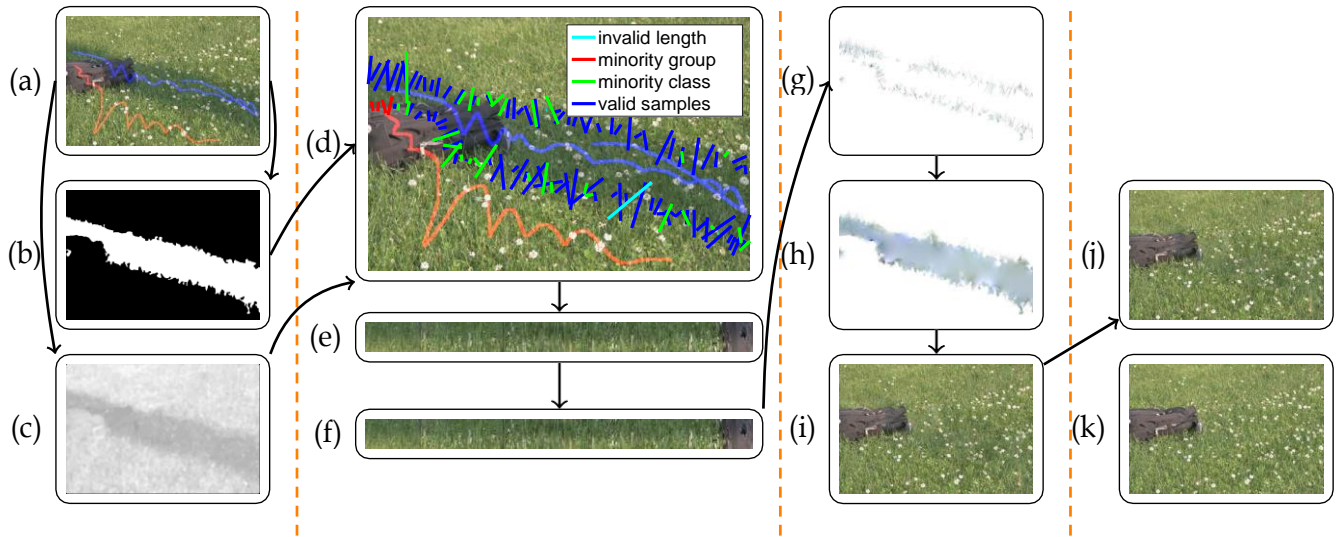


Fig. 1. Issues of shadow-free ground truth in the previous data set [2]. To easily examine the ground truth artifacts, we extract one half from a shadow image and another half from a shadow-free ground truth image and merge these two halves. For each sub-figure: top left segment – shadow-free image; bottom right segment – shadow image. (a) and (b) are taken from [2] which reflect the two annotated issues. An example from our data – which rejects image pairs with these properties is shown in (c).



1) Pre-Processing (§3.A) 2) Penumbra Unwrapping (§3.B) 3) Relighting (§3.C) 4) Color Correction (§3.D)

Fig. 2. Our shadow removal pipeline. (a) input: a shadow image and user strokes (blue for lit pixels and red for shadowed pixels); (b) detected shadow mask; (c) fusion image; (d) initial penumbra sampling (the actual density of samples are higher than the displayed samples’); (e) initial penumbra unwrap (only the shadow edges of the largest shadow segment is shown); (f) further aligned penumbra unwrap; (g) sparse shadow scale; (h) dense shadow scale; (i) initial shadow removal result; (j) color corrected shadow removal result; (k) ground truth.

where a_l is the positive fusing factor of C_l . The best fusing factors are derived by minimizing the following objective function E_b :

$$E_b(\mathbf{a}) = \mu(F_S) / \mu(F_L) + (\sigma(F_S) + \sigma(F_L)) / \sigma(F_{S \cup L}) \quad (3)$$

where \mathbf{a} is the vector of fusing factors and F_S and F_L are the two sets of shadow and lit pixels marked by user scribbles. In this paper, σ and μ are defined as functions that respectively compute the standard derivation and mean of a set of values. The first term ensures larger distinction between pixels of lit and shadow regions and the second term ensures smaller variation for pixels of the same lit or shadow regions.

2) Suppression of Texture The noise due to image texture is reduced by applying a median filter with a h_2 -by- h_2 neighborhood to F .

YCbCr color space offers perceptually meaningful information. Empirically, illumination appears dominantly in one of its channels. The illumination information in RGB channels are usually affected by texture noise. An example of comparison between fusing channels using YCbCr color space and RGB color space is shown in Fig.3.




Fig. 3. Comparison of fusion image using different color spaces. The same optimization scheme is applied to the image in Fig.2(a) but using different color spaces. The YCrCb fusion image presents more clean illumination information.

3.B. Penumbra Unwrapping

A shadow boundary generally has a noisy profile with a variable penumbra width. This can lead to inaccurate estimation of shadow scales and resulting artifacts. The penumbra is therefore unwrapped into a strip and its sampled columns of illumination change are aligned (e.g. Fig.2(f)). This improves the detection of

outliers and allows linearization of processing along the penumbra – leading to significant gains in efficiency and speed.

Algorithm 1: Penumbra Sample End Point Selection

input : boundary point (x_b, y_b) , fusion image F
output : two ends $(\mathbf{p}_s, \mathbf{p}_e)$ of a sampling line
 $\tilde{F} \leftarrow \nabla F$; $\mathbf{p}_s \leftarrow (x_b, y_b)$; $\mathbf{p}_e \leftarrow (x_b, y_b)$; $L \leftarrow |\tilde{F}(x_b, y_b)|$;
 $\Delta \mathbf{v} \leftarrow \tilde{F}(x_b, y_b) / L$;
repeat
 $\mathbf{v}_s \leftarrow \tilde{F}([\mathbf{p}_s])$; $\mathbf{v}_e \leftarrow \tilde{F}([\mathbf{p}_e])$;
 $L_s \leftarrow \mathbf{v}_s \cdot \Delta \mathbf{v}$; $L_e \leftarrow \mathbf{v}_e \cdot \Delta \mathbf{v}$;
 $\mathbf{p}_s \leftarrow \mathbf{p}_s - \Delta \mathbf{v}$; $\mathbf{p}_e \leftarrow \mathbf{p}_e + \Delta \mathbf{v}$;
until \mathbf{p}_s or \mathbf{p}_e is not within the range of F or $h_5 L_s > L$ or
 $h_5 L_e < L$;

Similar to prior work [1, 3], the intensity of sampling lines perpendicular to the shadow boundary (Fig.2(d)) are sampled. The length of a sampling line is determined by locating suitable start and end points guided by the fusion image F . A bi-directional search is initialized from each boundary point that extends the sampling line towards the lit area (end point) and the shadow area (start point) as described in Algorithm1. This extension is symmetric. The start and end points are initially set as the boundary point (x_b, y_b) and the direction vector $\Delta \mathbf{v}$ as the normalized gradient vector of (x_b, y_b) . To get the position for a start point, $\Delta \mathbf{v}$ is iteratively subtracted from the start point until its projected gradient is small enough (vice versa for the end point). Absolute gradient magnitude is not used in this algorithm because the gradient magnitudes of the soft penumbra edges can be very insignificant. The length of a sampling line thus depends on the starting gradient strength at the middle of penumbra.

