

# Evaluation of ADMM-based Poisson denoiser with different image priors on synthetic and real experimental images

Xiang Wu

**Abstract**—Image priors are powerful tools for solving many inverse problems such as imaging denoising. However, many existing priors have only been tested for Gaussian rather than Poisson noise. Even for those priors that have been tested for Poisson noise, usually synthetic data is used, which may not faithfully represent how these priors perform on real experimental images. Thus in this paper, I evaluated Poisson denoisers with different image priors based on an alternating direction method of multipliers (ADMM) framework for synthetic and real experimental images. Specifically, three different image priors, including L1-norm, total variation norm and Lysaker-Lundervold-Tai model, were implemented and tested. The results suggested that real experimental Poissonian images can be effectively denoised using ADMM-based denoisers with image priors, while images with different features prefer different image priors, as expected. As for the next step, different image priors may be combined into a joint regularization function, or applied in a sequential manner to further enhance the denoising performance.

**Index Terms**—Poisson Denoising, ADMM, Image Priors

## 1 INTRODUCTION

IN many photon-limited imaging systems, such as photon counting and fluorescence microscopy, Poisson noise is usually the dominant type of noise in the image. Unlike other noises such as thermal noise or read noise which can be greatly reduced via hardware engineering efforts, Poisson noise results from the quantum nature of light and thus always exists in images under low light conditions. Therefore, denoising Poissonian images necessitates post-processing algorithms, many of which use image priors as regularizers when solving the inverse problem.

Image priors are powerful tools for image denoising or solving other inverse problems as they generalize prior knowledge or models for a certain set of images. In brief, image priors summarize the like-hood information about the pixel distribution inside an image, thus allowing one to pick up the most likely solution among many feasible ones for the inverse problem. Although many different image priors have been proposed for image denoising, most of them have only been evaluated on Gaussian noise. Even for those existing papers working with Poisson noise, the noise distribution is usually simulated based on clean natural images and the denoisers were tested only on the synthetic data. Thus, it would be interesting to evaluate how the denoisers with different image priors/regularizers perform on real experimental Poissonian images with distinct features.

## 2 RELATED WORK

A great variety of image priors/regularizers have been proposed for image denoising. For example, total variation (TV), which is a very popular prior for natural images,

has been widely used for Gaussian and Poisson denoising [1,2]. However, TV prior is known to result in stair-casing artifacts [3], and thus other image priors such as the Lysaker-Lundervold-Tai (LLT) model [3], total generalized variation (TGA) [4], fractional-order total variation [5], and deep image prior [6] have been used to reduce the stair-casing artifacts for Gaussian or Poisson denoising. Furthermore, in other applications such as medical imaging or astronomy photography, the resulting images usually have different general features from natural images, and thus should benefit from other types of priors such as the L1-norm which promotes sparsity [7].

Specifically, in Ref [3], Lysaker et al. proposed the LLT model and applied the LLT-regularized denoisers for real magnetic resonance image (MRI) data. However, they assumed Gaussian noise rather than Poisson noise in this work. Furthermore, In Ref [5], Chowdbhury et al. generalized the TV and LLT priors to get fractional-order total variation prior, and tested their denoiser on Poissonian images. However, the Poisson noise in the image was simulated and no experiment images were tested in this work.

## 3 METHOD

In this paper, Poisson denoisers with different image priors based on an alternating direction method of multipliers (ADMM) framework were evaluated on both synthetic and real experimental images. Specifically, three different image priors were used as the regularizer, including L1-norm, TV norm and LLT norm. These three image priors generalize different set of images, with the L1-norm promoting sparsity in the image, the TV norm promoting sparsity in the first-order derivative and the LLT norm promoting sparsity in the second-order derivative. Furthermore, a diverse set of synthetic and experimental images with distinct features

- X. Wu is with the EE367 Winter 2022 course, Stanford University, Stanford, CA, 94305.  
E-mail: xiangwu@stanford.edu

were tested for the denoisers. One should expect the certain image prior matching the general features of the image should outperform the other two.

