

PRACTICAL DENOISING OF CLIPPED OR OVEREXPOSED NOISY IMAGES

Alessandro Foi

Department of Signal Processing, Tampere University of Technology
P.O. Box 553, 33101, Tampere, Finland
web: www.cs.tut.fi/~foi email: firstname.lastname@tut.fi

ABSTRACT

We study the denoising of signals from clipped noisy observations, such as digital images of an under- or over-exposed scene. From a precise mathematical formulation of the problem, we derive a set of homomorphic transformations that enable the use of existing denoising algorithms for non-clipped data (including arbitrary denoising filters for additive i.i.d. Gaussian noise). Our results have general applicability and can be “plugged” into current filtering implementations, to enable a more accurate and better processing of clipped data. Experiments with synthetic images and with real raw data from CCD sensor show the feasibility and accuracy of the approach.

1. INTRODUCTION

We consider the problem of recovering a signal from its noisy observations that have been clipped, i.e. observations whose range is limited by a lower and by an upper bound. A prominent example are images acquired with under- or over-exposed areas and particularly the raw-data images acquired by a digital imaging sensor (e.g., a CCD or CMOS imager). As a continuation of our previous work [6] on modeling and automatic noise parameter estimation for clipped images corrupted by Poissonian-Gaussian noise, in this paper we provide a complete methodology for the accurate denoising of such images. The goal of this paper is general and pragmatic: rather than developing a novel denoising algorithm specially designed for clipped noisy images, we derive a set of homomorphic transformations that enable the use of existing denoising algorithms for non-clipped data (including arbitrary denoising filters for additive i.i.d. Gaussian noise). Thus, our results can be “plugged” into current filtering implementations, allowing a more accurate and better processing of clipped data.

An interesting and important feature of clipped noisy signals is that they carry usable information about the signals’ values outside the acquisition range. By properly processing and denoising these signals, it is possible to extract this information and then produce images with a dynamic range that exceeds that of the clipped signal. Thus, we can partly overcome the problem of saturation of the imaging device when a scene is captured with overexposed regions. Furthermore, our procedure can be utilized for the accurate linearization of the sensor’s output with respect to the expectation of its input. Hence, we provide a practical solution which has direct application for a number of image processing problems, including deblurring/deconvolution, high dynamic range (HDR) imaging, sensor characterization, digital forensics, to name a few.

The paper is organized as follows. In the next section we introduce the general observation models for images corrupted by signal-dependent noise; the models are given for the observations both before and after the clipping. Next, we recall the relations and transformations which exist between the expectations and standard deviations of the clipped and the non-clipped random variables. The core of our contributions is given in Section 3, where we consider the denoising of a clipped signal using generic filters. In particular,

This work was supported by the Academy of Finland (application no. 213462, Finnish Programme for Centres of Excellence in Research 2006-2011, and application no. 118312, Finland Distinguished Professor Programme 2007-2010).

we discuss the use of conventional filters designed for additive i.i.d. Gaussian noise and derive specific homomorphic transformations to stabilize the variance of the clipped observations, to compensate the bias due to the clipped distribution in the variance-stabilized domain, and finally to compensate the estimation bias between the denoised clipped variables and the non-clipped true variables. Experiments with synthetic as well as with real raw data from a commercial CCD camera are presented in Section 4, showing the feasibility and accuracy of the developed approach. Relevant conclusions are given in the last section.

2. PRELIMINARIES

2.1 Signal-dependent noise model

Let $z(x)$, $x \in X \subset \mathbb{Z}^2$, be noisy observations with the expectations $E\{z(x)\} = y(x) \in Y \subseteq \mathbb{R}$, where the errors (noise) $\eta(x) = z(x) - y(x)$ are random independent and the standard deviation of these observations is modeled by $\text{std}\{\eta(x)\} = \sigma(y(x))$, $\sigma: \mathbb{R} \rightarrow \mathbb{R}^+$ being a deterministic function of y . Equivalently, we consider the generic signal-dependent noise observation model of the form

$$z(x) = y(x) + \sigma(y(x))\zeta(x), \quad x \in X, \quad (1)$$

where $\zeta: X \rightarrow \mathbb{R}$ is zero-mean independent random noise with standard deviation equal to 1. Clearly, $\eta(x) = \sigma(y(x))\zeta(x)$. Although $\text{std}\{\zeta(x)\} = 1$ for all $x \in X$, the probability distribution of ζ can be different at different samples (i.e. $\zeta(x_1) \approx \zeta(x_2)$ if $x_1 \neq x_2$). However, to allow a simple mathematical formulation, we approximate ζ as a standard normal random variable, $\zeta \sim \mathcal{N}(0, 1)$, which is formally equivalent to considering a heteroskedastic Gaussian noise $\eta(x) \sim \mathcal{N}(0, \sigma^2(y(x)))$ having signal-dependent variance. As discussed in [6], this can be a suitable approximation when dealing with the noise in the raw data from digital imaging sensors, for which the typical form of the function σ is $\sigma^2(y(x)) = ay(x) + b$, with the constants $a \in \mathbb{R}^+$ and $b \in \mathbb{R}$ depending on the sensor’s specific characteristics and on the particular acquisition settings (e.g., analog gain or ISO value, temperature, pedestal, etc.). Obviously, the trivial i.i.d. (signal-independent and homoskedastic) additive Gaussian noise case, where $\text{std}\{\eta(x)\} \equiv \sqrt{b}$, is obtained for $a = 0$.

