Algorithms for Poisson Phase Retrieval

Zongyu Li, *Student Member, IEEE,* Kenneth Lange and Jeffrey A. Fessler, *Fellow, IEEE*

***Abstract*—This paper discusses algorithms based on maximum likelihood (ML) estimation for phase retrieval where the measurements follow independent Poisson distributions. To optimize the log-likelihood for the Poisson ML model, we investigated and implemented several algorithms including a modified Wirtinger flow (WF), majorize minimize (MM) and alternating direction method of multipliers (ADMM), and compared them to the classical WF and Gerchberg Saxton (GS) methods for phase retrieval. Our modified WF approach uses a step size based on the observed Fisher information, eliminating all parameter tuning except the number of iterations. Simulation results using random Gaussian sensing matrix and discrete Fourier transform (DFT) matrix under Poisson measurement noise demonstrated that algorithms based on the Poisson ML model consistently produced higher quality reconstructions than algorithms (WF, GS) derived from Gaussian noise ML models when applied to such data. Moreover, the reconstruction quality can be further improved by adding regularizers that exploit assumed properties of the latent signal/image, such as sparsity of finite differences (anisotropic total variation (TV)) or of the coefficients of a discrete wavelet transform. In terms of the convergence speed, the WF using observed Fisher information for step size decreased NRMSE the fastest among all unregularized algorithms; the regularized WF approach also converged the fastest among all regularized algorithms with the TV regularizer approximated by the Huber function.**

***Index Terms*—Poisson phase retrieval, non-convex optimization.**

I. Introduction

Phase retrieval is an inverse problem that has many applications in engineering and applied physics [1, 2], including radar [3], X-ray crystallography [4], astronomical imaging [5] and speech processing [6]. The goal is to recover the signal from only the magnitude of linear measurements such as the Fourier transform [7]. This paper focuses on cases where the measurements follow independent Poisson distributions.

Zongyu Li and Jeffrey A. Fessler are with Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122 (e-mail: zonyul@umich.edu, fessler@umich.edu).

Kenneth Lange is with Departments of Computational Medicine, Human Genetics, and Statistics, University of California, Los Angeles, CA 90095 (e-mail: klange@ucla.edu).

Research supported in part by USPHS grants GM53275 and HG006139, and by NSF Grant IIS 1838179. Code for reproducing the results will be available on https://github.com/ZongyuLi-umich after the paper is accepted.

In most previous works, the measurement vector $\boldsymbol{y} \in \mathbb{R}^M$ was assumed to have statistically independent elements following Gaussian distributions:

$$y_i \sim \mathcal{N}(|\boldsymbol{a}_i' \boldsymbol{x}|^2 + b_i, \sigma^2), \tag{1}$$

where $\boldsymbol{a}_i' \in \mathbb{C}^N$ denotes the ith row of the system matrix $\boldsymbol{A} \in \mathbb{C}^{M \times N}$, where $i = 1, \ldots, M$. $\boldsymbol{x} \in \mathbb{F}^N$ denotes the true unknown signal, and b_i denotes a known mean background signal for the ith measurement, e.g., as arising from dark current [8]. Here the field $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$ depending on whether $\boldsymbol{x}$ is known to be real or complex. For this Gaussian noise model, the ML estimate of $\boldsymbol{x}$ corresponds to the following non-convex optimization problem

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x} \in \mathbb{F}^N}{\arg\min}\, g(\boldsymbol{x}), \;\; g(\boldsymbol{x}) \triangleq \sum_i \left| y_i - b_i - \left| \boldsymbol{a}_i' \boldsymbol{x} \right|^2 \right|^2. \tag{2}$$

To solve (2), numerous algorithms have been proposed, such as Wirtinger Flow [9], Gerchberg Saxton [10] and majorize-minimize (MM) methods [7].

However, in some low-photon count applications [11–16], a Poisson noise model is more appropriate:

$$y_i \sim \text{Poisson}(|\boldsymbol{a}_i' \boldsymbol{x}|^2 + b_i). \tag{3}$$

The ML estimation for (3) corresponds to the following optimization problem

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x} \in \mathbb{F}^N}{\arg\min}\, f(\boldsymbol{x}), \quad f(\boldsymbol{x}) \triangleq \sum_i \psi(\boldsymbol{a}_i' \boldsymbol{x}; y_i, b_i),$$
$$\psi(v; y, b) \triangleq (|v|^2 + b) - y \log(|v|^2 + b). \tag{4}$$

For (4), one can verify that the function

$$h(r; y, b) \triangleq (r^2 + b) - y \log(r^2 + b), \quad r \in \mathbb{R}, \tag{5}$$

is non-convex when $0 < b < y$. That property, combined with the modulus within the logarithm in (4), makes (4) a challenging optimization problem.

Similar problems for $b = 0$ have been considered previously [17–21]. Moreover, many optical sensors also have Gaussian readout noise [22, 23], so a more precise model would consider a sum of Gaussian and Poisson noise. However, the log likelihood for a Poisson plus Gaussian distribution is complicated, so a common approximation is to use a shifted Poisson model [24] that also leads

to the cost function in (4). An alternative to the shifted Poisson model could be to work with an unbiased inverse transformation of a generalized Anscombe transform approximation [25]. Such algorithms for a Poisson plus Gaussian noise model can be interesting directions for future work but will not be discussed in this paper.

In (4), we treat $0 \log 0$ as 0 because a Poisson random variable with zero mean can only take the value 0. In other words, if $|v| = b = 0$, then $\psi(v; y, b) = 0$. With this assumption, one can verify that ψ has a well-defined ascent direction (negative of descent direction [26]):

$$\dot{\psi}(v; y, b) = 2v \left(1 - \frac{y}{|v|^2 + b}\right), \quad v \in \mathbb{C}. \tag{6}$$

If $b > 0$, one can derive the following upper bound for the magnitude of the second-order ascent direction of ψ:

$$\ddot{\psi}(v; y, b) = \operatorname{sign}(v) \left(2 + 2y \frac{|v|^2 - b}{(|v|^2 + b)^2}\right),$$
$$|\ddot{\psi}(v; y, b)| \leq 2 + \frac{y}{4b}. \tag{7}$$

This bounded curvature property enables us to derive an MM algorithm [27] with a quadratic majorizer for (4).

For $b = 0$, one can verify that $\ddot{\psi}$ is unbounded so it is impossible to build a quadratic majorizer. Thus, we also investigated an alternating direction method of multipliers (ADMM) (Section III-C) adopted from [28]. The ADMM algorithm is mainly for completeness because in most real-world applications, background counts are rarely to be zero (e.g., considering naturally occurring background radiation).

The rest of this paper is organized as follows. Section II gives a brief overview of classical algorithms like WF and GS for phase retrieval. Section III introduces algorithms for Poisson phase retrieval, namely a modified Wirtinger flow method, a new MM approach, and a small variation of the ADMM algorithm of [28]. Section IV illustrates the implementation details of algorithms discussed in Section III. Section V provides numerical results using simulated data under different experimental settings. Section VI concludes this paper and provides future directions.

Notation: Bold upper/lower case letters (e.g., $\boldsymbol{A}$, $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{b}$) denote matrices and column vectors, respectively. Italics (e.g., μ, y, b) denote scalars. y_i and b_i denote the ith element in vector $\boldsymbol{y}$ and $\boldsymbol{b}$, respectively. $\mathbb{R}^N$ and $\mathbb{C}^N$ denote N-dimensional real/complex normed vector space, respectively. $(\cdot)^*$ denotes the complex conjugate and $(\cdot)'$ denotes Hermitian transpose. $\operatorname{diag}\{\cdot\}$ is a diagonal matrix constructed from a column vector. Unless otherwise defined, a subscript denotes outer iterations and superscript denotes the inner iterations, respectively. For example, $\boldsymbol{x}_k^i$ denotes the estimate of $\boldsymbol{x}$ at the ith inner iteration of the kth outer iteration. For gradients associated with complex numbers/vectors, the notation $\dot{\psi}(\cdot)$ and $\nabla(\cdot)$, should be considered as an ascent direction, not as a derivative.

II. Classic Methods for Gaussian model

A. Wirtinger Flow

In the classic Wirtinger flow (WF) algorithm [9], the gradient[1] for the Gaussian cost function (2) is

$$\nabla g(\boldsymbol{x}) = 4\boldsymbol{A}' \operatorname{diag}\{|\boldsymbol{A}\boldsymbol{x}|^2 - \boldsymbol{y} + \boldsymbol{b}\}\boldsymbol{A}\boldsymbol{x}. \tag{8}$$

To descend the cost function, reference [9] used a heuristic where the step size μ is rather small for the first few iterations and gradually becomes larger as the number of iterations increases. The intuition is that the gradient is noisy at the early iterations so a small step size is preferred. A drawback of this approach is that one needs to select hyper-parameters that control the growth of μ. An alternative approach is to perform backtracking for μ at each iteration [7], i.e., by reducing μ until the cost function decreases sufficiently. This approach guarantees decreasing the cost function monotonically but can increase the compute time of the algorithm due to the variable number of inner iterations. Section III will introduce a novel method to select the step size for the classical WF algorithm (8).

B. Gerchberg Saxton

An alternative to the intensity model (2) is the magnitude model that works with the square root of $\boldsymbol{y}$. In particular, by introducing a new variable $\boldsymbol{\theta}$ to represent the phase, one can derive the following optimization problem:

$$\hat{\boldsymbol{x}}, \hat{\boldsymbol{\theta}} = \underset{\boldsymbol{x} \in \mathbb{F}^N, \, \boldsymbol{\theta} \in \mathbb{C}^N}{\arg\min} \|\boldsymbol{A}\boldsymbol{x} - \operatorname{diag}\{\sqrt{\max(\boldsymbol{y} - \boldsymbol{b}, \boldsymbol{0})}\}\boldsymbol{\theta}\|_2^2,$$
$$\text{subject to} \quad |\theta_i| = 1, \ i = 1, ..., N. \tag{9}$$

To solve (9), the Gerchberg Saxton (GS) algorithm [10] was proposed that alternatively updates $\boldsymbol{x}$ and $\boldsymbol{\theta}$ in each iteration. The square root in (9) is reminiscent of the Anscombe transform that converts a Poisson random variable into another random variable that approximately has a standard Gaussian distribution. However, that approximation is accurate when the Poisson mean is

[1] If $\boldsymbol{x} \in \mathbb{R}^N$, then all gradients w.r.t. $\boldsymbol{x}$ in this paper should be real and hence use only the real part of expressions like (8).

sufficiently large (e.g., above 5), whereas this paper focuses on the lower-count regime.

WF and GS are two classic methods for phase retrieval under the Gaussian noise model. This paper uses them as baselines for comparison with the proposed methods.