Instead of processing unaligned and unselected samples individually [1, 3], we transform these samples into unified columns of the initial penumbra strip to enable fast batch processing. And, only the good samples are kept for shadow scale estimation. To avoid outliers, e.g. sampling lines at occlusion boundaries, invalid samples are filtered based on an assumption of similar shadow scales. A scale vector $Y_c = T_l - T_s$ is first computed where T_l and T_s are the average Log-domain RGB intensities of the lit and shadow halves of a sampling line. Y_c is then converted to spherical coordinates and considered as feature vector Y_s . DBSCAN clustering [29] (radius: h_3) is applied to Y_s for all samples, and samples that belong to the largest cluster are stored as valid ones with valid illumination. For finer scale estimation, valid clusters are further divided into a few sub-groups using mean-shift [30] (band width: h_4) and the samples of invalid sub-groups, whose total members are less than 10% of the largest sub-group’s, are discarded. Fig.2(d) shows an example of the above outlier detection.

To achieve an efficient batch shadow scale estimation, we need to cancel out the affection of variable penumbra width. The lengths of samples are normalized by re-sizing all the samples to a unique length n_a which is the maximum length of all valid samples. The normalized samples are then concatenated as columns to form the initial penumbra strip. Although our previous adaptive sampling already provides the intensity profiles with roughly aligned illumination changes, some minor errors may still exist. It is assumed that, after the previous intensity outlier filtering and the sample length normalization, the trends of intensity changes are similar and the dissimilarity only ap-

pears on the level of noise (e.g. background texture) and some minor alignment errors. This is resolved by a fine-scale alignment by optimization. For each column (intensity profile), the alignment process vertically shifts the column’s center and then stretches the column about its center by shifting its two ends. An illustration of this alignment is shown in Fig.4. The parameters




Fig. 4. Alignment of penumbra strip: The orange circles are the centers of columns in the penumbra strip. The orange dashed lines indicate the desired column length for a strip. The blue and green lines indicate the shifts required. In each iteration of the optimization, the alignment is in two steps: (left to middle) alignment of center; (middle to right) alignment of illumination change.

of this fine-scale alignment for each column are estimated by minimizing the following energy function E_a :

$$L_n = \Gamma(A_s, A_k, L_o) \quad (4)$$

$$E_a = \text{MSE}(L_n - L_a) \quad (5)$$

where A_s and A_k are the stretching shift and the center shift in the fine-scale alignment respectively, L_o is the scales of original column, L_a is the reference of alignment which is the average scale values of all valid columns (i.e. column-wise mean of the penumbra unwrap), L_n is the aligned unwrap, Γ is a function that aligns L_o according to the estimated alignment parameters, MSE is a function that computes mean squared error. The minimization is solved using a sequential quadratic programming algorithm [31].

3.C.Relighting

Using the aligned penumbra unwrap, a fast multi-scale shadow scale estimation is adopted for each shadow boundary. Compared with [4, 5, 14], our shadow scale estimation is fast and adaptive, which neither requires computational-costly pixel-wise optimization nor assumes any constrained data fitting models of illumination change, e.g. cubic curves or surface models. The non-linear image post-processing can significantly distort the original shadow scale change. Also, complex lighting conditions make the penumbra shadow scale change unpredictable. Instead, we only assume that illumination change is smooth and surface material change causes sharp intensity variation. Our recovery does not constrain the shape of the smooth shadow scale change of a sampling line. This means that our penumbra recovery is compatible with a wider range of shadow scale changes (e.g. ours can remove unconventional colored shadows) as long as the shadow scale changes of neighboring sampling lines are not too dissimilar. Our previous sample length normalization and alignment make it possible to estimate the shadow scale change by a simple and adaptive horizontal smoothing.

A pyramid (e.g. Fig.5(b)) of horizontally filtered penumbra unwraps using 5 averaging kernels in different sizes are computed so that texture noise can be canceled. The sizes of averaging kernels are specified as 1-by- $2^{\tilde{n}}$ where $\tilde{n} \in \{2, 3, 4, 5, 6\}$. The filtered intensities of the pyramid are then converted to shadow scales. For each RGB channel layer of each pyramid layer, the




Fig. 5. Pipeline of multi-scale shadow scale estimation. The aligned penumbra unwrap (a) is filtered using average kernels in exponentially increasing sizes to build a pyramid of shadow scales (b). The roughness of each column of each pyramid layer is measured and visualized in (c). Brighter colors indicate higher roughness. The horizontal and vertical dimensions in (c) refer to column index and layer index respectively. (d) is the visualized corresponding selections of layer index for each column (in white) after thresholding. (e) is the finale shadow scale composed by the shadow scales from the different layers of (b). (f) is the relit penumbra unwrap using (e).

estimated scales can be computed by dividing the intensities of each column by the intensity of the last element of each column (*i.e.* lit end). The optimum shadow scales for each column are selected from different layers of the pyramid. Column intensity with higher localness (*i.e.* filtered by a smaller kernel) and lower roughness are preferred. However, higher localness leads to higher roughness, so an optimum solution should balance these two properties. The roughness of intensity change $E_s(\tilde{c}, \tilde{n})$ (visualized in Fig.5(c)) is measured as follows:

$$E_s(\tilde{c}, \tilde{n}) = \frac{\partial^2 U(\tilde{r}, \tilde{c}, \tilde{n})}{\partial \tilde{r}^2} dr \quad (6)$$

where U is the penumbra unwrap, \tilde{n} is the layer index of pyramid, \tilde{c} and \tilde{r} are the column and row coordinates of the penumbra unwrap respectively. The optimum scales for each column are selected using a threshold of roughness T_s which is computed as the mean of all values in E_s . The column of one of the layers which has the lowest roughness above T_s is selected (visualized in Fig.5(d)). A shadow scale image of the penumbra unwrap (*e.g.* Fig.5(e)) can thus be formed by picking columns from different pyramid layers according to the selected layer index of each column. As the intensity samples, *i.e.* columns, have previously been aligned during the alignment of the unwrap, the estimated scales of each sampling are mapped back by using an reverse operation of Γ so that the estimated shadow scales are corresponding to the original unaligned intensities of the penumbra unwrap. The mapped-back shadow scales are then registered to their 2D positions in the image that a sparse shadow scale field is formed (*e.g.* Fig.2(g)). To obtain a dense scale field (*e.g.* Fig.2(h)), we propagate the sparse scales in the penumbra region by smoothly interpolating and extrapolating the scales in other regions using image in-painting [32]. The shadow-free image can be obtained by inverse scaling according to Eq.1.

