Variance stabilization in Poisson image deblurring

VARIANCE STABILIZATION IN POISSON IMAGE DEBLURRING

Lucio Azzari and Alessandro Foi

Tampere University of Technology

ABSTRACT

We consider the restoration of blurred images corrupted by Poisson noise using variance-stabilizing transformations (VST). Although VST are an established tool used extensively for denoising, their adoption in deconvolution problems is problematic because VST are necessarily nonlinear operators, and thus break the linear image-formation model typically adopted in deconvolution. We propose a deblurring framework where the image is 1) deconvolved by a linear regularized inverse filter, 2) transformed by VST into an image which can be treated as corrupted by strong spatially correlated noise with constant variance and known power spectrum, 3) denoised by a filter for additive colored Gaussian noise, 4) returned to the original range via inverse VST. We particularly analyze the stabilization of Poisson variates after linear filtering and characterize the noise power spectrum before and after application of VST. We present an efficient implementation of this original deblurring framework using the BM3D denoising filter, demonstrating state-of-the-art results which are especially appealing in low SNR imaging conditions.

***Index Terms*—** Poisson image deblurring, deconvolution, signal-dependent noise, variance stabilization, photon-limited imaging.

1. INTRODUCTION

We propose an entirely modular approach to Poisson image deblurring. It tackles individually the problems of deconvolution and denoising by first deconvolving the image using a linear regularized deconvolution and then by treating the deconvolved image as corrupted by spatially correlated signal-dependent noise. Using suitably scaled variance-stabilizing transformations (VST), the noise becomes spatially correlated signal-independent, ultimately leading to a denoising problem that can be solved using off-the-shelf filters for stationary colored Gaussian noise. To this end, we provide an explicit model for the power spectrum of the colored noise which follows from the deconvolution filter and the action of the VST.

It is crucial that we deconvolve the image before applying the VST, and not after: because VSTs are nonlinear, they turn any linear spatially invariant blur into a generic nonlinear spatially varying blur, which if filtered as in a deconvolution problem results in significant distortions around edges [1].

To deal effectively with low SNR observations typical of photon-limited image capture, we implement the proposed approach within the iterative VST denoising framework [2], which we hereby extend to image deblurring. This framework progressively improves the effectiveness of the VST by applying the transformation to a linear combination of the deconvolved image with the previous estimate and by adopting pixel binning. The BM3D filter for colored noise [3] is used for denoising the stabilized image.

The validity of the proposed method is supported by its excellent performance in the comparison against state-of-the-art methods for Poisson deblurring reported in the experimental section of the manuscript.

2. PRELIMINARIES AND RATIONALE

Let y be an unknown deterministic noise-free image, v the known point spread function (PSF) of the blur, and let $g = y \circledast v$, where $\circledast$ denotes the convolution. We model the observed blurry and noisy image z as composed of pixels $z(x)$, $x \in \Omega$[1], that are independent realizations of a Poisson process with parameter $g(x) \geq 0$:

$$z(x) \sim \mathcal{P}(g(x)), \quad \mathrm{P}(z(x)\,|g(x)) = \begin{cases} \frac{g(x)^{z(x)} e^{-g(x)}}{z(x)!} & z \in \mathbb{N} \cup \{0\} \\ 0 & \text{elsewhere.} \end{cases}$$

The goal of deblurring is to restore y from the observed z.

The expectation and variance of the observed signal z are

$$\mathrm{E}\{z|y\} = \mathrm{var}\{z|y\} = y \circledast v = g. \tag{1}$$

The Anscombe forward transformation [5] is a well-known VST for Poisson data in particular, and for data satisfying the expectation-variance identity (1) in general [6]. Applying it to z yields an image $a(z) = 2\sqrt{z + 3/8}$ which can be treated as corrupted by addtive white Gaussian noise (AWGN) with unit variance, but subject to a blur which is no longer convolutional. Thus, the advantage of attaining signal-independent homoskedastic noise is offset by a complicated nonlinear blur which requires sophisticated iterative inversion schemes [7, 8]. We are interested in extending to the Poisson case popular direct deblurring methods for AWGN observations [9, 3] that combine simple linear regularized deconvolution with shrinkage in a transform domain, while maintaining the simplicity of VST-based Poisson denoising methods such as [10, 2].

