Automatic Estimation and Removal of Noise from a Single Image

Ce Liu, *Student Member*, *IEEE*, Richard Szeliski, *Fellow*, *IEEE*, Sing Bing Kang, *Senior Member*, *IEEE*,
C. Lawrence Zitnick, *Member*, *IEEE*, and William T. Freeman, *Senior Member*, *IEEE*

Abstract—Image denoising algorithms often assume an additive white Gaussian noise (AWGN) process that is independent of the actual RGB values. Such approaches cannot effectively remove color noise produced by today's CCD digital camera. In this paper, we propose a unified framework for two tasks: automatic estimation and removal of color noise from a single image using piecewise smooth image models. We introduce the noise level function (NLF), which is a continuous function describing the noise level as a function of image brightness. We then estimate an upper bound of the real NLF by fitting a lower envelope to the standard deviations of per-segment image variances. For denoising, the chrominance of color noise is significantly removed by projecting pixel values onto a line fit to the RGB values in each segment. Then, a Gaussian conditional random field (GCRF) is constructed to obtain the underlying clean image from the noisy input. Extensive experiments are conducted to test the proposed algorithm, which is shown to outperform state-of-the-art denoising algorithms.

Index Terms—Image denoising, piecewise smooth image model, segmentation-based computer vision algorithms, noise estimation, Gaussian conditional random field, automatic vision system.

✦

---

1 INTRODUCTION

IMAGE denoising has been studied for decades in computer vision, image processing and statistical signal processing. This problem not only provides a good platform to examine natural image models and signal separation algorithms, but it is also important part of image enhancement for digital image acquisition systems. These two directions are both important and will be explored in this paper.

Most of the existing image denoising work assumes additive white Gaussian noise (AWGN) and removes the noise independently of the RGB image data. However, the type and level of the noise generated by digital cameras are unknown if the series and brand of the camera, as well as the camera settings (ISO, shutter speed, aperture, and flash on/off), are unknown. For instance, the exchangeable image file format (EXIF) metadata attached with each picture can be lost in image format conversion and image file transferring. Meanwhile, the statistics of the color noise is not independent of the RGB channels because of the demosaic process embedded in cameras. Therefore, most current denoising approaches are *not* truly automatic and *cannot* effectively remove color noise. This fact prevents the noise removal techniques from being practically applied to digital image denoising and enhancing applications.

In some image denoising software, the user is required to specify a number of smooth image regions to estimate the noise level. This motivated us to adopt a segmentation-based approach to automatically estimate the noise level from a single image. Since the noise level is dependent on the image brightness, we propose to estimate an upper bound of the *noise level function* (NLF) from the image. The image is partitioned into piecewise smooth regions in which the mean is the estimate of brightness, and the standard deviation is an overestimate of noise level. The prior probability of the noise level functions is learned by simulating the digital camera imaging process and are used to help estimate the curve correctly where there is missing data.

Since separating signal and noise from a single input is underconstrained, it is in theory impossible to completely recover the original image from the noise contaminated observation. The goal of image denoising is to preserve image features as much as possible while eliminating noise. There are a number of goals we want to meet in designing image denoising algorithms.

1. The perceptually *flat regions* should be as smooth as possible. Noise should be completely removed from these regions.
2. *Image boundaries* should be well preserved. This means that the boundary should not be blurred or sharpened.
3. *Texture detail* should not be lost. This is one of the hardest criteria to match. Since image denoising algorithms tend to smooth the image, it is easy to lose texture detail during smoothing.
4. The *global contrast* should be preserved (that is, the *low frequencies* of the denoised and input images should be identical).
5. No *artifacts* should appear in the denoised image.

The global contrast is probably the easiest to meet, whereas some of the other principles are nearly incompatible. For instance, goals 1 and 3 are difficult to adjust

---

- *C. Liu and W.T. Freeman are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 32 Vassar Street, Cambridge, MA 02139. E-mail: {celiu, billf}@mit.edu.*
- *R. Szeliski, S.B. Kang, and C.L. Zitnick are with Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399.*
  E-mail: {szeliski, SingBing.Kang, larryz}@microsoft.com.

together since most denoising algorithms cannot distinguish flat from textured regions within a single input image. Satisfying goal 5 is important but not always easy, for example, wavelet-based denoising algorithms tend to generate ringing artifacts.

Ideally, the same image model should be used for both noise estimation and denoising. We found that a segmentation-based approach is well suited to both tasks. After a natural image is oversegmented into piecewise smooth regions, the pixel values within each segment approximately lie on a 1D line in RGB space due to the physics of image formation [26], [24], [20]. This important fact can help to significantly reduce color noise. We further improve the results by constructing a Gaussian conditional random field (GCRF) to estimate the clean image (signal) from the noisy image.

Experiments are conducted, with both quantitatively convincing and visually pleasing results to demonstrate that our segmentation-based denoising algorithm outperforms the state of the art. Our approach is distinctively automatic since the noise level is automatically estimated. Automatically estimating the noise level can benefit other computer vision algorithms as well. For example, the parameters of stereo, motion estimation, edge detection, and super resolution algorithms can be set as a function of the noise level so that we can avoid tweaking the parameters for different noise levels [30].

The paper is organized as follows: After reviewing relevant work in Section 2, we introduce our piecewise smooth image model in Section 3. In Section 4, we propose the method for noise estimation from a single image. Our segmentation-based image denoising algorithm is presented in detail in Section 5, with results shown in Section 6. We discuss issues of color noise, modeling, and automation in Section 7 and provide concluding remarks in Section 8.

2   RELATED WORK

In this section, we briefly review previous work on image denoising and noise estimation. Image denoising techniques differ in the choice of image prior models, and many noise estimation techniques assume white Gaussian noise (AWGN).

2.1   Image Denoising

In the past three decades, a variety of denoising methods have been developed in the image processing and computer vision communities. Although seemingly very different, they all share the same property: to keep the meaningful edges and remove less meaningful ones. We categorize the existing image denoising work by their different natural image prior models and the corresponding representation of natural image statistics.

Wavelets. When a natural image is decomposed into multiscale-oriented subbands [31], we observe highly kurtotic marginal distributions [16]. To enforce the marginal distribution to have high kurtosis, we can simply suppress low-amplitude values while retaining high-amplitude values, a technique known as coring [40], [44].

In [43], the joint distribution of wavelets were found to be dependent. A joint coring technique is developed to infer the wavelet coefficients in a small neighborhood across different orientation and scale subbands simultaneously.

The typical joint distribution for denoising is a Gaussian scale mixture (GSM) model [38]. In addition, wavelet-domain hidden Markov models have been applied to image denoising with promising results [9], [14].

Although the wavelet-based method is popular and dominant in denoising, it is hard to remove the ringing artifacts of wavelet reconstruction. In other words, wavelet-based methods tend to introduce additional edges or structures in the denoised image.

Anisotropic diffusion. The simplest method for noise removal is Gaussian filtering, which is equivalent to solving an isotropic heat diffusion equation [47], a second-order linear PDE. To keep sharp edges, anisotropic diffusion can be performed using $I_t = div(c(x, y, t)\nabla I)$ [35], where $c(x, y, t) = g(\|\nabla I(x, y, t)\|)$, and g is a monotonically decreasing function. As a result, for high gradient pixels, $c(x, y, t)$ is small and therefore gets less diffused. For low gradient pixels, $c(x, y, t)$ has a higher value, and these pixels get blurred with neighboring pixels. A more sophisticated way of choosing $g(\cdot)$ is discussed in [4]. Compared to simple Gaussian filtering, anisotropic diffusion smooths out noise while keeping edges. However, it tends to overblur the image and sharpen the boundary with many texture details lost.

More advanced partial differential equations (PDEs) have been developed so that a specific regularization process is designed for a given (user-defined) underlying local smoothing geometry [53], preserving more texture details than the classical anisotropic diffusion methods.

FRAME and FOE. As an alternative to measuring marginal or joint distributions on wavelet coefficients, a complete prior model over the whole image can be learned from marginal distributions [19], [56]. Thus, it is natural to use a Bayesian inference for denoising or restoration [55], [41]. The processed image I typically takes the iterative form

$$I_t = I_{t-1} + \eta \left[\sum_{i=1}^{n} F_i^{-1} * \lambda_i'(F_i * I_{t-1}) + \frac{1}{\sigma^2(I^{obs} - I_{t-1})}^2 \right],$$

where $\{F_i\}$ are linear filters (F_i^{-1} is the filter obtained by mirroring F_i around its center pixel), $\{\lambda_i\}$ are the corresponding Gibbs potential functions, σ^2 is the variance of noise, and t is the index of iteration. Because the derivative λ_i' typically has high values close to zero and low values at high amplitude, the above PDE is very similar to anisotropic diffusion if the F_is are regarded as derivative filters at different directions [55].

Learning a Gibbs distribution using MCMC can be inefficient. Furthermore, these methods can have the same drawbacks as anisotropic diffusion: over smoothing and edge sharpening.

Bilateral filtering. An alternative way of adapting Gaussian filtering to preserve edges is bilateral filtering [52], where both space and range distances are taken into account. The essential relationship between bilateral filtering and anisotropic diffusion is derived in [2]. A fast bilateral filtering algorithm is also proposed in [11], [34].

Bilateral filtering has been widely adopted as a simple algorithm for denoising, for example, video denoising in [3]. However, it cannot handle speckle noise, and it also has the tendency to oversmooth and to sharpen edges.

Nonlocal methods. If both the scene and camera are static, we can simply take multiple pictures and use the mean to remove the noise. This method is impractical for a single

image, but a temporal mean can be computed from a spatial mean—as long as there are enough similar patterns in the single image. We can find the similar patterns to a query patch and take the mean or other statistics to estimate the true pixel value, for example, in [1], [6]. A more rigorous formulation of this approach is through sparse coding of the noisy input [12].

Nonlocal methods are an exciting innovation and work well for texture-like images containing many repeated patterns. However, compared to other denoising algorithms that have $O(n^2)$ complexity, where n is the image width, these algorithms have $O(n^4)$ time complexity, which is prohibitive for real-world applications.

Conditional random fields (CRFs). Recently, CRFs [27] have been a promising model for statistical inference. Without an explicit prior model on the signal, CRFs are flexible at modeling short and long range constraints and statistics. Since the noisy input and the clean image are well aligned at image features, CRFs, in particular, GCRFs can be well applied to image denoising. Preliminary success has been shown in the denoising work in [48]. Learning GCRFs is also addressed in [50].

