

Removal of Image Artifacts Due to Sensor Dust

Changyin Zhou*
Fudan University

Stephen Lin
Microsoft Research Asia

Abstract

Image artifacts that result from sensor dust are a common but annoying problem for many photographers. To reduce the appearance of dust in an image, we first formulate a model of artifact formation due to sensor dust. With this artifact formation model, we make use of contextual information in the image and a color consistency constraint on dust to remove these artifacts. When multiple images are available from the same camera, even under different camera settings, this approach can also be used to reliably detect dust regions on the sensor.

In contrast to image inpainting or other hole-filling methods, the proposed technique utilizes image information within a dust region to guide the use of contextual data. Joint use of these multiple cues leads to image recovery results that are not only visually pleasing, but also faithful to the actual scene. The effectiveness of this method is demonstrated in experiments with various cameras.

1. Introduction

Sensor dust is a common problem in digital photography, especially for digital single-lens reflex cameras (DSLRs) with interchangeable lenses. During lens exchange, the sensor becomes exposed to the environment, and dust may enter the camera before another lens is mounted. This accumulated dust blocks incoming light before it reaches the sensor, and appears as dark spots in images. Not only are these dust artifacts visually objectionable to photographers, but their alterations of scene appearance may impair computer vision algorithms.

Various tools have been designed for physically removing dust from a sensor, including cleaning liquids, brushes, vacuums, and dust adhesives. However, sensor cleaning remains a difficult task even for a professional. Most camera manufacturers explicitly warn users not to touch the sensor even for the purpose of removing dust.

Rather than eliminate dust by physical means, photographers often prefer to process images using photo software

to remove the appearance of sensor dust. This approach becomes necessary when dust particles are not noticed until viewing the pictures. Dust removal in software using, for example, a cloning tool requires skillful image manipulation to obtain satisfactory results. While manual processing of images can be effective when the background of a dust region is uniform or highly regular, it is often the case that a dust background does not have such regularity. An example of this is shown in Fig. 1(a,b). Although the dust artifacts are visually obvious, they are challenging to remove by hand, since their irregular backgrounds are difficult to reproduce by cloning.


To automatically recover a dust-free image, a cloning-style process may be simulated to fill in dust regions. Previous techniques along this direction include image inpainting and texture synthesis. As with manual cloning, these approaches have difficulties with irregular dust backgrounds, due to insufficient contextual information.

In contrast to the hole-filling applications addressed by these previous methods, dust regions generally contain some information about the obscured background area. As shown in Fig. 1, despite the attenuation of scene radiance by dust, the veins and surface texture on the pictured leaf are still partially observable. This source of information provides an important constraint on dust artifact removal that enables recovery of irregular backgrounds common in natural imagery.

To better utilize the information inside a dust region, we study and model the formation of image artifacts by sensor dust. Since sensor dust actually accumulates on a low-pass filter that rests upon the sensor array, the attenuation of light as seen from the sensor is related not only to the dust itself, but also to the imaging geometry specified by camera settings. The presented dust imaging model defines a function space of image artifacts with respect to camera settings, which together with the information from the dust region implicitly describes a space of possible solutions for the dust background. Further constraints based on contextual data in the image and color consistency of dust lead to image recovery results that are true to the scene.

Given a single image, there exists ambiguity on whether a darkened region indicates a dust artifact or is actually a

*This work was done while Changyin Zhou was a visiting student at Microsoft Research Asia.



(a) Image Thumbnail (b) Dust Regions (c) Texture Synthesis (d) Our Results (e) Ground Truth

Figure 1. A leaf photographed with a Canon 30D camera. In (b), it can be seen that some background information is available inside the dust regions. By utilizing this information in conjunction with the artifact formation model, our technique produces results more similar to the ground truth than texture synthesis. The ground truth was obtained by slightly shifting the camera to displace the locations of artifacts in the image.

part of the scene. Our method provides an ordered list of candidate dust regions in an image, based on adherence to image context and the dust artifact model, and then a user can specify which regions on this list actually contain dust. In many instances, a photographer will have a collection of images taken from the same camera at approximately the same time, possibly with various camera settings. Since the configuration of dust on the sensor can be considered fixed in such an image set, the proposed method can more reliably detect dust artifacts based on their consistency among the images with respect to camera settings.

