PoGaIN: Poisson-Gaussian Image Noise Modeling from Paired Samples

Nicolas Bähler[1] Majed El Helou [1,2] Étienne Objois [1] Kaan Okumuş [1] Sabine Süsstrunk [1]

[1]Ecole Polytechnique Fédérale de Lausanne (EPFL) [2]Swiss Federal Institute of Technology in Zurich (ETHZ)

Abstract

Image noise can often be accurately fitted to a Poisson-Gaussian distribution. However, estimating the distribution parameters from a noisy image only is a challenging task. Here, we study the case when paired noisy and noise-free samples are accessible. No method is currently available to exploit the noise-free information, which may help to achieve more accurate estimations. To fill this gap, we derive a novel, cumulant-based, approach for Poisson-Gaussian noise modeling from paired image samples.

Motivation

Denoising in special case

- Restrict problem to subspace where paired samples are obtainable
- Additional noisy-free information may improve estimation performance

Figure 1. Noise-free and noisy image from BSD300 [2].

Problem at hand

- Perform Poisson-Gaussian noise estimation, often required prior to denoising
- Restricted problem setting: paired samples

$\implies$ Design noise modeling approach which is taking advantage of special setting!

Applications:

- Dataset modelling, deep learning
- Imaging system analysis

Results

Noise modeling:

- Statistical approach based on cumulant expansion
- Exploiting paired samples
- Outperforming baseline methods

Understanding of the problem:

- Derivation of log-likelihood function

Explored approaches

Using noisy images only:

- *FOI*, algorithm introduced by Foi *et al.* [1]
- *CNN*, Convolutional Neural Network

Using paired samples:

- Maximum log-likelihood
- *VAR*, method based on variance
- *OURS*, approach using the cumulant expansion

Noise Model

Inspired by Foi *et al.* [1], where x and y are the noise-free and noisy images respectively:

$$y = \eta_p + \eta_g, \tag{1}$$

$$\eta_p \sim \frac{1}{a}\mathcal{P}(ax), \quad \eta_g \sim \mathcal{N}(0, b^2), \quad a \in (0, 100]. \tag{2}$$

Here, a is equal to the quantum efficiency expressed as a percentage of the imaging pipeline at hand. Further, one can see in formula (3) that the expected value of the model is independent of the parameters but that the variance strongly depends on the values of a and b:

$$\mathbb{E}[y] = x, \quad \mathbb{V}[y] = \frac{x}{a} + b^2. \tag{3}$$

Log-Likelihood

We derive the log-likelihood for the given problem, but the optimization to find its maximum value is computationally demanding, rendering this approach unfeasible on a larger scale,

$$\mathcal{LL}(y|a,b,x) = \sum_n \log\left(\sum_{k=0}^{\infty} \frac{(ax_n)^k}{k! b\sqrt{2\pi}} \exp\left(-ax_n - \frac{(y_n - k/a)^2}{2b^2}\right)\right). \tag{4}$$

OURS

Given a random variable $X \sim \mathcal{X}$:

Definition

Cumulant-generating function:

$$K_{\mathcal{X}}(t) = \log(\mathbb{E}[e^{\mathcal{X}t}]). \tag{5}$$

Definition

r-th cumulant of $\mathcal{X}$:

$$\kappa_r[\mathcal{X}] = K_{\mathcal{X}}^{(r)}(0). \tag{6}$$

Unbiased estimator via k-statistics [3], given n samples:

$$\kappa_2[\mathcal{X}] = \frac{n}{n-1} m_2(x), \quad \kappa_3[\mathcal{X}] = \frac{n^2}{(n-1)(n-2)} m_3(x), \tag{7}$$

using the sample central moments [4]

$$m_2(x) = \frac{n-1}{n}\sum_i (x_i - \overline{x})^2, \quad m_3(x) = \frac{(n-1)(n-2)}{n^2}\sum_i (x_i - \overline{x})^3, \tag{8}$$

where $\overline{x}$ denotes the mean.

