

Noise Suppression in Low-light Images through Joint Denoising and Demosaicing

Priyam Chatterjee*
Univ. of California, Santa Cruz
priyam@soe.ucsc.edu

Neel Joshi Sing Bing Kang
Microsoft Research
{neel, sbkang}@microsoft.com

Yasuyuki Matsushita
Microsoft Research Asia
yasumat@microsoft.com

Abstract

We address the effects of noise in low-light images in this paper. Color images are usually captured by a sensor with a color filter array (CFA). This requires a demosaicing process to generate a full color image. The captured images typically have low signal-to-noise ratio, and the demosaicing step further corrupts the image, which we show to be the leading cause of visually objectionable random noise patterns (splotches). To avoid this problem, we propose a combined framework of denoising and demosaicing, where we use information about the image inferred in the denoising step to perform demosaicing. Our experiments show that such a framework results in sharper low-light images that are devoid of splotches and other noise artifacts.

1. Introduction

Low-light noise is a significant problem in photography. Most consumer cameras have poor low-light characteristics, which typically result in images with noticeable noise artifacts. Taking longer exposures reduces noise; however, this is only useful for static cameras and scenes, otherwise a blurry image is obtained. Active lighting in the form of a flash is not always viable as it causes color aberrations and is effective only for nearby objects. Recently, camera manufacturers have started addressing the problem of low-light photography by introducing cameras with larger photon sensors to accumulate more light at shorter exposures. In this paper, we propose a software-based approach to suppress the random mid-frequency texture patterns (or *splotches*) and other noise effects in low-light images, making it possible to capture sharp as well as noise-free images with short exposures.

Most modern consumer cameras capture images on a single sensor with a color filter array (CFA), of which the Bayer pattern [2] (Fig. 5(a)) is the most popular. A demosaicing algorithm is then used to interpolate the missing colors (e.g., red and blue values wherever only green samples

*The work presented in this paper was done while the author was visiting Microsoft Research, Redmond.

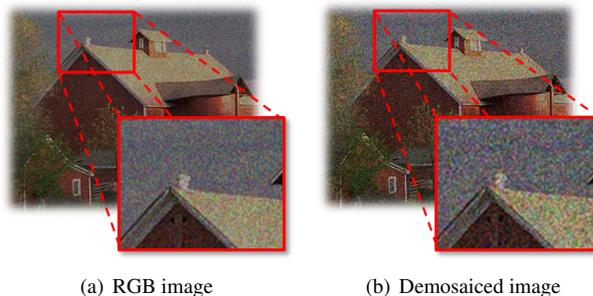


Figure 1. Effect of demosaicing on low-light noise characteristics: (a) RGB image with spatially independent Poisson noise; (b) demosaiced version of the noisy image. The simulated noisy RGB image was subsampled to form the Bayer pattern image. Notice how the demosaiced image demonstrates more splotches.

are observed) to form a full color image. With advances in optical and electronic systems, current cameras can produce high quality images when the images are well exposed. However, under low-light conditions, when the number of photons entering the imaging system is low, the raw image captured in the CFA has a low signal-to-noise ratio (SNR). Most cameras then produce quite noisy images. Cameras typically adapt to such low signal levels by increasing their sensitivity (ISO level) to the input signal, which amplifies the input noise. Such noise usually manifests itself in the form of splotches that are very salient to the human visual system (Fig. 1(b)). In our work, we concentrate on suppressing such visually unpleasant artifacts that typically corrupt light-limited images.

Our paper has two major contributions. First, we show, through simulations, that demosaicing is the main cause of the splotches commonly observed in low-light images. This finding leads us to our second contribution: a framework for combined denoising and demosaicing specifically targeting low-light images. To efficiently solve the joint problem, we pose the problem as that of vector upsampling in the presence of noise. This vector upsampling is carried out in a two-step framework that is motivated by Locally Linear Embedding (LLE) [23]. In the first step (denoising), embedding information is learned about the image from the noisy

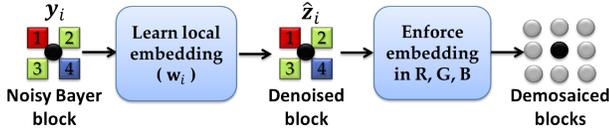


Figure 2. Outline of our denoising and demosaicing framework.

Bayer image, which is then enforced in an unknown low resolution image to perform the second step (demosaicing).

2. Related Work

Image denoising is a well-studied problem that has seen considerable performance improvement in recent years (see [6] for a survey). Broadly, denoising techniques can be classified as either transform-domain-based or spatial. Portilla *et al.* [21] proposed a wavelet-based method that was widely regarded as the state-of-the-art until recently. The recently proposed method of BM3D [8] employs a hybrid approach by first identifying similar patches within the noisy image. The similar patches are then denoised collaboratively in a transform-domain (*e.g.*, DCT) leading to remarkable denoising performance.

Spatial-domain methods can vary significantly in their approaches. Techniques range from denoising using image statistics [3, 7] or dictionaries [10] to using priors on the noise in a segmentation-based framework [17]. Another popular class of methods is based on directional filters [4, 24, 25] where denoising is performed by learning an embedding of the image in a local neighborhood. *Locality* can be defined either spatially [24, 25] or in the patch space [4], with the latter taking advantage of patch redundancy in the input image.