III. Methods for Poisson model

This section describes optimization algorithms for minimizing the unregularized Poisson log-likelihood, i.e., (4), as well as regularized versions of the form

$$\Phi(\boldsymbol{x}) = f(\boldsymbol{x}) + \beta R(\boldsymbol{x}) \tag{10}$$

where $R : \mathbb{F}^N \mapsto \mathbb{R}$ is a regularizer. The general methods in the paper are adaptable to many regularizers, but for simplicity we focus on regularizers that are based on the assumption that $\boldsymbol{T}\boldsymbol{x}$ is approximately sparse, for a $K \times N$ matrix $\boldsymbol{T}$. When $\boldsymbol{T}$ is prox-friendly, like the orthogonal discrete wavelet transform (ODWT), we consider the (non-smooth) regularizer $R(\boldsymbol{x}) = \|\boldsymbol{T}\boldsymbol{x}\|_1$, and we focus on algorithms based on proximal operators such as MM and ADMM. When $\boldsymbol{T}$ is not as prox-friendly, like the finite-difference matrix used in total variation (TV) regularization, we replace the 1-norm term with a Huber function regularizer of the form

$$\begin{aligned} R(\boldsymbol{x}) =& \mathbf{1}' h.(\boldsymbol{T}\boldsymbol{x}; \alpha) = \min_{\boldsymbol{z}} \frac{1}{2}\|\boldsymbol{T}\boldsymbol{x} - \boldsymbol{z}\|_2^2 + \alpha\|\boldsymbol{z}\|_1, \\ h(t;\alpha) =& \begin{cases} \frac{1}{2}|t|^2, & |t| < \alpha, \\ \alpha|t| - \frac{1}{2}\alpha^2, & \text{otherwise}, \end{cases} \end{aligned} \tag{11}$$

which involves solving for $\boldsymbol{z}$ analytically in terms of $\boldsymbol{x}$. Here the dot subscript notation $h.(\boldsymbol{z}; \alpha)$ denotes element-wise application of the function h to its first argument (as in the Julia language). This smooth regularizer is suitable for gradient-based methods like WF, as well as for versions of MM and ADMM. We refer to (11) as "TV regularizaton" even though it is technically (anisotropic) "corner rounded" TV.

A. Wirtinger Flow for Poisson Model

To generalize the Wirtinger flow algorithm to the Poisson cost function (4), the most direct approach simply replaces the gradient (8) by (6) in the WF framework [29] and performs backtracking to find the step-size μ, as in [7]. We propose a faster alternative next.

1) Observed Fisher information for step size: We propose to make a quadratic approximation along the gradient direction of the cost function at each iteration, and then apply one step of Newton's method to minimize that 1D quadratic. Because computing the Hessian can be computationally expensive in large-scale problems, we replace the Hessian by the observed Fisher information when applying Newton's method [30].

Specifically, we first approximate the 1D line search problem associated with (4) by the following Taylor series

$$\begin{aligned} \mu_k &= \arg\min_{\mu \in \mathbb{R}} f_k(\mu), \\ f_k(\mu) &\triangleq f(\boldsymbol{x}_k - \mu \nabla f(\boldsymbol{x}_k)) \approx f(\boldsymbol{x}_k) + \|\nabla f(\boldsymbol{x}_k)\|_2^2 \, \mu \\ &\quad + \frac{1}{2} \nabla f(\boldsymbol{x}_k)' \nabla^2 f(\boldsymbol{x}_k) \nabla f(\boldsymbol{x}_k) \mu^2, \end{aligned} \tag{12}$$

where one can verify that the (opposite direction of) minimizer is

$$\mu_k = \frac{\|\nabla f(\boldsymbol{x}_k)\|_2^2}{\operatorname{real}\{\nabla f(\boldsymbol{x}_k)' \nabla^2 f(\boldsymbol{x}_k) \nabla f(\boldsymbol{x}_k)\}}. \tag{13}$$

We next approximate the Hessian matrix $\nabla^2 f(\boldsymbol{x})$ using the observed Fisher information matrix associated with ML estimation:

$$\begin{aligned} \nabla^2 f(\boldsymbol{x}) &\approx \boldsymbol{I}(\boldsymbol{x}, \boldsymbol{b}) \\ &\triangleq \mathbb{E}_{\boldsymbol{y}} \Big[\big(\nabla f(\boldsymbol{x}; \boldsymbol{y}, \boldsymbol{b})\big) \big(\nabla f(\boldsymbol{x}; \boldsymbol{y}, \boldsymbol{b})\big)' \Big| \boldsymbol{x}, \boldsymbol{b} \Big] \\ &= \boldsymbol{A}' \mathbb{E}_{\boldsymbol{y}} \Big[\big(\dot{\psi}.(\boldsymbol{v}; \boldsymbol{y}, \boldsymbol{b})\big) \big(\dot{\psi}.(\boldsymbol{v}; \boldsymbol{y}, \boldsymbol{b})\big)' \Big| \boldsymbol{v}, \boldsymbol{b} \Big] \boldsymbol{A}, . \end{aligned} \tag{14}$$

One can verify that the marginal Fisher information for a single term $\psi(v; y, b)$ is

$$\begin{aligned} \bar{I}(v, b) &= \mathbb{E}_y \Big[\big|\dot{\psi}(v; y, b)\big|^2 \Big| v, b \Big] \\ &= \frac{4|v|^2}{|v|^2 + b}, \quad v \in \mathbb{C}, \ b > 0. \end{aligned} \tag{15}$$

Substituting (15) into (14) and then substituting (14) into (13) yields the simplified step-size expression

$$\mu_k \triangleq \frac{\|\nabla f(\boldsymbol{x}_k)\|_2^2}{\boldsymbol{d}_k' \, \boldsymbol{D}_1 \, \boldsymbol{d}_k} \in \mathbb{R}, \tag{16}$$

where $\boldsymbol{d}_k \triangleq \boldsymbol{A} \nabla f(\boldsymbol{x}_k)$ and $\boldsymbol{D}_1 \triangleq \operatorname{diag}\{\bar{I}.(\boldsymbol{A}\boldsymbol{x}_k, \boldsymbol{b})\}$. (Careful implementation avoids redundant matrix-vector products.)

This approach removes all tuning parameters other than number of iterations. In addition, using the observed Fisher information leads to a larger step size than using the best Lipschitz constant, i.e., $\max\{2 + \boldsymbol{y}/(4\boldsymbol{b})\}$ when $\boldsymbol{b} > 0$, hence accelerating convergence.

2) Fisher information for Gaussian case: Similarly, the marginal Fisher information for the scalar case of the Gaussian cost function (2) is

$$I(v, b) = 16|v|^2(|v|^2 + b), \quad v \in \mathbb{C}, \ b \geq 0. \tag{17}$$

Substituting (17) into (16), one can also derive a convenient step size μ_k for the WF algorithm for the Gaussian model (2) using its observed Fisher information to approximate the exact Hessian.

3) WF with regularization: For the smooth regularizer (11), we majorize the Huber function $h(t)$ using quadratic polynomials with the optimal curvature using the ratio $\dot{h}(z)/z$ [31, p. 184]. One can verify that the step size μ_k becomes

$$\begin{aligned}
&\mu_k \triangleq \frac{\|\nabla \tilde{f}(\boldsymbol{x}_k)\|_2^2}{\nabla \tilde{f}(\boldsymbol{x}_k)' \left(\boldsymbol{A}'\boldsymbol{D}_1\boldsymbol{A} + \beta \boldsymbol{T}'\boldsymbol{D}_2\boldsymbol{T}\right) \nabla \tilde{f}(\boldsymbol{x}_k)}, \\
&\nabla \tilde{f}(\boldsymbol{x}_k) \triangleq \nabla f(\boldsymbol{x}_k) + \beta \boldsymbol{T}' \dot{h}.(\boldsymbol{T}\boldsymbol{x}; \alpha), \\
&\boldsymbol{D}_2 \triangleq \operatorname{diag}\{\min.(\alpha \oslash |\boldsymbol{T}\boldsymbol{x}_k|, 1)\}, \qquad (18)
\end{aligned}$$

where $\oslash$ denotes element-wise division. Careful implementation avoids dividing by zeros and redundant matrix-vector products.

4) Truncated Wirtinger flow: To potentially reduce the error in gradient estimation due to noisy measurements, reference [20] proposed a truncated Wirtinger flow (TWF) approach that uses only those measurements satisfying a threshold criterion to calculate the Wirtinger flow gradient. In particular, the threshold criterion [18] is defined as

$$\left|y_i - |\boldsymbol{a}_i'\boldsymbol{x}|^2\right| \le a^h \frac{\left\|\boldsymbol{y} - |\boldsymbol{A}\boldsymbol{x}|^2\right\|_1}{M} \cdot \frac{|\boldsymbol{a}_i'\boldsymbol{x}|^2}{\|\boldsymbol{x}\|_2}, \qquad (19)$$

where a^h is a user-defined parameter that controls the threshold value. When a^h is chosen appropriately, y_i values that do not satisfy (19) will be truncated when calculating the gradient, to try to reduce noise.

Algorithm 1 summarizes the Wirtinger flow algorithm for the Poisson model that uses the observed Fisher information for step size and the optional gradient truncation for noise reduction. Note that the unregularized WF is just setting $\beta = 0$ in Algorithm 1. The results in Section V used Algorithm 1 but did not use gradient truncation, as will be explained in Section IV.

Algorithm 1: Wirtinger flow for the Poisson model

Input: $\boldsymbol{A}, \boldsymbol{y}, \boldsymbol{b}, \boldsymbol{x}_0$ and n (number of iterations)
for $k = 0, ..., n-1$ **do**
  if gradient is truncated **then**
    $\nabla \tilde{f}(\boldsymbol{x}_k) = \boldsymbol{A}_{\mathcal{T}}'\dot{\boldsymbol{\psi}}.([\boldsymbol{A}\boldsymbol{x}_k]_{\mathcal{T}}; \boldsymbol{y}_{\mathcal{T}}, \boldsymbol{b}_{\mathcal{T}})$
    $+\beta \boldsymbol{T}'\dot{h}.([\boldsymbol{T}\boldsymbol{x}_k]_{\mathcal{T}})$
  else
    $\nabla \tilde{f}(\boldsymbol{x}_k) = \boldsymbol{A}'\dot{\boldsymbol{\psi}}.(\boldsymbol{A}\boldsymbol{x}_k; \boldsymbol{y}, \boldsymbol{b}) + \beta \boldsymbol{T}'\dot{h}.(\boldsymbol{T}\boldsymbol{x}_k)$
  end
  if cost function is regularized **then**
    $\mu_k \leftarrow$ Computed by (18)
  else
    $\mu_k \leftarrow$ Computed by (16)
  end
  $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k - \mu_k \nabla f(\boldsymbol{x}_k)$
end
Output: $\boldsymbol{x}_n$

B. An MM Approach with a Quadratic Majorizer

A majorize-minimize (MM) algorithm [32] is a generalization of the expectation-maximization (EM) algorithm that solves an optimization problem by iteratively constructing and solving simpler surrogate optimization problems. Quadratic majorizers are very common in MM algorithms because they have closed-form solutions and are well-suited to conjugate gradient methods.