3.D.Color Correction

Images captured from popular imaging devices are often post-processed, *e.g.* gamma correction and JPEG compression, such that the linearity of photon intensity is not maintained. When the degree of post-processing is high, visible artifacts, *e.g.* differences in tone and contrast, may appear in shadow corrected areas as Eq.1 does not hold. This is because these post-processing

steps apply non-linear operations which break the linearity property that the intensity of a pixel is proportional to the amount of photons a sensor has received. A robust multi-scale color correction method is therefore proposed to address this issue. The improvement will only be significant for images which are over post-processed. Previous work has proposed global adjustments to align the intensity characteristics of the umbra and lit area [5, 14]. These assume that the surface around the penumbra has a similar texture and color but may lead to significant unnatural artifacts when they are dissimilar. To address this, we adopt an image detail alignment similar to Xiao *et al.* [15] that equates the spatially dependent variance of RGB intensities between the shadow and lit sides at different scales. Unlike [15], our method does not require a shadow removal for each filtered image level (which is computationally expensive). Instead, based on an initial shadow removal result, we only iteratively align the variances around shadow boundary on each scale. It is assumed that the average intensity of both sides of the shadow are accurate and that artifacts are due to the differences in intensity variance. Statistics are collected from the lit side pixels P_l and the umbra side pixels P_u both near the penumbra as the reference and source of color correction respectively. The algorithm for alignment is described in Algorithm2. where s is a scale, β is the maximum image dimensional of I^r , $bfilter$ is an operation that bilaterally filters [33] the input image (first parameter) using a standard deviation of the space (second parameter) and a range Gaussian (third parameter), I^h is an image of intensity variation, where c is the channel index, ς is a function which computes the median absolute deviation.

Finally, to smooth the color correction result, alpha blending is applied in RGB color space according to the shadow scale as shown in Eq.7.

$$I_c^f = I_c^r \circ S_c + I_c^{ra} \circ (1 - S_c) \quad (7)$$

where c is the channel index, S_c is the normalized scale field of S , I_c^f is the final shadow-free image. An illustration of the intermediate steps of color correction is shown in Fig6.

3.E.Parameter Learning

Our shadow removal approach includes the use of some standard algorithms that require the specification of parameters. To determine an appropriate set of parameters, we apply an optimization to learn these parameters from a subset of our ground truth data set. These parameters contain both integers and real numbers. It is therefore not possible to apply a gradient-based optimization method that requires the objective and constraint functions to be both continuous and have continuous first derivatives. To determine these parameters, we apply a mixed-integer genetic optimization method [34]. Let us define $H = h_1, h_2, \dots, h_6$ as a vector of the shadow removal parameters denoted throughout the paper. Our objective function E_p minimizes the sum of all error measurements as the follows:

$$E_p(H) = \sum_k e_k w_k \quad (8)$$

where e_k is the k^{th} error measurement, and w_k is the weight for e_k . We assume that the weights for all error measurements are the same (*i.e.* equally important), *e.g.*, $w_k = 1$. These error measurements are later introduced in §4.C (in Table4) and only all-pixel-error is used in our learning. Table2 shows the details of these parameters and their optimization configuration. In our

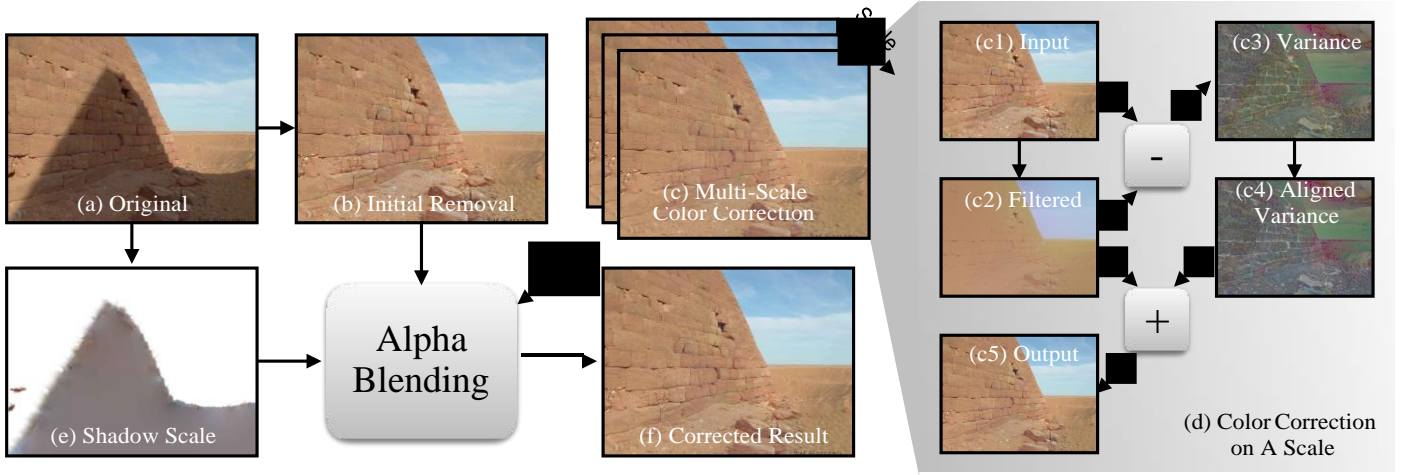


Fig. 6. Multi-scale color correction pipeline. The inconsistency in the initial shadow-free image (b) is fixed in the final output (f). The multi-scale color correction aligns the color variance at different scales from coarse to fine. On each single scale, the initial input image (c1) exhibits inconsistency of local variance between lit and shadow areas. The higher-frequency variation (c3) of shadow and lit areas are aligned in (c4). The corrected output (c5) can be obtained by adding (c4) to (c2).

Algorithm 2: Multi-Scale Color Correction

input : shadow removed image I^r , reference lit pixels P_l , source shadow pixels P_u , all shadow pixels P_s
output : color corrected image I_c^{ra}
 $I^{ra} \leftarrow I^r$;
for $s = 1$ **to** 3 **do** for each scale, e.g., Fig.6
 $I^l \leftarrow b f i l t e r (I^{ra}, \beta / 2^{s+1}, h_6)$ //apply bilateral filtering (e.g. Fig.6(c2))
 $I^h \leftarrow I^r - I^l$ //get local intensity variation image (e.g. Fig.6(c3))
for $c = 1$ **to** 3 **do** for each RGB channels
 $r_\sigma \leftarrow \zeta(I^h(P)) / \zeta(I^h(P_u))$ //get overall ratio of intensity variation
 $I_c^{re} \leftarrow r_\sigma I^h(P_s)$ //get aligned intensity variation image (e.g. Fig.6(c4))
 $I_c^{ra} \leftarrow I^r$ //copy intensities of lit pixels
 $I_c(P_s) \leftarrow I_c(P_s) + I_c^{re}$ //add aligned intensity variation back (e.g. Fig.6(c5))
end
end

Table 2. Parameter learning specification for the optimization.