### 3.1 Construction of ADMM-based Poisson denoiser

The Poisson denoising problem is first tangled by taking a Bayesian probabilistic perspective [8]. Assuming Poisson noise is the only source of noise in the image, the image formation model can be formulated as:

$$b = \mathcal{P}(x)$$

where  $x$  is the ground truth image and  $b$  is the noisy measurement. The probability for a certain measurement at a particular pixel  $i$  can be expressed as:

$$P(b_i|x_i) = \frac{(x_i)^{b_i} e^{-x_i}}{b_i!}$$

Then the joint probability of all measurements can be expressed as:

$$P(b|x) = \prod_{i=1}^M P(b_i|x_i) = \prod_{i=1}^M \frac{e^{\log(x_i)b_i} e^{-x_i}}{b_i!}$$

Apply Baye's rule, one can get the maximum-a-posterior (MAP) solution as:

$$x_{MAP} = \operatorname{argmin}_x x^T \mathbf{1} - \log(x)^T b + \lambda \Psi(x)$$

where  $\Psi(x)$  is the regularization term and  $\lambda$  is the relative weight of the regularization term. The term  $x^T \mathbf{1} - \log(x)^T b$  is the "data fidelity term" for Poisson noise. This problem may be difficult to solve directly, but can be solved in an ADMM framework by constructing the ADMM Augmented Lagrangian as:

$$L_\rho(x, z, u) = \sum_{i=1}^3 g_i(z_i) + \frac{\rho}{2} \|Kx - z + u\|_2^2 - \frac{\rho}{2} \|u\|_2^2$$

$$g_1(z_1) = z_1^T \mathbf{1} - \log(z_1)^T b$$

$$g_2(z_2) = \mathcal{I}_{\mathbb{R}_+}(z_2) = \begin{cases} 0 & z \in \mathbb{R}_+ \\ +\infty & z \notin \mathbb{R}_+ \end{cases}$$

$$g_3(z_3) = \|z_3\|_1$$

$$K = \begin{pmatrix} I \\ I \\ D \end{pmatrix}, \quad z = \begin{pmatrix} z_1 \\ z_2 \\ z_3 \end{pmatrix}, \quad u = \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}$$

where  $g_1(z_1)$  is the data fidelity term,  $g_2(z_2)$  is the indicator function for the non-negativity constraint, and  $g_3(z_3)$  is the regularizer. The matrix  $D$  takes different forms for different image priors, which are discussed in details in the next subsection.

#### 3.1.1 Image priors and regularizers

When using the L1-norm as the image prior, which promotes sparsity in the image, the matrix  $D$  is simply the identity matrix  $I$ . When using the TV-norm as the image prior, which promotes sparsity in the first-order derivative, the matrix  $D$  should be:

$$D = \begin{pmatrix} D_x \\ D_y \end{pmatrix}$$

Using the duality between the signal processing notation and linear algebra notation, we have:

$$D_x x = d_x * x, \quad d_x = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & 0 \end{pmatrix}$$

$$D_y x = d_y * x, \quad d_y = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

When using the LLT-norm as the image prior, which promotes sparsity in the second-order derivative, the matrix  $D$  should be:

$$D = \begin{pmatrix} D_{xx} \\ D_{yy} \end{pmatrix}$$

where  $D_x$  and  $D_y$  are:

$$D_{xx} x = d_{xx} * x, \quad d_{xx} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{pmatrix}$$

$$D_{yy} x = d_{yy} * x, \quad d_{yy} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & -2 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

#### 3.1.2 ADMM update rules

$x$ -update:

$$x \leftarrow \operatorname{prox}(v) = \operatorname{argmin}_x \frac{\rho}{2} \|Kx - v\|_2^2, \quad v = z - u$$

This is just a least-square problem and can be solved as follows:

$$x = (K^T K)^{-1} K^T v = (I + I + D^T D)^{-1} (v_1 + v_2 + D^T v_3)$$

Using convolution theorem, we can avoid inverting large matrix, but instead doing the calculation in the Fourier domain based on the image prior as follows:

L1-norm:

$$x = (v_1 + v_2 + v_3)/3$$

TV-norm:

$$x = \mathcal{F}^{-1} \left( \frac{\mathcal{F}(v_1) + \mathcal{F}(v_2) + \mathcal{F}(d_x) * \mathcal{F}(v_3) + \mathcal{F}(d_y) * \mathcal{F}(v_4)}{1 + 1 + \mathcal{F}(d_x) * \mathcal{F}(d_x) + \mathcal{F}(d_y) * \mathcal{F}(d_y)} \right)$$