3. PROPOSED METHOD

Algorithm 1 gives the proposed method, which we detail as follows. For clarity, binning operators are omitted from notation throughout Sections 3.1-3.5; binning is then presented separately in Section 3.6.

3.1. Regularized deconvolution

We compute a linear regularized deconvolution of z as

$$z^{\mathrm{RI}} = \mathcal{F}^{-1}(T^{\mathrm{RI}} Z) = t^{\mathrm{RI}} \circledast z, \qquad T^{\mathrm{RI}} = \frac{V^*}{|V|^2 + \mathcal{E}^2}, \tag{2}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ are, respectively, the forward and inverse Fourier transform, $Z = \mathcal{F}(z)$, $V = \mathcal{F}(v)$, V^* is the complex conjugate of V, and $\mathcal{E}^2 \geq 0$ is a regularization term. The deconvolved image z^{RI} is thus corrupted by nonstationary heteroskedastic spatially correlated noise whose variance and power spectrum Ψ^{RI} are determined by $g = y \circledast v$ and T^{RI}. In particular,

$$\mathrm{E}\{z^{\mathrm{RI}}|y\} = g \circledast t^{\mathrm{RI}}, \qquad \mathrm{var}\{z^{\mathrm{RI}}|y\} = g \circledast (t^{\mathrm{RI}})^2, \tag{3}$$

$$\Psi^{\mathrm{RI}} = \mathrm{var}\{Z^{\mathrm{RI}}|y\} = |T^{\mathrm{RI}}|^2 \,\mathrm{var}\{Z|y\} = |T^{\mathrm{RI}}|^2 \textstyle\sum_{x \in \Omega} g(x), \tag{4}$$

[1] For the sake of simplicity, we consider periodic boundary conditions and circulant convolution. Methods such as [4] (or references therein) can be adopted for dealing with unknown boundary in the practice.

Algorithm 1 Poisson Image Deconvolution via VST

$z^{\mathsf{RI}} = \mathcal{F}^{-1}(T^{\mathsf{RI}} Z)$
$\hat{y}_0 = z^{\mathsf{RI}}$
for $i = 1$ to K **do**
  $\bar{z}_i^{\mathsf{RI}} = \lambda_i z^{\mathsf{RI}} + (1-\lambda_i)\hat{y}_{i-1}$
  $\bar{\bar{z}}_i^{\mathsf{RI}} = f_i\left(\mathcal{B}_{h_i}\left[\bar{z}_i^{\mathsf{RI}}\right]\right)$
  $D_i = \Phi\left[\bar{\bar{z}}_i^{\mathsf{RI}}, \mathcal{B}_{h_i}\left[\tilde{\Psi}^{\mathsf{RI}}\right]\right]$
  $\hat{y}_i = \mathcal{B}_{h_i}^{-1}\left[\mathcal{I}_{f_i}^{\lambda_i}(D_i)\right]$
end for
return $\hat{y} = \hat{y}_K$

where $Z^{\mathsf{RI}} = \mathcal{F}(z^{\mathsf{RI}})$. Equalities (3) and (4) hold for any convolution filter t^{RI} applied to z, including regularized inverses typical of Gaussian deconvolution like Tikhonov (constant $\mathcal{E}$) and Wiener ($\mathcal{E} \propto 1/|Y|$) [9, 3]. The magnitude of $\mathcal{E}$ controls the trade-off between noise amplification and blur inversion. A small enough $\mathcal{E}$ makes the bias $\mathrm{E}\{z^{\mathsf{RI}}|y\} - y = g \circledast t^{\mathsf{RI}} - y$ practically negligible at the expense of a significant variance $\mathrm{var}\{z^{\mathsf{RI}}|y\}$, which thus requires powerful denoising. The following sections deal with the opportunity of stabilizing this variance via VST.