ID	Description	Value Range	Initial Value	Type
h_1	Gaussian filter Kernel size (§3.A)	[2, 15]	5	Int.
h_2	medium filter Kernel size (§3.A)	[2, 15]	10	Int.
h_3	DBScan Radius (§3.B)	[0.01, 0.5]	0.2	Real
h_4	meanshift Radius (§3.B)	[0.01, 0.5]	0.06	Real
h_5	gradient advance scale (Alg.1)	[2, 20]	10	Real
h_6	bilateral filter sampling spatial (Alg.2)	[0.1, 0.3]	0.2	Real

experiment, 5 test cases are randomly selected for computing each error measurement.

The optimum parameters for all error measurements are learned as 14 10 0.1124 0.0333 8.5195 0.2228. Our evaluation results in §4 are reported using these learned optimum parameters. Our experiment demonstrates that it is pos-

sible to learn an optimum set of shadow removal parameters from a ground truth data set. However, the learned parameters are also dependent on the amount and quality of training data. We further discover the choice of parameters for different types of shadows. An individual parameter learning process is performed for each shadow category measurement, i.e., $E_p(H) = e_k$ where k indicates the k^{th} error measurement. Table 3 shows our parameter learning results for the individual shadow categories (error measurement). These individual parameter learning results show that the optimum parameters can vary depending on type of shadow. It is therefore practical to provide some predefined parameter sets for different shadow removal tasks or a single parameter set that balances the shadow removal performance for all types of shadows.

4. EVALUATION

In this section, we first describe experiments that highlight our algorithms behavior given variable user inputs. The quality of our new ground truth versus existing state-of-the-art ground truth is then quantitatively evaluated. Finally, our algorithm is evaluated versus other state-of-the-art shadow removal methods based on our new dataset.

Table 3. Parameter learning results for individual error measurement.

ID	Texture			Softness			Brokenness			Colorfulness			Other
	Weak	Medium	Strong	Weak	Medium	Strong	Weak	Medium	Strong	Weak	Medium	Strong	
h_1	13	14	15	13	12	15	14	15	3	9	14	14	14
h_2	3	7	12	3	15	15	7	3	10	12	15	15	6
h_3	0.0269	0.3644	0.4418	0.1124	0.0193	0.0304	0.1815	0.0319	0.2733	0.2317	0.0304	0.0352	0.2775
h_4	0.0680	0.0193	0.0149	0.1049	0.1359	0.017	0.0426	0.0291	0.0203	0.0330	0.3464	0.0382	0.0307
h_5	5.3380	6.1115	2.2695	11.1833	9.5072	18.2173	8.5195	8.2637	5.0299	19.6150	19.8436	10.4065	7.6485
h_6	0.2274	0.1582	0.1547	0.2496	0.2392	0.2199	0.2204	0.2367	0.1658	0.2457	0.1612	0.2228	0.2360

4.A.P performance Stability Given Different User Inputs

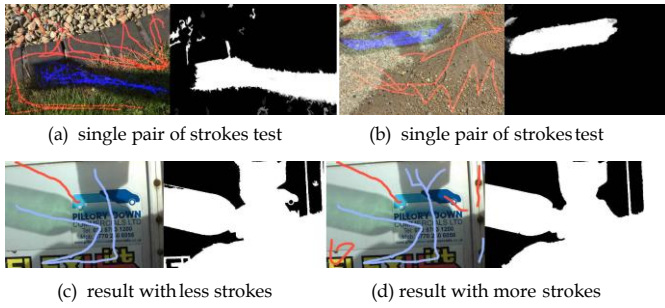


Fig. 7. Variable input behaviors: The top row shows two examples using single pairs of strokes. 10 examples of single strokes placed in different locations are supplied as input (red for lit and blue for shadow). The 2 gray-level images show the visualized probability of each pixel being marked in these 10 independent tests. Fewer gray pixels indicate higher stability, *i.e.* the image should only show black (0% probability) and white (100% probability) pixels when it is absolutely stable. The bottom row shows examples highlighting how additional strokes can improve the detection result (binary mask).

Given user-supplied single pairs of strokes of lit and shadow pixel samples, our shadow detection generates stable results in different conditions (*e.g.* Fig.7(a) and Fig.7(b)). In some cases, *e.g.* where the surface color is very shadow-like, the detection results can be improved by supplying more than one pair of strokes (*e.g.* Fig.7(c) and Fig.7(d)).

While our user input format is the same as the input format adopted in [15, 16], our method is found more robust to rougher (or fewer) scribbles since ours does not require an image matting process that is affected by sampled pixel location. Fig.8 shows an example.

4.B.Evaluation of Ground Truth Quality

Ideal pairs of ground truth images should have a minimum intensity difference in the common lit area - which will also indicate whether registration is poor (due to camera shake or scene movement - which should be rejected). This is utilized to assess the quality of ground truth candidates. The error image $\Delta I = I_s - I_g$ and the ratio image $I_r = \Phi(I_s) \oslash \Phi(I_g)$ are first computed, where I_s and I_g are the original shadow image and its shadow-free ground truth image (which differs from the processed shadow-free outputs I^{ϕ} in Eq.7) respectively, \oslash is element-wise division and Φ is a function that converts RGB image to gray-scale image. The set of pixels P_r of I_r that satisfies

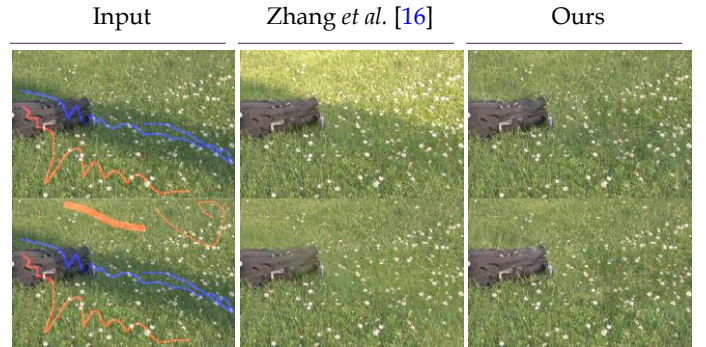


Fig. 8. Rougher stroke requirement: To generate a reasonable shadow removal result, our method requires less input strokes (in the same format) compared with [16].

$I_r(P_r) \geq 1$ are regarded as lit pixels. Due to some unavoidable minor global illumination changes and the inaccuracy in camera exposure control, the lit intensities in the shadow image can be higher than those in the shadow-free ground truth image. I_r can therefore be greater than 1. The ground truth error Q_d is computed as follows:

$$Q_d = \mu(|\Delta I(P_r)|) + \sigma(\Delta I(P_r)) \quad (9)$$

Ground truth pairs in our data set with $Q_d > 0.05$ are removed. Using this measure, our initial data capture of 195 test cases results in 186 test cases with stable illumination changes between the shadow and ground truth images. Comparing to the quality of other ground truth data sets, [2] (state of the art) results in mean error of 0.18 (leaving 28 out of 79 test cases) while ours is 0.02.