LLT-norm:

$$x = \mathcal{F}^{-1} \left( \frac{\mathcal{F}(v_1) + \mathcal{F}(v_2) + \mathcal{F}(d_{xx}) * \mathcal{F}(v_3) + \mathcal{F}(d_{yy}) * \mathcal{F}(v_4)}{1 + 1 + \mathcal{F}(d_{xx}) * \mathcal{F}(d_{xx}) + \mathcal{F}(d_{yy}) * \mathcal{F}(d_{yy})} \right)$$

(Note that for the TV-norm and LLT-norm, the original  $v_3$  was split into  $v_3$  and  $v_4$  for the ease of math and Python implementation.)

$z_1$ -update:

$$z_1 \leftarrow \text{prox}(v) = \underset{z_1}{\text{argmin}} (z_1)^T 1 - \log(z_1)^T b + \frac{\rho}{2} \|v - z_1\|_2^2$$

$$(v = x + u_1)$$

This can be solved by setting the gradient to be zero:

$$\nabla(J(z_1)) = 1 - \text{diag}(z_1)^{-1} b + \rho(z_1 - v) = 0$$

Note that the solution for the individual elements in  $z_1$  is independent from each other, thus this can be solved in an element-wise manner:

$$1 - b_i/(z_1)_i + \rho((z_1)_i - v_i) = 0$$

$$\rho(z_1)_i^2 + (1 - \rho v_i)(z_1)_i - b_i = 0$$

This is just a simple quadratic problem with the solution of:

$$(z_1)_i = -\frac{1 - \rho v_i}{2\rho} + \sqrt{\left(\frac{1 - \rho v_i}{2\rho}\right)^2 + \frac{b_i}{\rho}}$$

Note that only the positive solution is taken due to the non-negativity constraint.

$z_2$ -update:

$$z_2 \leftarrow \text{prox}(v) = \underset{z_2}{\text{argmin}}_{\mathbb{R}_+} \mathcal{I}_{\mathbb{R}_+}(z_2) + \frac{\rho}{2} \|v - z_2\|_2^2$$

$$(v = x + u_2)$$

As the indicator function goes to infinity whenever  $z_2 < 0$ , this is rather straight-forward:

$$(z_2)_i = \begin{cases} 0 & v_i < 0 \\ v_i & v_i \geq 0 \end{cases}$$

$z_3$ -update:

$$z_3 \leftarrow \text{prox}(v) = \underset{z_3}{\text{argmin}} \lambda \|z_3\| + \frac{\rho}{2} \|v - z_3\|_2^2 = S_{\lambda/\rho}(v)$$

$$(v = Dx + u_3)$$

Note that for ease of implementation of TV-norm and LLT-norm in Python, this element-wise soft thresholding  $z_3$  update was split into two parts corresponding to the two matrix components of the block matrix  $D$ .

$u$ -update:

$$z_3 \leftarrow u + Kx - z$$

### 3.2 Data preparation and collection

Three sets of synthetic images and four sets of experimental images were used in this project. For the synthetic images, the original clean image was first normalized against its possible maximum value (e.g. 255 for an 8-bit image), and then a noisy Poissonian image was simulated by assuming the peak intensity in the image is 5.

As for the experimental images, an electron-multiplying charge-coupled device (EMCCD) was used for data collection. The sensor chip of the EMCCD camera was deep-cooled to -70 °C to eliminate most of the thermal noise and read noise, leaving the Poisson noise the dominate noise in the images. For the noisy Poissonian image, the illumination condition was set to be very weak, and the exposure time

of the EMCCD camera was set to be 10 ms. Afterwards, another image with 10-s exposure time was taken under the same lighting condition, and was used as the "ground-truth" image. Specifically, the following sets of experimental images were collected: 1) Natural scenes of objects; 2) Microscopic image of United States Air Force (USAF) target in the transmission mode. 3) Microscopic image of microparticles in the dark-field mode. The microscopic images for 2) and 3) were acquired on a Leica DM 2700M microscope.