2.2 Clipping

In practice, the range of the acquisition system is always limited. Without loss of generality, we consider data given on the range $\tilde{Y} = [0, 1]$, where the extremes correspond to the maximum and minimum pixel values for the considered noisy image (e.g., raw data) format. Values above or below these bounds are replaced by the bounds themselves. This corresponds to the behavior of digital imaging sensors in the case of over- or under-exposure. Thus, we define the *clipped* observations \tilde{z} as

$$\tilde{z}(x) = \max\{0, \min\{z(x), 1\}\}, \quad x \in X, \quad (2)$$

where z is given by the (non clipped) signal-dependent noise model (1). In what follows, we use the tilde to denote variables directly related to clipped observations. The corresponding noise model for the clipped observations (2) is then

$$\tilde{z}(x) = \tilde{y}(x) + \tilde{\sigma}(\tilde{y}(x))\tilde{\zeta}(x), \quad (3)$$




Figure 1: Left and center: expectation $E\{\tilde{v}\}$ and standard deviation $\text{std}\{\tilde{v}\}$ of the clipped $\tilde{v} = \max(0, v)$ as functions \mathcal{E}_m and \mathcal{S}_m of μ , where $\mu = E\{v\}$ and $v \sim \mathcal{N}(\mu, 1)$. Right: the function $\mathcal{E}_r = \frac{\mu}{E\{\tilde{v}\}}$ of $\rho = \frac{E\{\tilde{v}\}}{\text{std}\{\tilde{v}\}}$. The small italic numbers indicate the corresponding value of μ .

where $\tilde{y}(x) = E\{\tilde{z}(x)\}$, the function $\tilde{\sigma} : [0, 1] \rightarrow \mathbb{R}^+$ is defined as $\tilde{\sigma}(\tilde{y}(x)) = \text{std}\{\tilde{z}(x)\}$, and $E\{\tilde{\zeta}(x)\} = 0$, $\text{var}\{\tilde{\zeta}(x)\} = 1$. Because of clipping, in general, we have that

$$\begin{aligned} \tilde{y}(x) &= E\{\tilde{z}(x)\} \neq E\{z(x)\} = y(x), \\ \tilde{\sigma}(\tilde{y}(x)) &= \text{std}\{\tilde{z}(x)\} \neq \text{std}\{z(x)\} = \sigma(y(x)). \end{aligned} \quad (4)$$

Thus, rewriting (3) as $\tilde{z}(x) = y(x) + [\tilde{y}(x) - y(x) + \tilde{\sigma}(\tilde{y}(x))\tilde{\zeta}(x)]$, we can see that, with respect to the true signal y , the clipped observations \tilde{z} are corrupted by an error (the term in square brackets) which has non-zero mean. Observe also that, even though $\text{var}\{\tilde{\zeta}(x)\} = \text{var}\{\zeta(x)\} = 1$, the distributions of ζ and $\tilde{\zeta}$ are different. In particular, assuming $\zeta(x) \sim \mathcal{N}(0, 1)$, we have that $\tilde{\zeta}(x)$ follows a censored normal distribution from above and from below [2].

2.3 Expectations and standard deviations of clipped variables and their transformations

A crucial point when working with clipped noisy signals is to understand how the variables and functions of the observation model (1) relate to those of the clipped observations' model (3). In particular, it is important to compute the functions \tilde{y} and $\tilde{\sigma}$ given σ and y , and vice versa.

As in [6], to simplify the calculations, we assume that the two clippings, the one from below ($z < 0$, $\tilde{z} = 0$) and the one from above ($z > 1$, $\tilde{z} = 1$), are not mixed by the randomness of the noise. In practice, this assumption is always verified, except for those extreme situations where the noise is dramatically strong (e.g., $\sigma \gg 0.2$ with respect to the range $\tilde{Y} = [0, 1]$).

Let $v \sim \mathcal{N}(\mu, 1)$ be a normally distributed random variable with mean $E\{v\} = \mu$ and unitary variance, and let $\tilde{v} = \max(0, v)$. It can be shown (see, e.g., [8], or [7], Chapter 20) that the expectation $E\{\tilde{v}\}$ and the variance $\text{var}\{\tilde{v}\}$ of the clipped (from below) \tilde{v} are

$$E\{\tilde{v}\} = \mathcal{E}_m(\mu) = \Phi(\mu)\mu + \phi(\mu), \quad (5)$$

$$\begin{aligned} \text{var}\{\tilde{v}\} &= \mathcal{S}_m^2(\mu) = \Phi(\mu) + \mathcal{E}_m(\mu)\mu - \mathcal{E}_m^2(\mu) = \\ &= \Phi(\mu) + \phi(\mu)\mu - \phi^2(\mu) + \Phi(\mu)\mu(\mu - \Phi(\mu)) - 2\phi(\mu)\mu, \end{aligned} \quad (6)$$

where ϕ and Φ are the probability density and cumulative distribution functions (p.d.f. and c.d.f.) of the standard normal $\mathcal{N}(0, 1)$. The plots of the expectation $E\{\tilde{v}\} = \mathcal{E}_m(\mu)$ and of the standard deviation $\text{std}\{\tilde{v}\} = \mathcal{S}_m(\mu)$ are shown, as functions of μ , in Figure 1.