2.2 Noise Estimation

Image-dependent noise can be estimated from multiple images or a single image. Estimation from multiple images is an overconstrained problem and was addressed in [25]. Estimation from a single image, however, is an underconstrained problem and further assumptions have to be made for the noise. In the image denoising literature, noise is often assumed to be additive white Gaussian noise (AWGN). A widely used estimation method is based on the mean absolute deviation [10]. In [17], the noise level is estimated from the gradient of smooth or small textured regions, and the signal-dependent noise level is estimated for each intensity interval. In [46], Stefano et al. proposed three methods to estimate noise levels based on training samples and the statistics (Laplacian) of natural images. In [37], a generalized expectation maximization algorithm is proposed to estimate the spectral features of a noise source corrupting an observed image.

Techniques for noise estimation followed by noise reduction have been proposed, but they tend to be heuristic. For example, in [22], a set of statistics of film grain noise are used to estimate and remove the noise produced from scanning the photographic element under uniform exposures. In [45], signal-dependent noise is estimated from the smooth regions of the image by segmenting the image gradient with an adaptive threshold. The estimated signal-dependent noise is applied to the whole image for noise reduction. This work was further extended in [21] by associating a default film-related noise model to the image based on its source identification tag. The noise model is then adjusted using the image statistics. In certain commercially available image enhancement software, such as Neat Image™,[1] the noise level can be semiautomatically estimated by specifying featureless areas to profile noise. Neat Image™ also provides calibration tools to estimate the amount of noise for a specific camera and camera setting; precalibrated noise profiles for various cameras are also available to directly denoise images.

By comparison, our technique avoids the tedious noise measurement process for each camera used. Furthermore,

---

1. http://www.neatimage.com.

our technique provides a principled way for estimating a continuous NLF from a single image under the Bayesian inference framework.

3 PIECEWISE SMOOTH IMAGE MODEL

The piecewise smooth image model was used by Terzopoulos [51] to account for the regularization of the reconstructed image. The concept of piecewise smooth (or continuous) was elaborated by Blake and Zisserman [5]. In this section, we discuss the reconstruction of piecewise smooth image model from an image and some important properties of this model.

3.1 Image Segmentation

Image segmentation algorithms are designed based on piecewise smooth image prior to partition pixels into regions with both similar spatial coordinates and RGB pixel values. There are a variety of segmentation algorithms, including mean shift [8] and graph-based methods [15]. Since the focus of this paper is not on segmentation, we choose a simple K-Means clustering method for grouping pixels into regions, as described in [57]. Each segment is represented by a mean color and spatial extent. The spatial extent is computed so that the shape of the segment is biased toward convex shapes and that all segments have similar size.

3.2 Segment Statistics and Affine Reconstruction

Let the image lattice be $\mathcal{L}$. It is completely partitioned to a number of regions $\{\Omega_i\}$, where $\mathcal{L} = \bigcup_i \Omega_i$ and $\Omega_i \cap \Omega_j = \emptyset$ for $i \neq j$. Let $v \in \mathbb{R}^2$ be the coordinate variable, and $I(v) \in \mathbb{R}^3$ be the RGB value of the pixel. Since in this section we focus on the statistics within each segment, we shall use Ω to represent a segment and $v \in \Omega$ to index pixels in segment Ω.

We can fit an affine model in segment Ω to minimize the squared error:

$$\mathbf{A}^* = \arg\min_{\mathbf{A}} \sum_{v \in \Omega} \left\| I(v) - \mathbf{A}[v^T \ 1]^T \right\|^2, \qquad (1)$$

where $\mathbf{A} \in \mathbb{R}^{3 \times 3}$ is the affine matrix. We call the reconstruction $f(v) = \mathbf{A}^*[v^T \ 1]^T$ the *affine reconstruction* of segment Ω. The residual is $r(v) = I(v) - f(v)$.

We assume that the residual consists of two parts: subtle texture variation $h(v)$, which is also part of signal, and noise $n(v)$, that is, $r(v) = h(v) + n(v)$. In other words, the observed image can be decomposed into $I(v) = f(v) + h(v) + n(v)$. The underlying clean image or signal is thus $s(v) = f(v) + h(v)$, which is to be estimated from the noisy input. $s(v)$, $h(v)$, and $n(v)$ are all 3D vectors in RGB space.

Let the covariance matrices of $I(v)$, $s(v)$, $h(v)$, and $n(v)$ be Σ_I, Σ_s, Σ_h, and Σ_n, respectively. We assume that $f(v)$ is a nonrandom process, and $r(v)$ and $n(v)$ are random variables. Therefore, $\Sigma_s = \Sigma_h$. Suppose signal $s(v)$, and noise $n(v)$ are independent, we have

$$\Sigma_r = \Sigma_s + \Sigma_n, \qquad (2)$$

which leads to

$$\Sigma_r \geq \Sigma_n \qquad (3)$$

that is, the variance of the residual bounds the noise variance.

Fig. 1. Illustration of piecewise smooth image model. (a) Original image. (b) Segmentation. (c) Per-segment affine reconstruction. (d) Affine reconstruction plus boundary blur. (e) The sorted eigenvalues in each segment. (f) RGB values projected onto the largest eigenvector.

3.3 Boundary Blur Estimation

If we merely use per-segment affine reconstruction, the reconstructed image has artificial boundaries, and the original boundaries would be artificially sharpened. To avoid that, we estimate the blur from the original image in the following way. For each hypothesized blur b from $b_{\min}(=0)$ to $b_{\max}(=2.5)$ in steps of $\Delta b(=0.25)$, we compute the blurred image $f_{\mathrm{blur}}(v;b) = f(v) * G(u;b)$, where $G(u;b)$ is a Gaussian kernel with sigma b. We then compute the error image I_{err} such that $I_{\mathrm{err}}(v;b) = \|I(v) - f_{\mathrm{blur}}(v;b)\|^2$. We dilate each boundary curve $\mathcal{C}_{ij}$ five times into regions Ω_i and Ω_j to obtain a mask Γ_{ij}. The best blur b_{ij}^* for $\mathcal{C}_{ij}$ corresponds to the minimum aggregate error $I_{err}(v;b)$ over Γ_{ij} or $b_{ij}^* = \arg\min_b \sum_{v \in \Gamma_{ij}} I_{err}(v,b)$.

To reinstate the blur in the transition region Γ_{ij}, we simply replace $f(v)$ with $f_{\mathrm{blur}}(v;b_{ij}^*)$. Note that this assumes that the amount of blur in Ω_i and Ω_j is the same, which is strictly not true in general. However, we found that this approximation generates satisfactory results. After this process is done for every pair of regions, we obtain boundary blurred piecewise affine reconstruction $f_{\mathrm{blur}}(v)$.

The piecewise smooth image model is illustrated in Figs. 1a, 1b, 1c, and 1d. The example image (Fig. 1a) taken from Berkeley image segmentation database [32] is partitioned to piecewise smooth regions (Fig. 1b) by the segmentation algorithm. The per-segment affine reconstruction is shown in Fig. 1c, where we can see artificial boundaries between regions and the true boundaries are sharpened. After blur estimation and reinstatement in Fig. 1d, the boundaries become much smoother.

3.4 Important Properties of the Piecewise Smooth Image Model

There are three important properties of our piecewise smooth image model that led us to choose it as the model for both noise estimation and removal. They are

1. the piecewise smooth image model is consistent with a sparse image prior;
2. the color distribution per each segment can be well approximated by a line segment, due to the physics of image formation [26], [24], [20]; and
3. the standard deviation of residual per each segment is the upper bound of the noise level in that segment.

The last property follows from (3). For the first two properties, we again use the example image in Fig. 1a to examine them. For the reconstructed image (Fig. 1d), we compute the log histograms of the horizontal and vertical derivatives and plotted them in Fig. 2. The long tails clearly show that the piecewise smooth reconstruction match the high-kurtosis statistics of natural images [33]. This image model also shares some similarity with the so-called dead leaves model [29].

For the second property, we compute the eigenvalues and eigenvectors of the RGB values $\{I(v)\}$ in each region. The eigenvalues are sorted in descending order and displayed in Fig. 1e. Obviously, the red channel accounts for the majority of the RGB channels, a fact that proves the first eigenvalue of each segment is significantly larger than the second eigenvalue. Therefore, when we project the pixel values onto the first


Fig. 2. The log histograms of the (a) horizontal and (b) vertical derivative filter responses of the reconstruction in Fig. 1d.
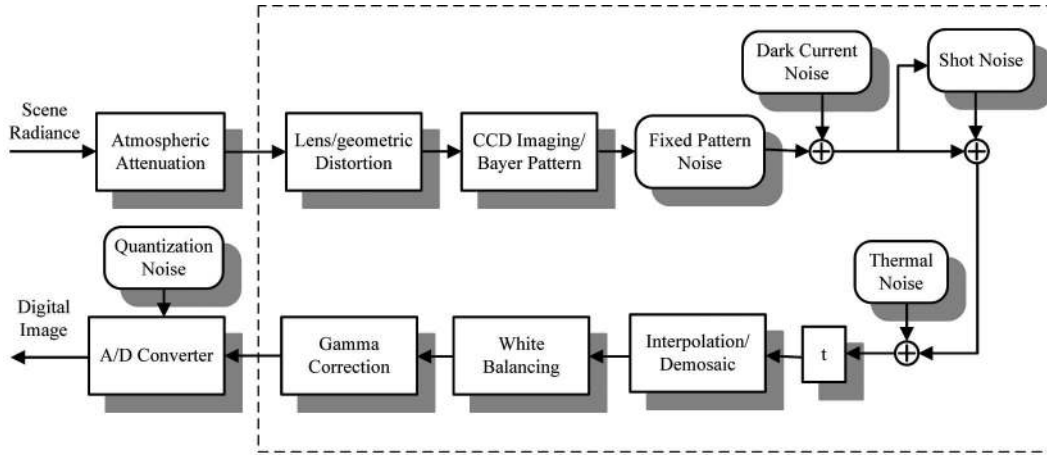
Fig. 3. CCD camera imaging pipeline, redrawn from [54].

eigenvalue while ignoring the other two, we get an almost perfect reconstruction in Fig. 1f. The mean square error (MSE) between the original image (Fig. 1a) and projected (Fig. 1f) is 5.31×10^{-18} or a peak signal to noise ratio (PSNR) of 35.12 dB. These numbers demonstrate that the RGB values in each segment lie in a line.

Having demonstrated these properties of our piecewise smooth image model, we are ready to develop models for both noise estimation and removal.