Let x_i be n pixels of a noise-free image, and X and Y be two random variables such that:

$$X \sim \mathcal{X}, \quad \mathbb{P}[X = x_i] = \frac{|\{k : x_k = x_i\}|}{n}, \quad Y \sim \mathcal{Y} = \frac{\mathcal{P}(a\mathcal{X})}{a} + \mathcal{N}(0, b^2). \tag{9}$$

We show that:

$$\begin{cases} \kappa_2[\mathcal{Y}] = \frac{\overline{x}}{a} + \overline{x^2} - \overline{x}^2 + b^2 \\ \kappa_3[\mathcal{Y}] = \overline{x^3} - 3\overline{x^2}\,\overline{x} + 2\overline{x}^3 + 3\frac{\overline{x^2}}{a} - 3\frac{\overline{x}^2}{a} + \frac{\overline{x}}{a^2} \end{cases}. \tag{10}$$

Hence, by using formula (7) one can estimate both $\kappa_2[\mathcal{Y}]$ and $\kappa_3[\mathcal{Y}]$ and insert it into formula (10) which leads to a system of two equations and two unknowns, a and b.

Estimation error

We obtain the mean squared error (MSE) shown below by picking 10 images from the validation set of BSD300 [2], 25 different values of a and b respectively, and also 10 different seeds to synthesize 62500 image pairs for validation. Note that the maximum log-likelihood approach is not evaluated here due to its computational complexity. Further, we show the MSE for $\hat{a}^{-1}$ and $\hat{b}^2$ because those are the values which are effectively estimated using our implementation before then either the inverse or the square root is taken, leading to a more direct evaluation of the estimation.

Figure 2. MSE for each method as a function of a (top) and b (bottom).