3.2. Convex combination

VSTs become typically inaccurate at low counts (i.e. low-intensity signals). As in [2], we can progressively improve the accuracy of a VST by increasing the SNR of its input combining it with an estimate of its expectation. Because $g \circledast t^{\mathsf{RI}} - y$ shall be negligible, we treat estimates of z^{RI} as direct estimates of y, which we denote by $\hat{y}$. Thus, we define

$$\bar{z}_i^{\mathsf{RI}} = \lambda_i z^{\mathsf{RI}} + (1-\lambda_i)\hat{y}_{i-1}, \tag{5}$$

where $0 < \lambda_i \leq 1$ is the combination coefficient and the subscript indices denote a symbol's instance at a particular iteration. The noise power spectrum of $\bar{z}_i^{\mathsf{RI}}$ is $\lambda_i^2 \Psi^{\mathsf{RI}}$, by trivial scaling of (4).

3.3. Variance stabilization

We stabilize $\bar{z}_i^{\mathsf{RI}}$ by the VST f_i, obtained by suitable scaling of a [5]:

$$f_i\left(\bar{z}_i^{\mathsf{RI}}\right) = a\left(c_i \bar{z}_i^{\mathsf{RI}}\right) = 2\sqrt{c_i \bar{z}_i^{\mathsf{RI}} + 3/8}, \quad c_i = \frac{\sum_{x\in\Omega} t^{\mathsf{RI}}(x)}{\lambda_i^2 \sum_{x\in\Omega} \left(t^{\mathsf{RI}}(x)\right)^2}. \tag{6}$$

The constant $c_i > 0$ in (6) is used to make the expectation and variance of $c_i \bar{z}_i^{\mathsf{RI}}$ approximately equal to each other, and consequently to make the Anscombe VST a stabilize the signal $c_i \bar{z}_i^{\mathsf{RI}}$ [6]. To understand c_i and appreciate the degree of approximation, let us consider

$$\mathrm{E}\left\{c_i \bar{z}_i^{\mathsf{RI}} | y\right\} = c_i \lambda_i \mathrm{E}\{z^{\mathsf{RI}}|y\} + c_i(1-\lambda_i)\, g \circledast t^{\mathsf{RI}} = c_i g \circledast t^{\mathsf{RI}}, \tag{7}$$

$$\mathrm{var}\left\{c_i \bar{z}_i^{\mathsf{RI}} | y\right\} = c_i^2 \lambda_i^2 \mathrm{var}\{t^{\mathsf{RI}}|y\} = c_i^2 \lambda_i^2\, g \circledast (t^{\mathsf{RI}})^2 . \tag{8}$$

We can evaluate (7) at any $x_0 \in \Omega$ by taking the Maclaurin expansion of $g(x_0 - \cdot)$:

$$\mathrm{E}\left\{c_i \bar{z}_i^{\mathsf{RI}} | y\right\}(x_0) = c_i \sum_{x\in\Omega} g(x_0 - x)\, t^{\mathsf{RI}}(x) = \tag{9}$$

$$= c_i \sum_{x\in\Omega} \sum_{k=0}^{+\infty} \frac{\partial^k g(x_0)\, x^k}{(-1)^k\, k!} t^{\mathsf{RI}}(x) = \sum_{k=0}^{+\infty} c_i \frac{\partial^k g(x_0)}{(-1)^k\, k!} \sum_{x\in\Omega} x^k t^{\mathsf{RI}}(x) .$$

Similarly, for (8):

$$\mathrm{var}\left\{c_i \bar{z}_i^{\mathsf{RI}} | y\right\}(x_0) = \sum_{k=0}^{+\infty} c_i^2 \lambda_i^2 \frac{\partial^k g(x_0)}{(-1)^k\, k!} \sum_{x\in\Omega} x^k \left(t^{\mathsf{RI}}(x)\right)^2 . \tag{10}$$

Comparing the k-th summand in (9) with the corresponding summand in (10) we observe the following: if c_i is as in (6), then the first summands (i.e. $k = 0$) coincide; if t^{RI} is even symmetric, then the summands are zero for every odd k thanks to the odd symmetry of x^k. Thus, for any even-symmetric PSF v, $\mathrm{E}\{c_i \bar{z}_i^{\mathsf{RI}}|y\}(x_0)$ and $\mathrm{var}\{c_i \bar{z}_i^{\mathsf{RI}}|y\}(x_0)$ differ only for approximation terms of $g(x_0 - \cdot)$ of even order 2 or greater. Based on this analysis, we can expect effective stabilization of variance from applying f_i (6) to $\bar{z}_i^{\mathsf{RI}}$ (5), particularly where g is smooth and for symmetric PSFs.