2. Related Work

The image appearance of dust that lies on a transparent lens cover was examined in [13]. While such protective covers may be important in certain scenarios, e.g., on the 2003 Mars Exploration Rovers in [13], digital cameras generally do not have them. If dust were to fall directly on a camera lens, the image effects would differ significantly from that of sensor dust, since the scene radiance would be uniformly attenuated over the entire sensor array. This is due to the fact that each sensor element receives light that is focused through the lens, and therefore the resultant artifact is not visually apparent. In [13], the problems of dust artifact removal and detection are not addressed.

Image completion and image inpainting methods have been used to recover missing portions of images. These techniques assume that there is no significant information within the regions to be filled, and exploit contextual information in the surrounding image. Bertalmio et al. [1] filled in missing regions by propagating image Laplacians from the region boundaries along isophote directions. Many techniques have since been proposed to take greater advan-

tage of contextual information. To handle curved structures, Chan and Shen [2] introduced the use of Euler’s elastica. Levin et al. [9] employed an image-specific prior to guide inpainting. Many example-based approaches [5][10][4] have also been proposed for image completion by interpolating pixels using texture synthesis techniques.

Although these previous approaches may generate visually plausible results, they do not consider available information inside the regions, and therefore may produce solutions that do not conform with the actual scene. This is exemplified in Fig. 1, where the texture synthesis result in (c) appears unnatural and different from the ground truth in (d). Some image information within a dust region may be preserved by band-limited texture synthesis [6], but the low-pass effects of dust vary over a dust region and would thus complicate such an approach. To better guide image completion, user specification of region structure has been utilized in [11][12]. In our work, this structural information is obtained directly from the dust region, and interpreted using the presented artifact formation model. Unlike the previous techniques, our work also can infer the locations of dust regions, based on both image context and the artifact formation model.

3. Formation of Dust Artifacts

To better utilize the image information within a dust region, an understanding of dust artifact formation is needed. We assume dust particles to be light attenuators that alter an image according to $I(x) = U(x) \cdot \alpha(x)$, where $I(x)$ is the recorded image intensity at pixel x , U represents the dust-free image that we wish to recover, and α models the attenuation due to sensor dust. Although dust artifacts may partly result from light diffraction, we consider this little ef-




Figure 2. The artifact produced by a given dust appears differently under different camera settings. From left to right, the F -number is set to $F/22$, $F/19$, $F/16$ and $F/11$, respectively. These artifact functions α are estimated from a uniform scene and contain image noise.

fect to be negligible, since dust regions do not exhibit interference effects associated with diffraction around the edges of a barrier.

Since the dust lies on a low-pass filter instead of directly on the sensor, the artifact function α represents a projection of the dust’s transparency function β . The geometry of this projection is dependent upon camera settings, which can substantially affect α for a given β as shown in Fig. 2. We note that light that passes through a dust particle may undergo some amount of scattering that is not modelled by a transparency function. However, we have found such scattering effects to be small enough to disregard.

3.1. Projection of Dust Effects

The projection of dust transparency from the low-pass filter to the camera sensor is illustrated in Fig. 3. The dust transparency function $\beta : L^2 \rightarrow [0, 1]$ manifests itself as an artifact function $\alpha : L^2 \rightarrow [0, 1]$ in the image. Essentially, $\alpha(x)$ denotes the proportion of scene radiance that passes through the dust and reaches pixel x .

It can be assumed that the camera focal length f is much larger than the offset d of the low-pass filter from the sensor, and that the dust size is much smaller than the aperture size A . From the projection geometry, the effect of an arbitrary point on the dust will be spread over the sensor according to the rect function $\Pi(\cdot)$, defined as

$$\Pi_\tau(x) = \begin{cases} 0 & \text{if } |x/\tau| > 1/2 \\ 1/2 & \text{if } |x/\tau| = 1/2 \\ 1 & \text{if } |x/\tau| < 1/2. \end{cases} \quad (1)$$

Based on triangle similarity as seen in Fig. 3, the scale τ of the rect pulse can be calculated as $\tau = d \cdot A/f$. The value of f/A is equal to the camera’s F -number setting, which is conventionally recorded in the EXIF data of an image file.

The projection from β to α can then be formulated as the following convolution operation:

$$\alpha = \beta * \Pi_\tau.$$

We note that since sensor dust lies on a low-pass filter, blurring from the filter will also have some slight effect on the formation of α . As discussed in the Appendix, this effect is small and will not be included in our artifact formation model.