Mathematically, denoising in Non-Local Means (NLM) [4] is performed by representing each patch in the image as $\hat{z}_i = \mathbf{Y}w_i$, where \mathbf{Y} is a matrix (or dictionary) of patches created from the noisy image, with the coefficients w_i being a measure of similarity between the reference noisy patch y_i and column elements in \mathbf{Y} . The self-similarity based framework has also been extended to demosaicing [5]. Instead of using a fixed \mathbf{Y} , Elad *et al.* [10] proposed to iteratively update the over-complete dictionary \mathbf{Y} and sparse coefficients w_i to improve performance. This framework was later adapted for different restoration problems by Mairal *et al.* [18, 19].

Although dictionary-based methods have been applied to a variety of image restoration tasks, they rely on the presence of the ground truth to act as a prior on the desired output. This is especially true for image demosaicing [18, 19]. They are, thus, not applicable to our problem where such ground truth cannot be obtained without making simplistic assumptions about low-light image generation. Moreover, restoration can be adversely affected if the input image is not described well by the dictionary.

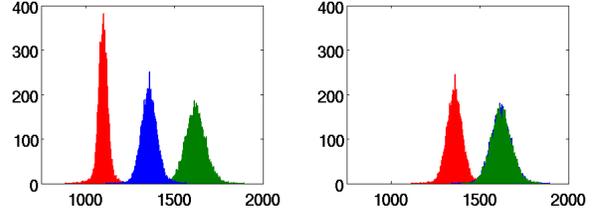


Figure 3. Histograms of noise values in different color channels of two different flat regions of the Macbeth color chart captured as Bayer pattern. The image was captured under low-light conditions at shutter speed 0.1 sec. and ISO 1600. In the plot to the right, the histograms for the green and blue channels are very similar.

Prior techniques do not explicitly target the removal of splotchy artifacts in low-light images. To remove such artifacts, it is essential to account for Poisson nature of the noise in the CFA (typically Bayer) image. The methods discussed thus far, and most CFA image denoising methods [11, 27], assume the noise to be Gaussian, and not signal-dependent. Deledalle *et al.* [9] recently extended NLM for Poisson noise by redefining the similarity measure between noisy grayscale patches. However, it is not clear how that work can be extended to Bayer images with non-homogeneous noise in different subsampled color channels.

In the class of denoising approaches that use color information, Hirakawa *et al.* [15] proposed an interpolation filter with the assumption that the color channels have similar higher frequencies. Zhang *et al.* [28] proposed a PCA-based approach for color reproduction from noisy raw data. Bennett *et al.*'s approach [3] assumes the true color at each pixel to be a combination of two colors which are estimated from a neighborhood. Although noise is accounted for in [3, 28], learning the basis colors and PCA vectors respectively can be challenging under low SNR. Another space-varying method for joint denoising and demosaicing was proposed by Menon *et al.* [20]. However, the noise there, as in [3, 28], is assumed to be independent and identically distributed Gaussian. As a result, they are not well tuned to handle low-light images.

In contrast to prior approaches, we are specifically interested in suppressing splotchy effects that are typical in low-light images. In order to design an effective technique to remove these artifacts, we first need to understand how they are generated. The next section describes our findings.

3. Splotch Artifacts in Low-Light Images

Images are typically captured using a CFA (*e.g.*, Bayer pattern), where each sensor element is capable of capturing light within specific wavelengths. The light (or photons) incident on a sensor element is accumulated over a period of time controlled by the shutter speed setting. Due to this photonic nature of light, the voltage resulting from the photon counting process is noisy. This noise is called shot noise.

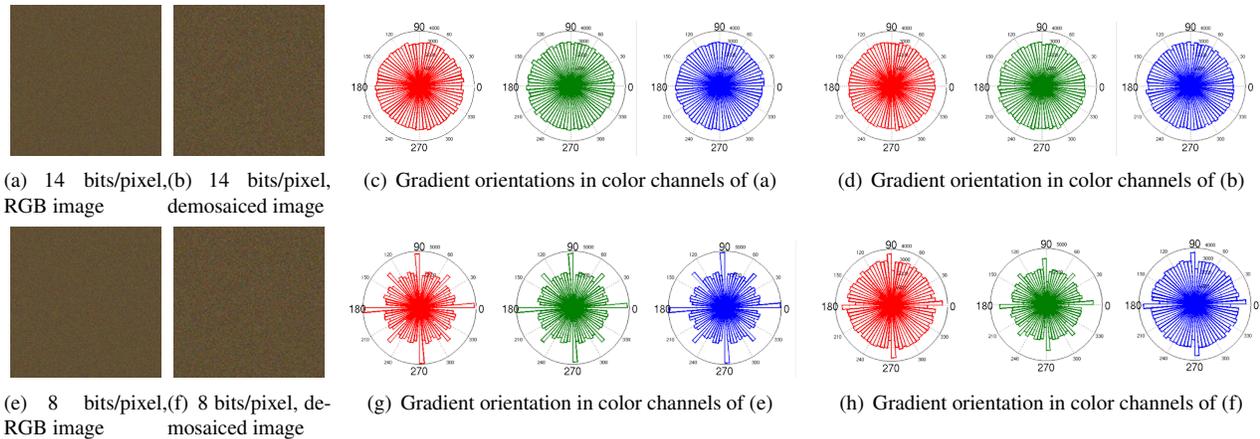


Figure 4. Directionality of Poisson noise visualized with gradient orientations histograms for images with various quantization levels.