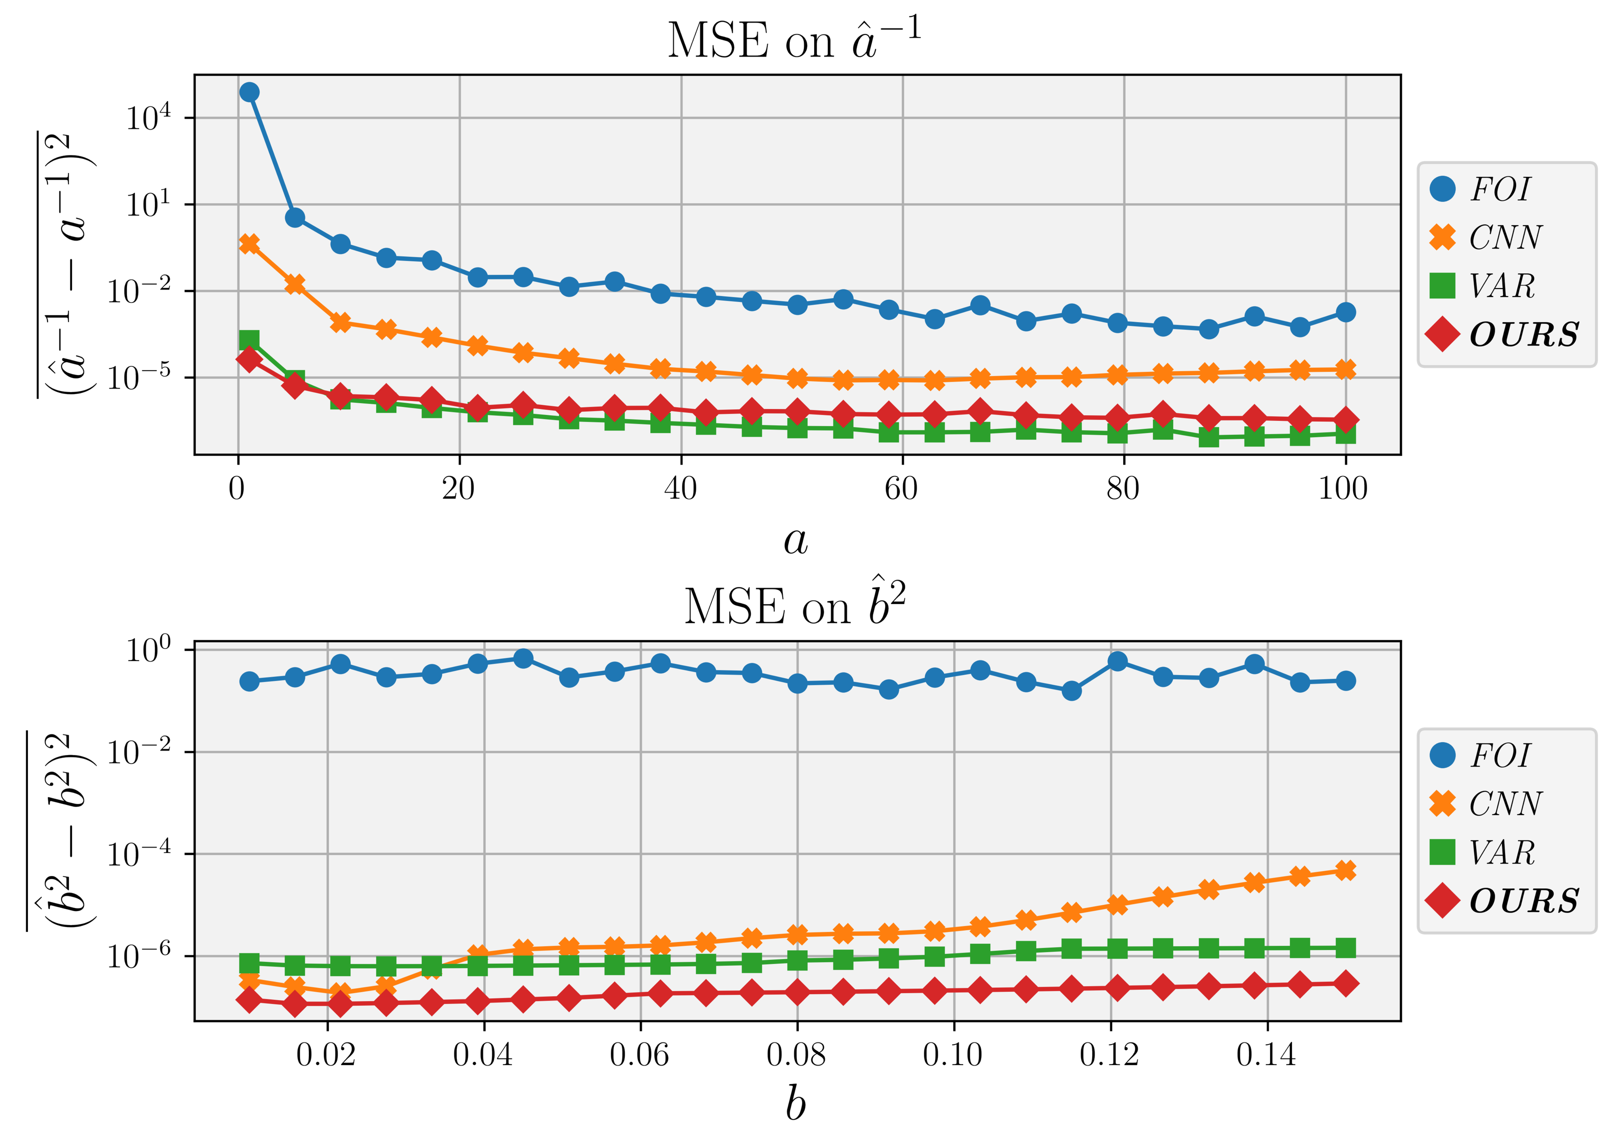

Table 1. Statistics about the MSE error on $\hat{a}^{-1}$ for various methods.

Method	Mean	Standard Dev.	75%-Quantile	Maximum
FOI	3.15×10^{3}	7.46×10^{5}	5.64×10^{-4}	1.86×10^{8}
CNN	1.78×10^{-2}	8.67×10^{-2}	7.40×10^{-5}	6.34×10^{-1}
VAR	8.00×10^{-6}	8.50×10^{-5}	≈ 0	3.54×10^{-3}
OURS	$\mathbf{3.00 \times 10^{-6}}$	$\mathbf{1.40 \times 10^{-5}}$	1.00×10^{-6}	$\mathbf{2.77 \times 10^{-4}}$

Table 2. Statistics about the MSE error on $\hat{b}^2$ for various methods.

Method	Mean	Standard Dev.	75%-Quantile	Maximum
FOI	3.46×10^{-1}	6.67	9.40×10^{-5}	6.16×10^{2}
CNN	8.00×10^{-6}	2.30×10^{-5}	5.00×10^{-6}	3.87×10^{-4}
VAR	1.00×10^{-6}	1.10×10^{-5}	≈ 0	4.45×10^{-4}
OURS	≈ 0	$\mathbf{1.00 \times 10^{-6}}$	≈ 0	$\mathbf{3.30 \times 10^{-5}}$

References

[1] Alessandro Foi, Mejdi Trimeche, Vladimir Katkovnik, and Karen Egiazarian. Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data. *IEEE Transactions on Image Processing*, 17(10):1737–1754, 2008.

[2] 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. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.

[3] Eric W. Weisstein. k-statistic from mathworld–a wolfram web resource.

[4] Eric W. Weisstein. Sample central moment. from mathworld–a wolfram web resource.