4 NOISE ESTIMATION FROM A SINGLE IMAGE

Although the standard deviation of each segment of the image is the upper bound of noise, as shown in (3), it is not guaranteed that the means of the segments cover the full range of image intensities and may not give a bound on the noise for all image intensities. Besides, the estimate of standard deviation itself is also a random variable, which has variance as well. Therefore, a rigorous statistical framework is needed for the inference. In this section, we introduce the noise level functions (NLFs) and a simulation approach to learn the priors. A Bayesian approach is proposed to infer the upper bound of the noise level function from a single input.

4.1 Learning a Prior for Noise Level Functions

The noise standard deviation as function of brightness is called the *noise level function* (NLF). For a particular brand of camera and a fixed parameter setting, the NLF can be estimated by fixing the camera on a tripod, taking multiple shots toward a static scene and then computing the mean as the estimate of the brightness and the standard deviation as the noise level for each pixel of every RGB channel. The standard deviation as a function of the mean intensity is the desired NLF. We shall use such an experimentally measured NLF as the reference method to test our algorithm, although it is expensive and time consuming.

As an alternative, we propose a simulation-based approach to obtain NLFs. We build a model for the noise level functions of (CCD) cameras. We introduce the terms of our camera noise model, showing the dependence of the NLF on the camera response function (CR function, the image brightness as function of scene irradiance). Given a CR function, we can synthesize realistic camera noise. Thus, from a parameterized set of CR functions, we derive the set of possible NLFs. This restricted class of NLFs allows us to

accurately estimate the NLF for image intensities not observed in the image.

4.1.1 Noise Model of a CCD Camera

A CCD digital camera converts the irradiance, the photons coming into the imaging sensor, to electrons and, finally, to bits. See Fig. 3 for the imaging pipeline of a CCD camera. There are five primary noise sources as stated in [25], namely, *fixed pattern noise*, *dark current noise*, *shot noise*, *amplifier noise*, and *quantization noise*. These noise terms are simplified in [54]. Following the imaging equation in [54], we propose the following noise model of a CCD camera:

$$I = f(L + n_s + n_c) + n_q, \qquad (4)$$

where I is the observed image brightness, $f(\cdot)$ is CRF, n_s accounts for all the noise components that are dependent on irradiance L, n_c accounts for the independent noise before gamma correction, and n_q is additional quantization and amplification noise. Since n_q is the minimum noise attached to the camera, and most cameras can achieve very low noise, n_q will be ignored in our model. We assume noise statistics $\mathrm{E}(n_s) = 0$, $\mathrm{Var}(n_s) = L\sigma_s^2$, and $\mathrm{E}(n_c) = 0$, $\mathrm{Var}(n_c) = \sigma_c^2$. Note the linear dependence of the variance of n_s on the irradiance L [54].

4.1.2 Camera Response Function

The camera response function models the nonlinear processes in a CCD camera that perform tonescale (gamma) and white balance correction [42]. There are many ways to estimate CR functions given a set of images taken under different exposures. To explore the common properties of many different CR functions, we downloaded 201 real-world response functions from http://www.cs.columbia.edu/ CAVE [23]. Note that we chose only 190 saturated CR functions since the unsaturated curves are mostly synthetic. Each CR function is a 1,024-dimensional vector that represents the discretized $[0, 1] \rightarrow [0, 1]$ function, where both irradiance L and brightness I are normalized to be in the range [0, 1]. We use the notation crf(i) to represent the ith function in the database.

4.1.3 Synthetic CCD Noise

In principle, we could set up optical experiments to measure precisely for each camera how the noise level changes with image brightness. However, this would be time consuming
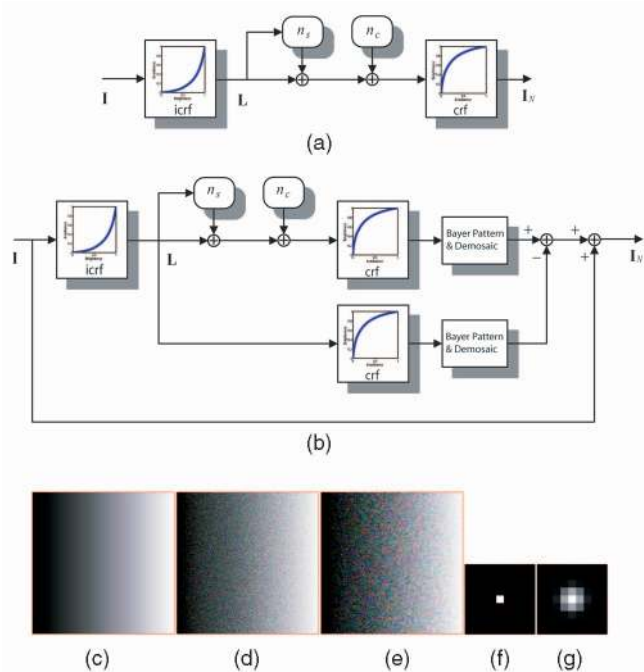
Fig. 4. Block diagrams showing noise simulations for color camera images. (a) shows independent white noise synthesis and (b) adds CCD color filter pattern sensing and demosaicing to model spatial correlations in the camera noise [28]. (c) Test pattern. (d) and (e) The synthesized images of (a) and (b). (f) and (g) The corresponding autocorrelation.

and might still not adequately sample the space of camera response functions. Instead, we use numerical simulation to estimate the noise function. The basic idea is to transform the image I by the inverse camera response function f^{-1} to obtain an irradiance plane L. We then take L through the processing blocks in Fig. 3 to obtain the noisy image I_N.

A direct way from (4) is to reverse transform I to irradiance L, add noise independently to each pixel, and transform to brightness to obtain I_N. This process is shown in Fig. 4a. The synthesized noise image, for the test pattern (Fig. 4c), is shown in Fig. 4d.

Real CCD noise is not white; however, there are spatial correlations introduced by "demosaicing" [39], that is, the reconstruction of three colors at every pixel from the single-color samples measured under the color filter array of the CCD. We simulate this effect for a common color filter pattern (Bayer) and demosaicing algorithm (gradient-based interpolation [28]); we expect that other filter patterns and demosaicing algorithms will give comparable noise spatial correlations. We synthesized CCD camera noise in accordance with Fig. 4b and took the difference between the demosaiced images with and without noise, adding that to the original image to synthesize CCD noise. The synthesized noise is shown in Fig. 4e. The autocorrelation functions for noise images (Figs. 4d and 4e) are shown in Figs. 4f and 4g, respectively, showing that the simulated CCD noise exhibits spatial correlations after taking into account the effects of demosaicing.

4.1.4 The Space of Noise Level Functions

Recall that NLF is defined as the variation of the standard deviation of noise with respect to image intensity. This function can be computed as
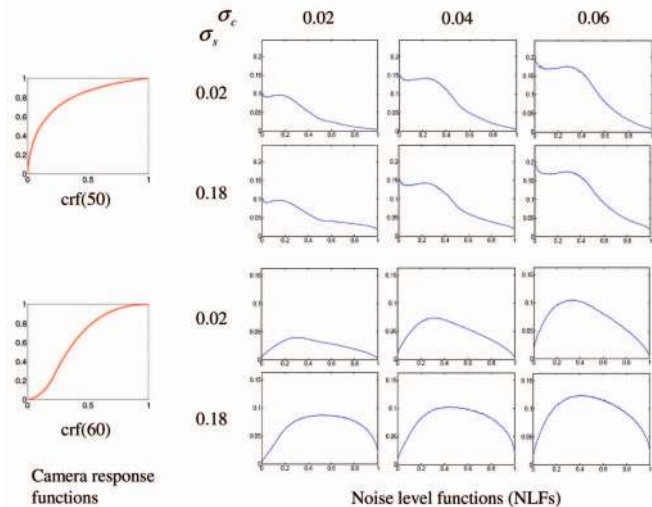

Fig. 5. The procedure of sampling noise level functions.

$$\tau(I) = \sqrt{\mathbb{E}[(I_N - I)^2]}, \qquad (5)$$

where I_N is the noise contaminated observation, and $I = \mathbb{E}(I_N)$. $\mathbb{E}(\cdot)$ is the expectation over underlying noise process.

Based on the CCD camera noise model (4) and noise synthesis process, I_N is a random variable dependent on the CRF f and noise parameters σ_s and σ_c. Because $L = f^{-1}(I)$, the NLF can also be written as

$$\tau(I; f, \sigma_s, \sigma_c) = \sqrt{\mathbb{E}[(I_N(f^{-1}(I), f, \sigma_s, \sigma_c) - I)^2]}, \qquad (6)$$

where $I_N(\cdot)$ is defined as noise synthesis process.

We use numerical simulation to estimate the noise function given f, σ_s, and σ_c, for each of red, green, and blue channels. This procedure is shown in Fig. 5. The smoothly changing pattern in Fig. 4c is used to estimate (6). To reduce statistical fluctuations, we use an image of dimension $1,024 \times 1,024$ and take the mean of 20 samples for each estimate.

To represent the whole space of noise level functions, we draw samples of $\tau(\cdot; f, \sigma_s, \sigma_c)$ from the space of f, σ_s, and σ_c. The downloaded 190 CR functions are used to represent the space of f. We found that $\sigma_s = 0.16$ and $\sigma_c = 0.06$ result in very high noise, so these two values are set as the maximum of the two parameters. We sample σ_s from 0.00 to 0.16 with step size 0.02 and sample σ_c from 0.01 to 0.06 with step size 0.01. We get a dense set of samples $\{\tau_i\}_{i=1}^{K}$ of NLFs, where $K = 190 \times 9 \times 6 = 10,260$. Using principal component analysis (PCA), we obtain the mean NLF $\overline{\tau}$, eigenvectors $\{w_i\}_{i=1}^{m}$, and the corresponding eigenvalues $\{v_i\}_{i=1}^{m}$. Thus, a noise function can be represented as

$$\tau = \overline{\tau} + \sum_{i=1}^{m} \beta_i w_i, \qquad (7)$$

where coefficient β_i is Gaussian distributed $\beta_i \sim \mathcal{N}(0, v_i)$, and the function must be positive everywhere, that is,

$$\overline{\tau} + \sum_{i=1}^{m} \beta_i w_i \geq 0, \qquad (8)$$

where $\overline{\tau}$, $w_i \in \mathbb{R}^d$, and $d = 256$. This inequality constraint implies that noise functions lie inside a cone in β space. The
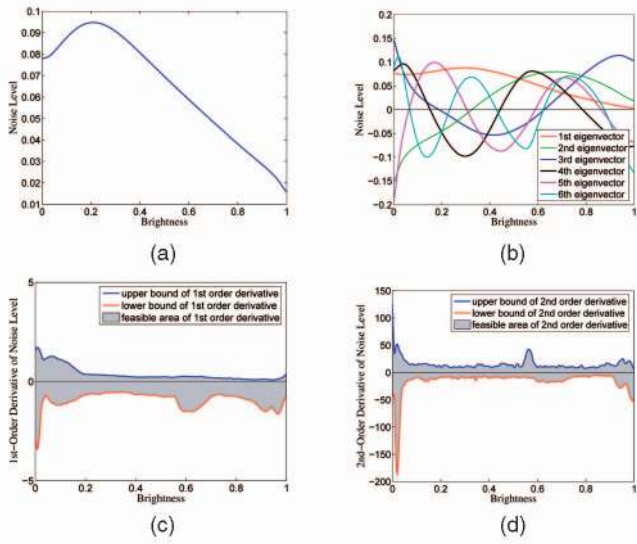
Fig. 6. The prior of noise level functions. (a) Mean of noise functions. (b) Eigenvectors of noise functions. (c) The bounds of first-order derivatives. (d) The bounds of second-order derivatives.

mean noise function and eigenvectors are displayed in Figs. 6a and 6b, respectively.