The voltage in each sensor element is usually magnified by a gain factor based on the ISO setting of the camera. While this increases sensitivity to the input signal, it also results in amplification of the noise. The voltage is then quantized to obtain the raw intensity values. Although imperfect electronic components also contribute to the noise affecting the image, the dominant noise at lower intensities has been shown to be shot and quantization noise [1].

Next, the camera demosaics the noisy Bayer image to produce a color image. Since most demosaicing algorithms do not account for noisiness of the raw data, a denoising process is often applied to the resulting image. This is followed by other non-linear operations such as gamma correction, color tone-mapping, and even sharpening before the final color image is stored in the camera. Here we are interested in identifying the chief causes of the splotches in low-light images.

We start by studying the noise characteristics of such images. As mentioned earlier, shot noise is a result of photonic nature of light. Since the intensity values are a result of a counting process, the noise is Poisson distributed, and can be assumed to be spatially independent. We confirmed this experimentally by capturing multiple images of a static scene (*e.g.*, Macbeth color chart) under various shutter speeds and ISO settings. We found that in any given flat region, the variance of the noise is proportional to the mean intensity of the flat region obtained by averaging the intensities of each color channel, the constant of proportionality being dependent on the camera's ISO setting (gain factor). Moreover, the histogram of the noise in each color channel in each flat area fits a Poisson density function. This is shown in Fig. 3 for two different flat regions of the color chart captured under low-light conditions with shutter speed 0.1 second and ISO 1600. This Poisson distribution is attributed to the shot noise [12, 14]. The quantization noise, on the other hand, is due to the digitization of the voltage.

Although modern cameras are capable of storing raw intensities with high precision (bits per pixel), we still study the effects of such quantization in low-light images.

In our study, we ran simulations to study the combined effects of different processes in the image formation pipeline, namely, quantization and demosaicing. Our simulator is designed to study the effects of different levels of Poisson distributed noise on different quantization levels of mainly constant-intensity images. The effect was studied when noise was added independently in each color channel of color images, as well as in the demosaiced image obtained from subsampling the noisy color image according to the Bayer pattern. The synthetically generated noisy images were visually inspected and analyzed (by computing gradient and orientation histograms).

In our simulations over multiple flat images, the effect of quantization level on the splotches was visually indiscernible for the full color as well as demosaiced images (Figures 4(a)-(d)). However, the histogram of gradient orientations (Figures 4(e)-(h)) exhibit stronger directionality along vertical and horizontal directions as quantization is reduced to 8 bits per pixel. Considering that even lower-end commercial cameras record information in at least 8 bits per pixel and the fact that this does not lead to any visually detectable increase in splotches, we eliminate quantization as a significant factor of image noise artifacts. Henceforth, we consider Poisson distributed shot noise as the major source of low-light noise.

Our simulations with varying levels of Poisson noise led to two important conclusions: (1) even when the noise is independent in each channel, patterns are visible in the noisy image (gradient orientation histograms do not show any bias in edge direction); (2) the visual patterns (splotches) are more pronounced in demosaiced images, irrespective of the demosaicing method employed. Interestingly, demosaicing does not seem to alter the approximately uniform distribu-

tion of gradient orientations. This leads us to believe that the noisy patterns in the image are a result of humans being adept at identifying patterns in any image [26], and that such patterns are noticeably enhanced by the demosaicing process. An example of such enhanced splotches is illustrated in Fig. 1 showing a house image with simulated Poisson noise added independently in each color channel and its demosaiced version obtained from the simulated CFA noisy image sampled according to the Bayer pattern.

Although modern demosaicing techniques (see [13, 16] for comprehensive reviews) take into account inter-channel dependencies, most do not account for noise in the observed Bayer samples. Interpolating across noisy color channels corrupts the noise characteristics in the demosaiced image, enhancing splotches in the process. Such artifacts are then propagated further in the imaging pipeline. Hence, it is important to suppress the Bayer domain Poisson noise *before* demosaicing. Most Bayer domain denoising methods [11, 27] do not handle Poisson noise. Moreover, the demosaicing method usually has no information about the denoising process and errors in denoising can be propagated easily. Motivated by this, we propose a method based on LLE [23] that we adapt to perform combined denoising and demosaicing in a vector upsampling framework.

4. Denoising–Demosaicing as Upsampling

In our proposed framework, we pose the problem of denoising and demosaicing as that of vector upsampling (Fig. 5). The task of demosaicing is to estimate the unknown color information at each pixel location. We consider this problem as estimating the unobserved blocks (Fig. 5(b)), at exact one pixel shifts in horizontal, vertical and diagonal directions. For example, consider the numbered regions in Fig. 5(a), estimating the horizontally shifted image results in estimating the green channel information at locations where only a red value was observed. In a similar manner, vertical and diagonally shifted image estimates result in obtaining estimates of different color channels. As a result, red, blue and two green color estimates are accumulated at each pixel location, from which a full color RGB image is obtained. Thus, we show that vector upsampling with an appropriate choice of vectors ($[R, G, G, B]$ patterns shown in the figure) can lead to demosaicing. As such, any method that performs 2-D upsampling of vector-valued functions can be used for this purpose. However, such interpolation methods must account for noise contamination of the samples, as well as the correlation between the observations ($[R, G, G, B]$ vector elements).