1) MM for Poisson ML problem: When $b = 0$, one can verify $\ddot{\psi}$ has a singularity at $v = 0$, which means ψ is not Lipschitz differentiable, so we need to pursue other algorithms for this case, as discussed in a subsequent section. In contrast, if $b > 0$, one can construct a quadratic majorizer on $\mathbb{R}$ for (5), as illustrated in the supplement. With a bit more work to generalize to $\mathbb{C}^N$, a quadratic majorizer for the Poisson ML cost function (4) has the form

$$\begin{aligned}
q(\boldsymbol{x}; \boldsymbol{x}_k) \triangleq\ & f(\boldsymbol{x}_k) + \operatorname{real}\Big\{(\boldsymbol{x} - \boldsymbol{x}_k)'\boldsymbol{A}'\dot{\boldsymbol{\psi}}.(\boldsymbol{A}\boldsymbol{x}_k; \boldsymbol{y}, \boldsymbol{b})\Big\} \\
&+ \frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}_k)'\boldsymbol{A}'\boldsymbol{W}\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}_k), \qquad (20)
\end{aligned}$$

where $\boldsymbol{W}$ denotes a diagonal curvature matrix. From (7), one choice of $\boldsymbol{W}$ uses the maximum of $\ddot{\psi}$:

$$\boldsymbol{W}_{\max} \triangleq \operatorname{diag}\{2 + \boldsymbol{y}/(4\boldsymbol{b})\} \in \mathbb{R}^{M \times M}. \qquad (21)$$

However, $\boldsymbol{W}_{\max}$ is suboptimal because the curvature of a quadratic majorizer of $\psi(v; \cdot)$ varies with $v = [\boldsymbol{A}\boldsymbol{x}_k]_i$. For example, when $|v| \to \infty$, then (4) is dominated by the quadratic term having curvature = 2; so if y is large and b is small, then $\boldsymbol{W}_{\max}$ can be much greater than the optimal curvature 2. Thus, instead of using $\boldsymbol{W}_{\max}$ to build majorizers, we propose to use the following improved curvature:

$$\begin{aligned}
&\boldsymbol{W}_{\text{imp}} \triangleq \operatorname{diag}\{\boldsymbol{c}.(\boldsymbol{A}\boldsymbol{x}_k; \boldsymbol{y}, \boldsymbol{b})\} \in \mathbb{R}^{M \times M}, \\
&c(s; y, b) \triangleq \begin{cases} \ddot{\psi}\Big(\frac{b + \sqrt{b^2 + b|s|^2}}{|s|}; y, b\Big), & s \neq 0, \\ 2, & s = 0. \end{cases} \qquad (22)
\end{aligned}$$

One can verify $\lim_{s \to 0} c(s; y, b) = 2$ so (22) is continuous over $s \in \mathbb{C}$. The Appendix proves that (22) provides a majorizer in (20) and is an improved curvature compared to $\boldsymbol{W}_{\max}$, though it is not necessarily the sharpest possible [33].

For the ML case where constraints or regularizers are absent, the quadratic majorizer (20) associated with (21) or (22) leads to the following MM update:

$$\begin{aligned}\boldsymbol{x}_{k+1} &= \underset{\boldsymbol{x}\in\mathbb{F}^N}{\arg\min}\, q(\boldsymbol{x};\boldsymbol{x}_k) \\ &= \boldsymbol{x}_k - (\boldsymbol{A}'\boldsymbol{W}\boldsymbol{A})^{-1}\boldsymbol{A}'\dot{\psi}.(\boldsymbol{A}\boldsymbol{x}_k;\boldsymbol{y},\boldsymbol{b}). \end{aligned} \tag{23}$$

If $\boldsymbol{x} \in \mathbb{R}^N$, then the MM update for $\boldsymbol{x}_{k+1}$ is

$$\boldsymbol{x}_k - (\mathrm{real}\{\boldsymbol{A}'\boldsymbol{W}\boldsymbol{A}\})^{-1}\mathrm{real}\{\boldsymbol{A}'\dot{\psi}.(\boldsymbol{A}\boldsymbol{x}_k;\boldsymbol{y},\boldsymbol{b})\}.$$

When N is large, the matrix inverse operation in (23) is impractical, so one alternative way is to run a few inner iterations of conjugate gradient (CG) to descend the quadratic majorizer and hence descend the original cost function.

2) Regularized MM: For the regularized cost function (10), one can use the quadratic majorizer (20) as a starting point. If the regularizer is prox-friendly, then the minimization step of an MM algorithm for the regularized optimization problem is

$$\boldsymbol{x}_{k+1} = \underset{\boldsymbol{x}\in\mathbb{F}^N}{\arg\min}\, q(\boldsymbol{x};\boldsymbol{x}_k) + \beta\|\boldsymbol{T}\boldsymbol{x}\|_1. \tag{24}$$

In this case we apply a proximal gradient method [34–36] to solve (24). Specifically we use the proximal optimized gradient method (POGM) with adaptive restart [36] that provides faster worst-case convergence bound than the fast iterative shrinkage-thresholding algorithm (FISTA) [35].

For the smooth regularizer (11), we replace (24) with

$$\boldsymbol{x}_{k+1} = \underset{\boldsymbol{x}\in\mathbb{F}^N}{\arg\min}\, q(\boldsymbol{x};\boldsymbol{x}_k) + \beta 1'h.(\boldsymbol{T}\boldsymbol{x};\alpha), \tag{25}$$

and we use CG for this minimization.

Algorithm 2 summarizes our MM algorithm with quadratic majorizer using the improved curvature (22).

Algorithm 2: MM algorithm for the Poisson model

Input: $\boldsymbol{A}, \boldsymbol{y}, \boldsymbol{b}, \boldsymbol{x}_0$ and n (number of iterations)
for $k = 0, ..., n-1$ **do**
  Build $q(\boldsymbol{x};\boldsymbol{x}_k)$ (20) using $\boldsymbol{W}_{\mathrm{imp}}$ (22)
  if cost function is regularized **then**
    if $\boldsymbol{T}$ is prox-friendly **then**
      Update $\boldsymbol{x}_k$ by (24) using POGM
    else
      Update $\boldsymbol{x}_k$ by (25) using CG
    end
  else
    Update $\boldsymbol{x}_k$ by (23) or CG
  end
end
Output: $\boldsymbol{x}_n$

C. ADMM

To handle the case when $b_i = 0$, we investigated an ADMM algorithm with variable splitting $v_i = \boldsymbol{a}_i'\boldsymbol{x}$. This ADMM algorithm generalizes to $b_i > 0$ by a simple modification of the v_i update. Previous work [28] also applied ADMM to a Poisson ML cost function (for $b_i = 0$) for a total-variation (TV) regularizer, with separate variable splitting for both the log-likelihood and for the gradient operator. Compared to the ADMM in [28], our ADMM variant also considers the $b_i > 0$ case and is slightly simpler, but is otherwise somewhat similar to [28]. The ADMM is included here primarily for completeness i.e., as an alternative to MM in the $b_i = 0$ case.

1) ADMM for Poisson ML problem: With variable splitting $v_i = \boldsymbol{a}_i'\boldsymbol{x}$, an augmented linearized Lagrangian of (4) when $b_i = 0$ is given by

$$\begin{aligned} L_\rho(\boldsymbol{v},\boldsymbol{x};\boldsymbol{\eta},\rho) = &\sum_i \left[(|v_i|^2) - y_i \log(|v_i|^2)\right] \\ &+ \frac{\rho}{2}\sum_i \left[|v_i - \boldsymbol{a}_i'\boldsymbol{x} + \eta_i|^2 - |\eta_i|^2\right], \end{aligned} \tag{26}$$

where $\boldsymbol{\eta}$ is the dual variable and $\rho > 0$ denotes the AL penalty parameter.

We followed the order of first updating $\boldsymbol{v}$, then $\boldsymbol{x}$ and finally $\boldsymbol{\eta}$. Because the updates for the phase and magnitude of $\boldsymbol{v}$ are separable; in particular, the update for the phase of $\boldsymbol{v}$ is

$$\mathrm{sign}(\boldsymbol{v}_{k+1}) = \mathrm{sign}(\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{\eta}_k). \tag{27}$$

To update the magnitude of $\boldsymbol{v}$, similar to [28], set $t_i = |\boldsymbol{a}_i'\boldsymbol{x}_k - \eta_i|$; then the update of a single component $|v_i|$ is given by the following positive solution of a quadratic polynomial:

$$|v_i| = \frac{\rho t_i + \sqrt{\rho^2 t_i^2 + 8y_i(2+\rho)}}{2(2+\rho)}. \tag{28}$$

Again, in the unregularized case, the $\boldsymbol{x}$ update is a simple least square (LS) problem that can be optimized by CG or the following operation (29) involving matrix inverse.

$$\boldsymbol{x}_{k+1} = (\boldsymbol{A}'\boldsymbol{A})^{-1}\boldsymbol{A}'(\boldsymbol{v}_{k+1} + \boldsymbol{\eta}_k), \tag{29}$$

Again, if $\boldsymbol{x} \in \mathbb{R}^N$, then the $\boldsymbol{x}$ update is

$$\boldsymbol{x}_{k+1} = (\mathrm{real}\{\boldsymbol{A}'\boldsymbol{A}\})^{-1}\mathrm{real}\{\boldsymbol{A}'(\boldsymbol{v}_{k+1} + \boldsymbol{\eta}_k)\}.$$

2) Regularized ADMM: For the regularized case (10), the $\boldsymbol{x}$ update becomes

$$\boldsymbol{x}_{k+1} = \underset{\boldsymbol{x}\in\mathbb{F}^N}{\arg\min} \frac{\rho}{2}\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{v}_{k+1} - \boldsymbol{\eta}_k\|_2^2 + \beta R(\boldsymbol{x}). \tag{30}$$

We solve this using CG or POGM, depending on whether the regularizer R is smooth or not.

The dual variable $\boldsymbol{\eta}$ ascent update [37] is simply

$$\boldsymbol{\eta}_{k+1} = \boldsymbol{\eta}_k + (\boldsymbol{v}_{k+1} - \boldsymbol{A}\boldsymbol{x}_{k+1}). \tag{31}$$

For the case $b_i > 0$, everything is the same as the case $b_i = 0$ except the update for $|v_i|$, for which one can verify the $|v_i|$ is instead a positive root of the following cubic polynomial

$$0 = (2+\rho)|v_i|^3 - \rho t_i|v_i|^2 + (2b_i - 2y_i + \rho b_i)|v_i| - \rho b_i t_i. \tag{32}$$

Owing to the intermediate value theorem and the non-negativity of ρ, b_i, t_i, this cubic (32) can only possess one or three positive real roots. If the cubic has one positive root, then the update of $|v_i|$ is simply to assign the single positive root. If the cubic has three positive roots, we choose the root that minimizes the following Lagrangian term based on (33):

$$(|v_i|^2 + b_i) - y_i \log(|v_i|^2 + b_i) + \frac{\rho}{2}(|v_i| - t_i)^2. \tag{33}$$

A natural extension is to vary AL penalty parameter along with the variable update every few iterations. Reference [37] considered the magnitude of primal residual $\boldsymbol{r}_{k+1} = \boldsymbol{A}\boldsymbol{x}_{k+1} - \boldsymbol{v}_{k+1}$ and dual residual $\boldsymbol{s}_{k+1} = \rho\boldsymbol{A}'(\boldsymbol{v}_{k+1} - \boldsymbol{v}_k)$, as a principle to select penalty parameter to potentially improve convergence for convex optimization problems. However, for non-convex problems, the penalty parameter ρ is preferred to be sufficiently large to enable the convergence of ADMM algorithm [38]. Thus, we used the following heuristic to update ρ every 10 iterations:

$$\rho_{k+1} = \begin{cases} 2\rho_k, & \|r_k\| > 10\|s_k\| \\ \rho_k/2, & \|s_k\| > 100\rho_k\|r_k\| \\ \rho_k, & \text{otherwise.} \end{cases} \tag{34}$$

Algorithm 3 summarizes the ADMM algorithm described above.

Algorithm 3: ADMM algorithm for the Poisson model

Input: $\boldsymbol{A}, \boldsymbol{y}, \boldsymbol{b}, \boldsymbol{x}_0$ and n (number of iterations)
Initialize:
$\boldsymbol{v}_0 \leftarrow \boldsymbol{A}\boldsymbol{x}_0$
$\boldsymbol{\eta}_0 \leftarrow \boldsymbol{v}_0 - \boldsymbol{A}\boldsymbol{x}_0$
for $k = 0, ..., n-1$ **do**
  Update $\text{sign}(\boldsymbol{v}_k)$ by (27)
  if $b_i = 0$ **then**
    Update $|\boldsymbol{v}_k|$ by (28)
  else
    Update $|\boldsymbol{v}_k|$ by selecting root based on (33)
  end
  if cost function is regularized **then**
    Update $\boldsymbol{x}_k$ by (30) using CG or POGM
  else
    Update $\boldsymbol{x}_k$ by (29) or CG
  end
  Update $\boldsymbol{\eta}_k$ by (31)
end
Output: $\boldsymbol{x}_n$

IV. Implementation Details

This section introduces the implementation details of algorithms discussed in the previous section and our experimental setup for the numerical simulation (Section V). Due to the page limit, Sections IV-D and IV-E directly state conclusions without supporting figures. See the supplemental materials for these figures.