Figure 3. Projection geometry of dust effects.

3.2. Function Space of α

Because of dust projection, a function space of possible dust artifacts α can be defined with respect to the projection geometry factor τ and the bounds of dust transparency β . As illustrated in Fig. 2, larger values of τ lead to dust artifacts that are essentially more blurred, because of a broader convolution kernel in Eq. 1. Since $\beta : L^2 \rightarrow [0, 1]$, the function space \mathcal{D}_τ of valid artifacts α with respect to τ can be expressed as

$$\mathcal{D}_\tau = \{\alpha : \beta * \Pi_\tau = \alpha, 0 \leq \beta \leq 1\}. \quad (2)$$

Together with the image information within a dust region, this function space implicitly defines a space of possible dust backgrounds, which serves as a constraint on the dust removal solution.

4. Dust Removal in a Single Image

For a given dust region $D \in \Omega$, the goal of dust removal is to recover its dust-free appearance U_D . According to our artifact formation model, U_D must satisfy

$$I_D = U_D \cdot (\beta * \Pi_\tau) + n$$

provided that image noise n is additive. The dust transparency function β and projection geometry factor τ are both unknown, since the low-pass filter displacement d is information generally not provided with a camera. As a result, the dust removal problem is ill-posed. To better condition the image recovery process, we utilize contextual in-

formation $I_{\Omega/D}$ from outside the dust region in conjunction with the artifact formation model.

4.1. Formulation

We formulate this recovery as the following MAP estimation problem¹:

$$\begin{aligned}\hat{U}_D &= \arg \max P(U_D, \alpha | I) \\ &= \arg \max P(I_D | U_D, \alpha) \cdot P(U_D | I_{\Omega/D}) \cdot P(\alpha).\end{aligned}$$

Expressed as a logarithmic energy function, the problem becomes the minimization of

$$E(U_D, \alpha | I) = \lambda_1 E_1(I_D | U_D, \alpha) + \lambda_2 E_2(U_D | I_{\Omega/D}) + \lambda_3 E_3(\alpha). \quad (3)$$

The energy term E_1 uses the artifact formation model to capitalize on the image information within the dust region. E_1 is defined as the L_2 distance between I_D and $U_D \cdot \alpha$:

$$E_1(I_D | U_D, \alpha) = \frac{1}{2} \|I_D - U_D \cdot \alpha\|^2. \quad (4)$$

As explained in [3], the coefficient λ_1 should ideally be related to the noise level of n , and should be large if the signal-to-noise ratio is large.

The term E_2 utilizes contextual constraints from the surrounding image. The recovered region U_D is expected to have local appearance features consistent with those in neighboring areas, and image transitions should be smooth at the region boundaries. Similar to [7], we define this energy to be

$$E_2(U_D | I_{\Omega/D}) = \sum_{D_i \in D} \text{dist}(U_{D_i}, I_{\Omega/D})^2, \quad (5)$$

where D_i is a set of overlapped sub-blocks in D , and

$$\text{dist}(U_{D_i}, I_{\Omega/D}) = \frac{1}{2} \|U_{D_i} - I_{B_i}\|^2$$

with $B_i = \arg \min_{B \in \Omega/D} \|U_{D_i} - I_B\|^2$.

E_3 represents the constraint from the function space of α . For an estimated α , we define $E_3(\alpha)$ as the minimal L_2 distance to the function space \mathcal{D}_τ :

$$E_3(\alpha) = \frac{1}{2} \|\alpha - \mathcal{P}_{\mathcal{D}_\tau}(\alpha)\|^2, \quad (6)$$

where $\mathcal{P}_{\mathcal{D}_\tau}(\alpha)$ is the projection of α onto the space \mathcal{D}_τ :

$$\mathcal{P}_{\mathcal{D}_\tau}(\alpha) = \arg \min_{u \in \mathcal{D}_\tau} \|\alpha - u\|.$$

¹ $\arg \max P(U_D, \alpha | I)$
 $= \arg \max P(U_D, \alpha, I_D, I_{\Omega/D})$
 $= \arg \max P(I_D | U_D, \alpha, I_{\Omega/D}) \cdot P(U_D, \alpha | I_{\Omega/D}) \cdot P(I_{\Omega/D})$
 $= \arg \max P(I_D | U_D, \alpha, I_{\Omega/D}) \cdot P(U_D, \alpha | I_{\Omega/D}).$

Since α is independent of U_D and $I_{\Omega/D}$, and I_D is conditionally independent of $I_{\Omega/D}$ given U_D and α , we have $\arg \max P(I_D | U_D, \alpha) \cdot P(U_D | I_{\Omega/D}) \cdot P(\alpha)$.