Here we propose to perform vector interpolation by learning *local* embeddings of each Bayer block (see Fig. 2), where locality is in terms of similarity of block intensities (vector magnitudes). This is performed by representing each observed block (black nodes) as a linear combination

of similar blocks to obtain a denoised Bayer image. This *embedding* is then enforced in the demosaicing step through an optimization framework. This implicitly enforces the cross-color correlation in demosaicing as the embeddings are learned simultaneously from all color channels.

4.1. Denoising Step

As mentioned earlier, denoising of the Bayer blocks is performed by representing each noisy block as a linear combination of similar blocks. To learn such an embedding, we first need to identify similar blocks. However, under low-light conditions, the weak signal corrupted by noise leads to low signal-to-noise ratios (SNR), making it difficult to identify similar blocks. To add robustness, we consider larger regions consisting of multiple spatially collocated blocks to form patches. Blocks with similar neighborhoods to any given reference patch are identified and the linear embedding weights are learned. The denoising step thus consists of learning a *local* embedding of a *patch* in the Bayer domain and using it to perform denoising. Below we provide an outline for the denoising process (Fig. 6):

1. Form patches centered at each $[R, G, G, B]$ node. For any given patch centered at \mathbf{y}_i identify the K -most similar patches (centers denoted by \mathbf{y}_j) within the image subject to noise corruption.
2. Represent the center node of the patch $\hat{\mathbf{z}}_i$ under consideration of similarity to \mathbf{y}_i as a linear combination of the similar patches to be of the form

$$\hat{\mathbf{z}}_i = \sum_{j=1}^K w_{ij} \mathbf{y}_j = \mathbf{Y}_i \mathbf{w}_i, \quad (1)$$

where \mathbf{Y}_i is the matrix formed by concatenating the vectorized center nodes of the similar patches \mathbf{y}_j , w_{ij} are the weights that measure the confidence of a patch centered at \mathbf{y}_j being similar to that centered at \mathbf{y}_i , and \mathbf{w}_i is a vector containing the w_{ij} entries. The weights are calculated in an optimization framework as

$$\hat{\mathbf{w}}_i = \arg \min_{\mathbf{w}_i} \|\mathbf{y}_i - \mathbf{Y}_i \mathbf{w}_i\|_p, \quad (2)$$

with p chosen to fit the error distribution.

3. Repeat the process for blocks with different patterns, namely, $[G, R, B, G]$, $[G, B, R, G]$, & $[B, G, G, R]$.

Note that with blocks of a certain pattern, we match blocks with the same pattern only. The embedding learning process is carried out for each such pattern to obtain multiple denoised estimates of the Bayer image. The final estimate is obtained as an average of the multiple estimates. This secondary averaging process reduces the residual noise in the initial estimates to produce the final denoised raw image. More importantly, doing so allows us to obtain multiple embeddings of each pixel, which provides more constraints for the optimization in the demosaicing step.

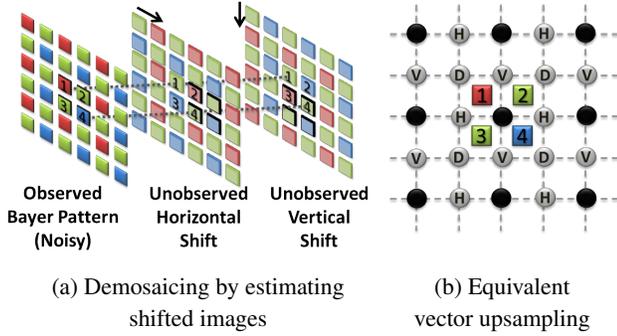


Figure 5. Demosaicing of noisy Bayer data as vector upsampling: (a) single pixel shifts in horizontal, vertical, and diagonal (not shown) directions fill-in missing color information at different locations; (b) corresponding visualization of vector upsampling where the black nodes are observed noisy Bayer patterns and gray nodes marked H, V, and D are filled-in through horizontal, vertical and diagonal shifted image estimation respectively.

In order to apply this scheme, we first need to identify similar patches, given only their noisy observations. Under low-light conditions, the noise has been shown to be Poisson distributed which is dependent on the (unknown) true intensity of the Bayer pattern image. Similarity measures based on such Poisson noise has been proposed by Alter *et al.* [1]. However, in our case we found that using an ℓ_2 distance measure between noisy patches

$$d(\mathbf{y}_i, \mathbf{y}_j) = \|\mathbf{y}_i - \mathbf{y}_j\|_2^2, \quad (3)$$

to identify similar patches works reasonably well for our purposes. This is because we are merely interested in ranking noisy patches based on similarity to a reference patch to identify the $K = 30$ most similar patches. The weight calculation mechanism of Eq. (2) then implicitly rejects dissimilar patches by giving them low weights.

The patch-based embedding of each block allows us to perform denoising of the raw data as outlined in Fig. 6. Demosaicing this image using any demosaicing method results in a denoised color image. However, further noise suppression is achieved when the demosaicing step makes use of the local linear embedding information learned in the denoising step.