4.C.Quantitative evaluation of shadow removal

In previous work [2, 6], the quality of shadow removal is measured by directly using the per-pixel error between the shadow removal result and shadow-free ground truth. However, a shadow in a smaller size or a lighter shadow can result in a smaller initial error between the original shadow image and its shadow-free ground truth. It is thus unfair to judge that a method is better only because the error between the shadow-removed image and its shadow-free ground truth is smaller. In our work, we cancel out the affects of the size and darkness of the shadow. We therefore compute the error ratio E_r as our quality measurement:

$$E_r = E_n/E_o \quad (10)$$

where E_n is the error between the ground truth (no shadow) and shadow removal result, and E_o is the error between the ground

truth (no shadow) and the original shadow image. This normalized measure better reflects removal improvements towards the ground truth independent of original shadow intensity and size. We assess E_n and E_o using Root-Mean-Square-Error (RMSE) of RGB intensity. To test robustness, the standard deviation for each measurement is also computed. Unlike previous un-categorized test [2, 6], our removal test is based on our data set of 186 cases, which contains challenging soft, broken and color shadows and shadows cast on strong textured surfaces as well as simpler shadows, plus 28 remaining cases from [2] – resulting in 214 test cases in total. Each case is rated according to 4 attributes, which are *texture*, *brokenness*, *colorfulness* and *softness*, in 3 perceptual degrees from weak to strong which were aggregated by 5 users. In Table 4 (left side), the combined shadow removal error results from both automatic and semi-automatic shadow removal algorithms (all 214 cases) are shown. We separate the error measurement for all pixels and only shadow pixels. In our experiments, our method shows significant leading performance across all metrics. According to Eq. 10, E_o for all pixels is lower than E_o for the shadow pixels only because the intensity errors of the lit pixels are close to 0 and E_o measures the RMSE. In addition, the new RMSE E_n for both all pixels and shadow pixels only are very close after shadow removal, the error ratio E_r for the shadow pixels only is therefore generally lower.

To encourage open comparison in the community, we provide an online benchmark¹ for quantitative evaluation of shadow removal.

4.D. Analysis of Shadow Categories and Attributes

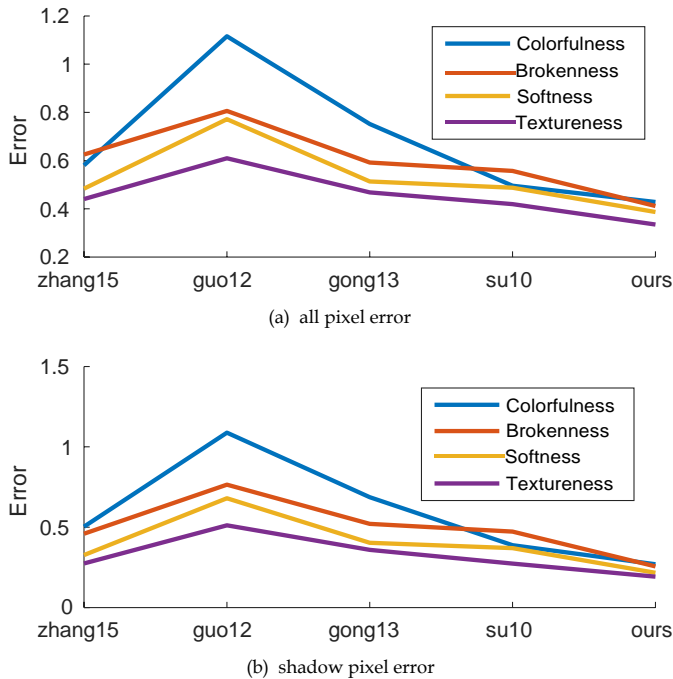


Fig. 9. Parallel coordinate charts of the quantitative results in Table 4. The ticks zhang15, guo12, gong13, and su10 refer to [16], [2], [14], and [4] accordingly. The scores presented here are the average scores of all three degrees for each attribute.

To investigate the affects of different shadow categories and attributes, the quantitative result in Table 4 is summarized by vi-

ualizing the result using the parallel coordinate charts in Fig. 9. Such a visualization is insightful as strong performance of one method could direct practitioners to favor one algorithm over another in some problem cases. Overall, colored shadows are shown to be significantly the most difficult shadows to remove and shadows cast on high texture the easiest challenge. Broken shadows are slightly more difficult to process than soft shadows, although both of them are in the range of medium difficulty. [2] and [14] show relatively significant disadvantages in processing color shadows, while [4] demonstrates obvious difficulty in processing broken shadows. The trend of the other methods and attributes are otherwise similar. In our tests, our method overall demonstrates the best performance for all types of shadows analyzed.

4.E. Visual Comparison

Fig. 10 shows some typical visual results of shadow removal on various scenarios from our data set.

4.F. Efficiency Comparison

Table 5 shows the required time for processing 0.3 mega-pixel color images shown in Figure 10 on a 3.1GHz machine. Our MATLAB implementation generally requires less system processing time than the other two MATLAB implementations of [2, 14], one MATLAB+C implementation [16] and one C++ implementation [4]. Compared with the other user-assisted methods [4, 14], our method also generally requires less time for user-interaction. The slower performance of [16] is majorly caused by its slow image matting pre-processing step.

4.G. Limitation And Future Work

As is the case with all current shadow removal methods, our algorithm has most difficulty in extreme cases, e.g. Fig. 11, where shadows are highly soft, broken, colorful (mixed by at least two different colors). Although the shadow effects can be significantly reduced by our method, the artifacts are still noticeable. A shadow image may also contain a mixture of more than one strong shadow attributes (e.g. the last column of Fig. 11). We have challenged the multi-category shadow removal for the first time in our community (with a leading performance), but have not resolved the extreme cases. This highlights a direction for future work.

Our method requires users to supply reasonable inputs. We have not considered its tolerance for very careless user inputs (e.g. mistakenly marking many shadow pixels as lit samples). Besides, insufficient user inputs may result in a poor shadow detection. Since a sufficiently trained shadow classifier may be robust to this issue, another future work could be improving shadow detection by combining user inputs with the shadow masks generated from an automatic shadow classifier.

5. CONCLUSIONS

We have presented an interactive method for fast shadow removal together with a state-of-the-art ground truth. Our method balances the complexity of user input with robust shadow removal performance. Our quantitatively-verified ground truth data set overcomes issues of mismatched illumination and registration in existing data sets. We have evaluated our method against several state-of-the-art methods using a thorough quantitative test and shown leading state of the art performance. Besides the opportunities for improving shadow removal quality for the categorized shadows in our dataset, the detection

¹ http://cs.bath.ac.uk/%7Ehmg299/shadow_eval/

Table 4. Shadow removal errors for test cases according to four attributes. The left and right sides of the table show the error scores where all pixels in the image are used, and just shadow area pixels respectively. For each score of each attribute, the images with other predominant attributes (strong) are not used. Hence, test cases have a strong single bias towards one of the attributes. “Other” refers to a set of shadow cases showing no markedly predominant attributes. “Mean” refers to the average score for each category. Standard derivations are shown in brackets. In our ordering, the average error is compared before comparing the standard derivation. Method [2] is trained using a large shadow detection data set from [19]. The user input for Method [16] is a combination of the simple input for our method and some additional strokes for accommodating the sensitive shadow detection of [16]. The best scores are made bold.