3.4. Colored Noise Power Spectrum and Denoising

The variance of the stabilized $f_i(\bar{z}_i^{\mathsf{RI}})$ can be treated as approximately unitary. However, because VST is a pointwise transformation, the noise retains much of the spatial correlation given by $\lambda_i^2 \Psi^{\mathsf{RI}}$. We model the noise power spectrum after stabilization as

$$\tilde{\Psi}^{\mathsf{RI}} = |\Omega|^2 \Psi^{\mathsf{RI}} \| \Psi^{\mathsf{RI}} \|_1^{-1}, \tag{11}$$

i.e. as the scaling of $\lambda_i^2 \Psi^{\mathsf{RI}}$ that attains unitary spatial domain variance. We then denoise $\bar{\bar{z}}_i^{\mathsf{RI}} = f_i(\bar{z}_i^{\mathsf{RI}})$ with a filter Φ for colored noise:

$$D_i = \Phi\left[\bar{\bar{z}}_i^{\mathsf{RI}}, \tilde{\Psi}^{\mathsf{RI}}\right]. \tag{12}$$

When using transform-domain filters such as [9, 3], the power spectrum $\tilde{\Psi}^{\mathsf{RI}}$ determines the internal shrinkage thresholds for each transform coefficient. In our experiments we use the BM3D collaborative filter Φ for colored noise [3] to filter the stabilized data. The inputs of Φ are the stabilized image $\bar{\bar{z}}_i^{\mathsf{RI}}$ and the noise power spectrum $\tilde{\Psi}^{\mathsf{RI}}$.

Figure 1 provides a one-dimensional illustration of the analysis described in Sections 3.3 and 3.4. Observe that neither the identity between (7) and (8) nor the stabilization $\mathrm{var}\left\{f_i(\bar{z}_i^{\mathsf{RI}})\right\} = 1$ hold exactly although the prediction of the noise power spectrum can be quite accurate. These discrepancies are due to the fact that the stabilized noise follows in fact a more complex model than what is explained by the power spectrum $\tilde{\Psi}^{\mathsf{RI}}$ alone. In most cases we found these modeling imprecision to be acceptable, as reflected in the goodness of the experimental results for a broad set of natural images and peak values.

3.5. VST Inversion

After denoising, we apply the "exact unbiased" inverse $\mathcal{I}_a$ of the Anscombe transformation [10] scaled by the coefficient c_i^{-1}, which we denote by $\mathcal{I}_{f_i}^{\lambda_i}$. This gives an estimate of the blur- and noise-free image y as $\hat{y}_i = \mathcal{I}_{f_i}^{\lambda_i}(D_i) = c_i^{-1} \mathcal{I}_a(D_i)$. The inverse $\mathcal{I}_a$ had been designed for Poisson data and, even upon scaling, is *not* an exact unbiased inverse for data that underwent deconvolution (2) and convex combination (5). In fact, the expectation mapping $\mathrm{E}\left\{f_i(\bar{z}_i^{\mathsf{RI}})\middle| y\right\} \mapsto y$ is only approximate, even if we assume $\mathrm{E}\{\bar{z}_i^{\mathsf{RI}}|y\} = y$, and it is subject to approximation errors comparable in nature to those analyzed in Section 3.3 with regard to the forward VST.