4.2. Demosaicing step

Once the raw image is denoised, the missing color information needs to be estimated to form the final image. This process is known as demosaicing. One can use any existing method to perform demosaicing of the denoised raw image. However, demosaicing being a special form of interpolation inherently leads to reduced sharpness in the final image. Moreover, for images with low SNR as in our case, the noise may not be sufficiently removed by denoising. One way to further suppress the noise is to use any existing demosaicing algorithm to obtain an initial color estimate and

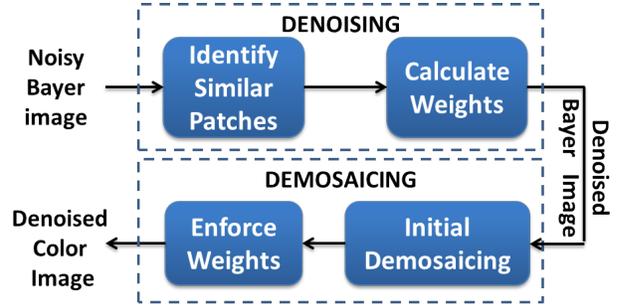


Figure 6. Our combined denoising and demosaicing framework.

enforce the embedding on it to obtain the final output. However, such a method relies heavily on the performance of the initial demosaicing. To counter that effect, we propose to perform demosaicing in an optimization framework.

The motivation of our optimization framework arises from the relation between a high resolution (HR) image and its corresponding low resolution (LR) version. Consider the Bayer image as an incomplete observation of an unknown HR image that we wish to estimate. A $2 \times$ downsampled LR image could then be formed by first averaging each color channel of the HR image in a 2×2 region and then subsampling the smoothed image. Thus, the color information at any pixel at location i of the LR image is obtained from averaging of the colors at the i -th block of the HR image, only some of which are observed in the Bayer image. Without loss of generality, consider the red color channel in Fig. 5(a). We can see that the red value in the LR image corresponding to the numbered block (indexed by i) is obtained from the observed r_{i1} and unobserved r_{il}^{un} for $l = 2, 3, 4$. For notational convenience, we denote r_{i1} as r_i . The embeddings for each block in the Bayer domain can thus be thought to be similar to the embedding for the corresponding pixel in the (unknown) LR image. This is in keeping with the spirit of LLE [23], where dimensionality reduction is performed by enforcing local embeddings learned in the higher dimension to data in the lower dimensional space. In our context, we learn the local embeddings in the partially observed HR (Bayer) image and enforce them in the downsampled LR image. Thus, we can write

$$r_i + \sum_{l=2}^4 r_{il}^{un} = \sum_j w_{ij} (r_j + \sum_{l=2}^4 r_{jl}^{un}) \quad \forall i. \quad (4)$$

Rewriting the above equation in terms of all the observed and unobserved samples of the HR image in a vector form (\mathbf{r} and \mathbf{r}^{un} respectively) we obtain

$$\mathbf{r} + \mathbf{D}\mathbf{r}^{un} = \mathbf{W}(\mathbf{r} + \mathbf{D}\mathbf{r}^{un}) \quad (5)$$

$$\Rightarrow \tilde{\mathbf{r}} = (\mathbf{W} - \mathbf{I})\mathbf{D}\mathbf{r}^{un} = \tilde{\mathbf{W}}\mathbf{r}^{un}, \quad (6)$$

where \mathbf{D} is a matrix operator where each row implements a summation of the unobserved samples corresponding to the observed samples \mathbf{r} , \mathbf{W} is a sparse matrix with w_{ij} as its

i, j -th entry, $\tilde{\mathbf{r}} = \mathbf{r} - \mathbf{W}\mathbf{r}$ is a function of observed samples only, and $\tilde{\mathbf{W}} = (\mathbf{W} - \mathbf{I})\mathbf{D}$ with \mathbf{I} as the identity matrix. In the demosaicing step, we are interested in estimating \mathbf{r}^{un} given the estimated $\tilde{\mathbf{W}}$ and $\tilde{\mathbf{r}}$. Note that there are many more unobserved red values than observed ones making the problem ill-posed. Hence, we employ a regularization framework to reduce the solution space. We use Eq. (6) as our data fitting term and an initial demosaicing of the denoised Bayer image as prior information (\mathbf{r}_0^{un}). The final estimate for the demosaiced image is obtained by solving for the unknown intensities as

$$\hat{\mathbf{r}}^{un} = \arg \min_{\mathbf{r}^{un}} \|\tilde{\mathbf{r}} - \tilde{\mathbf{W}}\mathbf{r}^{un}\|_p + \|\mathbf{\Lambda}(\mathbf{r}^{un} - \mathbf{r}_0^{un})\|_q, \quad (7)$$

where $\mathbf{\Lambda}$ is a diagonal matrix with varying diagonal entries that allow us to control the penalty of deviating from the initial estimate, and p, q are chosen according to the distribution of the error for the data fitting and regularization terms. We use $p = q = 0.65$, since the histogram of the data fitting error fits a hyper-Laplace distribution for our simulated experiments (see the following section). This also allows us to avoid over-fitting at the higher intensities and over-smoothing in the low-intensity regions [22].

As mentioned earlier, noise may still be present in the denoised Bayer image. Such noise is more perceptible in the smoother regions. Thus, allowing the final estimate in the smoother regions to deviate from the initial estimate can be advantageous. On the other hand, enforcing the prior can be good in the textured regions where there may be fewer similar patches. The diagonal entries of $\mathbf{\Lambda}$ can thus be chosen to be inversely related to the number of similar blocks (large w_{ij} values) found in the embedding stage which serves as a good measure of the reliability of the learned embedding. However, we found that simply using a fixed $\mathbf{\Lambda} = 0.1 \mathbf{I}$ in the optimization framework leads to satisfactory suppression of noise and splotches in low-light images, as demonstrated in Sec. 5.