## 4 ANALYSIS AND EVALUATION

The Poissonian synthetic or experimental images were used as the input for the ADMM-based Poisson denoisers with three different images priors. The output images were then compared with the ground-truth image both qualitatively and quantitatively. On one hand, the denoised images based on different images priors were visually compared with the noisy images and the ground-truth images for evaluation. On the other hand, the peak signal-to-noise-ratio (PSNR) values between the denoised images and the ground-truth images were computed to quantitatively evaluate the denoiser performance on different types of images. The PSNR value was calculated using the mean squared error (MSE) as follows:

$$MSE = \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n [I_{gt}(i, j) - I_{output}(i, j)]^2$$

$$PSNR = 10 \log_{10} \left( \frac{\max(I_{gt})^2}{MSE} \right)$$

where  $I_{gt}$  is the ground-truth image,  $I_{output}$  is the denoised image, and  $i$  and  $j$  represent the column and row index of the image, respectively.

Furthermore, I also evaluated how fast each denoiser converge by calculating the PSNR values as a function of iteration number. Additionally, two parameters,  $\lambda$  (relative weight of the image prior) and  $\rho$  (relative weight of the ADMM penalty term), were manually tuned for each image prior on each set of synthetic and experimental image so that the PSNR value of the denoised image sits near the local (if not global) maxima in the parameter space. The iteration number was fixed at 75 for all images and denoisers.

## 5 EXPERIMENTAL RESULTS

### 5.1 Evaluation on synthetic data

The three sets of synthetic images, including the original clean image, the simulated Poisson image, and three denoised images using different image priors are shown in **Fig. 1** below. These denoised results for simulated Poisson noise revealed several key findings:

First, the denoised images using L1-norm as the image prior (**Fig. 1**, the third column) performed the worst among all three priors, as evidenced by the incremental improvement in the image qualities and PSNR values for all three synthetic data sets. Specifically, the L1-norm regularized images just look like another Poissonian image with slightly higher photon counts. This actually makes sense, since none of these three synthetic images is sparse, thus violating the assumption behind the L1-norm.