Degree	All Pixels Errors					Shadow Pixels Errors				
	Zhang [16]	Guo [2]	Gong [14]	Su [4]	Ours	Zhang [16]	Guo [2]	Gong [14]	Su [4]	Ours
Textureness										
Weak	0.35 (0.17)	0.53 (0.50)	0.32 (0.19)	0.35 (0.24)	0.26 (0.16)	0.16 (0.20)	0.42 (0.57)	0.16 (0.18)	0.17 (0.34)	0.10 (0.09)
Medium	0.39 (0.25)	0.59 (1.09)	0.38 (0.33)	0.36 (0.16)	0.26 (0.11)	0.28 (0.25)	0.47 (1.15)	0.27 (0.35)	0.21 (0.27)	0.12 (0.09)
Strong	0.58 (0.38)	0.71 (0.60)	0.70 (0.42)	0.60 (0.41)	0.49 (0.40)	0.39 (0.50)	0.64 (1.03)	0.65 (0.55)	0.49 (0.60)	0.36 (0.44)
Mean	0.44 (0.27)	0.61 (0.73)	0.47 (0.31)	0.43 (0.27)	0.34 (0.22)	0.27 (0.38)	0.51 (0.92)	0.36 (0.36)	0.29 (0.40)	0.19 (0.21)
Softness										
Weak	0.37 (0.24)	0.52 (1.08)	0.33 (0.31)	0.33 (0.21)	0.23 (0.10)	0.24 (0.42)	0.39 (1.13)	0.21 (0.32)	0.18 (0.33)	0.10 (0.09)
Medium	0.40 (0.20)	0.70 (0.36)	0.44 (0.21)	0.42 (0.11)	0.34 (0.15)	0.25 (0.26)	0.64 (0.43)	0.29 (0.25)	0.24 (0.14)	0.15 (0.10)
Strong	0.69 (0.49)	1.09 (0.75)	0.76 (0.35)	0.72 (0.21)	0.60 (0.27)	0.49 (0.62)	1.01 (0.97)	0.71 (0.53)	0.69 (0.54)	0.40 (0.25)
Mean	0.48 (0.31)	0.77 (0.73)	0.51 (0.29)	0.49 (0.18)	0.39 (0.18)	0.33 (0.43)	0.68 (0.84)	0.40 (0.37)	0.37 (0.34)	0.22 (0.15)
Brokenness										
Weak	0.37 (0.23)	0.59 (0.98)	0.36 (0.29)	0.34 (0.15)	0.25 (0.13)	0.24 (0.40)	0.48 (1.04)	0.23 (0.31)	0.18 (0.24)	0.11 (0.09)
Medium	0.43 (0.22)	0.42 (0.29)	0.44 (0.25)	0.49 (0.38)	0.29 (0.14)	0.27 (0.27)	0.27 (0.35)	0.29 (0.28)	0.38 (0.58)	0.14 (0.11)
Strong	1.07 (0.47)	1.42 (1.06)	0.98 (0.31)	0.85 (0.25)	0.69 (0.30)	0.88 (0.72)	1.55 (1.84)	1.05 (0.50)	0.86 (0.65)	0.52 (0.32)
Mean	0.63 (0.31)	0.81 (0.78)	0.59 (0.29)	0.56 (0.26)	0.41 (0.19)	0.46 (0.46)	0.76 (1.08)	0.52 (0.36)	0.47 (0.49)	0.26 (0.17)
Colorfulness										
Weak	0.36 (0.18)	0.48 (0.64)	0.32 (0.18)	0.34 (0.18)	0.24 (0.11)	0.21 (0.24)	0.36 (0.78)	0.19 (0.20)	0.18 (0.23)	0.10 (0.08)
Medium	0.60 (0.50)	1.67 (2.29)	0.83 (0.67)	0.52 (0.24)	0.48 (0.18)	0.57 (1.06)	1.56 (2.07)	0.67 (0.73)	0.45 (0.65)	0.24 (0.14)
Strong	0.78 (0.57)	1.20 (0.99)	1.10 (0.68)	0.63 (0.49)	0.56 (0.31)	0.72 (1.00)	1.34 (2.33)	1.20 (1.18)	0.54 (0.84)	0.46 (0.48)
Mean	0.58 (0.41)	1.12 (1.31)	0.75 (0.51)	0.50 (0.30)	0.43 (0.20)	0.50 (0.77)	1.09 (1.73)	0.69 (0.70)	0.39 (0.57)	0.27 (0.23)
Other										
	0.35 (0.16)	0.38 (0.52)	0.27 (0.17)	0.26 (0.08)	0.19 (0.06)	0.16 (0.22)	0.25 (0.58)	0.14 (0.20)	0.09 (0.05)	0.06 (0.02)

and removal for highly-complicated shadows, such as overlapping shadows caused multiple light sources with different light colors, and shadows caused by transparent objects with complicated inner structure and color, is still an open problem for the community.

ACKNOWLEDGMENT

This work was funded by EPSRC Engineering and Physical Sciences Research Council (EP/M023281/1 and EP/M001768/1).

REFERENCES

1. E. Arbel and H. Hel-Or, “Shadow removal using intensity surfaces and texture anchor points,” *Transaction on Pattern Analysis and Machine Intelligence* **33**, 1202–1216 (2011).
2. R. Guo, Q. Dai, and D. Hoiem, “Paired regions for shadow detection and removal,” *Transaction on Pattern Analysis and Machine Intelligence* **PP**, 1–1 (2012).
3. S. H. Khan, M. Bennamoun, F. Sohel, and R. Togneri, “Automatic shadow detection and removal from a single image,” *transactions on pattern analysis and machine intelligence* **38**, 431–446 (2016).
4. Y.-F. Su and H. H. Chen, “A three-stage approach to shadow field estimation from partial boundary information,” *Transaction on Image Processing* **19**, 2749–2760 (2010).
5. F. Liu and M. Gleicher, “Texture-consistent shadow removal,” in “European Conference on Computer Vision,” (Springer, 2008), pp. 437–450.
6. Y. Shor and D. Lischinski, “The shadow meets the mask: Pyramid-based shadow removal,” *Computer Graphics Forum* **27**, 577–586 (2008).
7. H. Barrow and J. Tenenbaum, “Recovering intrinsic scene characteristics,” *Computer Vision System, A Hanson & E. Riseman (Eds.)* pp. 3–26 (1978).
8. G. D. Finlayson, M. S. Drew, and C. Lu, “Entropy minimization for shadow removal,” *International Journal of Computer Vision* **85**, 35–57 (2009).
9. C. Fredembach and G. Finlayson, “Simple shadow removal,” in “International Conference on Pattern Recognition,” , vol. 1 (IEEE, 2006), vol. 1, pp. 832–835.
10. Q. Yang, K.-H. Tan, and N. Ahuja, “Shadow removal using bilateral filtering,” *Transaction on Image Processing* **21**, 4361–4368 (2012).
11. M. Drew, C. Lu, and G. Finlayson, “Removing shadows using