A. Initialization

Reference [9] proposed to use $\tilde{\boldsymbol{x}}_0$, the leading eigenvector of $\boldsymbol{A}'\,\text{diag}\{\boldsymbol{y} - \boldsymbol{b}\}\boldsymbol{A}$, as an initial estimate of $\boldsymbol{x}$. To accommodate signals of arbitrary scale, we scaled that leading eigenvector by

$$\hat{\alpha} = \underset{\alpha\in\mathbb{R}}{\arg\min} \|\boldsymbol{y} - \boldsymbol{b} - |\alpha\boldsymbol{A}\tilde{\boldsymbol{x}}_0|^2\|_2^2 = \frac{\sqrt{(\boldsymbol{y}-\boldsymbol{b})'|\boldsymbol{A}\tilde{\boldsymbol{x}}_0|^2}}{\|\boldsymbol{A}\tilde{\boldsymbol{x}}_0\|_4^2}. \tag{35}$$

Then, our initial estimate $\boldsymbol{x}_0$ is $\hat{\alpha}\tilde{\boldsymbol{x}}_0$ or $|\hat{\alpha}\tilde{\boldsymbol{x}}_0|$, depending on whether $\boldsymbol{x}$ is complex or real.

B. Ambiguities

To handle the phase ambiguity (all the algorithms can recover the signal only to within a constant phase shift due to the loss of global phase information), we used the following normalized root mean square error (NRMSE) metric, similar to [7], to evaluate all the algorithms:

$$\text{NRMSE} = \frac{\|\hat{\boldsymbol{x}} - \boldsymbol{x}_{\text{true}}\,\mathrm{e}^{\imath\phi}\|_2}{\|\boldsymbol{x}_{\text{true}}\|_2}, \quad \mathrm{e}^{\imath\phi} = \text{sign}(\boldsymbol{x}_{\text{true}}'\hat{\boldsymbol{x}}), \tag{36}$$

where $\hat{\boldsymbol{x}}$ and $\boldsymbol{x}_{\text{true}}$ denote the reconstructed signal/image and the true signal/image, respectively.

In the DFT matrix setting, more ambiguities such as circular shift and conjugate inversion [39] are involved. To circumvent these ambiguities, one method is to use several different masks that introduce redundancy in the measurement vector $\boldsymbol{y}$. In particular, in the Fourier phase

retrieval problem setting [40], the measurement vector $\boldsymbol{y}$ has elements given by

$$y[\tilde{n}] = \Big| \sum_{n=0}^{N-1} x[n] e^{-\imath 2\pi n \tilde{n}/\tilde{N}} \Big|^2 + b[\tilde{n}], \qquad (37)$$

where $\tilde{N} = 2N - 1$ (here we consider the over-sampled case), and $\tilde{n} = 0, ..., \tilde{N} - 1$. After introducing redundant masks, the measurement model becomes

$$y_l[\tilde{n}] = \Big| \sum_{n=0}^{N-1} x[n] D_l[n] e^{-\imath 2\pi n \tilde{n}/\tilde{N}} \Big|^2 + b_l[\tilde{n}], \qquad (38)$$

where $\boldsymbol{y}_l \in \mathbb{R}^{\tilde{N}}$ for $i = 1, \ldots, L$ and D_l denotes the lth of L masks. Our experiment used $L = 21$ masks to define the overall system matrix $\boldsymbol{A} \in \mathbb{C}^{L\tilde{N} \times N}$, where the first mask has full sampling and the remaining 20 have sampling rate 0.5 with random sampling patterns.

C. MM for DFT Matrix Setting

When $\boldsymbol{A}$ is a DFT matrix, we conjectured that the dynamic range of the y_i values can be very wide, leading to an ill-conditioned curvature matrix $\boldsymbol{W}$. In particular, the curvature changes dramatically across different coordinates, and the MM algorithm based on a quadratic majorizer using the maximum or improved curvature can converge very slowly. Thus, for the DFT case we imposed a threshold on the elements of the curvature matrix $\boldsymbol{W}$ as follows: $\tilde{\boldsymbol{W}} = \min(\boldsymbol{W}, \boldsymbol{W}_\mu + 3\,\boldsymbol{W}_\sigma)$, where $\boldsymbol{W}_\mu$ and $\boldsymbol{W}_\sigma$ denote the mean and the standard deviation of the diagonal elements of the original $\boldsymbol{W}$, respectively. We used $\tilde{\boldsymbol{W}}$ to construct quadratic "majorizers" for the DFT case. In addition, we found that it could be computationally expensive to directly compute the best Lipschitz constant ($\mathcal{L} = \|\boldsymbol{A}'\tilde{\boldsymbol{W}}\boldsymbol{A}\|_2$) every time $\tilde{\boldsymbol{W}}$ changes, so we used $\mathcal{L} = \|\boldsymbol{A}'\boldsymbol{A}\|_2 \|\tilde{\boldsymbol{W}}\|_2$ in MM algorithms, where one can verify

$$\|\boldsymbol{A}'\boldsymbol{A}\|_2 = c_{\text{DFT}}^2 \tilde{N}^2 \max\Big\{ \sum_{l=1}^{L} D_l \Big\}, \qquad (39)$$

where c_{DFT} is the scaling factor applied to $\boldsymbol{A}$.

D. Solving Quadratic Sub-problems

As discussed in the previous section, performing matrix inverse is impractical in large scale problems. Moreover, interestingly, even in small/medium scale problems when $\boldsymbol{A}$ is random Gaussian matrix, we found that running few iterations of CG was somewhat (or much) faster than performing matrix inverse (back slash in Julia) yet could give almost the same reconstructed quality. Thus, we used the CG algorithm instead of performing matrix inverse whenever solving a least squares problem, i.e., in (23), (29), (30) shown in Section III. In particular, we ran 3 iterations of CG in the first 10 outer iterations (out of 200 iterations) and only 1 CG update (equivalent to steepest descent) after that.

When $\boldsymbol{A}$ is a DFT matrix, including over-sampling and masks like in (38), one can verify that $\boldsymbol{A}'\boldsymbol{A}$ is a diagonal matrix that is trivial to invert. In this case we used matrix inverse rather than CG in the ADMM and Gerchberg Saxton algorithms.

E. Truncated Wirtinger Flow

TWF was commonly used in phase retrieval applications where the measurements follow independent Gaussian distribution. In our experiments with Poisson noise, however, we did not apply truncation because firstly, we did not observe any improved reconstruction quality or noise reduction for various setting of a^h compared to WF. We found that the reconstructed error (NRMSE) of TWF decreases almost monotonically as the truncatation threshold parameter a^h in (19) becomes larger, which is consistent with results in [18]. Furthermore, to achieve a comparable reconstruction quality with the non-truncated WF, one needs to keep about 95% of indices, indicating that the computational efficiency can also be hardly improved by TWF under the Poisson noise model, without sacrificing reconstruction quality. In contrast, TWF can instead be computationally inefficient because it requires computing the truncated indices in each iteration, especially when both the iteration number and M are large.

F. Summary

Considering all the factors discussed above, the algorithms compared in our experiments are introduced and abbreviated as follows.

- WF-Gaussian and WF-Poisson: non-truncated Wirtinger flow algorithms based on the Gaussian and Poisson ML models, using the observed Fisher information to select the step size. Regularized WF-Poisson used Huber version (11) if the original regularizer is not prox-friendly, as discussed in Section III.
- GS: the Gerchberg Saxton algorithm based on the Gaussian model (9).
- LSMM: Unregularized MM algorithm with quadratic majorizer using the improved curvature.

- ADMM: Unregularized ADMM algorithm with adaptive AL penalty parameter ρ that is initialized as 16.
- Regularized MM and ADMM: These algorithms use either the regularizer $R(\boldsymbol{x}) = \|\boldsymbol{T}\boldsymbol{x}\|_1$, where $\boldsymbol{T}$ denotes the ODWT matrix, or (11), where $\boldsymbol{T}$ denotes the finite difference matrix. The regularized ADMM uses the same AL penalty parameter strategy as the unregularized ADMM.

We ran all algorithms for 200 iterations on a server with Ubuntu 16.04 LTS operating system having Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz and 251 GB memory. We investigated 2 different system matrix $\boldsymbol{A}$: complex random Gaussian matrix and discrete Fourier transform (DFT) matrix. All elements in the measurement vector $\boldsymbol{y}$ were simulated to follow independent Poisson distributions per (3).

V. Numerical Simulations Results

This section reports numerical simulations that compare the convergence rates of the various algorithms and that demonstrate that using the Poisson log-likelihood improves the quality of the estimates under various experimental settings.

A. Random Gaussian Sensing Matrix

First we consider the case when all the elements in the system matrix $\boldsymbol{A}$ are independent random variables following a complex Gaussian distribution. We scaled $\boldsymbol{A}$ by a constant such that the average of $|\boldsymbol{a}_i'\boldsymbol{x}_{\text{true}}|^2$ is 2. The mean background counts $\boldsymbol{b}$ were set to 0.1 (a constant vector). The number of measurements M ranged from 2e3 to 1e4 with an interval of 1e3. We repeated each experiment 50 times and averaged the results to reduce the effects of statistical randomness.

1) True signals: Fig. 1 shows the real and imaginary parts of the piece-wise constant, complex true signal ($N = 100$). Here we exploit the piece-wise uniformity of the true signal, for which a natural choice for $\boldsymbol{T}$ is finite differences, i.e., a discrete TV regularizer [28]. We denote the WF Poisson, LSMM and ADMM algorithms using the "corner rounded" TV regularizer (11) as WF Poisson-TV, LSMM-TV and ADMM-TV, respectively. We empirically chose the regularizer parameters $\beta = 8$ and $\alpha = 0.5$ in (25) and (30). We ran 3 iterations of CG to update $\boldsymbol{x}$ in (25) and (30).

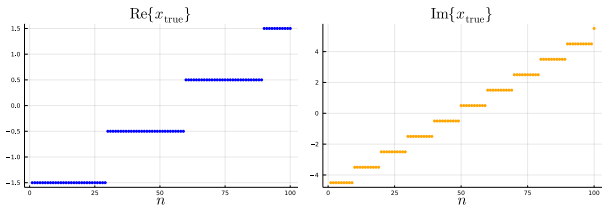

(a) Real part of $\boldsymbol{x}_{\text{true}}$. (b) Imaginary part of $\boldsymbol{x}_{\text{true}}$.

Fig. 1: The true signal used in simulations when $\boldsymbol{A}$ is random Gaussian matrix.

2) Accuracy comparison: Fig. 2 compares the NRMSE of all algorithms for $b_i > 0$ and $b_i = 0$, respectively. As expected, algorithms based on the Poisson model produce consistently better quality results (lower NRMSE) than algorithms based on the Gaussian model. In particular, the curves for LSMM, ADMM and WF Poisson overlap; perhaps they all converged to a global minimizer. Furthermore, regularized algorithms (WF Poisson-TV, LSMM-TV and ADMM-TV) consistently led to the lowest NRMSE for all M tested in the experiments; and were less sensitive to the change of M than the unregularized algorithms. These results illustrate, as expected, that regularizers that exploit properties of the true signal can guide the algorithms to better recover $\boldsymbol{x}$, when the measurements are noisy or corrupted, and especially when the number of measurements is low, as evident in Fig. 2, where the largest improvement of adding regularization was observed at the smallest M ($M = 2000$).