Computing $\mathcal{P}_{\mathcal{D}_\tau}(\alpha) \in \mathcal{D}_\tau$ is equivalent to finding a β^* , ($0 < \beta^* < 1$) that minimizes $F = \frac{1}{2} \|\alpha - \beta^* \bar{\Gamma}_\tau\|^2$, and can be solved using the partial derivative $\frac{\partial F}{\partial \beta} = (\alpha - \beta^* \bar{\Gamma}_\tau) * \bar{\Gamma}_\tau$, where $\bar{\Gamma}_\tau(x, y)$ is defined as $\Gamma_\tau(-x, -y)$. Determination of the value of τ used here will be explained in Sec. 4.4.

Intuitively, recovery of a dust-free image using the energy function in Eq. 3 may be viewed as a constrained texture synthesis or inpainting. The term E_2 could be any texture synthesis or inpainting objective function to propagate contextual information from the surrounding image. In this paper, we use patch-based texture synthesis [4] in conjunction with the global optimization procedure of [7]. This synthesis is regulated by physical knowledge on artifact formation (E_3) together with observed image information within the dust region (E_1). These cues jointly suggest a recovery solution that tends toward the ground truth.

4.2. Optimization

In the optimization of Eq. 3, estimation of α relies on a good prediction of U_D , whereas U_D depends on an accurate function α . For this chicken-and-egg problem, we use the alternating minimization approach [3] to optimize $E(U_D, \alpha | I)$.

We start from an initialization $U_D^{(0)}$ obtained by applying conventional texture synthesis on a set of candidate dust regions, which are determined simply by fitting 2D Gaussians on small overlapping image patches and evaluating their concavity. An alternating sequence of conditional minimizers is then computed:

$$u_D^{(0)} \rightarrow \alpha^{(0)} \rightarrow u_D^{(1)} \rightarrow \alpha^{(1)} \rightarrow \dots \rightarrow u_D^{(n)} \rightarrow \alpha^{(n)} \rightarrow \dots$$

according to

$$\begin{aligned}\alpha^{(n)} &= \arg \min E(\alpha | I, U_D^{(n)}) \\ U_D^{(n+1)} &= \arg \min E(U | I, \alpha^{(n)}),\end{aligned}$$

where the partial derivatives of E can be computed from Eqs. 3, 4, 5, and 6 as

$$\begin{aligned}\partial E / \partial U_D &= \lambda_1 \alpha' (U_D \cdot \alpha - I_D) + \lambda_2 \sum_{D_i \in D} (U_{D_i} - I_{B_i}) \\ \partial E / \partial \alpha &= \lambda_1 U_D' (U_D \cdot \alpha - I_D) + \lambda_3 (\alpha - \mathcal{P}_{\mathcal{D}_\tau}(\alpha)).\end{aligned}$$

The sequence is iteratively computed until $\|\alpha^{(n)} - \alpha^{(n+1)}\| < \varepsilon$.

4.3. Color Consistency

The preceding algorithm may be applied independently for each color channel of an image; however, certain correlations between channels may be exploited as an additional constraint for radiometrically calibrated cameras. Since a dust particle is typically composed of a single material, we

can generally assume that it is monochromatic, i.e., of a uniform color. After normalization, we expect α to be proportional in the three channels, such that $\alpha_r(x)/\sum_x \alpha_r(x) = \alpha_g(x)/\sum_x \alpha_g(x) = \alpha_b(x)/\sum_x \alpha_b(x)$. The normalized α functions in the three channels should then have minimal Kullback-Leibler (K-L) divergence [8] to an optimized α^* :

$$\alpha^* = \arg \min_{\sum_{\alpha(x)=1} \sum_{i=r,g,b} KL(\alpha || (\alpha_i / \sum_x \alpha_i))}, \quad (7)$$

where $KL(\cdot || \cdot)$ is the K-L divergence:

$$KL(P || Q) = \sum_x P(x) \log \frac{P(x)}{Q(x)}.$$

From Eq. 7, we can define an energy function that represents a color consistency constraint on dust:

$$E_4(\alpha_r, \alpha_g, \alpha_b) = \frac{1}{2} \sum_{c=r,g,b} \|\alpha_c - \alpha^* \cdot \sum_x \alpha_c\|^2.$$