4.3. Poisson Noise Model and Error Norm

The form of the error norm in our LLE framework is dependent on the image noise distribution. While it is known that low-light noise is Poisson distributed, without further analysis, it is not obvious what the error distributions would be on the terms in Equations 2 and 7. We empirically measured the distributions, by generating Poisson distributed samples in a low-intensity range from a set of known samples with known coefficients. We fit a generalized Gaussian to this distribution and found it to be hyper-Laplacian ($p = 0.65 < 1$), with the variance proportional to the mean value of the samples. This distribution is consistent with a mean subtracted Poisson distribution as it is sparser than a Gaussian, which is to be expected for a Poisson distribution at low-intensity levels, and its variance is proportional to the original mean of the distribution. We minimized this norm using iterative re-weighted least-squares.

5. Results

To validate our method, we perform experiments with real-world images captured with different cameras, under different low-light conditions. In Fig. 7, we compare our method to straight demosaicing results – we use the ‘demosaic’ function in Matlab. The noise artifacts are significantly reduced by our approach across a large range of SNR, image content, and lighting conditions (indoor, outdoor, daytime, nighttime, and flash), without any parameter tuning or user input to handle the change in the noise levels and image content. The images in these results were shot with a Canon 1Ds Mark III and Rebel Xti, which have very different image sensors, showing that our method generalizes across camera types. Note how the extremely high-textured splotch patterns seen in the straight demosaiced results do not appear in our results – the smooth areas appear smooth, yet textures are largely retained.

In Fig. 8 we compare our results to several existing techniques. For our comparisons, we choose two methods that perform denoising after demosaicing, as is commonly done. The methods used are Neat Image and BM3D [8] (state-of-the-art denoiser [6]). Since our method performs denoising on Bayer images, we also compare to a recent Bayer denoising method [27]. Finally, we show results from a joint denoising and demosaicing method proposed by Menon *et al.* [20]. For fair comparisons, we tuned the parameters of each method to obtain the most visually pleasing results.

From Fig. 8, it can be seen that our method maintains a good balance of denoising smooth regions while retaining textured regions. For previous methods, a trade-off must be made. Either the support-region needs to be large to remove the splotches, which then blurs the image detail, or vice versa. Moreover, with many of the previous methods, even when the high and mid-frequency splotches are removed, lower frequency components of the artifacts remain, where smooth regions appear to have a low-frequency low-intensity color pattern. The benefits of our method are best seen when viewing the images on a computer screen.

6. Conclusions

We proposed a framework for combined denoising and demosaicing of low-light images, specifically targeted at reducing the splotches and noise effects in such images. For this we make use of a local linear embedding framework where a linear embedding is learned in the denoising step. This local embedding is then enforced within an optimization framework to perform refinement of an initial demosaicing of the denoised image. We showed that the proposed method is able to considerably improve the quality of the final output image over the traditional approaches of demosaicing and then denoising and also performs better