3) Speed comparison: Fig. 3 compares how quickly different algorithms decrease the NRMSE when the number of measurements is $M = 5000$. Again, as expected, WF Poisson decreases NRMSE faster than WF Gaussian for Poisson data. Furthermore, the quadratic approximation fits better to (4) than (2) because (4) is more similar to a quadratic function, i.e., has bounded curvature when $b_i > 0$, whereas (2) is a quartic with unbounded curvature. Among unregularized algorithms, the WF Poisson and LSMM have comparable speed yet both are much faster than ADMM. For regularized algorithms, the WF Poisson-TV decreases NRMSE the fastest; both WF Poisson-TV and LSMM-TV are much faster than ADMM-TV and even have comparable speed performance as the unregularized algorithms, but with notably improved reconstruction accuracy (lower NRMSE).

B. Discrete Fourier Transform Matrix

In many optical imaging applications, the system matrix $\boldsymbol{A}$ is often modeled with a DFT matrix. The measurement model and the corresponding system matrix $\boldsymbol{A}$ are

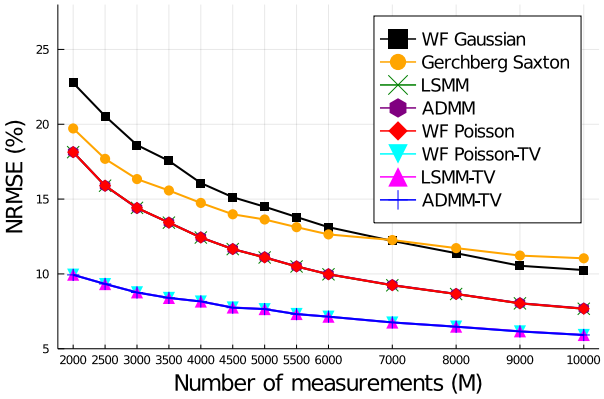

(a)

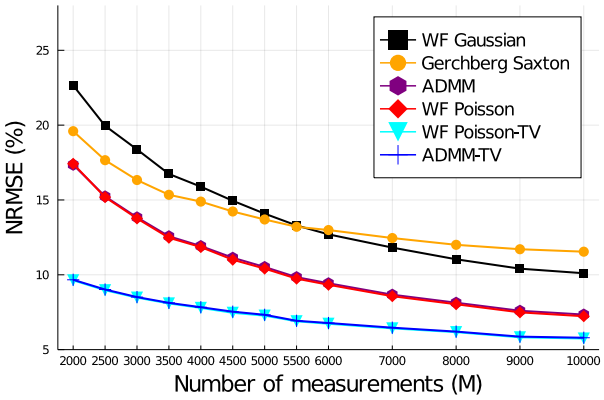

(b)

Fig. 2: NRMSE vs. number of measurements (M) when $\boldsymbol{A}$ is random Gaussian, where (a) and (b) correspond to $b_i = 0.1$ and $b_i = 0$, respectively.

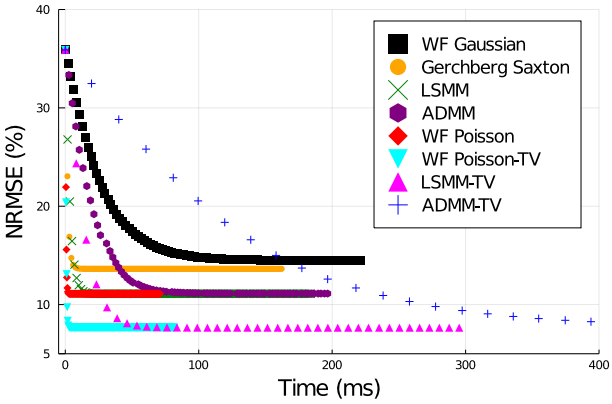

(a)

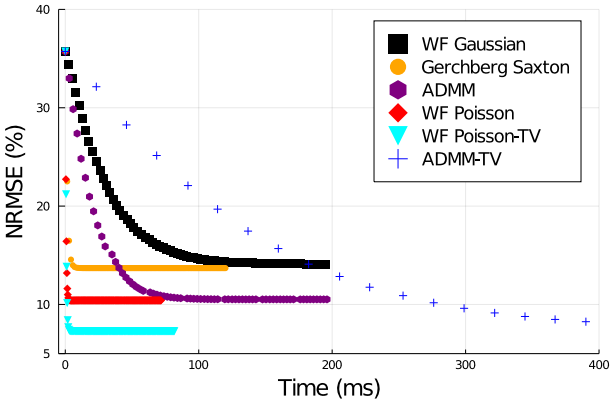

(b)

Fig. 3: NRMSE vs. time (ms) when $\boldsymbol{A}$ is random Gaussian where $M = 5000$. Subfigure (a) and (b) correspond to $b_i = 0.1$ and $b_i = 0$, respectively.

illustrated in (38). Here we scaled the (masked) DFT matrix $\boldsymbol{A}$ by a constant c_{DFT} so that the average of $|\boldsymbol{a}_i' \boldsymbol{x}_{\text{true}}|^2$ is 1 and the background counts vector $\boldsymbol{b}$ was set to 0.1. Because the pixel values in an image cannot be negative, we initialized $\boldsymbol{x}$ as $|\hat{\alpha}\tilde{\boldsymbol{x}}_0|$, as discussed in Section IV-A.

1) True signals: Our true image is a Shepp-Logan phantom of size 256×256. See Fig. 4. Pixel values range from 0 to 1. Different from the previous section, here we exploited the fact that the detail coefficients of the discrete wavelet transform of the true image are sparse, so we chose the matrix $\boldsymbol{T}$ in the regularizer to be an ODWT matrix. Here, the WF algorithm was not considered because it cannot accommodate proximal gradient methods. We empirically chose the regularization parameter $\beta = 32$ in (24) and (30). We ran 10 inner iterations of POGM with adaptive restart [36] for each outer iteration. Due to the page limit, results tested on more images are shown in the supplement.

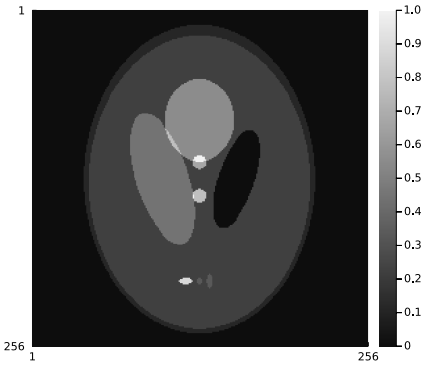

Fig. 4: Test true image for DFT matrix setting (Section V-B).

2) Accuracy comparison: Fig. 5 compares the images reconstructed by the WF, GS, MM and ADMM algorithms. Again, as expected, algorithms based on the Poisson model had better reconstruction quality than algorithms derived from Gaussian model. The regularized algorithms that exploit the assumed properties of the true

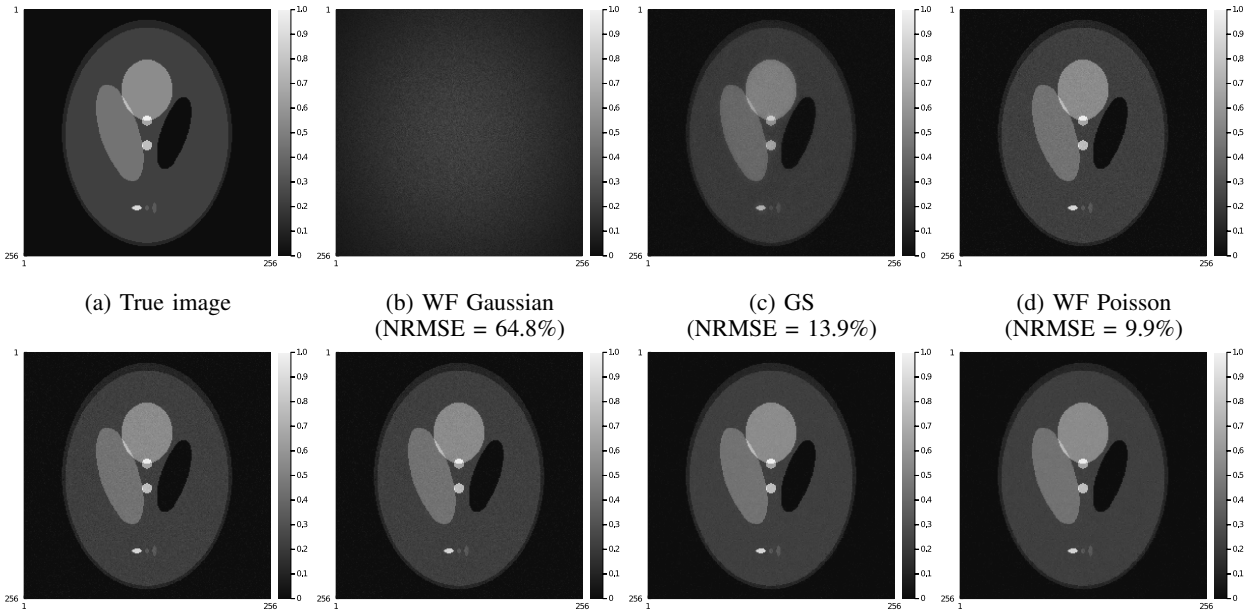

(a) True image
(b) WF Gaussian (NRMSE = 64.8%)
(c) GS (NRMSE = 13.9%)
(d) WF Poisson (NRMSE = 9.9%)
(e) LSMM (NRMSE = 9.9%)
(f) ADMM (NRMSE = 9.6%)
(g) LSMM-ODWT (NRMSE = 6.1%)
(h) ADMM-ODWT (NRMSE = 5.8%)

Fig. 5: Reconstructed image and corresponding NRMSE compared to the true image (Shepp-Logan phantom of size 256×256), for a sensing system with $L = 21$ masked DFT matrices. The average of $|\boldsymbol{a}_i'\boldsymbol{x}|^2$ and b_i were set to 1 and 0.1, respectively.

image can further improve the reconstruction quality. Reference [7] found that the WF-Gaussian algorithm had a very low recovery probability when $\boldsymbol{A}$ is a DFT matrix for Gaussian noise (Fig. 9 in [7]), and similarly we found that WF Gaussian algorithm failed to reconstruct a useful image for Poisson noise. Exploring why WF Gaussian fails could be an interesting future work.

3) Speed comparison: Fig. 6 shows how quickly the algorithms decrease NRMSE when $\boldsymbol{A}$ is a DFT matrix equipped with 21 random masks. Among unregularized algorithms, WF Poisson converged the fastest. Perhaps due to the simple inverse operation of the diagonal matrix $\boldsymbol{A}'\boldsymbol{A}$, ADMM converged faster than LSMM, for both unregularized and regularized cases. Thus, based on the convergence rate of algorithms shown in Fig. 6, an efficient strategy would be to first run few iterations of WF Poisson then switch to the ADMM-ODWT method.

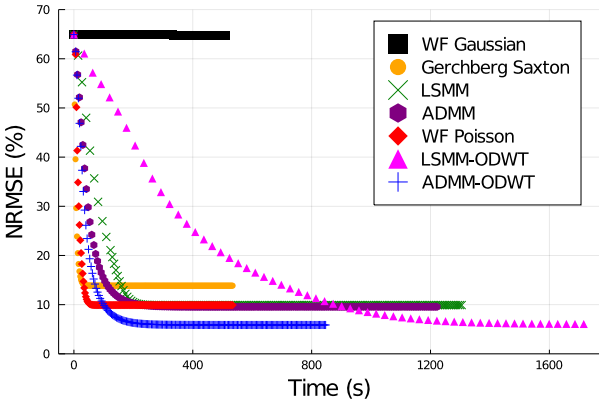

Fig. 6: NRMSE vs. time (s) when $\boldsymbol{A}$ is a DFT matrix.