This energy term can be integrated into Eq. 3 to obtain $E = \lambda_1 E_1 + \lambda_2 E_2 + \lambda_3 E_3 + \lambda_4 E_4$. Within the alternating minimization algorithm, the partial derivative of E with respect to α for each channel $c = r, g, b$ is given by

$$\begin{aligned} \partial E / \partial \alpha_c &= \lambda_1 \cdot \sum U'_D(U_D \cdot \alpha_c - I_D) \\ &+ \lambda_3 \cdot (\alpha_c - \mathcal{P}_{\mathcal{D}_\tau}(\alpha_c)) \\ &+ \lambda_4 \cdot (\alpha_c - \alpha^* \cdot \sum_x \alpha_c). \end{aligned}$$

4.4. Estimation of τ

As mentioned previously, the value of τ is generally unknown, because the low-pass filter displacement d from the sensor is typically unavailable. Since τ defines the function space \mathcal{D}_τ of α , it affects the energy term E_3 defined in Eq. 6.

The volume of space \mathcal{D}_τ is negatively related to the value of τ . As seen from Eq. 2, when τ approaches zero, \mathcal{D}_τ will accommodate any function $0 \leq \alpha \leq 1$; as τ increases, the space of valid artifact functions becomes constricted. Consequently, values of τ that are greater than the actual value τ^* will have a corresponding space \mathcal{D}_τ that may not admit true artifact functions in an image. As a result, the optimization energy E computed for $\tau > \tau^*$ tends to increase sharply with respect to τ . On the other hand, a space \mathcal{D}_τ for $\tau \leq \tau^*$ will admit valid artifacts, and hence have optimization energies that decrease less rapidly with respect to τ . Some decrease in energy still exists because a broader range of artifact functions allows more flexible use of contextual data.

Based on this characteristic, τ may be estimated by examining the optimization energy E computed for sampled values of τ , as shown in the leftmost graph of Fig. 5. Since E changes rapidly for $\tau > \tau^*$ and slowly for $\tau \leq \tau^*$, τ can be estimated as $\arg \max_\tau \partial^2 E / \partial \tau^2$. In practice, we




Figure 4. Candidate dust regions can be detected based on optimization energy. The only true dust regions are $E=0.07$ in (a), and $E=0.06$ in (b). The other candidates in (b) are part of the image texture, and have less color consistency in their estimated artifact functions.

have found that small discrepancies between τ and τ^* do not lead to appreciable differences in results, likely because valid artifacts still lie in \mathcal{D}_τ for an underestimated τ , and dusts without sharp edges are still admissible for a slightly overestimated τ .

We note that d is a fixed value for a given camera model. If this distance has been accurately recovered for one camera, e.g., using multiple images as will be explained in Sec. 5.2, then other users of this camera model can utilize this value to determine τ as $d \cdot A/f$, where f/A is the F-number setting.

4.5. Detection of Candidate Dust Regions

When τ is set, the proposed method can determine optimization energy values for regions in an image. Since lower energies indicate a recovery result that is more consistent with image context and the dust artifact model, their corresponding regions are more likely candidates to contain dust. Some results of candidate detection are shown in Fig. 4. Dust artifacts can often be distinguished from dark scene spots because of dust projection characteristics and color consistency. A list of candidate regions ordered by increasing energy could be presented to the user to expedite dust detection.