Exploiting these functions and omitting the spatial coordinate x from notation, the *direct* and *inverse* transformations which link σ and y to \tilde{y} and $\tilde{\sigma}$ can be expressed as follows [6].

2.3.1 Direct transformation (\tilde{y} and $\tilde{\sigma}$ from y and σ)

Provided that $y = E\{z\}$ and $\sigma(y) = \text{std}\{z\}$ from the basic model (1) are known, the expectation $\tilde{y} = E\{\tilde{z}\}$ and the standard deviation $\tilde{\sigma}(\tilde{y}) = \text{std}\{\tilde{z}\}$ from the observation model (3) are obtained as

$$\tilde{y} = \mathcal{A}(y, \sigma(y)) = \sigma(y)\mathcal{E}_m\left(\frac{y}{\sigma(y)}\right) - y + 1 - \sigma(y)\mathcal{E}_m\left(\frac{1-y}{\sigma(y)}\right), \quad (7)$$

$$\tilde{\sigma}(\tilde{y}) = \mathcal{B}(y, \sigma(y)) = \sigma(y)\mathcal{S}_m\left(\frac{y}{\sigma(y)}\right)\mathcal{S}_m\left(\frac{1-y}{\sigma(y)}\right). \quad (8)$$

In particular, (7) and (8) define transformations that bring the standard deviation curve $(y, \sigma(y))$ to its clipped counterpart $(\tilde{y}, \tilde{\sigma}(\tilde{y}))$.

2.3.2 Inverse transformation (y from $\tilde{\sigma}$ and \tilde{y})

The non-clipped y from the model (1) can be calculated from the clipped \tilde{y} and $\tilde{\sigma}(\tilde{y})$ as

$$y = \mathcal{C}(\tilde{y}, \tilde{\sigma}(\tilde{y})) = \tilde{y}\mathcal{E}_r\left(\frac{\tilde{y}}{\tilde{\sigma}(\tilde{y})}\right) - \tilde{y} + 1 - (1 - \tilde{y})\mathcal{E}_r\left(\frac{1 - \tilde{y}}{\tilde{\sigma}(\tilde{y})}\right), \quad (9)$$

where \mathcal{E}_r is defined implicitly as function of $\rho = \frac{\mathcal{E}_m(\mu)}{\mathcal{S}_m(\mu)} = \frac{E\{\tilde{v}\}}{\text{std}\{\tilde{v}\}}$ by $\mathcal{E}_r(\rho) = \frac{\mu}{\mathcal{E}_m(\mu)}$. Figure 1 shows the plot of \mathcal{E}_r as function of ρ .

3. DENOISING CLIPPED SIGNALS

A generic denoising algorithm can be modeled as an operator whose output is an estimate of the expectation of the noisy input. Formally, let $\mathbf{D} : \mathbb{R}^{|\mathcal{X}|} \rightarrow \mathbb{R}^{|\mathcal{X}|}$ be the denoising operator, then

$$\mathbf{D}(\tilde{z}) = \widehat{E\{\tilde{z}\}} \approx E\{\tilde{z}\} = \tilde{y}. \quad (10)$$

It means that when we denoise \tilde{z} , as output we do not get an estimate of y , but rather an estimate of \tilde{y} . Analogously, we may say that $\mathbf{D}(\tilde{z})$ is a biased estimator of y , in the sense that $E\{\mathbf{D}(\tilde{z})\} \approx \tilde{y} \neq y$; in such a case, the bias error can be expressed as $\tilde{y} - y$.

For the sake of simplification, we may even assume that \mathbf{D} is an ideal operator that can always accurately recover the expectation of a given (non clipped) input image corrupted by additive i.i.d. Gaussian noise. This is an appealing assumption, but there remain a number of important issues that need to be considered when we apply \mathbf{D} for the denoising of clipped data.

Mainly, the clipping noise is signal-dependent and as such it requires special care: first, when the unknown noise parameters are estimated from the image; and second, when we aim at suppressing the noise. Furthermore, the noise samples follow essentially non-Gaussian and asymmetrical distributions. Finally, the estimation bias $y - \tilde{y}$ caused by the clipping needs to be compensated.