3.6. Binning

At the first few iterations of the algorithm, we may not have accurate estimates of y. Following [2], we combine the convex combination (5) with binning, as a means to increase the SNR towards effective stabilization of variance. The binning operator $\mathcal{B}_{h_i}$ replaces a block of $h_i \times h_i$ pixels with their sum; therefore the binned image $\mathcal{B}_{h_i}[\bar{z}_i^{\mathsf{RI}}]$ has lower resolution than $\bar{z}_i^{\mathsf{RI}}$. Binning amounts to a convolution with a uniform $h_i \times h_i$ kernel followed by downsampling by h_i in both dimensions. Hence, as convolution is associative, it can be composed with t^{RI} in order to calculate both the power spectrum Ψ^{RI} of

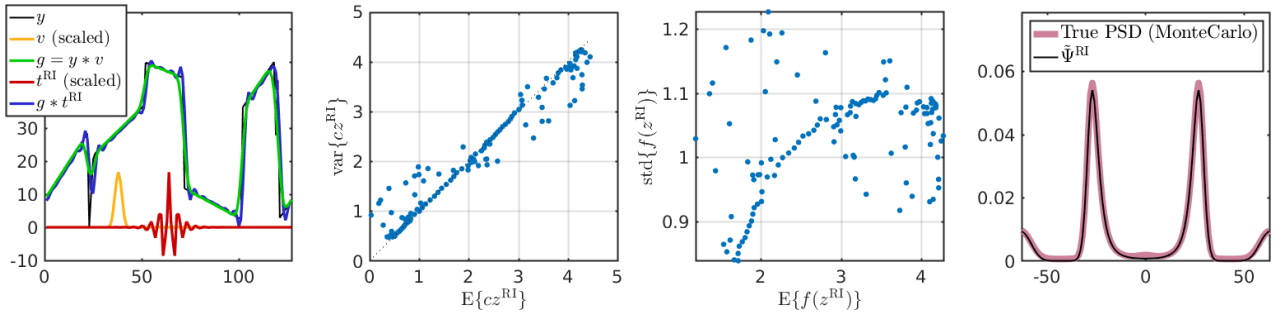

Fig. 1. One-dimensional illustration, from left to right: a piecewise-affine signal y undergoes blurring by a PSF v and Poisson noise, its expectation $g = y \circledast v$, the regularized deconvolution kernel t^{RI}, and $g \circledast t^{\mathrm{RI}} = \mathrm{E}\{z^{\mathrm{RI}}|y\}$ (3); expectation vs variance scatterplot of cz^{RI} (7),(8) (mostly distributed along the diagonal); expectation vs standard deviation scatterplot after VST (6) (showing approximate stabilization); model noise power spectrum $\tilde{\Psi}^{\mathrm{RI}}$ of $f(z^{\mathrm{RI}})$ showing good match with the empirical power spectral density (PSD) Monte Carlo estimate.

$\mathcal{B}_{h_i}[\bar{z}_i^{\mathrm{RI}}]$ (4),(5), and the suitable constant c_i for the VST to be applied to $\mathcal{B}_{h_i}[\bar{z}_i^{\mathrm{RI}}]$ (6). Aliasing caused by the downsampling must be accounted for when calculating Ψ^{RI}.

After denoising and inverse transformation, we upsample the image $\mathcal{I}_{f_i}^{\lambda_i}(D_i)$ via the debinning operator $\mathcal{B}_{h_i}^{-1}$, which yields a full-size image $\hat{y}_i$ such that $\mathcal{B}_{h_i}[\hat{y}_i] = \mathcal{I}_{f_i}^{\lambda_i}(D_i)$. Implementation details on the adopted binning and debinning operators can be found in [2].

4. EXPERIMENTS

4.1. Regularization term and second stage filtering

As a preliminary, let $\Gamma_\delta = \delta\,|V|\sqrt{|\Omega|\sum_{x\in\Omega} y(x)/\sum_{x\in\Omega} y(x)} - |V|^2$, where $\delta \geq 0$ is a fixed constant. If $\mathcal{E}^2 = \Gamma_\delta$, then by (4) $\Psi^{\mathrm{RI}} = \delta^{-2}|\Omega|^{-1}\left(\sum_{x\in\Omega} y(x)\right)^2$, which can be thought as the flat spectrum of white noise whose average spatial variance (possibly signal-dependent) is $\delta^{-2}|\Omega|^{-2}\left(\sum_{x\in\Omega} y(x)\right)^2$.