4.6. Results

The cropped photograph in Fig. 1 was taken with a Canon 30D DSLR camera. The energy function coefficients were set to $\lambda_1 = 0.9, \lambda_2 = 0.4, \lambda_3 = 0.9, \lambda_4 = 0.9$ for all experiments in this paper, and the value of τ was determined to be 10 in this case. From the ordered list of candidate dust regions output by our algorithm, the top two were selected for image recovery. We compare our technique with the texture synthesis method in [4] and ground truth obtained by slightly shifting the camera. Because of subpixel shifts in the ground truth image, we present a visual rather than a quantitative comparison of results. Without consideration




Figure 5. Optimization energy E_d vs. low-pass filter displacement d for twelve dust regions in a Kodak DCS 330. The red curves represent the means among the twelve dusts. From left to right: with a single input image at F/11; with two images at F/11 and F/22; with four images at F/11, F/16, F/19, F/22. For a single image, d can be estimated from second derivatives, and for multiple images, d can be determined from the minimum energy.

of image information within the dust region, texture synthesis generates leaf structures and texture details that are different from the ground truth. Moreover, some leaf details appear slightly blurred due to texture irregularity and patch blending. Our method preserves most of the original leaf appearance, with less susceptibility to blurring because of the structural constraints provided by the dust region and artifact formation model.

Fig. 7 shows an example from a Kodak DCS 330 camera. τ is determined to be 4. In the evaluation of dust candidates, 17 (indicated by green arrows) of the 20 actual dust regions in the image appear among the top 18 candidates in the ordered list. The actual dust regions are determined manually from a photo of a uniform scene. The three overlooked dust regions are marked by red arrows, and the one incorrect candidate of the 18 is labelled with a yellow arrow. Although the dust detection method is not fully accurate, it nevertheless provides a helpful tool to the user.

Because of the large number of dust regions, the ground truth was captured using a different camera that is free of sensor dust. The scene appears slightly different with the other camera, but it provides an idea of the actual dust-free appearance. In comparison to conventional texture synthesis, our method better recovers the true scene appearance within the dust regions.

5. Multiple-Image Input

Having multiple input images, even with different camera settings, can improve detection and removal of dust artifacts. The dust transparency function β can be considered fixed among images captured within a single session, and this additional constraint enables estimation of β as well as a more reliable determination of d .

5.1. Formulation

Suppose we have a set of images $I_i, i = 1, 2, \dots, m$ and their corresponding F-numbers, $\rho_1, \rho_2, \dots, \rho_m$. Similar to the formulation of color consistency in dust, α functions in




Figure 6. Two β functions derived from multiple-image input.

these images should closely correspond to a common β :

$$\beta = \arg \min_{\beta} \sum_i \|\beta * \nabla_{\tau_i} - \alpha_i\|^2, \quad (8)$$

where $\tau_i = d/\rho_i$, and d is obtained as described in Sec. 5.2. From Eq. (8), we can define an energy term that represents consistency of β among multiple images:

$$E_5 = E(\alpha_1, \alpha_2, \dots, \alpha_n, d) = \frac{1}{2} \sum_i \|\beta * \nabla_{\tau_i} - \alpha_i\|^2, \quad (9)$$

which has the derivative

$$\partial E_5 / \partial \alpha_i = \beta * \nabla_{\tau_i} - \alpha_i.$$

Thus, the overall energy function for multiple-image input is

$$E^{multi} = \sum_{i=1}^m E_d^i + \lambda_5 \cdot E_5, \quad (10)$$

where E_d^i denotes the single-image energy for each I_i .

5.2. Estimation of d and β

With a larger number of input images, a more reliable estimate of d can be obtained. Among m images, a set of α functions $\alpha_1, \dots, \alpha_m$ is computed for each dust region. Since these α functions arise from the same dust, they should satisfy m equations $\alpha_i = \beta * \nabla_{d/\rho_i}, i = 1, \dots, m$. The correct value of d should give the same β when deconvolving the functions α_i by ∇_{d/ρ_i} ; otherwise, the inconsistency in β will lead to a high energy in Eq. 9.

Greater reliability in estimating d is exemplified in Fig. 5. For sampled values of d , we denote the optimization energy as E_d , and plot E_d with respect to d . We set

$\lambda_5 = 0.6$ for all examples in this paper. With two images as input, a minimum in E_d becomes noticeable, and the minimum becomes clearer with four images.

Simultaneously, an estimate of β can be computed from Eq. 8. Given $\alpha_1, \alpha_2, \dots, \alpha_n$, the function β can be solved by setting the partial derivative of $\sum_i \|\beta * \Pi_{\tau_i} - \alpha_i\|^2$ with respect to β to zero. This gives us

$$\beta = \text{deconv}(\Sigma(\alpha_i * \Pi_{\tau_i}), \Sigma(\Pi_{\tau_i} * \Pi_{\tau_i})),$$

where $C = \text{deconv}(A, B)$ is the deconvolution operation to seek a C that satisfies $A = B * C$.