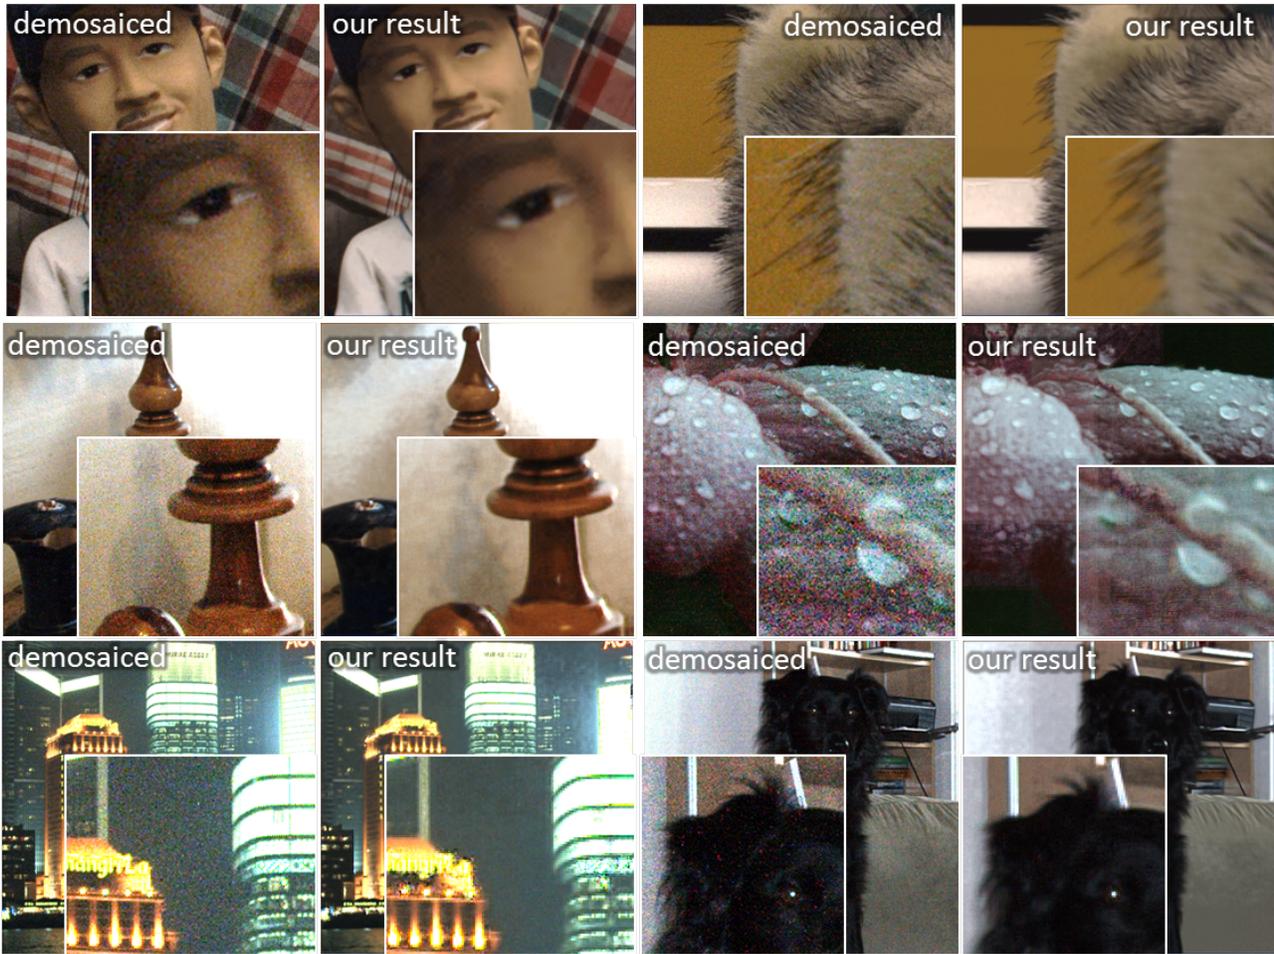


Figure 7. Comparisons of real images captured under various conditions to those obtained by straight demosaicing. The low-light images are shot with different cameras (Canon 1Ds Mark III and Rebel Xti)

than recently proposed joint denoising and demosaicing approaches.

Our method assumes access to the raw image data before it is corrupted by various stages in the image formation process. However, such raw data is not easily accessed for a considerable number of consumer cameras. For those cases and for existing images, it is necessary to invert the demosaicing and other in-camera processing to estimate the raw data before applying our algorithm. This is a direction of our future work.

References

- [1] F. Alter, Y. Matsushita, and X. Tang. An intensity similarity measure in low-light conditions. In *Proc. of ECCV*, volume 4, pages 267–280, Graz, Austria, May 2006.
- [2] B. E. Bayer. Color imaging array. *US Patent 3971065*, July 1976.
- [3] E. P. Bennett, M. Uyttendaele, C. L. Zitnick, R. Szeliski, and S. B. Kang. Video and image Bayesian demosaicing with a two color image prior. In *Proc. of ECCV*, volume 1, pages 508–521, Graz, Austria, May 2006.
- [4] A. Buades, B. Coll, and J. M. Morel. A review of image denoising methods, with a new one. *Multiscale Modeling and Simulation*, 4(2):490–530, 2005.
- [5] A. Buades, B. Coll, J.-M. Morel, and C. Sbert. Self-similarity driven color demosaicking. *IEEE T. IP*, 18(7):1192–1202, July 2009.
- [6] P. Chatterjee and P. Milanfar. Is denoising dead? *IEEE T. IP*, 19(4):895–911, Apr. 2010.
- [7] T. S. Cho, N. Joshi, C. L. Zitnick, S. B. Kang, R. Szeliski, and W. T. Freeman. A content-aware image prior. In *Proc. of IEEE CVPR*, San Francisco, CA, June 2010.
- [8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3d transform-domain collaborative filtering. *IEEE T. IP*, 16(8):2080–2095, Aug. 2007.
- [9] C.-A. Deledalle, F. Tupin, and L. Denis. Poisson NL means: Unsupervised non local means for Poisson noise. In *Proc. of IEEE ICIP*, pages 801–804, Hong Kong, Sept. 2010.