Based on the above, we adopt the following regularization term $\mathcal{E}$:

$$\mathcal{E}^2 = \max\left(0, \min\left(\epsilon_{\mathrm{RI}}^2 \sum_{x\in\Omega} y(x), \Gamma_\delta\right)\right), \qquad (13)$$

where $\epsilon_{\mathrm{RI}} > 0$ is a constant. For large enough δ, (13) yields the common Tikhonov-type regularized inverse used, *e.g.*, by [9, 3] for Gaussian deblurring. This inverse systematically introduces some small bias at every frequency. Up to a certain value of δ, the bias is regulated by the Tikhonov-type regularized inverse; for lower values of δ, variance and bias are then determined by δ; if $|V|$ is sufficiently large, the inverse at that particular frequency may be exactly unbiased. In practice, in the sums over Ω we can reliably use z as a surrogate for the true y. We fix $\delta = 3.2$, as this value was found to provide as much noise as the denoising filter could efficiently handle. To further improve the filtering results, we introduce a second filtering stage in which we perform the same denoising procedure, but we deconvolve z adopting the following regularization in (2):

$$\mathcal{E}^2 = \max\left(0, \min\left(|Y|^{-2}\,\epsilon_{\mathrm{RWI}}^2 \sum_{x\in\Omega} y(x), \Gamma_\delta\right)\right), \qquad (14)$$

where in practice we use the result obtained using (13) as a surrogate for the true y. For large δ, (14) yields the common regularized empirical-Wiener inverse [9, 3]. Like for (13), here too $\delta = 3.2$.

4.2. Algorithm parameters

Our current implementation[2] of Algorithm 1 is determined by six parameters: ϵ_{RI}, ϵ_{RWI} (regularization constants), K (number of iterations), λ_K, h_1, h_K (first and last bin sizes); other values of λ_i and h_i are defined as $\lambda_i = 1 - \frac{i-1}{K-1}(1-\lambda_K)$ and $h_i = \max\{h_K, h_1 - 2i + 2\}$. The image z and PSF v are the only input to our algorithm; similar to [2], all parameters are adaptively selected based on the quantiles of z, following a training over a set of four images (*Lena*, *Bridge*, *Man*, and *Peppers*) not included in the experiments shown below.

[2]Matlab software available at http://www.cs.tut.fi/~foi/invansc/

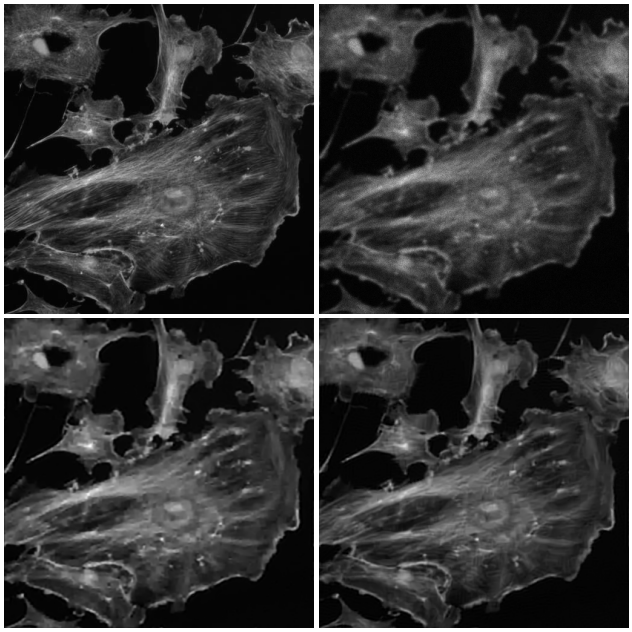

Fig. 2. Comparison of deblurring results. Left to right, top to bottom: *Fluocells* at peak 255, observed Poisson image blurred by a Gaussian PSF with variance 3 (PSNR 28.01 dB), restored by the PURE-LET [11] (31.42 dB), result of proposed method (31.67 dB).