Eigenvectors serve as basis functions to impose smoothness to the function. We also impose upper and lower bounds on first and second-order derivatives to further constrain noise functions. Let $\mathbf{T} \in \mathbb{R}^{(d-1)\times d}$ and $\mathbf{K} \in \mathbb{R}^{(d-2)\times d}$ be the matrix of first and second-order derivatives [47]. The constraints can be represented as

$$b_{min} \leq \mathbf{T}\tau \leq b_{max}, \quad h_{min} \leq \mathbf{K}\tau \leq h_{max}, \tag{9}$$

where b_{min}, $b_{max} \in \mathbb{R}^{d-1}$, h_{min}, and $h_{max} \in \mathbb{R}^{d-2}$ are estimated from the training data set $\{\tau_i\}_{i=1}^{K}$.

4.2 Likelihood Model
Since the estimated standard deviation of each segment is an overestimate of the noise level, we obtain an upper bound estimate of the noise level function by fitting a lower envelope to the samples of standard deviation versus mean of each RGB channel. The examples of these sample points are shown in Fig. 8. We could simply fit the noise function in the learned space to lie below all the sample points yet close to them. However, because the estimates of variance in each segment are noisy, extracting these estimates with hard constraints could result in bias due to a bad outlier. Instead, we follow a probabilistic inference framework to let every data point contribute to the estimation.

Let the estimated standard deviation of noise from k pixels be $\hat{\sigma}$, with σ being the true standard deviation. When k is large, the square root of the chi-square distribution is approximately $\mathcal{N}(0, \sigma^2/k)$ [13]. In addition, we assume a noninformative prior for large k and obtain the posterior of the true standard deviation σ given $\hat{\sigma}$:

$$p(\sigma|\hat{\sigma}) \propto \frac{1}{\sqrt{2\pi\sigma^2/k}}\exp\left\{-\frac{(\hat{\sigma}-\sigma)^2}{2\sigma^2/k}\right\}$$
$$\approx \frac{1}{\sqrt{2\pi\hat{\sigma}^2/k}}\exp\left\{-\frac{(\sigma-\hat{\sigma})^2}{2\hat{\sigma}^2/k}\right\}. \tag{10}$$
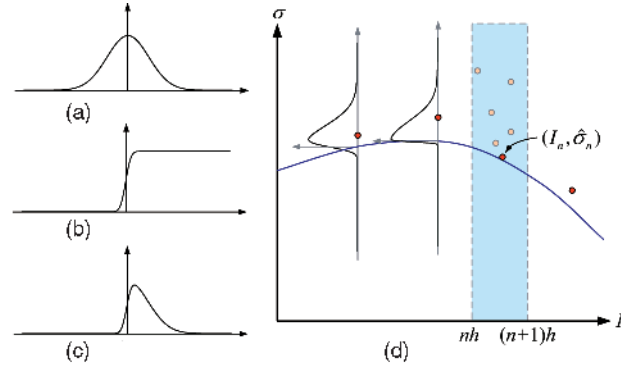

Fig. 7. The likelihood function of (12). Each single likelihood function (c) is a product of a Gaussian pdf (a) and Gaussian cdf (b).

Let the cumulative distribution function of a standard normal distribution be $\Phi(z)$. Then, given the estimate $(I, \hat{\sigma})$, the probability that the underlying standard deviation σ is larger than u is

$$\Pr[\sigma \geq u|\hat{\sigma}] = \int_u^\infty p(\sigma|\hat{\sigma})d\sigma = \Phi\left(\frac{\sqrt{k}(\hat{\sigma}-u)}{\hat{\sigma}}\right). \tag{11}$$

To fit the noise level function to the lower envelope of the samples, we discretize the range of brightness [0, 1] into uniform intervals $\{nh, (n+1)h\}_{n=0}^{\frac{1}{h}-1}$. We denote the set $\Omega_n = \{(I_i, \hat{\sigma}_i)|nh \leq I_i \leq (n+1)h\}$ and find the pair $(I_n, \hat{\sigma}_n)$ with the minimum variance $\hat{\sigma}_n = \min_{\Omega_n} \hat{\sigma}_i$. The fitted lower envelope should most probably be lower than all the estimates while being as close as possible to the samples. Mathematically, the likelihood function is the probability of seeing the observed image intensity and noise variance measurements given a particular noise level function. It is formulated as

$$\mathcal{L}(\tau(I)) = P(\{I_n, \hat{\sigma}_n\}|\tau(I))$$
$$\propto \prod_n \Pr[\sigma_n \geq \tau(I_n)|\hat{\sigma}_n]\exp\left\{-\frac{(\tau(I_n)-\hat{\sigma}_n)^2}{2s^2}\right\}$$
$$= \prod_n \Phi\left(\frac{\sqrt{k_n}(\hat{\sigma}_n-\tau(I_n))}{\hat{\sigma}_n}\right)\exp\left\{-\frac{(\tau(I_n)-\hat{\sigma}_n)^2}{2s^2}\right\}, \tag{12}$$

where s is the parameter to control how close the function should approach the samples. This likelihood function is illustrated in Fig. 7, where each term (Fig. 7c) is a product of a Gaussian probability density function (pdf) with variance s^2 (Fig. 7a) and a Gaussian cdf with variance $\hat{\sigma}_n^2$ (Fig. 7b). The red dots are the samples of the minimum in each interval. Given the function (blue curve), each red dot is probabilistically beyond but close to the curve with the pdf in Fig. 7c.

4.3 Bayesian MAP Inference
The parameters we want to infer are actually the coefficients on the eigenvectors $x_l = [\beta_1 \cdots \beta_m]^T \in \mathbb{R}^m$, $l = 1, 2, 3$ of the noise level function for RGB channels. Let the sample set to fit be $\{(I_{ln}, \hat{\sigma}_{ln}, k_{ln})\}$. Bayesian MAP inference, finding the most probable point from the posterior, is the optimization problem

$$\{x_l^*\} = \arg\min_{\{x_l\}} \sum_{l=1}^{3} \left\{ \sum_n \left[-\log \Phi\left(\frac{\sqrt{k_{ln}}}{\hat{\sigma}_n} (\hat{\sigma}_{ln} - e_n^T x_l - \overline{\tau}_n) \right) \right. \right.$$

$$\left. + \frac{(e_n^T x_l + \overline{\tau}_n - \hat{\sigma}_{ln})^2}{2s^2} \right] + x_l^T \Lambda^{-1} x_l$$

$$\left. + \sum_{j=1, j>l}^{3} (x_l - x_j)^T \mathbf{E}^T (\gamma_1 \mathbf{T}^T \mathbf{T} + \gamma_2 \mathbf{K}^T \mathbf{K}) \mathbf{E}(x_l - x_j) \right\}$$

$$(13)$$

subject to

$$\overline{\tau} + \mathbf{E}x_l \geq 0, \tag{14}$$

$$b_{min} \leq \mathbf{T}(\overline{\tau} + \mathbf{E}x_l) \leq b_{max}, \tag{15}$$

$$h_{min} \leq \mathbf{K}(\overline{\tau} + \mathbf{E}x_l) \leq h_{max}. \tag{16}$$

In the above formula, the matrix $\mathbf{E} = [w_1 \cdots w_m] \in \mathbb{R}^{d \times m}$ contains the principal components, e_n is the nth row of $\mathbf{E}$, and $\Lambda = \text{diag}(v_1, \cdots, v_m)$ is the diagonal eigenvalue matrix. The last term in the objective function accounts for the similarity of the NLF for RGB channels. Their similarity is defined as a distance on the first and second-order derivative. Since the dimensionality of the optimization is low, we use the Matlab standard nonlinear constrained optimization function `fmincon` for optimization. The function was able to find a local optimal solution for all the examples we tested.

4.4 Experimental Results on Noise Estimation

We have conducted experiments on both synthetic and real noisy images to test the proposed noise estimation algorithm. First, we applied our CCD noise synthesis algorithm in Section 3.3 to 17 randomly selected pictures from the Berkeley image segmentation database [32] to generate synthetic test

TABLE 1
The Statistics of the L^2 and L^∞ Norms between the
Estimated NLF and the Ground Truth

Norm	mean	std. deviation
L^2	0.0048	0.0033
L^∞	0.0110	0.0120

images. To generate the synthetic CCD noise, we specified a CR function and two parameters σ_s and σ_c. From this information, we also produced the ground truth NLF using the training database in Section 4.1.4. For this experiment, we selected crf(60), $\sigma_s = 0.10$, and $\sigma_c = 0.04$. Then, we applied our method to estimate the NLF from the synthesized noisy images. Both L^2 and L^∞ norms are used to measure the distance between the estimated NLF and the ground truth. The error statistics under the two norms are listed in Table 1, where the mean and maximum value of the ground truth are 0.0645 and 0.0932, respectively.

Some estimated NLFs are shown in Fig. 8. In Fig. 8a, we observe many texture regions especially at high-intensity values, which implies high-signal variance. The estimated curves (in red, green, and blue) do not tightly follow the lower envelope of the samples at high intensities, although they deviate from the true noise function (in gray) slightly. In Fig. 8b, the samples do not span the full intensity range, so our estimate is only reliable where the samples appear. This shows a limit of the prior model: the samples are assumed to be well distributed. The estimation is reliable if the color distribution spans the full range of the spectrum and there are textureless regions, as in Fig. 8c.