3.1 Noise estimation

In [6], we proposed an algorithm for the automatic estimation of the parameters of clipped signal-dependent noise models from a single noisy image. The algorithm utilizes a one-level wavelet decomposition of the noisy image and the transformations of Section 2.3 for the maximum-likelihood estimation of the noise parameters and hence of the curves $\tilde{\sigma}$ and \tilde{y} . We use this algorithm as the first step in the processing of clipped noisy images¹.

3.2 Noise removal

In general, when dealing with signal-dependent noise, we can either use filters specifically designed for such heteroskedastic noise (e.g., [5], [4]), or we can exploit a variance-stabilizing homomorphic transformation (e.g., [1], [11]) and then apply practically any filter for homoskedastic noise on the transformed noisy signal. Here, we discuss both alternative approaches and concentrate our attention on *transform-domain* filters designed for either heteroskedastic or homoskedastic *Gaussian* noise, as these are in practice the most powerful and most widely used ones (see, e.g., [10] and references

¹To simplify notation, here we use the symbols σ and $\tilde{\sigma}$ for the true as well as for the estimated curves. In [6] the latter are denoted by $\hat{\sigma}_{\text{fit}}$ and $\hat{\tilde{\sigma}}_{\text{fit}}$.

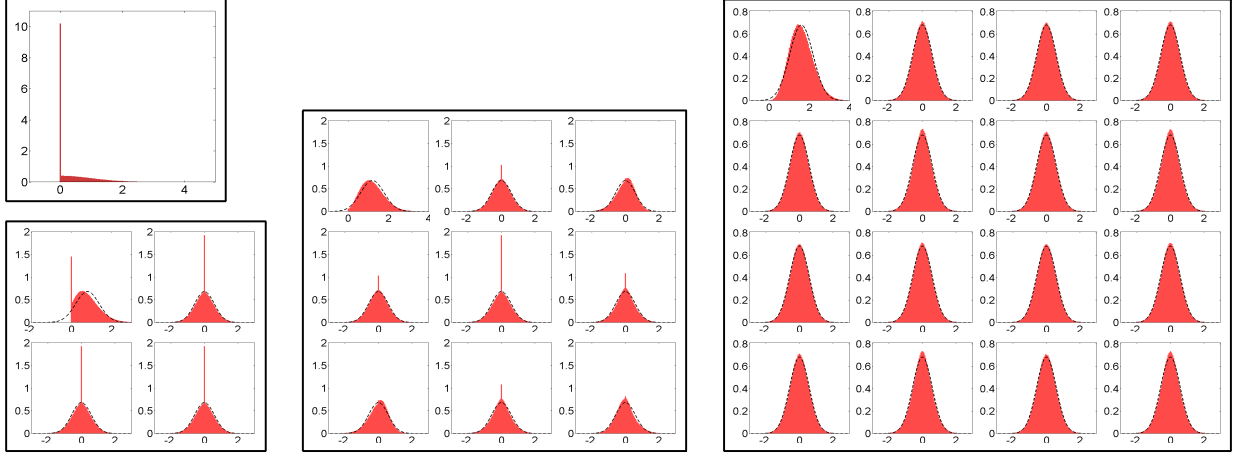


Figure 2: Normalized histograms of an i.i.d. clipped noise field \tilde{v} , $\tilde{v}(\cdot) = \max(0, v(\cdot))$ (where $v(\cdot) \sim \mathcal{N}(0, 1)$) and of each coefficient of its 2-D DCT transforms of size $n \times n$, $n=2,3,4$. The dashed lines show the probability distribution functions of $\mathcal{N}(n\mathcal{E}_m(0), \mathcal{S}_m(0))$ (DC terms) and $\mathcal{N}(0, \mathcal{S}_m(0))$ (AC terms). As the size of the transform increases, the distribution of the coefficients converges very quickly to these Gaussian bells. The histograms have equispaced bins of width 0.05 and are computed from $2 \cdot 10^7$ independent realizations of $\tilde{v}(\cdot)$.

therein). In what follows, we refer to the two approaches as *heteroskedastic filtering* and *homoskedastic filtering*, respectively, and denote the corresponding denoising filters as \mathbf{D}_{he} and \mathbf{D}_{ho} . Due to the essential non-Gaussianity of the clipped noise, the latter approach requires additional care.