4.3. Results

Table 1 compares the proposed method against four deblurring algorithms [11, 12, 13, 7] over a set of three images (*Cameraman*, *Fluocells*, and *Moon*) at six different peak values, blurred by a Gaussian PSF with variance 3. The PSNR (dB) results are averaged over 10 Poisson noise realizations. This experimental setup is reproduced from [11], where the PURE-LET algorithm recently demonstrated state-of-the-art results.

From the table we observe that the proposed method outperforms all other techniques, including [11], in all cases. The advantage is consistent throughout the table, and is especially signifi-

Method	Peak	*Cameraman*	*Moon*	*Fluocells*
Proposed	255	**24.54**	**27.75**	**31.66**
PURE-LET [11]		24.46	27.66	31.42
PoissonHessReg [12]		23.04	27.00	30.59
SPIRAL-TAP-TI [13]		24.06	25.61	30.46
PoissonDeconv [7]		22.78	25.03	30.96
Proposed	51	**23.55**	**26.68**	**29.84**
PURE-LET [11]		23.35	26.26	29.69
PoissonHessReg [12]		21.97	25.73	28.58
SPIRAL-TAP-TI [13]		23.12	24.15	28.59
PoissonDeconv [7]		22.07	24.77	27.75
Proposed	25.5	**23.03**	**26.06**	**29.03**
PURE-LET [11]		22.85	25.69	28.88
PoissonHessReg [12]		21.38	25.15	27.60
SPIRAL-TAP-TI [13]		22.22	24.93	28.05
PoissonDeconv [7]		21.57	24.62	27.19
Proposed	5.1	**21.67**	**24.89**	**27.07**
PURE-LET [11]		21.41	24.35	26.79
PoissonHessReg [12]		19.64	23.87	25.14
SPIRAL-TAP-TI [13]		20.91	22.01	26.39
PoissonDeconv [7]		18.65	20.66	21.86
Proposed	2.55	**21.15**	**24.24**	**26.11**
PURE-LET [11]		20.65	23.91	25.81
PoissonHessReg [12]		18.70	23.27	24.06
SPIRAL-TAP-TI [13]		20.30	21.68	25.17
PoissonDeconv [7]		15.03	15.28	18.51
Proposed	1.275	**20.32**	**23.79**	**25.15**
PURE-LET [11]		19.87	23.45	24.81
PoissonHessReg [12]		17.57	22.46	23.17
SPIRAL-TAP-TI [13]		18.26	19.11	23.18
PoissonDeconv [7]		10.86	11.95	15.90

Table 1. Poisson deblurring results (PSNR, dB, average over 10 noise realizations).

cant for the lower peak values, i.e. when noise is relatively stronger. Inspecting the visual comparison in Figure 2, we can see that the proposed method achieves higher PSNR by preserving more details. Figure 3 illustrates the effective restoration provided by the proposed method for a case with stronger blur and stronger noise.

With the proposed algorithm on a 3.4-GHz Intel i7 CPU, deblurring *Fluocells* (512×512) and *Cameraman* (256×256) at peak 25.5 takes, on average, 9.1 and 2.2 seconds, respectively. At peak 2.55 the times reduce to 1.4 and 0.4 seconds, respectively. The proposed algorithm is faster at low peaks because of the larger bin sizes $h_i, i = 1, \ldots, K$. We note that in all our experiments $K \leq 4$, which is drastically smaller than the number of iterations typically required by iterative deblurring schemes. These processing times can be roughly compared to those reported for PURE-LET [11], which is the fastest among the considered algorithms (7.15 seconds for *Fluocells* at peak 25.5 and 1.04 seconds for *Cameraman* at peak 2.55).

5. CONCLUSIONS

We presented a deblurring approach for Poisson data where a VST is applied between a preliminary linear deconvolution and a denoising step that filters the data according to a stationary spatially correlated noise model. To the best of our knowledge, this is the first attempt of using VST in this way for Poisson deblurring. Experimental results show that the proposed method achieves state-of-the-art results in Poisson deblurring, at a modest computational cost and using an off-the-shelf denoising filter for stationary Gaussian noise. The approach can be directly extended to 3D deconvolution by adopting the volumetric version of the filter [14]. Future research shall address the parameters selection with regard to the PSF type and further analysis of the approximation errors when inverting the VST.