We conducted a further experiment as a sanity check. We took 29 images of a static scene by Canon™ EOS 10D (ISO
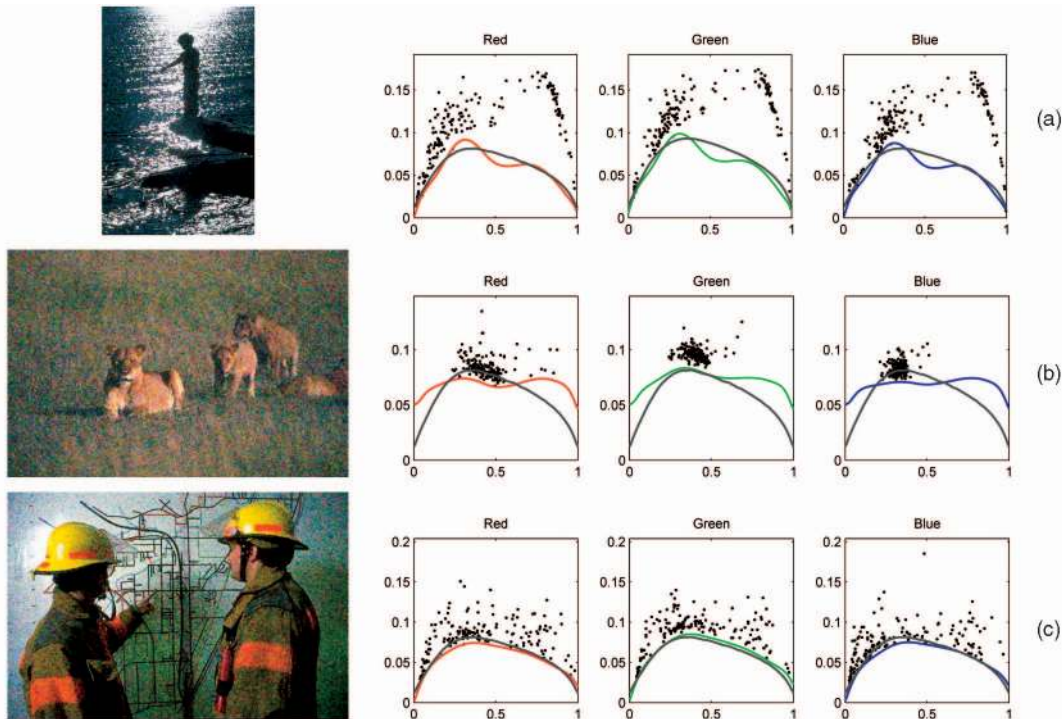

Fig. 8. Synthesized noisy images and their corresponding NLFs (noise standard deviation as a function of image brightness). The red, green, and blue curves are estimated using the proposed algorithm, whereas the gray curves are the true values for the synthetically generated noise.
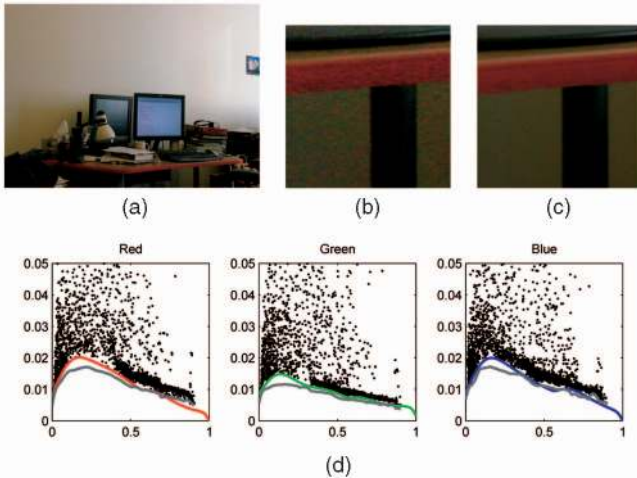
Fig. 9. Comparison of estimated camera noise with experimental measurement. (a) shows one of the 29 images taken with a Canon™ EOS 10D. An enlarged patch is shown for (b) a single image and (c) the mean image. (d) is the estimated NLF from a single image (color), showing good agreement with the ground truth (gray), measured from the noise variations over all 29 images.

1600, exposure time 1/30 second and aperture f/19) and computed the mean image. One sample is shown in Fig. 9a. A close-up view of sample (Fig. 9a) and the mean image is shown in Figs. 9b and 9c, respectively. Clearly, the noise is significantly reduced in the mean image. Using the sample variance over the 29 images as a function of mean intensity, we calculated the experimentally measured NLF and compared that to the NLF estimated by our method from only one image. The agreement between the NLFs in each color band is very good, see Fig. 9d.

We also applied the algorithm to estimating NLFs from the other images taken by a CCD camera. We evaluated our results based on repeatability: pictures taken by the same camera with the same setting on the same date should have the same NLF, independent of the image content. We collected two pictures taken by a Canon™ EOS DIGITAL REBEL and estimated the corresponding NLFs, as shown in Figs. 10a and 10b. Even though image Fig. 10a is missing high intensity values, the estimated NLFs estimated from the two images are similar, showing the ability of our Bayesian

method to extrapolate the NLF beyond observed image intensities.

5 SEGMENTATION-BASED DENOISING

Recall from Section 3.2 that the observation $I(v)$ is decomposed into signal $s(v)$ and noise $n(v)$. Given the characteristics of the noise that have been estimated from the previous section, we are now ready to extend the signal from the noisy observation.

5.1 Zeroth-Order Model

Let $\mu = [\mu_1 \mu_2 \mu_3]^T \in \mathbb{R}^3$ be the mean color for segment Ω after the piecewise smooth image reconstruction to the input image I. Suppose the noise is independent for RGB channels, and we obtain the covariance matrix of noise in this segment:

$$\hat{\mathbf{\Sigma}}_n = \text{diag}\big(\tau^2(\mu_1), \tau^2(\mu_2), \tau^2(\mu_3)\big). \quad (17)$$

From the independence assumption of the noise and signal, we obtain (from (2))

$$\hat{\mathbf{\Sigma}}_s = \mathbf{\Sigma}_r - \hat{\mathbf{\Sigma}}_n. \quad (18)$$

It is possible that the estimated $\hat{\Sigma}_s$ is not positive definite. For this case, we simply enforce the minimum eigenvalue of $\hat{\Sigma}_s$ to be a small value (0.0001).

We simply run Bayesian MAP estimation for each pixel to estimate the noise based on the obtained second-order statistics. Since

$$p(s(v)|I(v)) \propto p(I(v)|s(v))p(s(v))$$
$$\propto \exp\left\{-\frac{1}{2}[I(v) - s(v)]^T \hat{\mathbf{\Sigma}}_n^{-1}[I(v) - s(v)]\right\} \quad (19)$$
$$\exp\left\{-\frac{1}{2}[s(v) - f(v)]^T \hat{\mathbf{\Sigma}}_s^{-1}[s(v) - f(v)]\right\},$$

where $f(v)$ is the piecewise smooth reconstruction, the optimal estimation has a simple closed-form solution:

$$s^*(v) = \arg\max p(s(v)|I(v))$$
$$= (\hat{\mathbf{\Sigma}}_n^{-1} + \hat{\mathbf{\Sigma}}_s^{-1})^{-1}(\hat{\mathbf{\Sigma}}_n^{-1}I(v) + \hat{\mathbf{\Sigma}}_s^{-1}f(v)). \quad (20)$$
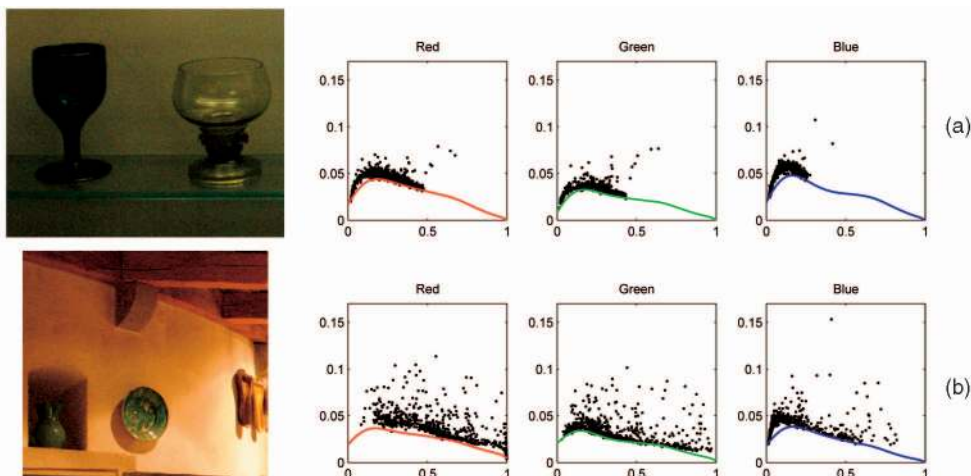

Fig. 10. The two images are taken by a Canon™ EOS DIGITAL REBEL and the estimated noise level functions. Very similar noise level functions are derived, even though the two images have very different tonescales.

Fig. 11. Noise contaminated images and the denoised results by the zeroth-order model. A patch at a fixed position marked by a red rectangle is zoomed-in and inset at the bottom right of each image. Clearly, the zeroth-order model significantly removes the chrominance component of color noise. For high noise level, the discontinuities between the neighboring segments are further removed by the first-order model (see Figs. 13 and 14 and Table 2). (a) Five percent AWGN. (b) Denoised by the zeroth-order model ($\mathrm{PSNR} = 32.04$). (c) Ten percent AWGN. (d) Denoised by the zeroth-order model ($\mathrm{PSNR} = 27.05$).

This simply down weighs the pixel values from $I(v)$ to $f(v)$ using the covariance matrices as weights. For a scaled identity $\hat{\Sigma}_n$, it is easy to show that the attenuation along each principal direction in the color covariance matrix is $\lambda_i/(\lambda_i + \sigma_n)$, where λ_i is the variance in the ith direction. Qualitatively, as this variance tends toward zero (either because the nondominant direction has low variance or the region is untextured), the cleaned up residual is progressively more attenuated.