3.2.1 Heteroskedastic filtering

Here, we assume that the denoising filter \mathbf{D}_{he} can handle the signal-dependent standard deviation $\tilde{\sigma}$. However, the noise distributions of different samples are all different and non-Gaussian. Fortunately, this does not constitute a serious problem for transform-based algorithms because, in practice, the transform coefficients of clipped data have distributions which are nearly Gaussians. In particular, let us consider a clipped noise field $\tilde{v} = \max(0, v)$, where $v(\cdot) \sim \mathcal{N}(\mu, 1)$. For any basis element ψ of an orthonormal transform, the distribution of the transform coefficient $\langle \psi, \tilde{v} \rangle$ can be approximated by $\mathcal{N}(\sum_j \psi(j) \mathcal{E}_m(\mu), \mathcal{S}_m(\mu))$. Figure 2 illustrates this approximation when $\mu = 0$ in the case of the 2-D discrete cosine transform DCT transform of size $n \times n$, $n=2,3,4$. As can be seen in the figure, the Gaussianization of the AC coefficients is faster. This is due to the positive and negative signs in the samples of the corresponding basis elements. The magnitude of the impulse at 0 (visible in some the histograms of smaller size DCT) decays exponentially with the number K of non-zero samples of the basis element, $\Phi(-\mu)^K$ [6]. For example, the nine basis elements $\{\psi_{i,j}\}_{i,j=1,2,3}$ of the 3×3 DCT have $K(i,j)=[3 \ 2 \ 3]^T [3 \ 2 \ 3]$ non-zero samples, respectively, and for $\mu=0$ we have $\Phi(-\mu)=\frac{1}{2}$; it means that the magnitude of the impulse at 0 in the histogram for $i=j=2$ ($K=4$) is approximately $2^{6-4}=4$ times larger than those for $i+j=3, 5$ ($K=6$), $2^{9-4}=32$ times larger than those for $i, j \in \{1, 3\}$ ($K=9$), and $2^{4-1}=8$ times smaller than the impulse in the histogram of \tilde{v} ($K=1$). Each basis element of the 4×4 DCT has $K=16$ non-zero samples, hence the magnitude of the impulse in their histograms is $2^{16-1}=32768$ times smaller than the impulse in the histogram of \tilde{v} and thus practically negligible.

Of course, because of the central-limit theorem, one always gets some sort of ‘‘Gaussianization’’, regardless of the distribution of the original samples. However, in the case of clipped normal samples, the convergence is remarkably fast, as can be seen in the figure.

We should also remark that, in the transform domain, the clipped noise is no longer independent and that instead some correlation exists in the noise of different coefficients. Nevertheless, it is well known that such correlation is in practice a secondary issue which is always encountered whenever the transform is redundant or overcomplete.

Therefore, we conclude that a transform-domain filter \mathbf{D}_{he} can be applied successfully on clipped data and that the approximation (10) holds.

3.2.2 Homoskedastic filtering

A. Variance-stabilizing transformation. A specifically designed homomorphic transformation $f: [0, 1] \rightarrow \mathbb{R}$ can be utilized in order to stabilize the signal-dependent standard deviation of the clipped observations \tilde{z} to a desired constant $c \in \mathbb{R}^+$ and thus apply a denoising algorithm \mathbf{D}_{ho} for homoskedastic noise on $f(\tilde{z})$. Following the simple approach which appears in many works (e.g., [11]), we use a first-order Taylor expansion for a monotonically increasing f , from which follows $\text{std}\{f(\tilde{z})\} \approx f'(E\{\tilde{z}\}) \text{std}\{\tilde{z}\} = f'(\tilde{y}) \tilde{\sigma}(\tilde{y})$, and then solve for $\text{std}\{f(\tilde{z})\} \equiv c$. Up to an arbitrary additive constant, this yields

$$f(t) = \int_{t_0}^t \frac{c}{\tilde{\sigma}(s)} ds, \quad t, t_0 \in [0, 1], \quad (11)$$

i.e. the c -stabilizing homomorphic transformation is the indefinite integral of $\frac{c}{\tilde{\sigma}(\tilde{y})}$, with the integration with respect to the argument \tilde{y} .

B. Boundedness of the transformation. Let us now discuss the convergence of the integral (11) (and thus the boundedness of f) when $\tilde{\sigma}(s) \rightarrow 0$. Without loss of generality, we restrict to the case of clipping from below, for which (7) and (8) reduce to

$$\tilde{y} = \sigma(y) \mathcal{E}_m\left(\frac{y}{\sigma(y)}\right), \quad \tilde{\sigma}(\tilde{y}) = \sigma(y) \mathcal{S}_m\left(\frac{y}{\sigma(y)}\right). \quad (12)$$

Let $\mu = \frac{y}{\sigma(y)}$; we have

$$\tilde{\sigma}(\tilde{y}) = \sigma(y) \sqrt{\Phi(\mu) + \mathcal{E}_m(\mu) \mu - \mathcal{E}_m^2(\mu)}. \quad (13)$$

We give proofs of divergence and convergence for the two main cases which are of practical interest. **Case 1.** Let $\sigma(y) \xrightarrow{y \rightarrow 0} 0$

and assume that $\mu = \frac{y}{\sigma(y)} \rightarrow 0$. Then, from (12)-(13) and (5)-(6), $\tilde{y} \simeq \sigma(y) \phi(0)$ and $\tilde{\sigma}(\tilde{y}) \simeq \sigma(y) \sqrt{\Phi(0) - \phi^2(0)}$, which implies that $\tilde{\sigma}(\tilde{y}) \simeq \tilde{y} \sqrt{\Phi(0) - \phi^2(0)}/\phi(0)$ and hence that (11) diverges logarithmically at 0. For example, $\sigma(y) = \sqrt{y}$ leads to such situation.