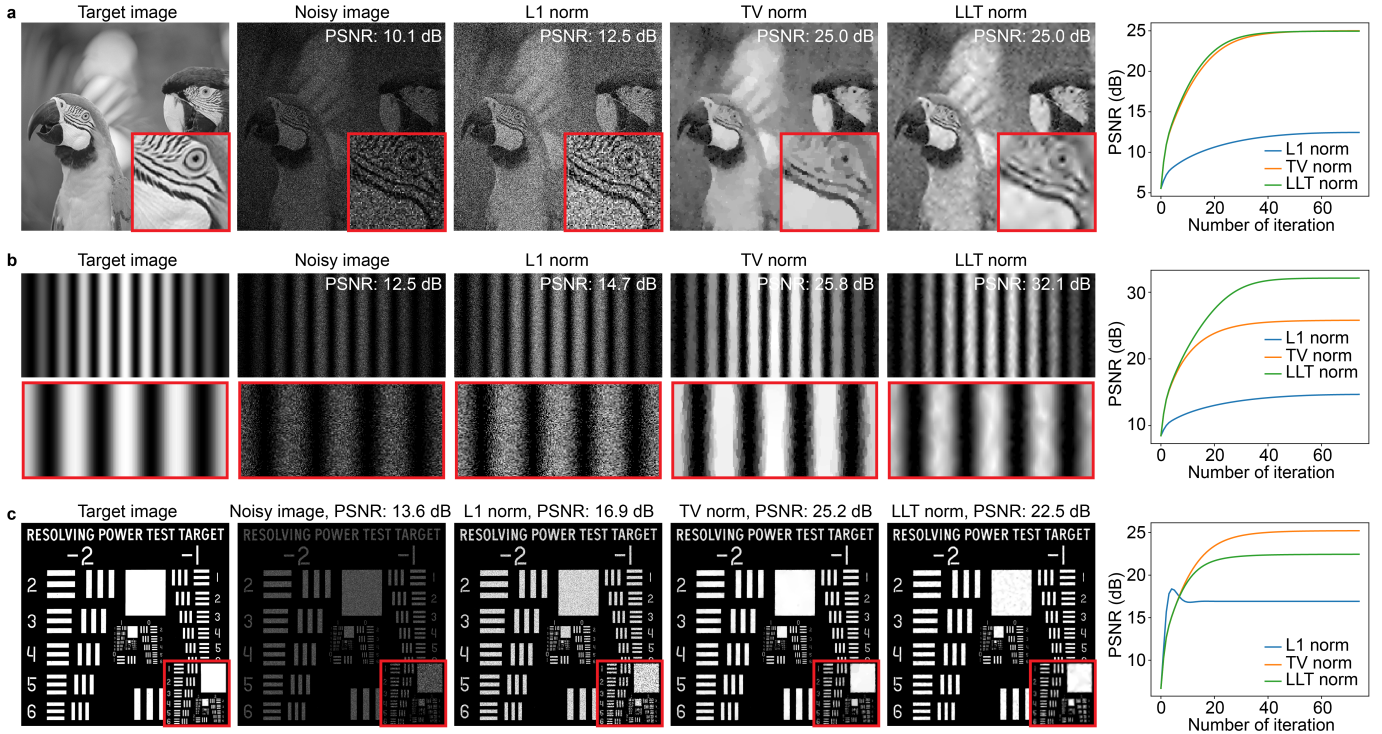


Fig. 1. Evaluation of ADMM-based Poisson denoisers on different synthetic images, including birds (a), interference pattern (b) and USAF targets (c). The Poisson noise was simulated on the normalized clean image by assuming a peak intensity of 5. The images from left to right are: the original clean image (the ground-truth image), the synthetic Poissonian image, the denoised image with L1-norm, the denoised image with TV norm and the denoised image using LLT norm. The insets with red boxes on the bottom right of each image show the zoom-in view at specific locations. The PSNR values for the noisy and denoised images with respect to the ground-truth images are labeled on the image title or in the right top corner of the image. The last column shows PSNR values of the denoised image as a function of iteration number for all three image priors (L1-norm: blue; TV norm: orange; LLT norm: green).

Second, both the TV norm and the LLT norm regularize the images relatively well, resulting in denoised images quite similar to the ground-truth images visually (Fig. 1, the fourth and fifth columns). The PSNR values also confirm the substantial improvement in the imaging quality after denoising with TV or LLT priors. This confirms the utilities of TV norm and LLT norm as regularizers for Poisson denoising problem.

Third, the TV norm and LLT norm introduce different image artifacts. This can be clearly seen in the birds image (Fig. 1a, the fourth and fifth columns), where the TV-denoised image shows obvious staircasing artifact in the face of the bird. This staircasing artifact from TV norm has been reported previously [3], and the LLT norm was proposed to address this issue. Indeed, the LLT-denoised image does not show any staircasing artifact and seems to reconstruct the strips and the eye on the bird face better. However, the LLT norm introduces speckle patterns, which are seen as blurry black dots.

Fourth, images with different features prefer different image priors. The birds image (Fig. 1a), which is a representative natural scene, seems to have no preference for TV or LLT prior. This is also shown by the almost identical PSNR values as a function of iteration numbers for these two priors. The interference pattern image (Fig. 1b) however, has very strong preference for LLT prior over TV prior, as evidenced by the much better visual image quality after denoising and also the much higher PSNR values for LLT norm.

This result is expected, since there is always gradual change of intensity for this interference pattern, which does not fit the piece-wise constant assumption of the TV norm very well. As a result, the denoised image based on TV norm has a quite large area with constant intensity at the interference peak. The LLT norm however, promotes linear variation of the intensity, and thus reconstructs the interference pattern much better. In contrast, the USAF synthetic image, which can be considered as an ideal case with only piece-wise constant components, prefers TV norm over LLT norm. The denoised USAF target image based on TV norm is almost identical to the ground-truth image visually, while the one based on LLT norm has some speckle artifacts that degrades the image quality. Furthermore, it is also interesting to see that despite this preference of different image priors for different images, the PSNR iteration curves for TV norm and LLT norm almost overlay with each other for small iteration numbers. This is probably because the ADMM-denoisiers primarily reconstruct the key features inside the image that are promoted by both the priors. It would then be interesting to compare these denoised images at early iteration stage before the convergence to see which features are reconstructed first, which is beyond the scope of this project and remains to be explored in the future.