Equation (20) is applied to every pixel, where $\hat{\Sigma}_n$ and $\hat{\Sigma}_s$ vary from segment to segment. Since there is no spatial relationship of pixels in this model, we call it *zeroth-order model*. An example of denoising using zeroth-order model is shown in Fig. 11, where the algorithm is tested by synthetic AWGN with noise levels of 5 percent and 10 percent. Clearly, the zeroth-order model significantly removes the chrominance component of color noise. The results are acceptable for 5 percent noise level, and we can see discontinuities between the neighboring segments for 10 percent noise level because the spatial correlation has been accounted for.

5.2 First-Order Gaussian Conditional Random Field

The values of the neighboring pixels are correlated in natural images. We chose to regularize with a conditional random field (CRF) [27], [48], where the spatial correlation is a function of the local patch of the input image, over the Markov random field (MRF) [19] to avoid having a complete prior model on images as in the approach of [41], Moreover, we model it as a Gaussian CRF since all the energy functions are quadratic. We call it *first-order model* because the spatial correlation is captured by the first-order derivative filters. Likewise, we can have a second-order model or even higher order. However, we found that first-order model is sufficient for the denoising task.

Let the estimated covariance matrices of signal and noise be $\hat{\Sigma}_s(i)$ and $\hat{\Sigma}_n(i)$ for segment Ω_i. The CRF is formulated as

$$p(s|I) =$$
$$\frac{1}{Z}\exp\left\{ -\frac{1}{2}\sum_i \sum_{v\in\Omega_i} \left[\left(s(v) - I(v)\right)^T \hat{\Sigma}_n^{-1}(i)\left(s(v) - I(v)\right) + \right.\right.$$
$$\left.\left. \left(s(v) - f(v)\right)^T \hat{\Sigma}_s^{-1}(i)\left(s(v) - f(v)\right) + \xi_i w(v) \sum_{j=1}^{m} F_j^2(v) \right] \right\}.$$

$$(21)$$

In the above equation, $F_j = \phi_j * s$, the convolution of s with filter ϕ_j. For this first-order GCRF, we choose horizontal and vertical derivative filters for ϕ_j (that is, $m = 2$). $w(v)$ and ξ_i are both weights to balance the importance of spatial correlation. ξ_i is the weight for each segment. We find that ξ_i can be a linear function of the mean noise level in segment Ω_i. $w(v)$ is derived from the filter responses of the original image. Intuitively, $w(v)$ should be small when there is clear boundary at v to weaken spatial correlation between pixels and be large when there is no boundary to strengthen that spatial correlation. Boundaries can be detected by Canny edge detection [7], but we found that the algorithm is more stable when $w(v)$ is set to be a function of local filter responses. We use orientationally elongated Gabor sine and cosine filters [18] to capture the boundary energy of the underlying noise-free image. The boundary energy is the sum over all the orientations and sin/cos phases. We then use a nonlinear function to map the energy to the local value of the weight matrix, that is, $y = (1 - \tanh(t_1 x))^{t_2}$, where $t_1 = 0.6$ and $t_2 = 12$ in our implementation. Solving (21) is equivalent to solving a linear system, which can be
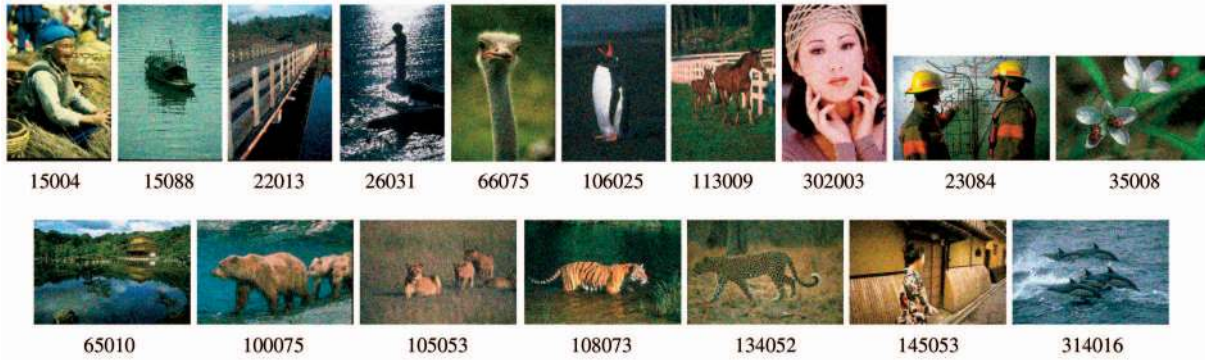
Fig. 12. Seventeen images are selected from Berkeley image segmentation database [32] to evaluate the proposed algorithm. The file names (numbers) are shown beneath each picture.

effectively computed by conjugate gradient method that only requires iterative filtering [48], [50].

6 EXPERIMENTAL RESULTS ON IMAGE DENOISING

Our automatic image denoising system consists of two parts, noise estimation and denoising. To have a fair comparison with other denoising algorithms, we first test our denoising algorithms using synthetic AWGN with constant and known noise level (σ). Then the whole system is tested with the images contaminated with real CCD camera noise.

6.1 Synthetic AWGN

We selected 17 images covering different types of objects and scenes from the Berkeley segmentation data set [32] and added AWGN with 5 percent and 10 percent noise level to test our denoising algorithm. The noise-contaminated images with 10 percent noise level are shown in Fig. 12. We also ran standard bilateral filtering [52] (our implementation), curvature preserving PDE [53] (publicly available implementation),[2] and wavelet joint coring, GSM [38] (publicly available implementation).[3] Default parameter settings are used for the downloaded code. For curvature preserving PDE, we tweaked the parameters and found that the best results can be obtained by setting $alpha = 1$ and $iter = 4$ for $\sigma = 10\%$ and $alpha = 0.5$, $iter = 7$ for $\sigma = 5\%$. We compare the other methods to our own using both visual inspection in Figs. 13 and 14, and peak signal to noise ratio (PSNR) statistics in Table 2.

It is clear that our technique consistently outperforms bilateral filtering, curvature preserving PDE, and wavelet joint coring. In terms of PSNR, our technique outperforms these algorithms by a significant margin. When $\sigma = 0.05$, that is, the noise level is low, even the zeroth-order model outperforms the state-of-the-art wavelet GSM. When $\sigma = 0.10$, that is, the noise level is high, the first-order model outperforms wavelet GSM by 1.3 PSNR on the average.

The results are also visually inspected in Figs. 13a, 13b, 13c, 13d, and 13e, corresponding to image 35008, 23084, 108073, 65010, and 66075 in Fig. 12, respectively. Some close-up views of the denoising results are shown in Fig. 14. The curvature preserving PDE method generates color-fringing artifacts around the strong edges. Wavelet coring tends to produce

2. http://www.greyc.ensicaen.fr/~dtschump/greycstoration/download.html.
3. http://decsai.ugr.es/~javier/denoise/.

color and blurring artifacts, especially in Figs. 14a and 14d. Our algorithm, however, is able to smooth out flat regions, preserve sharp edges, as well as keep subtle texture details. In Figs. 13 and 14a, our algorithm achieved sharper boundaries of the bug and preserved the texture of the flower. In Fig. 14b, many curves with a variety of width are well reconstructed, whereas the wavelet coring introduced color fringing artifacts around boundaries. In Fig. 14c, the whiskers of the tiger are sharper and clearer by our algorithm and so are the stems and leaves of the grasses. In Fig. 14d, the texture details of the leaves are preserved, whereas the clouds are well smoothed. The ostrich head in Fig. 14e is a failure example for our method, where the upper neck part is oversmoothed and some artificial boundaries are generated for the furry edges. Note that while our system does not always completely remove the noise for the texture regions, the results can still look visually pleasing since the chrominance component of the noise is removed. In addition, the remaining noise in the texture regions as in Fig. 14d is not noticeable.

Overall, our algorithm outperforms the state-of-the-art denoising algorithms on the synthetic noise case. It takes our unoptimized Matlab implementation less than one minute on the average to denoise one picture (with a typical resolution of 481×321) in the Berkeley database. Our experiments were run on a 2.8 GHz Pentium D PC.

6.2 Real CCD Noise

We further tested our automatic denoising system using the pictures taken by CCD cameras under high-noise conditions [36]. The picture in Fig. 15a was taken by Canon™ EOS DIGITAL REBEL, with intense noise for the dim pixels but less for the bright ones. The noise level functions are estimated and displayed in Fig. 17, which agree with the observation. To compare, we also run wavelet coring (GSM), with $\sigma = 10\%$ and $\sigma = 15\%$, and the results are shown in Figs. 15b and 15c, respectively. The denoising result automatically generated by our system is shown in Fig. 15d. The close-up inspections of these results are shown in Fig. 16. Clearly, with the constant noise level assumption, the wavelet-coring algorithm cannot balance the high- and low-noise areas. When $\sigma = 10\%$, it does a better job for the bright pixels with sharper edges, and when $\sigma = 15\%$, it does a better job for the dark pixels with smoother regions. However, overall, we can still see blocky color artifacts, overly smoothed boundaries and loss of texture details. The result produced by our system shows that our technique is able to overcome these problems. In the noisy image of a nearly uniform region, Fig. 16 row (1), our method outputs an almost uniform patch. In row (2), the boundary is

Fig. 13. Close-up view of the denoising results. See text for the explanation.

much sharper, whereas in row (3), many subtle texture details are preserved. Overall, we believe our algorithm generates visually more appealing result (we cannot compute PSNR since there is no ground-truth clean image).

We tested our algorithm on another challenging example shown in Fig. 18a. As shown in Fig. 18b, the wavelet coring cannot effectively remove the color noise because of the spatial correlation of the color noise. Even though the result generated by our automatic denoising system in Fig. 18c overly sharpens the edges to have cartoon style, the noise gets completely removed and the image looks visually more pleasing.

Since our segmentation algorithm generates bigger segments for flat regions and smaller segments for texture regions, the parameter setting of segmentation does not significantly influence the overall performance of the system if it is within a reasonable range. We used the same parameter

settings for the benchmark test and found that visually more pleasing results are achieved for the real CCD noise examples in Fig. 18 if bigger segments are allowed (unfortunately we could not measure the PSNR). Certainly, better segmentation will further improve the denoising system, but we believe our current segmentation algorithm is sufficient.

7   DISCUSSION

Having shown the success of our model using both synthetic and real noise, we want to comment on the denoising problem and our modeling.