Case 2. Let $\tilde{\sigma}(\tilde{y}) \xrightarrow{\tilde{y} \rightarrow 0} 0$ and assume that exists some $\varepsilon > 0$ such that $\sigma(y) \geq \varepsilon$. This implies that $\mu \rightarrow -\infty$ and hence that $y \rightarrow -\infty$. From (13), we obtain

$$\begin{aligned} \tilde{\sigma}(\tilde{y}) &= \sigma(y) \sqrt{\mathcal{E}_m(\mu) \left(\frac{\Phi(\mu)}{\mathcal{E}_m(\mu)} + \mu - \mathcal{E}_m(\mu) \right)} = \\ &= \sigma(y) \sqrt{\mathcal{E}_m(\mu) \left(\left(\frac{\Phi(\mu)\mu + \phi(\mu)}{\Phi(\mu)} \right)^{-1} + \mu - \mathcal{E}_m(\mu) \right)} = \end{aligned}$$




Figure 3: Noisy \tilde{z} (3), and denoised and debiased \hat{y} (15) for the *Testpat* and *Man* test images of size 1024×1024 .

$$= \sigma(y) \sqrt{\mathcal{E}_m(\mu) \left(\left[(\mu + R^{-1}(-\mu))^{-1} + \mu \right] - \mathcal{E}_m(\mu) \right)},$$
 where R is the Mill's ratio, $R(\mu) = \frac{1-\Phi(\mu)}{\phi(\mu)}$. Exploiting a classical asymptotic expansion [9], $R(\mu) = \frac{1}{\mu} - \frac{1}{\mu^3} + \frac{1 \cdot 3}{\mu^5} - \frac{1 \cdot 3 \cdot 5}{\mu^7} + \dots$, we can approximate the term in the square brackets above as $-\frac{2}{\mu} + O(\frac{1}{\mu^3})$. Moreover, $\mathcal{E}_m(\mu) \xrightarrow{\mu \rightarrow \infty} 0$ at exponential rate and, in particular, much faster than $(-\frac{2}{\mu})^2$. It means that $-\frac{2}{\mu} + O(\frac{1}{\mu^3})$ is much larger than $\sqrt{\mathcal{E}_m(\mu)}$. Hence, for large enough $-\mu$, we have

$$\begin{aligned}
 \tilde{\sigma}(\tilde{y}) &= \sigma(y) \sqrt{\mathcal{E}_m\left(\frac{y}{\sigma(y)}\right) \left(-\frac{2}{\mu} + O(\mu^{-3})\right)} > \\
 &> \sigma(y) \sqrt{\mathcal{E}_m\left(\frac{y}{\sigma(y)}\right)} = \sigma^{\frac{1}{4}}(y) \tilde{y}^{\frac{3}{4}} \geq \varepsilon^{\frac{1}{4}} \tilde{y}^{\frac{3}{4}}.
 \end{aligned}$$

This proves our result, because in a neighborhood of 0 the integrand in (11) is bounded from above by $c\varepsilon^{-\frac{1}{4}}\tilde{y}^{-\frac{3}{4}}$, which has itself a convergent indefinite integral at 0.

For obvious reasons of numerical stability, it is preferable and recommended to deal with bounded functions; therefore, in practice, for those cases (e.g., Case 1) where (11) is divergent, we shall use $f(t) = \int_{t_0=0}^t \frac{c}{\max\{\tilde{\sigma}(s), \varepsilon\}} ds$, where $\varepsilon > 0$ is a small regularization constant. This also ensures that f is Lipschitz. In what follows, we will always assume that f is a bounded, strictly increasing (hence invertible) function, and that $t_0=0, f(0)=0$.

C. Transform-domain filtering. In the light of the previous first-order approximation, $f(\tilde{z})$ can be treated as a clipped normal random variable, with clipping from below at $f(0) = 0$ and from above at $f(1)$, and variance (after clipping) equal to c^2 . Thus, qualitatively, the rapid Gaussianization discussed in Section 3.2.1 holds also after the variance stabilization, enabling the effectiveness of the homoskedastic denoising filter. Hence, in this case, the approximation (10) has the form $\mathbf{D}_{\text{ho}}(f(\tilde{z})) \approx E\{f(\tilde{z})\}$.

D. Debiasing and inverse transformation. Because of the non-linearity of f , the first-order approximation is never exact. While minor approximation errors on the variance are typically acceptable, the errors on the expectations are not, because they result in a systematic estimation bias. More precisely, we have that $E\{f(\tilde{z})\} \neq f(E\{\tilde{z}\})$; this discrepancy must be compensated before the inversion of f .