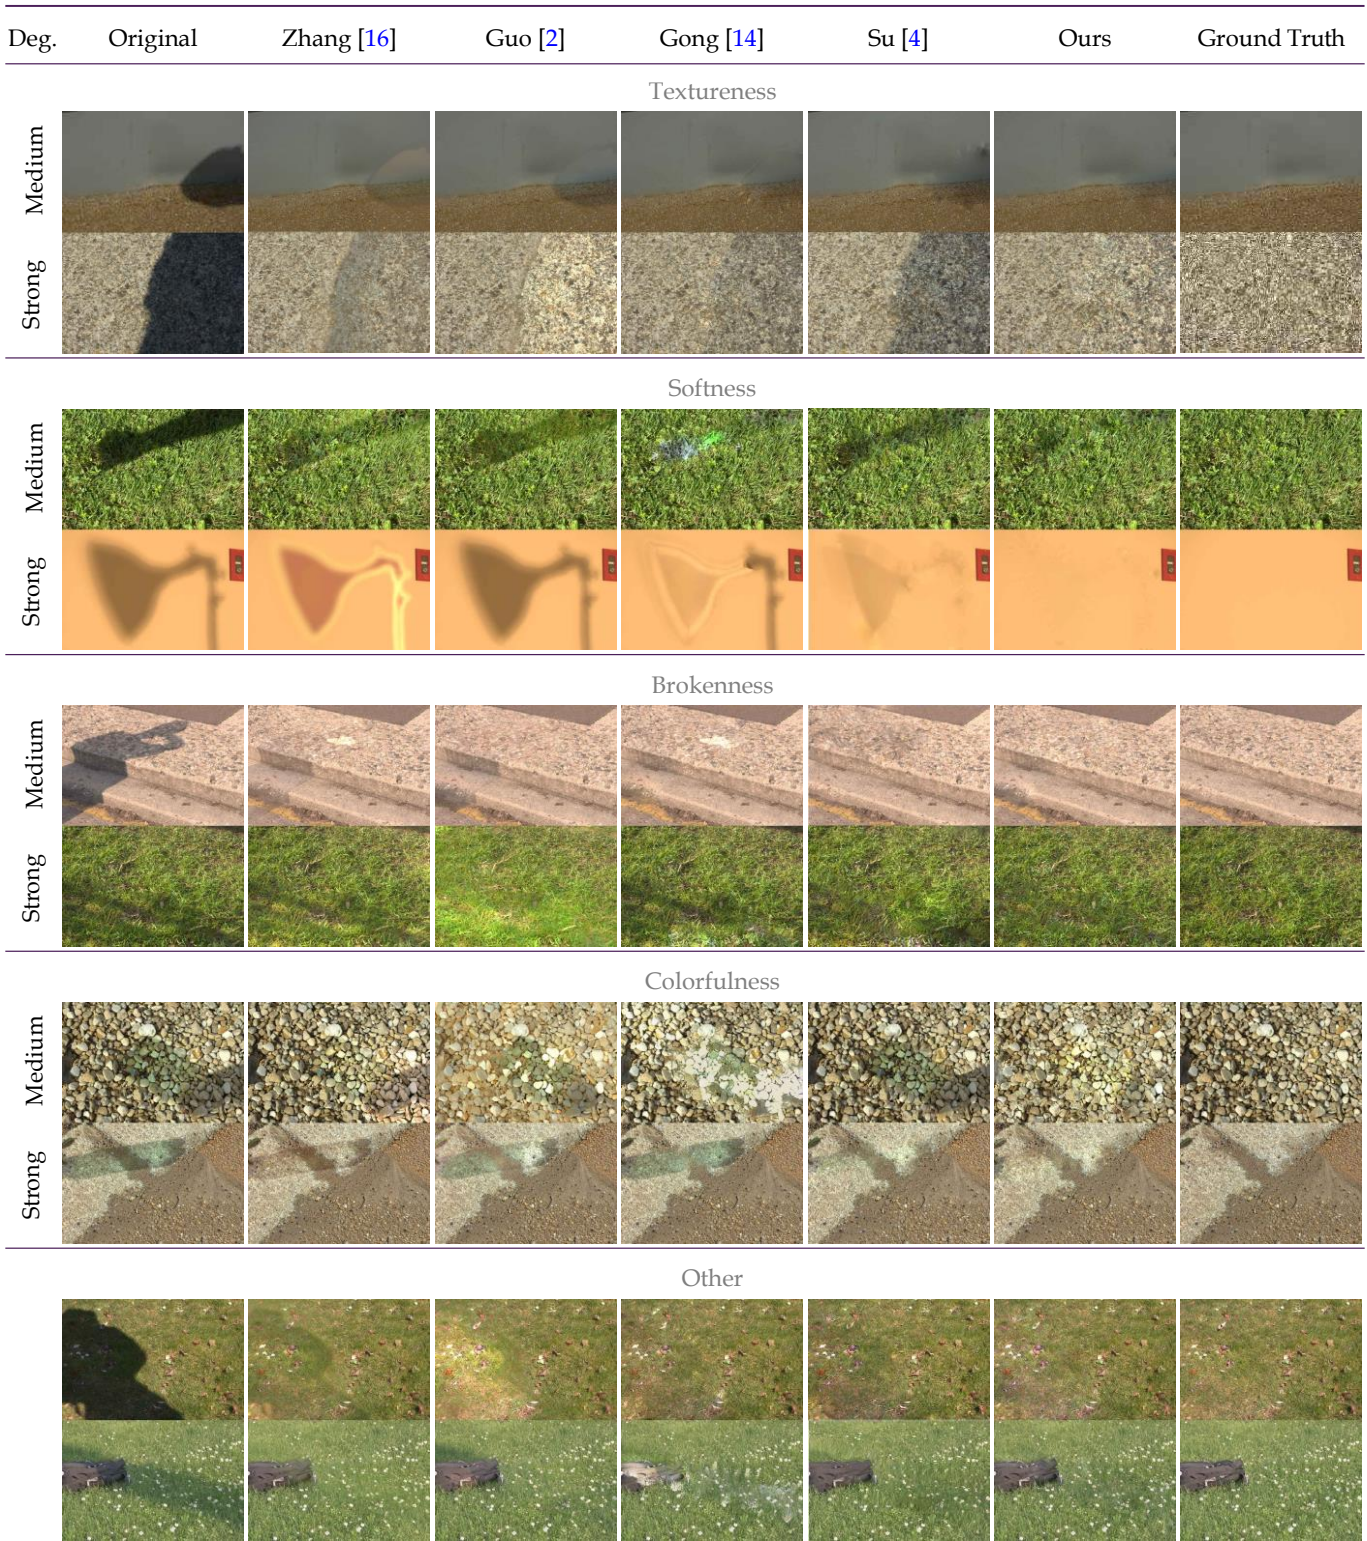


Fig. 10. Comparisons using images from our data set. The table shows our results given test cases with stronger degrees of the corresponding attribute except for "Other", which refers to cases where there is no predominantly strong attribute. Degree weak is not shown as weak means almost no attribute feature.