VI. Conclusion

This paper compares several algorithms including Wirtinger flow (WF), Gerchberg Saxton (GS), majorize-minimize (MM) and alternating direction method of multipliers (ADMM), for ML estimation and regularized ML estimation for phase retrieval from Poisson measurements. For the WF algorithm, instead of performing backtracking or using a heuristic step size, we proposed to use a quadratic approximation along the gradient direction and replace the Hessian by the observed Fisher information so that the step size can be computed efficiently without any tuning parameter. For the MM algorithm with quadratic majorizers, we proposed a curvature formula and proved that it is superior to the curvature built on the upper bound of the second derivative of the Poisson ML cost function, i.e., the Lipschitz constant of its derivative.

Simulation results showed that the ML estimates based on the Poisson model (WF Poisson, MM, ADMM) yielded consistently lower NRMSE than algorithms based on Gaussian ML model (WF Gaussian and GS), when tested on simulated signals/images with random complex Gaussian matrix and discrete Fourier transform (DFT) matrix, for measurements following Poisson distribution.

Furthermore, as expected, the regularized algorithms designed for the Poisson model that exploit assumed properties of the true signals/images, such as sparsity of finite differences (anisotropic total variation) or of the coefficients of a discrete wavelet transform, can further lower the NRMSE compared to unregularized algorithms. Regarding the computational efficiency in the unregularized case, the WF Poisson algorithm decreased the NRMSE fastest. For regularized algorithms, the regularized WF Poisson algorithm (WF Poisson-TV) decreased the NRMSE the fastest for a smooth regularizer. For an 1-norm regularizer, WF Poisson is not an option because the ineligibility to accommodate proximal gradient methods. In such case, regularized ADMM is faster than the regularized LSMM because it can benefit from the trivial inverse implementation of $\boldsymbol{A}'\boldsymbol{A}$, i.e., when $\boldsymbol{A}$ is (masked) DFT matrix, and hence faster in our experiments.

In summary, the proposed algorithms are all effective for both unregularized and regularized phase retrieval problems from Poisson measurements. Future work includes investigating Poisson plus Gaussian noise models, exploring why WF Gaussian fails under DFT matrix setting, establishing sufficient conditions for global convergence, investigating algorithms that can handle other kind of regularizers (e.g., deep learning methods), seeking the optimal curvature, and testing algorithms derived for the Poisson noise model under a wider variety of experimental settings.

APPENDIX

This appendix proves that the curvature formula (22) provides a majorizer for the negative log-likelihood of the Poisson model that always is bounded above by the maximum curvature (7).

For Poisson phase retrieval, we drop the subscript i and irrelevant constants and focus on the negative log-likelihood for real case for simplicity:

$$h(r) = (r^2 + b) - y \log(r^2 + b), \tag{40}$$

where $r \in \mathbb{R}, y \geq 0, b \geq 0$.

One can generalize the majorizer derived here for (40) to the complex case by taking the magnitude and some other minor modifications.

First, we consider some simple cases:

- If $y = 0$, then (40) is a quadratic function, so no quadratic majorizer is needed.
- If $b = 0$ and $y > 0$ then (40) has unbounded 2nd derivative so no quadratic majorizer exists.
- If $b = 0$ and $r = 0$, then y must be zero because a Poisson random variable with zero mean can only take the value 0. Thus again quadratic majorizer is not needed.

So hereafter, we assume that $y > 0,\ b > 0$. Under these assumptions, the derivatives of (40) are:

$$\dot{h}(r) = 2r\left(1 - \frac{y}{r^2 + b}\right), \tag{41}$$

$$\ddot{h}(r) = 2 + 2y\frac{r^2 - b}{(r^2 + b)^2}, \tag{42}$$

$$h^{(3)}(r) = \frac{2yr(3b - r^2)}{(r^2 + b)^3}, \tag{43}$$

where $h^{(3)}(r)$ denotes the third derivative. Clearly, $\dot{h}(r)$ is convex on $(-\infty, -\sqrt{3b}]$ and $[0, \sqrt{3b}]$, and concave on $[-\sqrt{3b}, 0]$ and $[\sqrt{3b}, +\infty)$, based on the sign of $h^{(3)}(r)$.

A quadratic majorizer of h at point s has the form:

$$H(r; s) = h(s) + \dot{h}(s)(r - s) + \frac{1}{2}c(s)(r - s)^2. \tag{44}$$

The derivative of this function (w.r.t. r) is:

$$\dot{H}(r; s) = c(s)(r - s) + \dot{h}(s). \tag{45}$$

By design, this kind of quadratic majorizer satisfies $H(s; s) = h(s)$ and $\dot{H}(s; s) = \dot{h}(s)$. From (43), we note that $r^2 = 3b$ is a maximizer of $\ddot{h}$ so the maximum curvature is:

$$\ddot{h}(r) \leq 2y\frac{2b}{(4b)^2} + 2 = 2 + \frac{y}{4b}. \tag{46}$$

Proposition: $H(r; s)$ defined in (44) is a majorizer of $h(r)$ when $c(s) = c_{\text{imp}}(s)$, where:

$$c_{\text{imp}}(s) \triangleq \begin{cases} \ddot{h}(g(s)), & s \neq 0, \\ \lim\limits_{s \to 0} \ddot{h}(g(s)), & s = 0, \end{cases} \tag{47}$$

where

$$g(s) \triangleq \frac{b + \sqrt{b^2 + bs^2}}{s}. \tag{48}$$

By construction, the proposed curvature $c(s)$ is at most the max curvature given in (46).

Proof: Because of the symmetry of $\ddot{h}(r)$, it suffices to prove the proposition for $s \geq 0$ without loss of generality. First we consider some trivial cases:

1) If $s = 0$, one can verify $\lim_{s\to 0} \ddot{h}(g(s)) = 2$. In this case, $H(r;s)$ is simply

$$\begin{aligned} H(r;0) &= h(0) + \frac{1}{2}c(0)r^2 \\ &= r^2 + b - y\log(b) \\ &\geq r^2 + b - y\log(r^2+b) = h(r). \end{aligned} \tag{49}$$

2) If $s = \sqrt{3b}$, one can verify

$$\ddot{h}(g(\sqrt{3b})) = 2 + \frac{y}{4b}. \tag{50}$$

So in this case $\ddot{h}(g(s))$ equals the maximum curvature.

Hereafter, we consider only $s > 0$ and $s \neq \sqrt{3b}$.

We note that it suffices to prove

$$\begin{aligned} &\forall r \in (-\infty, s],\ \dot{h}(r) \geq \dot{H}(r;s), \\ &\forall r \in [s, +\infty),\ \dot{h}(r) \leq \dot{H}(r;s), \end{aligned} \tag{51}$$

because if (51) holds, then $\forall \tilde{r} < s$:

$$\begin{aligned} H(s;s) - H(\tilde{r};s) &= \int_{\tilde{r}}^{s} \dot{H}(r;s)dr \\ &\leq \int_{\tilde{r}}^{s} \dot{h}(r)dr = h(s) - h(\tilde{r}), \end{aligned} \tag{52}$$

and $\forall \tilde{r} > s$:

$$\begin{aligned} H(\tilde{r};s) - H(s;s) &= \int_{s}^{\tilde{r}} \dot{H}(r;s)dr \\ &\geq \int_{s}^{\tilde{r}} \dot{h}(r)dr = h(\tilde{r}) - h(s). \end{aligned} \tag{53}$$

Together with $H(s;s) = h(s)$, we have shown that equation (51) implies $H(r;s) \geq h(r)$, $\forall r \in \mathbb{R}$.

Substituting $\dot{H}(r;s) = c(s)(r-s) + \dot{h}(s)$ into (51), one can verify that showing (51) becomes showing

$$c_{\text{imp}}(s) \geq \frac{\dot{h}(r) - \dot{h}(s)}{r - s}, \quad \forall r \in \mathbb{R},\ r \neq s. \tag{54}$$

Furthermore, when $s > 0$, the parabola $H(\cdot;s)$ is symmetric about its minimizer:

$$\begin{aligned} a = a(s) &\triangleq \arg\min_{r}\ H(r;s) = s - \frac{\dot{h}(s)}{c_{\text{imp}}(s)} \\ &= \frac{s\ddot{h}(g(s)) - \dot{h}(s)}{\ddot{h}(g(s))} \geq 0. \end{aligned} \tag{55}$$

This minimizer is nonnegative because $\dot{h}(s) \leq 2s$ and

$$\begin{aligned} c_{\text{imp}}(s) = \ddot{h}(g(s)) &= 2 + \frac{ys^2(b + \sqrt{b^2+bs^2})}{b(b + s^2 + \sqrt{b^2+bs^2})^2} \\ &\geq 2. \end{aligned} \tag{56}$$

Thus, if $h(r) \leq H(r;s)$ when $r \geq 0$, we have $h(-r) = h(r) \leq H(r;s) \leq H(-r;s) = H(r+2a;s)$, so it suffices to prove (54) only for $r \geq 0$, which simplifies (54) to showing

$$c_{\text{imp}}(s) \geq \frac{\dot{h}(r) - \dot{h}(s)}{r - s}, \quad \forall r \geq 0,\ r \neq s. \tag{57}$$

In short, if (57) holds, then $H(r;s) \geq h(r),\ \forall r \in \mathbb{R}$.

To prove (57), we exploit a useful property of $c_{\text{imp}}(s)$. Under geometric view, $c_{\text{imp}}(s)$ defines (the ratio of) an affine function connecting points $(g(s), \dot{h}(g(s)))$ and $(s, \dot{h}(s))$ is tangent to $\dot{h}(r)$ at point $r = g(s)$, so that one can verify

$$\ddot{h}(g(s)) = c_{\text{imp}}(s) = \frac{\dot{h}(g(s)) - \dot{h}(s)}{g(s) - s}, \quad g(s) \neq s. \tag{58}$$

Fig. 7 visualizes this tangent property.

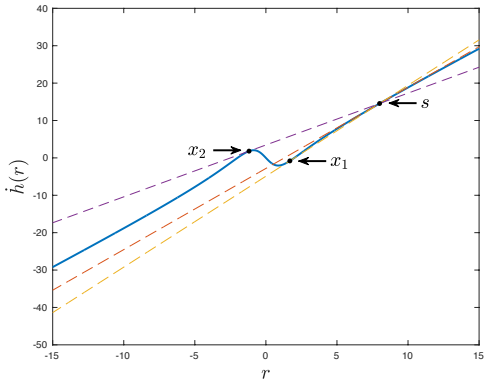

Fig. 7: Visualization of tangent property where $s = 8,\ y = 6,\ b = 2$. The blue curve is $\dot{h}(r)$, and the dashed lines are tangent to $\dot{h}(r)$ at $x_1 = g(s) = (b + \sqrt{b^2+bs^2})/s,\ x_2 = (b - \sqrt{b^2+bs^2})/s$ and s.

The reason why $g(s) \neq s$ is that one can verify $g(s) = s$ implies $s = \sqrt{3b}$ for $s > 0$ that has already been proved above.