First, let us observe that the p.d.f. of \tilde{z} (2) is $\wp_{\tilde{z}}(\zeta) = \Phi\left(\frac{-y}{\sigma(y)}\right) \delta_0(\zeta) + \frac{1}{\sigma(y)} \phi\left(\frac{\zeta-y}{\sigma(y)}\right) \chi_{[0,1]} + \left(1 - \Phi\left(\frac{1-y}{\sigma(y)}\right)\right) \delta_0(1-\zeta)$, where $\chi_{[0,1]}$ denotes the characteristic function of the interval $[0, 1]$ and δ_0 is the Dirac delta impulse at 0. Clearly, $f(E\{\tilde{z}\}) = f\left(\int_0^1 \zeta \wp_{\tilde{z}}(\zeta) d\zeta\right)$ and $E\{f(\tilde{z})\} = \int_0^1 f(\zeta) \wp_{\tilde{z}}(\zeta) d\zeta = \int_0^1 f(\zeta) \frac{1}{\sigma(y)} \phi\left(\frac{\zeta-y}{\sigma(y)}\right) d\zeta + f(1) \left(1 - \Phi\left(\frac{1-y}{\sigma(y)}\right)\right)$ (note that $f(0) = 0$). Let now $h : [0, f(1)] \rightarrow [0, f(1)]$ be the function defined (implicitly varying y) by $f(E\{\tilde{z}\}) \mapsto h(f(E\{\tilde{z}\})) = E\{f(\tilde{z})\}$. The invertibility of h follows easily from the monotonicity of f . Thus, we have that $h^{-1}(\mathbf{D}_{\text{ho}}(f(\tilde{z}))) \approx f(E\{\tilde{z}\})$ and hence that





Figure 4: Estimated standard-deviation functions σ and $\tilde{\sigma}$ for the *Testpat* and *Man* test images and for the raw-data image acquired by the CCD sensor of a Fujifilm FinePix S9600 camera. The true curves corresponding to the two test images are drawn (in a different color) beneath the respective estimated lines.

$$f^{-1}(h^{-1}(\mathbf{D}_{\text{ho}}(f(\tilde{z})))) \approx E\{\tilde{z}\}, \quad (14)$$

which is the final form of the approximation (10) for the case of homoskedastic filtering. Due to (14), we can define a denoising filter $\tilde{\mathbf{D}}_{\text{he}}$ for heteroskedastic noise as $\tilde{\mathbf{D}}_{\text{he}} = f \circ \mathbf{D}_{\text{ho}} \circ h^{-1} \circ f^{-1}$.

3.3 Output debiasing

As shown in Section 3.2, using either \mathbf{D}_{he} or $\tilde{\mathbf{D}}_{\text{he}}$ (from now on denoted collectively as \mathbf{D}), we can get an estimate of $E\{\tilde{z}\} = \tilde{y}$. However, our goal is to estimate the non-clipped y . Its estimate can be obtained exploiting the inverse transformation (9) as

$$\hat{y} = \mathcal{C}(\mathbf{D}(\tilde{z}), \tilde{\sigma}(\mathbf{D}(\tilde{z}))). \quad (15)$$

As can be intuited from Figure 1 and by the very definition of \mathcal{E}_r , the range of \hat{y} can in principle be the whole $(-\infty, \infty)$. Nevertheless, because \mathcal{E}_r has unbounded derivative and because of finite precision, the actual range that can be obtained is rather limited. Firstly, with double-precision floats (as, e.g., in Matlab) the limit is set at about 8σ beyond the 0 and 1 bounds, which roughly means that the range can be extended up to $(16\sigma+1)$ times (note that $\Phi(\mu)$ approaches the relative spacing between adjacent double-precision floats at $\mu = -8.1$). However, in practice, the achievable range can be much smaller, because of estimation errors. In this sense, it is important to emphasize that the inverse \mathcal{C} (9) does not consider the estimation errors in the estimation of y or $\sigma(y)$, i.e. it is not an optimal inverse (such as a MMSE or ML inverse). The fact that the extended range increases with σ follows directly from the form of (7)-(9). It means that by denoising we can “take advantage of the noise” to obtain an image with a wider range, since the range of \tilde{z} and \tilde{y} is always smaller than that of y . Qualitatively, the estimation accuracy of \hat{y} depends on the smoothness (e.g., bandwidth) of y .

4. EXPERIMENTAL RESULTS

Due to length limitation, we show experiments for the homoskedastic filtering only, as this is the approach that has more general applicability. As denoising filter \mathbf{D}_{ho} , we use the BM3D algorithm [3]. Similar to [6], for all experiments we follow the signal-dependent noise models (1)-(3), with the noise term in (1) composed of two mutually independent parts, a *Poissonian* signal-




Figure 7: Noisy raw-data image \tilde{z} (3) (Fujifilm FinePix S9600 camera), denoised $\mathbf{D}(\tilde{z})$ (14), and denoised and debiased \hat{y} (15).




Figure 5: Variance-stabilizing homomorphic transformation f (11) for the three standard-deviation functions shown in Figure 4. The curves are plotted with respect to a normalized vertical scale $[0, f(1)]$.




Figure 6: Cross-sections of the original y (1), noisy \tilde{z} (3), denoised $\mathbf{D}(\tilde{z})$ (14), and of the denoised and debiased \hat{y} (15), for the 171th line of the *Testpat* image (the line is indicated by the black triangle next to \tilde{z} in Figure 3).

dependent component η_p and a *Gaussian* signal-independent component η_g : $\sigma(y(x))\zeta(x) = \eta_p(y(x)) + \eta_g(x)$. In particular, $(y(x) + \eta_p(y(x)))\chi \sim \mathcal{P}(\chi y(x))$, where $\chi > 0$ and \mathcal{P} denotes the Poisson distribution, and $\eta_g(x) \sim \mathcal{N}(0, b)$. Thus, as discussed in Section 2.1, $\sigma^2(y(x)) = ay(x) + b$, with $a = \chi^{-1}$.