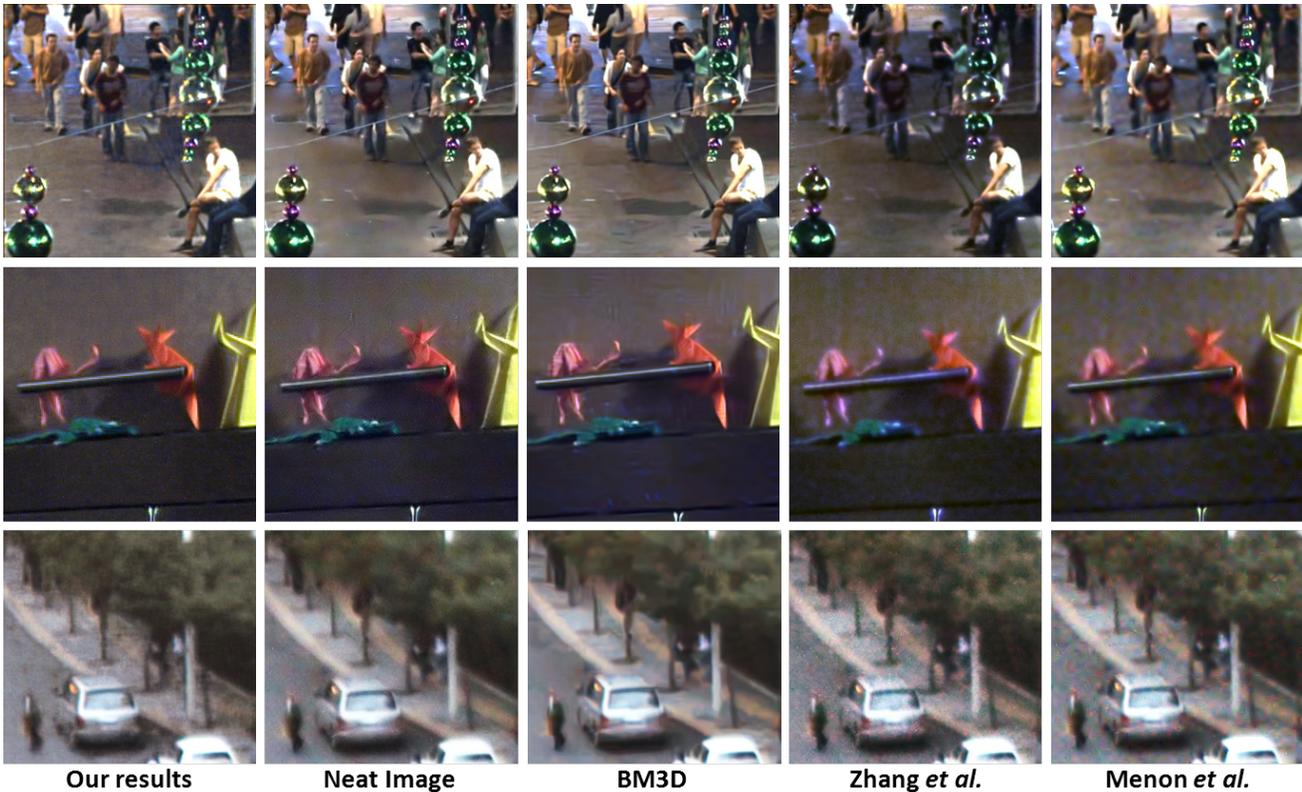


Figure 8. Our results compared to existing techniques - Neat ImageTM, BM3D [8], Zhang *et al.* [27], and Menon *et al.* [20].

- [10] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE T. IP*, 15(12):3736–3745, Dec. 2006.
- [11] H. Faraji and W. J. MacLean. CCD noise removal in digital images. *IEEE T. IP*, 15(9):2676–2685, Sept. 2006.
- [12] A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian. Practical Poissonian-Gaussian noise modeling and fitting for single-image raw data. *IEEE T. IP*, 17(10):1737–1754, Oct. 2008.
- [13] B. K. Gunturk, J. Glotzbach, Y. Altunbasak, R. W. Schafer, and R. M. Mersereau. Demosaicking: Color filter array interpolation. *IEEE Sig. Process. Mag.*, 22(1):44–54, Jan. 2005.
- [14] G. E. Healey and R. Kondepudy. Radiometric ccd camera calibration and noise estimation. *IEEE T. PAMI*, 16(3):267–276, Mar. 1994.
- [15] K. Hirakawa and T. Parks. Joint demosaicing and denoising. *IEEE T. IP*, 15(8):2146–2157, Aug. 2006.
- [16] X. Li, B. Gunturk, and L. Zhang. Image demosaicing: A systematic survey. In *Proc. of SPIE Electronic Imaging Conf.*, volume 6822, page 68221J, San Jose, CA, Jan. 2008.
- [17] C. Liu, R. Szeliski, S. B. Kang, C. L. Zitnick, and W. T. Freeman. Automatic estimation and removal of noise from a single image. *IEEE T. PAMI*, 30(2):299–314, Feb. 2008.
- [18] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Non-local sparse models for image restoration. In *Proc. of ICCV*, Kyoto, Japan, Sept. 2009.
- [19] J. Mairal, M. Elad, and G. Sapiro. Sparse representation for color image restoration. *IEEE T. IP*, 17(1):53–69, Jan. 2008.
- [20] D. Menon and G. Calvagno. Joint demosaicking and denoising with space-varying filters. In *Proc. of IEEE ICIP*, pages 477–480, Cairo, Nov. 2009.
- [21] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli. Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE T. IP*, 12(11):1338–1351, Nov. 2003.
- [22] M. Raginsky, R. M. Willett, Z. T. Harmany, and R. F. Marcia. Compressed sensing performance bounds under Poisson noise. *IEEE T. SP*, 58(8):3990–4002, Aug. 2010.
- [23] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
- [24] H. Takeda, S. Farsiu, and P. Milanfar. Kernel regression for image processing and reconstruction. *IEEE T. IP*, 16(2):349–366, Feb. 2007.
- [25] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. of IEEE ICCV*, pages 839–846, Jan. 1998.
- [26] J. Whitson and A. Galinsky. Lacking control increases illusory pattern perception. *Science*, 322(5898):115, 2008.
- [27] L. Zhang, R. Lukac, X. Wu, and D. Zhang. PCA-based spatially adaptive denoising of CFA images for single-sensor digital cameras. *IEEE T. IP*, 18(4):797–812, Apr. 2009.
- [28] L. Zhang, X. Wu, and D. Zhang. Color reproduction from noisy CFA data of single sensor digital cameras. *IEEE T. IP*, 16(9):2184–2197, Sept. 2007.