7.1   Color Noise

As shown in Section 3, the color of the pixels in a segment is approximately distributed along a 1D subspace of the three-dimensional RGB space. This agrees with the fact that the
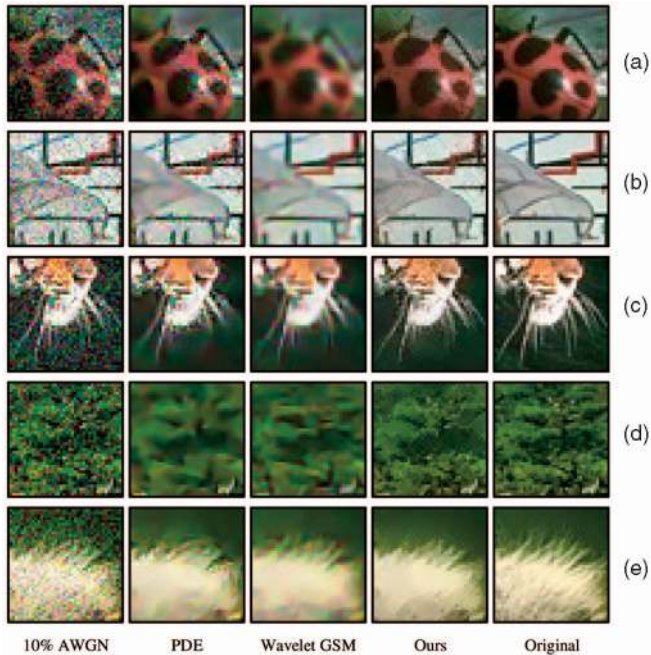
Fig. 14. Some close-up views of the denoising results in Fig. 13 are shown here. Our algorithm generated crispier images without color fringing artifacts as produced by PDE [53] and wavelet GSM [38] approaches.

Fig. 15. Comparison of denoising algorithms on a real CCD camera noise input. (a) Noisy input. (b) Wavelet GSM $\sigma = 10\%$t. (c) Wavelet GSM $\sigma = 15\%$. (d) Ours.

strong sharp boundaries are mainly produced by the change of materials (or reflectance), whereas the weak smooth boundaries are mainly produced by the change of lighting [49]. Since the human vision system is accustomed to these patterns, color noise, which breaks the 1D subspace rule, can appear visually annoying. Our denoising system was designed based on this 1D subspace rule to effectively

remove the chrominance component of color noise. The results of our zeroth-order model in Fig. 11 demonstrate that the images look significantly more pleasing when the chrominance component is removed.

7.2 Conditional Model versus Generative Model

In our system, we do segmentation only once to obtain a piecewise smooth model of the input image. If we treat the region partitions as a hidden variable that generates the noise

TABLE 2
PSNR for the Images in Berkeley Image Segmentation Database

PSNR	$\sigma = 5\%$					$\sigma = 10\%$				
File name	bilat	PDE	wavelet	0th	1st	bilat	PDE	wavelet	0th	1st
100075	29.32	29.76	31.27	31.69	31.68	26.47	27.72	28.31	28.14	28.96
105053	32.33	32.54	34.01	33.77	34.02	30.05	30.56	31.41	30.63	31.95
106025	34.47	34.29	36.13	35.75	36.44	30.94	31.58	32.57	32.03	34.22
108073	29.98	29.84	31.48	31.94	31.98	25.61	26.98	27.94	28.33	29.21
113009	30.73	30.72	32.89	32.31	32.61	27.38	27.80	29.91	28.89	30.19
134052	30.02	30.03	32.09	32.58	32.88	25.71	27.38	28.20	28.61	29.55
145053	29.51	29.12	31.72	31.88	32.26	23.84	25.83	27.23	27.46	28.71
15004	28.61	28.24	30.74	30.98	31.51	23.38	24.77	26.35	25.58	27.50
15088	29.55	29.19	33.36	32.41	32.74	25.02	26.64	28.83	27.71	28.76
22013	29.92	29.50	31.31	32.17	32.33	25.09	26.42	27.12	27.14	28.84
23084	30.31	29.76	32.14	32.04	32.64	24.63	26.34	27.24	27.05	29.23
26031	28.76	27.93	28.87	31.20	31.24	21.58	22.97	23.95	25.55	26.65
302003	31.29	30.93	33.70	32.94	33.96	26.85	27.43	29.47	27.91	30.84
314016	28.43	29.26	31.28	31.57	31.44	25.00	26.93	27.64	26.81	27.83
35008	33.28	33.40	35.74	34.84	35.97	29.25	30.85	31.23	30.24	33.27
65010	29.62	29.46	30.95	31.99	32.18	25.20	26.45	26.73	27.18	28.41
66075	32.57	32.46	33.36	35.02	35.03	28.33	29.78	29.69	29.79	31.80
mean	30.51	30.78	32.41	32.65	32.99	26.14	27.43	28.46	28.18	29.76

"bilat," "PDE," "wavelet," "zeroth," and "first" stand for bilateral filtering [52], curvature preserving PDE [53], wavelet (GSM) [38], zeroth-order and first-order model, respectively. The images with green are displayed, cropped, and enlarged in Fig. 13.

Fig. 16. Close-up view of the denoising results in Fig. 15. (a) Noisy input. (b) Wavelet GSM $\sigma = 10\%$. (c) Wavelet GSM $\sigma = 15\%$. (d) Ours.

image, then the conditional model becomes a generative model. Inference in the generative model would require integration over the region partitions. Intuitively, the segmentation of the noisy input image could be noisy and unreliable and could result in many possible segmentations.

One approach of this full Bayesian model is to sample partitions from the input image, obtain the denoised image for each segmentation, and compute the mean as the output. This approach would possibly improve the results by removing some of the boundary artifacts, but it is intractable in practice because of the huge space of partitions. Another approach is to treat the partition as missing data and use expectation-maximization (EM) algorithm to iterate between segmenting the image based on the denoised image (E-step), and estimating the denoised image based on the segmentation (M-step). This approach is also intractable in practice because many iterations are required. Nevertheless, these full

Bayesian approaches might be promising directions for future segmentation-based image processing systems with more powerful computation.

7.3 Automation of Computer Vision System

The performance of a computer vision system is sensitive to peripheral parameters, for example, noise level, blur level, resolution/image quality, lighting, and view point. For example, for image denoising the noise level is an important parameter to the system. Poor results may be produced with a wrong estimate of the noise level. Most existing computer vision algorithms focus on addressing the problems with known peripheral parameters, but the algorithms have to be tweaked to fit different imaging conditions. Therefore, an important direction is to make computer vision systems account for the important peripheral parameters fully automatically [30]. Our system is one first step: automatic estimation and removal of noise makes denoising algorithm robust to noise level.

8 CONCLUSION

Based on a simple piecewise-smooth image prior, we proposed a segmentation-based approach to automatically estimate and remove noise from color images. The NLF is obtained by estimating the lower envelope of the standard deviations of image variance per segment. The chrominance of the color noise is significantly removed by projecting the RGB pixel values to a line in color space fitted to each segment. The noise is removed by formulating and solving a Gaussian conditional random field. Experiments were conducted to test both the noise estimation and removal algorithms.

We verified that the estimated noise level is a tight upper bound of the true noise level in three ways: 1) by showing good agreement with experimentally measured noise from repeated exposures of the same image, 2) by repeatedly measuring the same NLF with the same camera for different image content, and 3) by accurately estimating known


Fig. 17. Estimated NLFs. (a) NLFs for the noisy sample in Fig. 15. (b) NLFs for the example in Fig. 18.


Fig. 18. The denoising results of a very challenging example. (a) Noisy input. (b) Denoised by wavelet GSM [38]. (c) Denoised by our algorithm.

synthetic noise functions. Our noise estimation algorithm can be applied to not only denoising algorithms but other computer vision applications to make them independent of noise level [30].

Our denoising algorithm outperforms the state-of-the-art wavelet denoising algorithms on both synthetic and real noise-contaminated images by generating shaper edges, producing smoother flat regions and preserving subtle texture details. These features match our original criteria that we proposed for a good denoising algorithm.

ACKNOWLEDGMENTS

REFERENCES

[1] S.P. Awate and R.T. Whitaker, "Higher-Order Image Statistics for Unsupervised, Information-Theoretic, Adaptive, Image Filtering," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2005.

[2] D. Barash, "A Fundamental Relationship between Bilateral Filtering, Adaptive Smoothing, and the Nonlinear Diffusion Equation," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 24, no. 6, pp. 844-847, June 2002.

[3] E.P. Bennett and L. McMillan, *Video Enhancement Using Per-Pixel Virtual Exposures, ACM SIGGRAPH,* pp. 845-852, 2005.

[4] M.J. Black, G. Sapiro, D.H. Marimont, and D. Heeger, "Robust Anisotropic Diffusion," *IEEE Trans. Image Processing,* vol. 7, no. 3, pp. 421-432, 1998.

[5] A. Blake and A. Zisserman, *Visual Reconstruction.* MIT Press, 1987.

[6] A. Buades, B. Coll, and J.-M. Morel, "A Non-Local Algorithm for Image Denoising," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2005.

[7] J. Canny, "A Computational Approach to Edge Detection," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 8, no. 6, pp. 679-698, 1986.

[8] D. Comaniciu and P. Meer, "Mean Shift: A Robust Approach Toward Feature Space Analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 24, pp. 603-619, 2002.

[9] M.S. Crouse, R.D. Nowak, and R.G. Baraniuk, "Wavelet-Based Statistical Signal Processing Using Hidden Markov Models," *IEEE Trans. Signal Processing,* vol. 46, no. 4, pp. 886-902, 1998.

[10] D. Donoho, "De-Noising by Soft-Thresholding," *IEEE Trans. Information Theory,* vol. 41, no. 3, pp. 613-627, 1995.

[11] F. Durand and J. Dorsey, "Fast Bilateral Filtering for the Display of High-Dynamic-Range Images," *Proc. ACM SIGGRAPH '02,* pp. 257-266, 2002.

[12] M. Elad and M. Aharon, "Image Denoising via Learned Dictionaries and Sparse Representation," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2006.

[13] M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions.* Wiley-Interscience, 2000.

[14] G. Fan and X.-G. Xia, "Image Denoising Using a Local Contextual Hidden Markov Model in the Wavelet Domain," *IEEE Signal Processing Letters,* vol. 8, no. 5, pp. 125-128, 2001.

[15] P.F. Felzenszwalb and D.P. Huttenlocher, "Efficient Graph-Based Image Segmentation," *Int'l J. Computer Vision,* vol. 59, pp. 167-181, 2004.

[16] D. Field, "Relations between the Statistics of Natural Images and the Response Properties of Cortial Cells," *J. Optical Soc. Am. A,* vol. 4, no. 12, pp. 2379-2394, Dec. 1987.