First, we show results for two 1024×1024 test images [12] with simulated noise: *Testpat* ($a=0, b=0.2^2$) and *Man* ($\chi=30, b=0.1^2$). The noisy images are shown in Figure 3. The standard-deviation functions σ and $\tilde{\sigma}$ estimated with the algorithm [6] are shown in Figure 4. The estimated curves match with the respective true curves. Figure 5 presents the computed variance-stabilizing transformations (11)². The final debiased and denoised estimates \hat{y} (15) are shown in Figure 3. The PSNR (dB) of \tilde{z} , $\mathbf{D}(\tilde{z})$ (14), and \hat{y} for *Testpat* are 4.85, 26.49, and 29.35, respectively. The corresponding numbers for *Man* are 7.64, 27.69, and 27.96. These numbers demonstrate the substantial improvement obtained by compensating the bias due to the clipped noise. This improvement is well illustrated by the cross-sections of *Testpat* shown in Figure 6.

We also show results of processing a real raw-data 1224×922 image (green subcomponent) from the CCD sensor of a Fujifilm FinePix S9600 camera at ISO 1600. Figure 7 shows the raw-data \tilde{z} , $\mathbf{D}(\tilde{z})$, and \hat{y} . The latter image has an extended range $[-0.03, 1.33]$

and it is thus shown after scaling to $[0, 1]$. The extension of the range is clearly visible in the overexposed highlights of the scene.

5. CONCLUSIONS

We identified conditions and derived transformations for the rigorous processing of clipped noisy signals using generic filters. Experiments demonstrate the success of the proposed approach. The shown techniques can also be used for the accurate linearization of the sensor response, including the correct pedestal (black) level measurement. They are thus relevant for high dynamic range (HDR) imaging applications, where a number of differently exposed images, all of which present substantial over- or underexposed regions, are composed to yield one image where all regions are visible, and for all fields where accuracy in the estimation of the sensor response is crucial, such as inverse imaging (allowing to use of linear deconvolution for deblurring over-exposed raw-data), digital forensics (sensor identification), etc. Related Matlab software is provided online at www.cs.tut.fi/~foi/sensornoise.html

REFERENCES

- [1] Anscombe, F.J., “The Transformation of Poisson, Binomial and Negative-Binomial Data”, *Biometrika*, vol. 35, no. 3/4, pp. 246-254, 1948.
- [2] Cohen, A., *Truncated and censored samples*, CRC Press, 1991.
- [3] Dabov, K., A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising by sparse 3D transform-domain collaborative filtering”, *IEEE Trans. Image Process.*, vol. 16, no. 8, Aug. 2007. <http://www.cs.tut.fi/~foi/GCF-BM3D>
- [4] Foi, A., R. Bilcu, V. Katkovnik, and K. Egiazarian, “Adaptive-Size Block Transforms for Signal-Dependent Noise Removal”, *Proc. 7th Nordic Signal Processing Symposium, NORSIG 2006*, Reykjavik, Iceland, June 2006.
- [5] Foi, A., V. Katkovnik, and K. Egiazarian, “Signal-dependent noise removal in Pointwise Shape-Adaptive DCT domain with locally adaptive variance”, *Proc. 15th Eur. Signal Process. Conf., EUSIPCO 2007*, Poznań, September 2007.
- [6] Foi, A., M. Trimeche, V. Katkovnik, and K. Egiazarian, “Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data”, to appear in *IEEE Trans. Image Process.*
- [7] Greene, W., *Econometric Analysis*, 4th ed., Prentice Hall, 2000.
- [8] Johnson, N.L., S. Kotz, and N. Balakrishnan, *Continuous univariate distributions*, 2nd ed., Wiley, 1994.
- [9] Kendall, M., and A. Stuart, *The advanced theory of statistics, Volume 1*, 2nd ed., Charles Griffin and co., 1963.
- [10] Lansel, S., D. Donoho, and T. Weissman, “DenoiseLab: A Standard Test Set and Evaluation Method to Compare Denoising Algorithms”, benchmarks and MATLAB software online at <http://www.stanford.edu/~slansel/DenoiseLab/>.
- [11] Prucnal, P.R., and B.E.A. Saleh, “Transformation of image-signal-dependent noise into image signal-independent noise”, *Optics Letters*, vol. 6, no. 7, July 1981.
- [12] The USC-SIPI Image Database, University of Southern California, <http://sipi.usc.edu/database/>.

14.85
17.06

²The integral (11) converges and f is bounded for any $a \geq 0$ and $b > 0$, as proved in Case 2 of Section 3.2.2.B. Thus, we can set $t_0=0$ and have $f(0)=0$.