flash/noflash image edges," in "International Conference on Multimedia and Expo," (IEEE, 2006), pp. 257–260.
 12. N. Salamaty, A. Gernain, and S. Susstrunk, "Removing shadows from images using color and near-infrared," in "International Conference on

Image Processing," (IEEE, 2011), pp. 1713–1716.
 13. T.-P. Wu, C.-K. Tang, M. S. Brown, and H.-Y. Shum, "Natural shadow matting," *ACM Transactions on Graphics* **26**, 8 (2007).
 14. H. Gong, D. Cosker, C. Li, and M. Brown, "User-aided single image

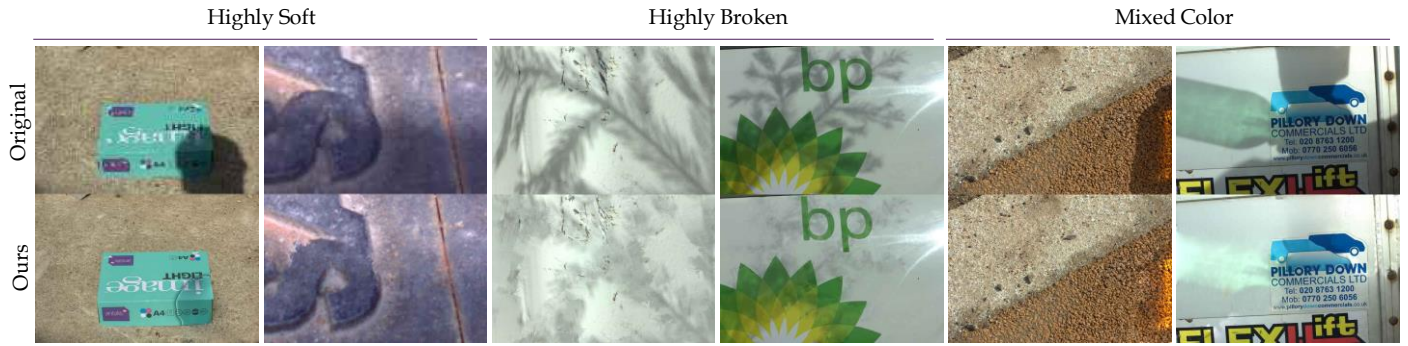


Fig. 11. Failure cases – where all the leading methods tested result in visible image artifacts or fail to remove the shadow. These images are also included in our shadow removal dataset published online.

Table 5. Time comparison of shadow removal: T_u , T_s refer to time (in seconds) for user interaction and time for system processing respectively.

Degree	Zhang [16] MATLAB+C		Guo [2] MATLAB		Gong [14] MATLAB		Su [4] C++		Ours MATLAB	
	T_u	T_s	T_u	T_s	T_u	T_s	T_u	T_s	T_u	T_s
Textureiness										
Medium	6	40	0	157	2	35	22	44	4	7
Strong	4	38	0	74	2	81	20	44	4	10
Softness										
Medium	6	35	0	159	4	267	18	44	4	13
Strong	7	33	0	23	9	31	109	43	3	21
Brokenness										
Medium	6	35	0	61	8	88	37	43	4	14
Strong	8	35	0	158	10	241	61	51	8	54
Colorfulness										
Medium	6	34	0	160	3	385	30	45	5	51
Strong	7	35	0	56	2	132	19	46	5	11
Other										
	6	35	0	74	4	301	33	42	4	17
	6	34	0	80	2	114	28	42	4	13

shadow removal," in "International Conference on Multimedia and Expo," (IEEE, 2013), pp. 1–6.

15. C. Xiao, R. She, D. Xiao, and K.-L. Ma, "Fast shadow removal using adaptive multi-scale illumination transfer," in "Computer Graphics Forum," , vol. 32 (Wiley Online Library, 2013), vol. 32, pp. 207–218.

16. L. Zhang, Q. Zhang, and C. Xiao, "Shadow remover: Image shadow removal based on illumination recovering optimization," Transactions on Image Processing **24**, 4623–4636 (2015).

17. C. Fredembach and G. Finlayson, "Hamiltonian path-based shadow removal," in "British Machine Vision Conference," , vol. 2 (BMVA, 2005), vol. 2, pp. 502–511.

18. J.-B. Huang and C.-S. Chen, "Moving cast shadow detection using physics-based features," in "Conference on Computer Vision and Pattern Analysis," (IEEE, 2009), pp. 2310–2317.

19. J. Zhu, K. G. G. Samuel, S. Masood, and M. F. Tappen, "Learning to recognize shadows in monochromatic natural images," in "Conference on Computer Vision and Pattern Analysis," (IEEE, 2010).

20. J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan, "Detecting ground shadows in outdoor consumer photographs," in "European Conference

on Computer Vision," (Springer, 2010), pp. 322–335.

21. T. F. Y. Vicente, C.-P. Yu, and D. Samaras, "Single image shadow detection using multiple cues in a supermodular mrf," in "British Machine Vision Conference," (BMVA, 2013).

22. X. Huang, G. Hua, J. Tumblin, and L. Williams, "What characterizes a shadow boundary under the sun and sky?" in "International Conference on Computer Vision," (IEEE, 2011), pp. 898–905.

23. L. Shen, T. Wee Chua, and K. Leman, "Shadow optimization from structured deep edge detection," in "Conference on Computer Vision and Pattern Recognition," (IEEE, 2015), pp. 2067–2074.

24. G. Finlayson, C. Fredembach, and M. Drew, "Detecting illumination in images," in "International Conference on Computer Vision," (IEEE, 2007), pp. 1–8.

25. T. M. Mitchell, "Machine learning. wcb," (1997).

26. G. D. Finlayson, S. D. Hordley, and M. S. Drew, "Removing shadows from images using retinex," in "Color Imaging Conference," (2002), pp. 73–79.

27. G. D. Finlayson, S. D. Hordley, C. Lu, and M. S. Drew, "On the removal of shadows from images," Transactions on Pattern Analysis and Machine Intelligence **28**, 59–68 (2006).

28. G. D. Finlayson, S. D. Hordley, and M. S. Drew, "Removing shadows from images," in "European Conference on Computer Vision," (Springer, 2002), pp. 823–836.

29. M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise." in "KDD," , vol. 96 (ACM, 1996), vol. 96, pp. 226–231.

30. D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," Transaction on Pattern Analysis and Machine Intelligence **24**, 603–619 (2002).

31. J. Nocedal and S. Wright, "Numerical optimization," (Springer, 2006), Springer series in operations research, chap. 18, 2nd ed.

32. M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," (ACM, 2000), pp. 417–424.

33. S. Paris and F. Durand, "A fast approximation of the bilateral filter using a signal processing approach," International Journal of Computer Vision **81**, 24–52 (2009).

34. K. Deep, K. P. Singh, M. Kansal, and C. Mohan, "A real coded genetic algorithm for solving integer and mixed integer optimization problems," Applied Mathematics and Computation **212**, 505–518 (2009).