Recovered β functions for two of the dust regions are displayed in Fig. 6. We note that with multiple input images, candidate dust regions can be more reliably detected because of the powerful constraint that they have similar β functions in each image.

5.3. Results

Fig. 8 exhibits results for three complex input images taken by a Canon 20D camera at F/16, F/14 and F/11. There exist two actual dust regions (indicated by red rectangles) on the sensor. Since these photographs each contain numerous small dark regions, detection of candidate dust regions becomes difficult when processing each image independently. When applying our single-image method on the top image of (a), the two true dust regions appear among the top 11 candidates (marked by green arrows). Similar results are obtained when processing the other two photographs individually. When using all three images as input, the two dust regions become the top two candidates.

In the close-up images, the quality of our recovery results is seen to be higher than that of conventional texture synthesis. We observe that although our single-image algorithm is less effective at detecting dust candidates, it can achieve recovery results that are comparable to our multiple-image technique, which employs much more image information.

6. Conclusion

In this paper, we introduced an approach to sensor dust removal that, unlike inpainting or texture synthesis methods, exploits the image information inside the damaged region. To effectively utilize this data, we presented an artifact formation model for sensor dust. The dust region information together with the artifact model jointly provides a constraint on the possible dust removal solutions. Image context and color consistency of dust are additionally utilized to recover a visually pleasing image that closely matches the actual scene. When multiple images are available, this approach can further offer the user a reliable list of candidate dust regions to choose from. Experiments demonstrate effective results with different cameras and complex scenes.

In future work, we plan to examine various cameras and determine the displacements d between the low-pass filter and imaging sensor. Although low-pass filters in digital cameras exist in different forms, such as birefringent plates or molded plastic sheets, and may be supplemented with an infrared cutoff filter, the number or ranges of typical d values may possibly be limited. Such a prior model on d could provide a useful constraint in our method for single-image input.

Although our method targets sensor dust removal, the presented framework could potentially be used in other applications. For example, similar artifacts often appear in scanned images due to fingerprints and smudges on the scanner glass. With a suitable model for artifact formation, the constrained use of contextual information in our dust removal technique might then be applied for scanner images.

Acknowledgements

We thank Hover Xue for providing several photographs with sensor dust artifacts, and Dong Xu for early discussions on this work.

Appendix: Effects of Low-pass Filtering

The low-pass filter in a camera is designed to blur a photograph by one or two pixels in order to prevent aliasing and moiré effects. Since sensor dust lies on this filter, the blurring it induces will also influence the formation of a dust artifact. For simplicity, we will discuss this effect in 1D.

For a low-pass filter that blurs an image by one pixel, we obtain an image $I = (U \cdot \alpha) * \Pi_1$, where U is the dust-free image prior to filtering. If the dust were not present, then the low-pass filtered image would instead be $I' = U * \Pi_1$. The original artifact function α without low-pass filtering would thus be transformed to $\alpha' = I/I'$ by the low-pass filter. In the discrete case, $(f * \Pi_1)(x) = [f(x) + f(x+1)]/2$, such that the artifact function can be expressed as

$$\alpha'(x) = \alpha(x) + \frac{U(x+1)}{U(x) + U(x+1)} \cdot [\alpha(x+1) - \alpha(x)].$$

The transformed artifact $\alpha'(x)$ is therefore bounded by $\alpha(x)$ and $\alpha(x+1)$. Because of image noise and the typical smoothness of $\alpha = \beta * \Pi_\tau$, we approximate α' with α and ignore the effect of the low-pass filter on dust artifacts.

References

- [1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *ACM SIGGRAPH*, pages 417–424, 2000. 2
- [2] T. Chan and J. Shen. Non-texture inpaintings by curvature-driven diffusions. *J. Visual Comm. Image Rep.*, 12(4):436–449, 2001. 2
- [3] T. F. Chan and C.-K. Wong. Total variation blind deconvolution. *IEEE Trans. Image Proc.*, 7(3):370–375, 1998. 4




Figure 7. Dust removal and candidate detection for a single image taken with a Kodak DCS 330 camera at $F/18$. In (a), green arrows indicate 17 actual dust regions among the top 18 candidates; red arrows show actual dust regions missing from the candidate list; and the yellow arrow shows a non-dust region that was included among the candidates. This candidate list can aid a user in dust detection. The ground truth was taken using another camera that is free of sensor dust. One dust region is enlarged for examination of recovery quality in comparison to texture synthesis and the ground truth.