Let $f(r) = (\dot{h}(r) - \dot{h}(s))/(r-s)$, where $r \geq 0$ and $r \neq s$, plugging in $\dot{h}(r)$ and $\dot{h}(s)$ yields:

$$f(r) = 2 + \frac{2y(sr - b)}{(s^2+b)(r^2+b)}. \tag{59}$$

Differentiating w.r.t. r leads to:

$$\dot{f}(r) = \frac{2y}{s^2+b} \cdot \frac{-sr^2 + 2br + bs}{(r^2+b)^2}, \tag{60}$$

where one can verify the positive root of $-sr^2 + 2br + bs = 0$ is $g(s)$ that is given by (48).

Together with $\dot{f}(r) > 0$ when $r \in (0, g(s))$ and $\dot{f}(r) < 0$ when $r \in (g(s), \infty)$, we have (57) holds because $f(r)$ achieves its maximum at $f(g(s))$:

$$f(r) \leq f(g(s)) = c_{\text{imp}}(s). \quad \blacksquare \tag{61}$$

REFERENCES

[1] K. Jaganathan, Y. C. Eldar, and B. Hassibi. *Phase retrieval: an overview of recent developments*. 2015.

[2] P. Grohs, S. Koppensteiner, and M. Rathmair. "Phase retrieval: uniqueness and stability". In: *SIAM Review* 62.2 (2020), 301–50.

[3] P. Jaming. "Phase retrieval techniques for radar ambiguity problems." In: *J. Four. Anal. Appl.* 5.4 (1999), pp. 309–329.

[4] R. P. Millane. "Phase retrieval in crystallography and optics". In: *J. Opt. Soc. Am. A* 7.3 (Mar. 1990), 394–411.

[5] J. Dainty and J. Fienup. "Phase retrieval and image reconstruction for astronomy". In: *Imag. Recov. Theory Appl.* 13 (Jan. 1987), pp. 231–275.

[6] L. R. Rabiner and B.-H. Juang. *Fundamentals of Speech Recognition*. Eng. Clif., NJ: Pren. Hall, 1993.

[7] T. Qiu, P. Babu, and D. P. Palomar. "PRIME: phase retrieval via majorization-minimization". In: *IEEE Trans. Sig. Proc.* 64.19 (Oct. 2016), 5174–86.

[8] D. L. Snyder, C. W. Helstrom, A. D. Lanterman, M. Faisal, and R. L. White. "Compensation for readout noise in CCD images". In: *J. Opt. Soc. Am. A* 12.2 (Feb. 1995), 272–83.

[9] E. Candes, X. Li, and M. Soltanolkotabi. "Phase Retrieval via Wirtinger Flow: Theory and Algorithms". In: *IEEE Trans. Info. Theory* 61.4 (Apr. 2015), pp. 1985–2007.

[10] R. W. Gerchberg and W. O. Saxton. "Practical Algorithm for Determination of Phase from Image and Diffraction Plane Pictures". In: *OPTIK* 35.2 (1972), 237–246.

[11] P. Thibault and M. Guizar-Sicairos. "Maximum-likelihood refinement for coherent diffractive imaging". In: *New J. of Phys.* 14.6 (June 2012), p. 063004.

[12] A. Goy, K. Arthur, S. Li, and G. Barbastathis. "Low photon count phase retrieval using deep learning". In: *Phys. Rev. Lett.* 121.24 (Dec. 2018), p. 243902.

[13] R. Xu, M. Soltanolkotabi, J. P. Haldar, W. Unglaub, J. Zusman, A. F. J. Levi, and R. M. Leahy. *Accelerated Wirtinger Flow: A fast algorithm for ptychography*. 2018.

[14] D. A. Barmherzig and J. Sun. "Low-photon holographic phase retrieval". In: *OSA: Computational Optical Sensing and Imaging*. 2020.

[15] I. Vazquez, I. E. Harmon, J. C. R. Luna, and M. Das. "Quantitative phase retrieval with low photon counts using an energy resolving quantum detector". In: *J. Opt. Soc. Am. A* 38.1 (Jan. 2021), 71–9.

[16] H. Lawrence, D. A. Barmherzig, H. Li, M. Eickenberg, and M. Gabrie. *Phase retrieval with holography and untrained priors: Tackling the challenges of low-photon nanoscale imaging*. 2020.

[17] K. Choi and A. D. Lanterman. "Phase retrieval from noisy data based on minimization of penalized I-divergence". In: *J. Opt. Soc. Am. A* 24.1 (Jan. 2007), 34–49.

[18] L. Bian, J. Suo, J. Chung, X. Ou, C. Yang, F. Chen, and Q. Dai. "Fourier ptychographic reconstruction using Poisson maximum likelihood and truncated Wirtinger gradient". In: *Nature Sci. Rep.* 6.1 (2016).

[19] E. J. Candes, Y. C. Eldar, T. Strohmer, and V. Voroninski. "Phase retrieval via matrix completion". In: *SIAM J. Imaging Sci.* 6.1 (2013), 199–225.

[20] Y. Chen and E. J. Candes. "Solving random quadratic systems of equations is nearly as easy as solving linear systems". In: *Comm. Pure Appl. Math.* 70.5 (May 2017), 822–83.

[21] H. Chang and S. Marchesini. "Denoising Poisson phaseless measurements via orthogonal dictionary learning". In: *Optics Express* 26.16 (Aug. 2018), 19773–96.

[22] Y. Zhang, P. Song, and Q. Dai. "Fourier ptychographic microscopy using a generalized Anscombe transform approximation of the mixed Poisson-Gaussian likelihood". In: *Optics Express* 25.1 (Jan. 2017), 168–79.

[23] I. Kang, F. Zhang, and G. Barbastathis. "Phase extraction neural network (PhENN) with coherent modulation imaging (CMI) for phase retrieval at low photon counts". In: *Optics Express* 28.15 (July 2020), 21578–600.

[24] D. L. Snyder, A. M. Hammoud, and R. L. White. "Image recovery from data acquired with a charge-coupled-device camera". In: *J. Opt. Soc. Am. A* 10.5 (May 1993), 1014–23.

[25] M. Makitalo and A. Foi. "Optimal Inversion of the Generalized Anscombe Transformation for Poisson-Gaussian Noise". In: *IEEE Transactions on Image Processing* 22.1 (2013), pp. 91–103.

[26] H. Zhang and D. P. Mandic. "Is a complex-valued stepsize advantageous in complex-valued gradient learning algorithms?" In: *IEEE Trans. Neural Net. Learn. Sys.* 27.12 (Dec. 2016), 2730–5.

[27] D. Böhning and B. G. Lindsay. "Monotonicity of quadratic approximation algorithms". In: *Ann. Inst. Stat. Math.* 40.4 (Dec. 1988), 641–63.

[28] H. Chang, Y. Lou, Y. Duan, and S. Marchesini. "Total variation–based phase retrieval for poisson noise removal". In: *SIAM journal on imaging sciences* 11.1 (2018), pp. 24–55.

[29] H. Zhang, Y. Liang, and Y. Chi. "A nonconvex approach for phase retrieval: Reshaped Wirtinger flow and incremental algorithms". In: *J. Mach. Learning Res.* 18.141 (2017), 1–35.

[30] K. Lange. "A gradient algorithm locally equivalent to the EM Algorithm". In: *J. Royal Stat. Soc. Ser. B* 57.2 (1995), 425–37.

[31] P. J. Huber. *Robust statistics*. New York: Wiley, 1981.

[32] D. R. Hunter and K. Lange. "A tutorial on MM algorithms". In: *American Statistician* 58.1 (Feb. 2004), 30–7.

[33] J. de Leeuw and K. Lange. "Sharp quadratic majorization in one dimension". In: *Comp. Stat. Data Anal.* 53.7 (May 2009), 2471–84.

[34] I. Daubechies, M. Defrise, and C. De Mol. "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint". In: *Comm. Pure Appl. Math.* 57.11 (Nov. 2004), 1413–57.

[35] A. Beck and M. Teboulle. "A fast iterative shrinkage-thresholding algorithm for linear inverse problems". In: *SIAM J. Imaging Sci.* 2.1 (2009), 183–202.

[36] D. Kim and J. A. Fessler. "Adaptive restart of the optimized gradient method for convex optimization". In: *J. Optim. Theory Appl.* 178.1 (July 2018), 240–63.

[37] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers". In: *Found. Trends Mach. Learn.* 3.1 (Jan. 2011), pp. 1–122.

[38] Y. Wang, W. Yin, and J. Zeng. "Global convergence of ADMM in nonconvex nonsmooth optimization". In: *J. Sci. Comp.* 78.1 (2019), 29–63.

[39] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev. "Phase Retrieval with Application to Optical Imaging: A contemporary overview". In: *IEEE Sig. Proc. Mag.* 32.3 (2015), pp. 87–109.

[40] T. Bendory, R. Beinert, and Y. C. Eldar. "Fourier Phase Retrieval: Uniqueness and Algorithms". In: *Compressed Sensing and its Applications*. Applied and Numerical Harmonic Analysis. Springer Intl. Publishing, 2018, pp. 55–91.

[41] R. Chandra, Z. Zhong, J. Hontz, V. McCulloch, C. Studer, and T. Goldstein. "PhasePack: A Phase Retrieval Library". In: *Asil. Conf. Sig. Sys. Comp.* (2017), pp. 1617–1621.

[42] C. A. Metzler, M. K. Sharma, S. Nagesh, R. G. Baraniuk, O. Cossairt, and A. Veeraraghavan. "Coherent inverse scattering via transmission matrices: Efficient phase retrieval algorithms

and a public dataset". In: *Proc. Intl. Conf. Comp. Photography*. 2017, 1–16.

SUPPLEMENT

A. Illustration of Quadratic Majorizers for Poisson ML Cost Function

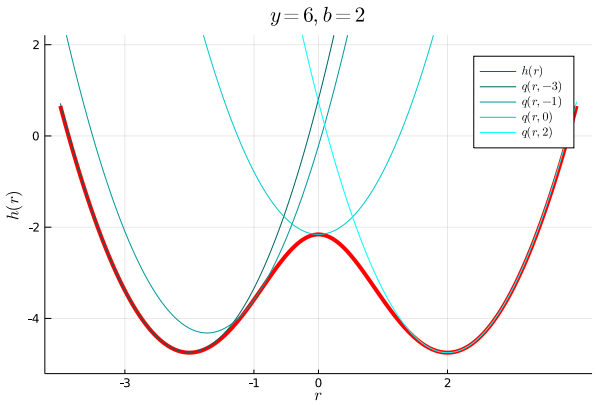

Fig. S.1: Quadratic majorizers for the non-convex Poisson log-likelihood function $h(r; y, b)$ when $y = 6$ and $b = 2$.

B. Illustration of the Improved Curvature

Recall that the cost function of a real single term for Poisson noise model is

$$h(r) = (r^2 + b) - y \log(r^2 + b), \quad r \in \mathbb{R}, \tag{S.1}$$

where y and b are known measurement and background count, respectively. A sufficient and necessary condition to be a quadratic majorizer for (S.1) is

$$\frac{1}{2} c(s)(r - s)^2 + \dot{h}(s)(r - s) + h(s) \geq h(r), \quad \forall r \in \mathbb{R}. \tag{S.2}$$

Following (S.2), one can express the optimal curvature [33] as

$$c_{\text{opt}}(s) = \sup_{r \neq s} \frac{2\Big(h(r) - h(s) - \dot{h}(s)(r - s)\Big)}{(r - s)^2}, \tag{S.3}$$

However, because $h(r)$ involves logarithm plus the derivative of (S.3) is complicated, we found it difficult to derive an analytical solution for $c_{\text{opt}}(s)$, i.e., by selecting an appropriate root after setting derivative to zero. In this paper, alternatively, we derived a sub-optimal yet simpler curvature (aka the improved curvature as proved in the Appendix)

$$c_{\text{imp}}(s) = \sup_{r \neq s} \frac{\dot{h}(r) - \dot{h}(s)}{r - s}, \tag{S.4}$$

where one can verify c_{imp} is mathematically close to c_{opt} if the supremum is achieved near s.