[17] W. Förstner, *Image Preprocessing for Feature Extraction in Digital Intensity, Color and Range Images.* Springer Lecture Notes on Earth Sciences, 1998.

[18] D. Gabor, "Theory of Communication," *J. IEE,* vol. 93, no. 26, pp. 429-457, 1946.

[19] S. Geman and D. Geman, "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 6, pp. 721-741, 1984.

[20] T. Gevers and A.W.M. Smeulders, "Color-Based Object Recognition," *Pattern Recognition,* vol. 32, no. 3, pp. 453-464, 1999.

[21] E.B. Gindele and N. Serrano, "Estimating Noise for a Digital Image Utilizing Updated Statistics," US patent 7,054,501, 2006.

[22] R.T. Gray and D.R. Cok, "Adjusting Film Grain Properties in Digital Images," US patent 5,641,596, 1997.

[23] M.D. Grossberg and S.K. Nayar, "Modeling the Space of Camera Response Functions," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 26, no. 10, pp. 1272-1282, Oct. 2004.

[24] G. Healey, "Segmenting Images Using Normalized Color," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 22, no. 1, pp. 64-73, 1992.

[25] G.E. Healey and R. Kondepudy, "Radiometric CCD Camera Calibration and Noise Estimation," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 16, no. 3, pp. 267-276, Mar. 1994.

[26] G. Klinker, S. Shafer, and T. Kanade, "A Physical Approach to Color Image Understanding," *Int'l J. Computer Vision,* vol. 4, no. 1, pp. 7-38, Jan. 1990.

[27] J. Lafferty, A. McCallum, and F. Pereira, "Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data," *Proc. Int'l Conf. Machine Learning,* 2001.

[28] C. Laroche and M. Prescott, "Apparatus and Methods for Adaptively Interpolating a Full Color Image Utilizing Chrominance Gradients," U.S. patent 5,373,322, 1994.

[29] A.B. Lee, D. Mumford, and J. Huang, "Occlusion Models for Natural Images: A Statistical Study of a Scale-Invariant Dead Leaves Model," *Int'l J. Compter Vision,* vol. 41, nos. 1/2, pp. 35-59, 2001.

[30] C. Liu, W.T. Freeman, R. Szeliski, and S.B. Kang, "Noise Estimation from a Single Image," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* pp. 901-908, June 2006.

[31] S.G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 11, no. 7, pp. 674-693, July 1989.

[32] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A Database of Human Segmented Natural Images and Its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics," *Proc. IEEE Int'l Conf. Computer Vision,* vol. 2, pp. 416-423, July 2001.

[33] D. Mumford and B. Gidas, "Stochastic Models for Generic Images," *Quarterly J. Applied Math.,* vol. LIX, no. 1, pp. 85-111, 2001.

[34] S. Paris and F. Durand, "A Fast Approximation of the Bilateral Filter Using a Signal Processing Approach," *Proc. European Conf. Computer Vision,* 2006.

[35] P. Perona and J. Malik, "Scale-Space and Edge Detection Using Anisotropic Diffusion," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 12, no. 7, pp. 629-639, July 1990.

[36] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama, "Digital Photography with Flash and No-Flash Image Pairs," *ACM Trans. Graphics,* vol. 23, no. 3, pp. 664-672, 2004.

[37] J. Portilla, "Full Blind Denoising through Noise Covariance Estimation Using Gaussian Scale Mixtures in the Wavelet Domain," *Proc. IEEE Int'l Conf. Image Processing,* pp. 1217-1220, 2004.

[38] J. Portilla, V. Strela, M.J. Wainwright, and E.P. Simoncelli, "Image Denoising Using Scale Mixtures of Gaussians in the Wavelet Domain," *IEEE Trans. Image Processing,* vol. 12, no. 11, pp. 1338-1351, Nov. 2003.

[39] R. Ramanath, W.E. Snyder, G.L. Bilbro, and W.A. Sander, "Demosaicking Methods for Bayer Color Arrays," *J. Electronic Imaging,* vol. 11, no. 3, pp. 306-315, July 2002.

[40] J. Rossi, "Digital Techniques for Reducing Television Noise," *J. Soc. Motion Picture and Television Eng.,* vol. 87, no. 134-140, 1978.

[41] S. Roth and M.J. Black, "Fields of Experts: A Framework for Learning Image Priors," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2005.

[42] W.F. Schreiber, *Fundamentals of Electronic Imaging Systems.* Springer, 1986.

[43] E.P. Simoncelli, "Statistical Models for Images: Compression, Restoration and Synthesis," *Proc. 31st Asilomar Conf. Signals, Systems and Computers,* pp. 673-678, 1997.

[44] E.P. Simoncelli and E.H. Adelson, "Noise Removal via Bayesian Wavelet Coring," *Proc. IEEE Int'l Conf. Image Processing,* vol. I, pp. 379-382, 1996.

[45] P.D. Snyder, T.F. Pawlicki, and R.S. Gaborski, "Apparatus and Method for Signal Dependent Noise Estimation and Reduction in Digital Images," US patent 5,923,775, 1999.

[46] A. Stefano, P. White, and W. Collis, "Training Methods for Image Noise Level Estimation on Wavelet Components," *EURASIP J. Applied Signal Processing,* vol. 16, pp. 2400-2407, 2004.

[47] G. Strang, *Introduction to Applied Math.* Wellesley-Cambridge Press, 1986.

[48] M.F. Tappen, E.H. Adelson, and W.T. Freeman, "Estimating Intrinsic Component Images Using Non-Linear Regression," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2006.

[49] M.F. Tappen, W.T. Freeman, and E.H. Adelson, "Recovering Intrinsic Images from a Single Image," *Advances in Neural Information Processing System,* pp. 1343-1350, 2003.

[50] M.F. Tappen, C. Liu, E.H. Adelson, and W.T. Freeman, "Learning Gaussian Conditional Random Fields for Low-Level Vision," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2007.

[51] D. Terzopoulos, "Regularization of Inverse Visual Problems Involving Discontinuities," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 8, no. 4, pp. 413-424, July 1986.

[52] C. Tomasi and R. Manduchi, "Bilateral Filtering for Gray and Color Images," *Proc. Sixth Int'l Conf. Computer Vision,* pp. 839-846, 1998.

[53] D. Tschumperlé, "Fast Anisotropic Smoothing of Multi-Valued Images Using Curvature-Preserving PDE's," *Int'l J. Computer Vision,* vol. 68, no. 1, pp. 65-82, 2006.

[54] Y. Tsin, V. Ramesh, and T. Kanade, "Statistical Calibration of CCD Imaging Process," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 480-487, 2001.

[55] S.C. Zhu and D. Mumford, "Prior Learning and Gibbs Reaction-Diffusion," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 19, no. 11, pp. 1236-1250, Nov. 1997.

[56] S.C. Zhu, Y. Wu, and D. Mumford, "Filters, Random Fields and Maximum Entropy (FRAME): Towards a Unified Theory for Texture Modeling," *Int'l J. Computer Vision,* vol. 27, no. 2, pp. 107-126, 1998.

[57] C.L. Zitnick, N. Jojic, and S.B. Kang, "Consistent Segmentation for Optical Flow Estimation," *Proc. IEEE Int'l Conf. Computer Vision,* 2005.

Ce Liu received the BS degree in automation and the ME degree in pattern recognition from the Department of Automation, Tsinghua University in 1999 and 2002, respectively. From 2002 to 2003, he worked at Microsoft Research Asia as an assistant researcher. He is currently a doctoral student at Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Lab (CSAIL). His research interests include computer vision, computer graphics, and machine learning. He received a Microsoft Fellowship in 2005 and the outstanding student paper award at the Advances in Neural Information Processing Systems (NIPS) in 2006. He is a student member of the IEEE.

Richard Szeliski received the PhD degree in computer science from Carnegie Mellon University, Pittsburgh, in 1988. He leads the Interactive Visual Media Group at Microsoft Research, which does research in digital and computational photography, video scene analysis, 3D computer vision, and image-based rendering. He joined Microsoft Research in 1995. Prior to Microsoft, he worked at Bell-Northern Research, Schlumberger Palo Alto Research, the Artificial Intelligence Center of the Stanford Research Institute (SRI) International, and the Cambridge Research Lab of Digital Equipment Corp. He has published more than 100 research papers in computer vision, computer graphics, medical imaging, and neural nets, as well as the book *Bayesian Modeling of Uncertainty in Low-Level Vision.* He was a program committee chair for the International Conference on Computer Vision (ICCV '01) and the 1999 Vision Algorithms Workshop. He served as an associate editor of the *IEEE Transactions on Pattern Analysis and Machine Intelligence* and on the editorial board of the *International Journal of Computer Vision* and is also a founding editor of the *Foundations and Trends in Computer Graphics and Vision.* He is a fellow of the IEEE.

Sing Bing Kang received the PhD degree in robotics from Carnegie Mellon University in 1994. He is currently a senior researcher at Microsoft Corp., working on image-based modeling, as well as image and video enhancement. He has coedited two books in computer vision (*Panoramic Vision* and *Emerging Topics in Computer Vision*) and coauthored a book on image-based rendering. He is a senior member of the IEEE.

C. Lawrence Zitnick received the PhD degree in robotics from Carnegie Mellon University in 2003. His thesis focused on algorithms for efficiently computing conditional probabilities in large-problem domains. Previously, his work centered on stereo vision, including cooperative and parallel algorithms, as well as developing a portable 3D camera. Currently, he is a researcher at the Interactive Visual Media group at Microsoft Research. His latest work includes spatial and temporal video interpolation, object recognition, and computational photography. He is a member of the IEEE.

William T. Freeman received the BS degree in physics and the MS degree in electrical engineering from Stanford University in 1979, the MS degree in applied physics from Cornell University in 1981, and the PhD degree in computer vision from the Massachusetts Institute of Technology in 1992. He is a professor of electrical engineering and computer science at the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology (MIT), joining the faculty in 2001. From 1992 to 2001, he worked at Mitsubishi Electric Research Labs (MERL), in Cambridge, Massachsetts, most recently, as senior research scientist and associate director. His current research interests include machine learning applied to computer vision, Bayesian models of visual perception, and computational photography. Previous research topics include steerable filters and pyramids, the generic viewpoint assumption, color constancy, separating "style and content," and computer vision for computer games. He holds over 25 patents. From 1981 to 1987, he worked at the Polaroid Corp., and during 1987-1988, he was a foreign expert at the Taiyuan University of Technology, China. He is a senior member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.