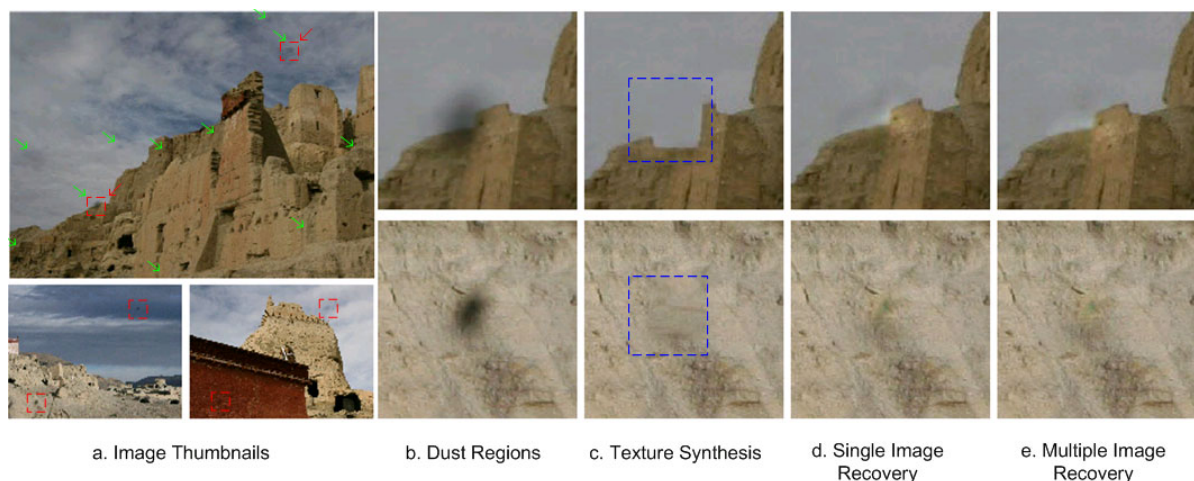


Figure 8. A complex example with numerous dust-like regions. (a) Three pictures taken by a Canon 20D at F -numbers $F/16$, $F/14$ and $F/11$. The top image shows candidate detection results (green arrows) using only that image. There exist only two actual dust regions on the sensor (red rectangles). Using all three photographs yields the two actual dust regions as the top two candidates. (b) Close-ups of the dust regions in the left images. (c) Texture synthesis result, where the synthesis region is outlined in blue. (d) Result from single-image dust removal. (e) Result of dust removal using all three images together.

- [4] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. *ACM SIGGRAPH*, pages 341–346, 2001. 2, 4, 5
- [5] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *Proc. IEEE International Conference on Computer Vision*, pages 1033–1038, 1999. 2
- [6] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. In *ACM SIGGRAPH*, pages 229–238, 1995. 2
- [7] N. Komodakis and G. Tziritas. Image completion using global optimization. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages 442–452, 2006. 4
- [8] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals Math. Stat.*, 22:79–86, 1951. 5
- [9] A. Levin, A. Zomet, and Y. Weiss. Learning how to inpaint from global image statistics. In *Proc. IEEE International Conference on Computer Vision*, pages 305–312, 2003. 2
- [10] L. Liang, C. Liu, Y.-Q. Xu, B. Guo, and H.-Y. Shum. Real-time texture synthesis by patch-based sampling. *ACM Trans. on Graphics*, 20(3):127–150, 2001. 2
- [11] P. Perez, M. Gagnat, and A. Blake. Patchworks: example-based region tiling for image editing. *Technical Report, MSR-TR-2004-04*, 2004. 2
- [12] J. Sun, L. Yuan, J. Jia, and H.-Y. Shum. Image completion with structure propagation. *ACM Trans. on Graphics*, 24(3):861–868, 2005. 2
- [13] R. G. Willson, M. W. Maimone, A. E. Johnson, and L. M. Scherr. An optical model for image artifacts produced by dust particles on lenses. In *Proc. ISAIRAS*, 2005. 2