## 5.2 Evaluation on real experimental data

Having validated the performance of the constructed Poisson denoisers on synthetic data, I next sought to evaluate

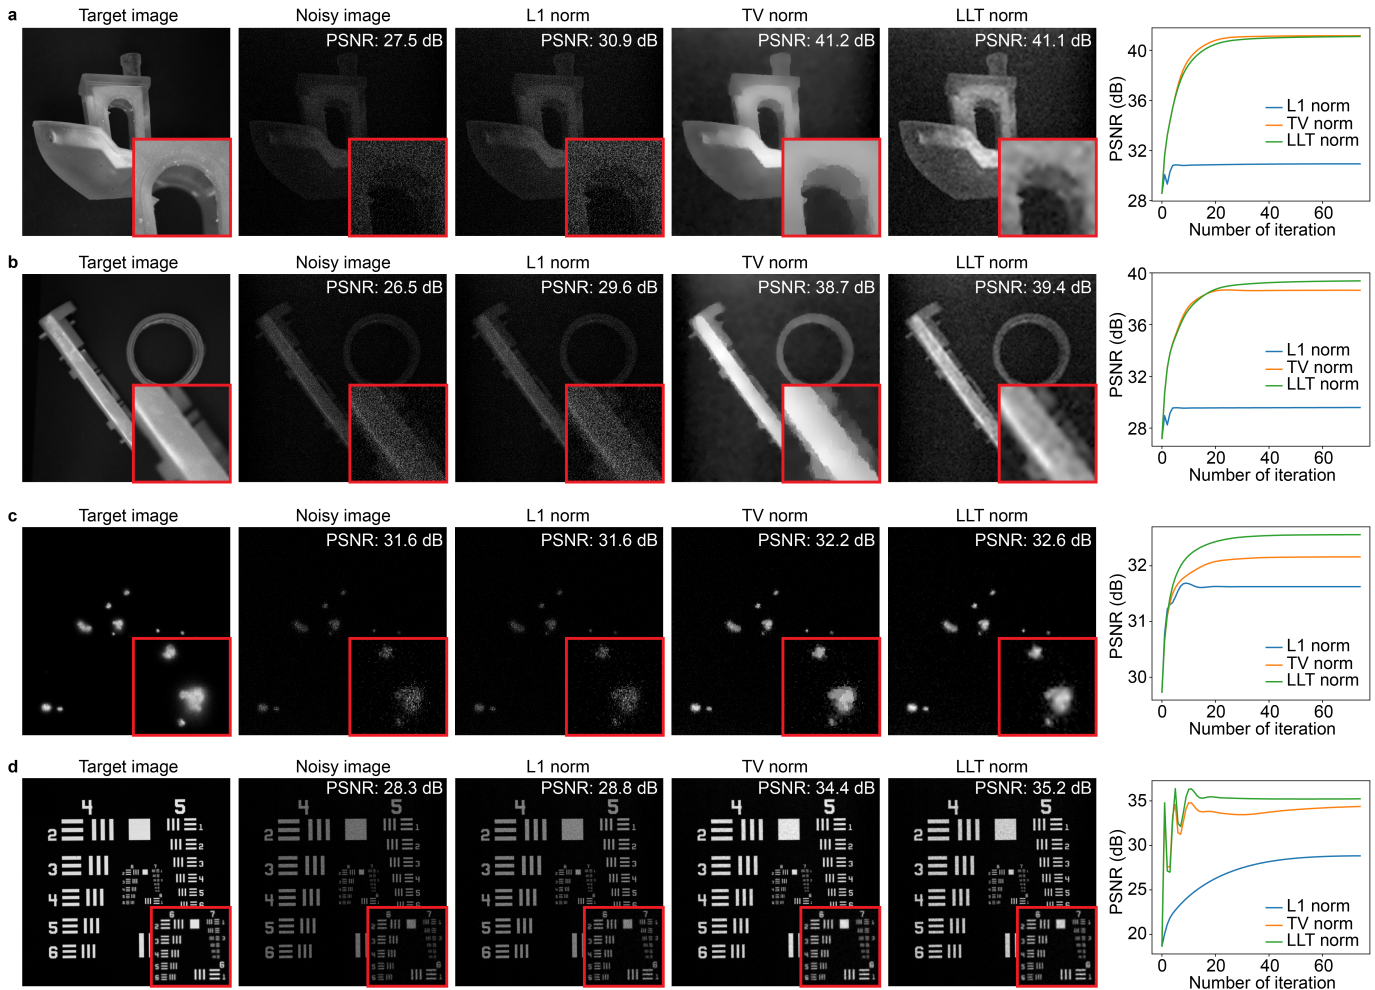


Fig. 2. Evaluation of ADMM-based Poisson denoisers on different experimental images, including 3D-printed boat (a), toolbox level and tape (b), dark-field microscopy image of microparticles (c) and transmission microscopy image of USAF target (d). The images from left to right are: the image with 10-s exposure time (the ground-truth image), the Poissonian image with 10-ms exposure time, the denoised image with L1-norm, the denoised image with TV norm and the denoised image using LLT norm. The insets with red boxes on the bottom right of each image show the zoom-in view at specific locations. The PSNR values for the noisy and denoised images with respect to the ground-truth images are labeled in the right top corner of the image. The last column shows PSNR values of the denoised image as a function of iteration number for all three image priors (L1-norm: blue; TV norm: orange; LLT norm: green).

their denoising performance on real experimental data. The four sets of experimental images, including the long-exposure ground-truth image, the short-exposure noisy image, and three denoised images using different image priors are shown in Fig. 2 above. These denoised results for experimental Poissonian images revealed several key findings:

First, similar to the results on synthetic images, the ADMM denoiser using L1-norm as the prior does not denoise the image very well (Fig. 2, the third column). This is a little bit surprising for the particle dark-field image (Fig. 2c), which I hoped to capture a sparse image to fit with the assumption behind the L1-norm. This unexpectedly poor denoising performance of the L1-norm prior on the particle dark-field image suggests that this image is not truly sparse. Indeed, each individual particles seem to be comprised of many pixels.

Second, different images prefer different priors, as we have seen earlier for the synthetic images. Specifically, the denoiser with TV prior performs slightly better than the one with LLT prior on the 3D-printed boat image (Fig.