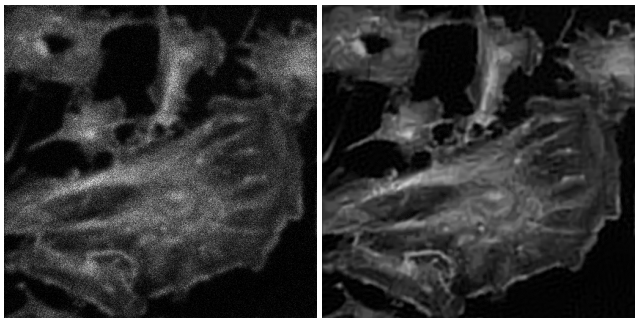

Fig. 3. Deblurring of *Fluocells* at peak 25.5. Left: observed Poisson image blurred by a Gaussian PSF with variance 9 (PSNR 20.52 dB). Right: restored by the proposed method (27.18 dB).

6. REFERENCES

[1] Y.-W. Tai, X. Chen, S. Kim, S. J. Kim, F. Li, J. Yang, J. Yu, Y. Matsushita, and M. S. Brown, "Nonlinear camera response functions and image deblurring: Theoretical analysis and practice," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 10, pp. 2498–2512, 2013.

[2] L. Azzari and A. Foi, "Variance Stabilization for Noisy+Estimate Combination in Iterative Poisson Denoising," *IEEE Sig. Process. Lett.*, vol. 23, no. 8, pp. 1086–1090, Aug. 2016.

[3] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image restoration by sparse 3D transform-domain collaborative filtering," in *Proc. SPIE Electronic Imaging 2008*, no. 681207, Jan. 2008.

[4] M. Almeida and M. Figueiredo, "Deconvolving images with unknown boundaries using the alternating direction method of multipliers," *IEEE Trans. Image Process.*, vol. 22, no. 8, pp. 3074–3086, 2013.

[5] F. J. Anscombe, "The transformation of Poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3/4, pp. 246–254, Dec. 1948.

[6] S. Bar-Lev and P. Enis, "On the construction of classes of variance stabilizing transformations," *Statistics & Probability Letters*, vol. 10, no. 2, pp. 95–100, 1990.

[7] F.-X. Dupé, J. M. Fadili, and J.-L. Starck, "A proximal iteration for deconvolving Poisson noisy images using sparse representations," *IEEE Trans. Image Process.*, vol. 18, no. 2, pp. 310–321, 2009.

[8] S. Harizanov, J. Pesquet, and G. Steidl, "Epigraphical projection for solving least squares anscombe transformed constrained optimization problems," in *Int. Conf. Scale Space and Variational Meth. Comp. Vis.* Springer, 2013, pp. 125–136.

[9] R. Neelamani, H. Choi, and R. Baraniuk, "ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems," *IEEE Trans. Signal Process.*, vol. 52, no. 2, pp. 418–433, 2004.

[10] M. Mäkitalo and A. Foi, "Optimal Inversion of the Anscombe Transformation in Low-Count Poisson Image Denoising," *IEEE Trans. Image Process.*, vol. 20, no. 1, pp. 99–109, Jan. 2011.

[11] J. Li, F. Luisier, and T. Blu, "Deconvolution of Poissonian images with the PURE-LET approach," in *2016 IEEE Int. Conf. Image Process.*, Sept. 2016, pp. 2708–2712.

[12] S. Lefkimmiatis and M. Unser, "Poisson image reconstruction with Hessian Schatten-norm regularization," *IEEE Trans. Image Process.*, vol. 22, no. 11, pp. 4314–4327, 2013.

[13] Z. Harmany, R. Marcia, and R. Willett, "This is SPIRAL-TAP: Sparse Poisson intensity reconstruction algorithms—theory and practice," *IEEE Trans. Image Process.*, vol. 21, no. 3, pp. 1084–1096, 2012.

[14] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi, "Nonlocal transform-domain filter for volumetric data denoising and reconstruction," *IEEE Trans. Image Process.*, vol. 22, no. 1, pp. 119–133, 2013.