2a) as it converges faster and the final image looks cleaner visually. This is expected since there are many piece-wise constant components in this boat image. The toolbox level and tape image (Fig. 2b) however, prefers LLT norm over TV norm. This is mainly because there are many specular reflection features on the toolbox level, which involve large variation of the intensity within a few pixels. Such strong spatial intensity variation cannot be reconstructed by the TV norm, which just smooths out these features into a large piece of constant value, but fits well with the linear variation assumption behind the LLT norm, which results in better denoising performance. Similar phenomenon is also observed for the particle dark-field image (Fig. 2c), which prefers LLT norm over TV norm. The staircasing artifact in the TV-denoised image is very obvious, as a single particle is split into a few different components, each with a relatively constant value. The LLT norm on the contrary, results in a much smoother particle appearance.

Third, the LLT norm outperforms TV norm for denoising

the experimental images of USAF target (Fig. 2d), which is the opposite for the synthetic data in Fig. 1c and also quite unexpected. At a closer look, one can actually see speckle artifacts in the TV-denoised image, although the amplitude is smaller than that in the LLT-denoised image. The PSNR iteration curves also look weird, which experience a few sharp peaks as the iteration proceeds before final convergence. It would be interesting to take a look at these denoised images at the early stage of iteration, which might be helpful for understanding this unusual behavior.

## 6 DISCUSSION AND CONCLUSION

In this project, I constructed and evaluated ADMM-based Poisson denoisers with three different image priors on both synthetic and experimental images. The results suggest that both TV norm and LLT norm are good image priors for denoising Poissonian images in an ADMM framework, while each of them have their particular advantages and disadvantages. The TV norm promotes piece-wise constant, thus performs better on images with large area of constant values, such as the synthetic USAF target image (Fig. 1c) and the experimental image of 3D-printed boat (Fig. 2a). This assumption behind the TV norm, however, also causes the staircasing artifacts and smear out some gradually varying features in the image. The LLT norm on the contrary, promotes linear variation of the intensity and thus outperforms TV norm on the images with matching features such as the synthetic interference pattern (Fig. 1b) and the experimental image of the toolbox level and tape (Fig. 2b). Although LLT norm helps eliminate the staircasing artifact, it introduces speckle artifacts.

Furthermore, the denoiser with L1-norm as the image prior does not perform well on any set of synthetic or experimental image. This is quite surprisingly for experimental dark-field image of microparticles (Fig. 2c), as it was intended to be used as a sparse image. This is probably because Fig. 2c is still not sparse enough, as each individual particle comprise of multiple pixels. As a future work, I would like to combine L1-norm with the TV norm or LLT norm as a joint regularier and see if that enables better denoising performance for this type of "pseudo-sparse" images.

Despite the fact that the L1-norm prior fails to denoise the images well, it indeed results in some improvement in the PSNR values, and the output image looks like another Poissonian image with higher photon counts (Fig. 1,2, the third column). This leads me to think that it could be interesting to use these L1-norm-denoised image as the input for the other ADMM-based Poisson denoisers with TV or LLT prior. Maybe this sequential denoising algorithm can result in better denoising performance.

Additionally, I also identified another interesting future direction when working on this project, which is to stop the iteration process at some early stages and look at the denoised images. This was inspired by the almost identical PSNR iteration curves for the TV norm and LLT norm for most of the images at the early iteration stage, suggesting that some key common features were reconstructed first by both the image priors. This early iteration study may also be helpful for understanding the unusual fluctuation in the

PSNR iteration curve for the experimental image of USAF target (Fig. 2d).

Lastly, one may also notice that in general the noisy experimental images have a much higher PSNR values than the noisy synthetic images, although some of the experimental images look much more noisier (such as the 3D-printed boat image in Fig. 2a). This might be because the background for the experimental images is always dark, which might seem to fit well with the ground-truth image even without any denoising. This can also be seen in the PSNR iteration curve, where even the initial guess of all-zeros give higher PSNR values for the experimental images than for the synthetic images.

To conclude, I have evaluated three different image priors within an ADMM framework for denoising synthetic and experimental Poissonian images. TV-norm and LLT-norm are both good priors for Poisson denoising problem, each with its particular advantages but also introducing different types of artifacts. For the future direction, I may combine different image priors as a joint regularizer or apply these denoisers sequentially to hopefully improve the denosing performance. It would also be interesting to stop the iteration earlier to get some insight about the primary common reconstructed features by different image priors.

## REFERENCES

- [1] S. V. Venkatakrishnan, C. A. Bouman and B. Wohlberg, "Plug-and-Play priors for model based reconstruction," in *2013 IEEE Global Conference on Signal and Information Processing*, 2013, pp. 945-948.
- [2] M. A. T. Figueiredo and J. M. Bioucas-Dias, "Restoration of Poissonian Images Using Alternating Direction Optimization," in *IEEE Transactions on Image Processing*, vol. 19, no. 12, pp. 3133-3145, 2010.
- [3] M. Lysaker, A. Lundervold and Xue-Cheng Tai, "Noise removal using fourth-order partial differential equation with applications to medical magnetic resonance images in space and time," in *IEEE Transactions on Image Processing*, vol. 12, no. 12, pp. 1579-1590, 2003.
- [4] K. Bredies, K. Kunisch and T. Pock, "Total Generalized Variation," in *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 492-526, 2009.
- [5] M. R. Chowdhury, J. Zhang, J. Qin, and Y. Lou, "Poisson image denoising based on fractional-order total variation," in *Inverse Problems and Imaging*, vol. 14, no. 1, pp. 77-96, 2020.
- [6] Y. Jo, S. Y. Chun, J. Choi, "Rethinking Deep Image Prior for Denoising," in *arXiv*, 2021.
- [7] G. Wetzstein, "Lecture 10: Image Deconvolution with the Half-quadratic Splitting (HQS) Method," in *Stanford EE367 Lecture Notes*, 2022.
- [8] G. Wetzstein, "Noise, Denoising, and Image Reconstruction with Noise," in *Stanford EE367 Course Notes*, 2022.