Letting $t = r - s$, one can empirically find $c_{\text{opt}}(s)$ by searching all possible t for any s. For Fig. S.2, we search from $t = -1000 : 0.0011 : 1000$ for all $s = -20 : 0.009 : 20$.

Fig. S.2 shows that the proposed curvature lies between the maximum curvature and the optimal curvature (found empirically). In particular, when s is near zero, the proposed curvature and the optimal curvature are close to each other, which partially illustrates the near optimality of c_{imp} when s is near zero. We also noticed that there still remains space for improvement when $|s|$ is large, which can be an interesting future work. Fig. S.3 compares majorizers built using the maximum curvature, improved curvature and the optimal curvature as shown in Fig. S.2a with the original cost function $h(r)$.

C. Computation Time Comparison Between CG and BS

Fig. S.4 compares the convergence rates of the Poisson phase retrieval algorithms (GS, MM and ADMM) that involve solving least squares problems. In Fig. S.4, $M = 3000$ and $N = 100$ are both small enough so that both backsubstitution (BS) (backslash in Julia) and CG are viable options for solving the inner quadratic optimization problems. In every case the CG version of 250 iterations decreased NRMSE faster than the BS version.

D. Fisher Information vs. Line Search in WF

Fig. S.5 compares the WF algorithm using the observed Fisher information with WF using line search for step size. We found that WF Fisher decreased NRMSE faster in time than WF line search, while WF line search can decrease more NRMSE within a single iteration. Considering the computational efficiency (NRMSE vs. time), we decided to use the WF Fisher in this paper.

E. Truncated vs. Non-truncated WF

Fig. S.6 compares truncated WF (TWF) vs. non-truncated WF under Poisson noise model. We found that, similar to results in [18], the reconstruction error (NRMSE) for TWF is almost monotonically decreasing as the truncate threshold parameter a^h goes larger. To achieve a comparable NRMSE with non-truncated WF, one needs to set $a^h > 100$, however, under this setting,

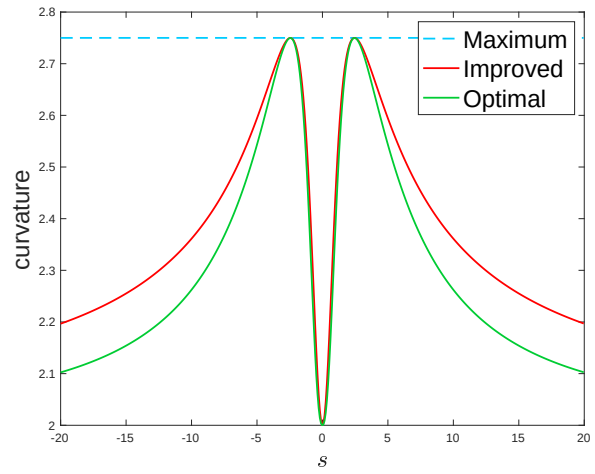

(a) $y = 6,\ b = 2$.

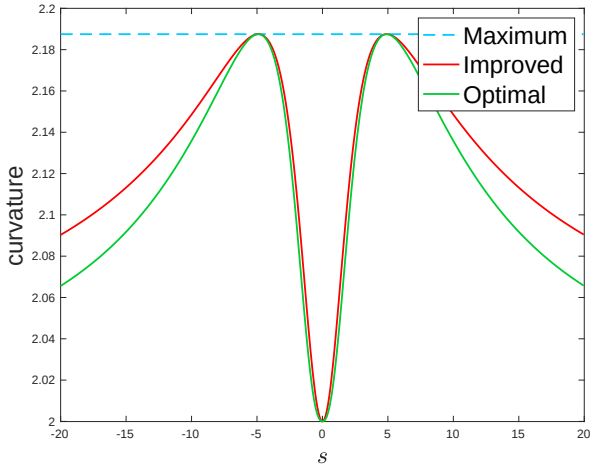

(b) $y = 6,\ b = 8$.

Fig. S.2: Plot of proposed curvature and max and optimal curvature, for the case $y = 6,\ b = 2$ and $y = 6,\ b = 8$, respectively.

less than 5% indices of $\boldsymbol{y}$ are truncated when calculating the gradient. This implies that almost all the measurements are useful for the reconstruction and therefore should not be truncated. In addition, calculating the truncated indices in each iteration can be computationally inefficient, especially when M is large. Considering all reasons described above, we did not use gradient truncation (TWF) in this paper.

F. Huber Function vs. Alternating Minimization

For non prox-friendly regularizers, other than the Huber function, an alternative can be introducing another variable and applying alternating minimization. In particular,

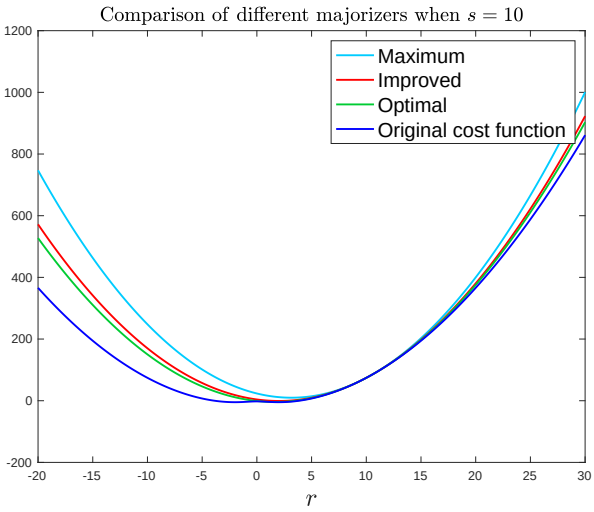

(a) Visualization of majorizers vs. $h(r)$.

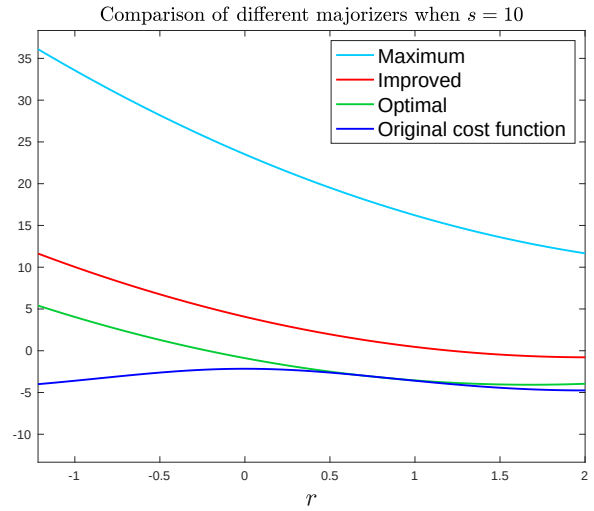

(b) Enlarged around $r = 0.5$.

Fig. S.3: Comparison of quadratic majorizers with different curvatures, for $y = 6$ and $b = 2$.

The update of $\boldsymbol{x}$ in LSMM becomes

$$\boldsymbol{x}_{k+1}, \boldsymbol{z}_{k+1} = \underset{\boldsymbol{x}\in\mathbb{F}^N, \boldsymbol{z}\in\mathbb{C}^K}{\arg\min}\ Q_k(\boldsymbol{x}, \boldsymbol{z}),$$

$$Q_k(\boldsymbol{x}, \boldsymbol{z}) \triangleq q(\boldsymbol{x}; \boldsymbol{x}_k) + \beta\Big(\frac{1}{2}\|\boldsymbol{T}\boldsymbol{x} - \boldsymbol{z}\|_2^2 + \alpha\|\boldsymbol{z}\|_1\Big),$$

where one can alternatively update $\boldsymbol{x}$ and $\boldsymbol{z}$. The $\boldsymbol{x}$ update uses the closed-form solution that involves matrix inverse or conjugate gradient. The $\boldsymbol{z}$ update is simply a soft-thresholding operation.

Similarly, the update of $\boldsymbol{x}$ in ADMM is

$$\boldsymbol{x}_{k+1}, \boldsymbol{z}_{k+1} = \underset{\boldsymbol{x}\in\mathbb{F}^N, \boldsymbol{z}\in\mathbb{C}^K}{\arg\min}\ \frac{\rho}{2}\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{v}_{k+1} - \boldsymbol{\eta}_k\|_2^2 + \beta\Big(\frac{1}{2}\|\boldsymbol{T}\boldsymbol{x} - \boldsymbol{z}\|_2^2 + \alpha\|\boldsymbol{z}\|_1\Big).$$

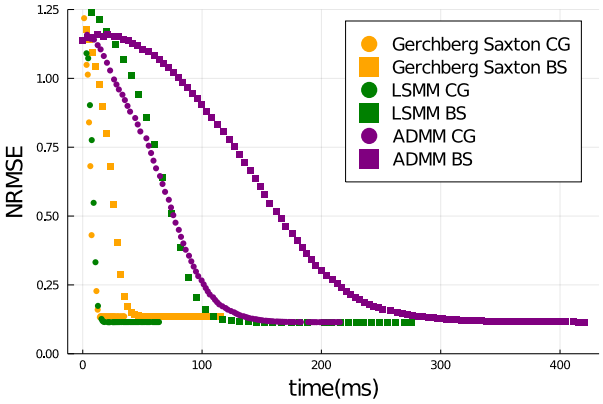

(a)

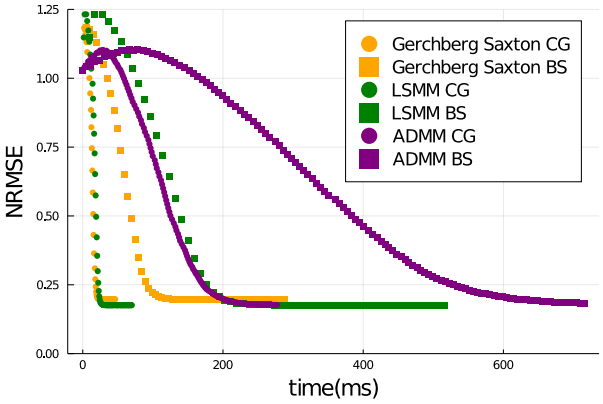

(b)

Fig. S.4: NRMSE vs. time (ms) using different update strategy. Here $\boldsymbol{A}$ is random Gaussian with the average of $\boldsymbol{a}_i'\boldsymbol{x}$ equals to 2, $M = 3000$ and $b_i = 0.5$. Subfigure (a) and (b) correspond to experiments on two true signals, namely $\boldsymbol{x}_{\text{true-A}}$ and $\boldsymbol{x}_{\text{true-B}}$, respectively. $\boldsymbol{x}_{\text{true-B}}$ is identical to the $\boldsymbol{x}_{\text{true}}$ defined in our paper; and $\boldsymbol{x}_{\text{true-A}}$ is the real part of $\boldsymbol{x}_{\text{true-B}}$.

We compared these two approaches (Huber function vs. alternating minimization) and found that using Huber function was more efficient than alternating minimization, as evident in Fig. S.7.

G. Another Test Image for DFT Matrix Setting

We also tested our proposed algorithms on a binary "UM1817" image of size 64×64, as shown in the upper-left corner of Fig. S.9. Here the experimental settings are identical to those for the Shepp-Logan phantom shown in Section V-B.

1) Accuracy comparison: Fig. S.9 showed the reconstructed results among different algorithms. Again, similar to Fig. 5, algorithms based on Poisson ML model yielded consistently better reconstruction quality than

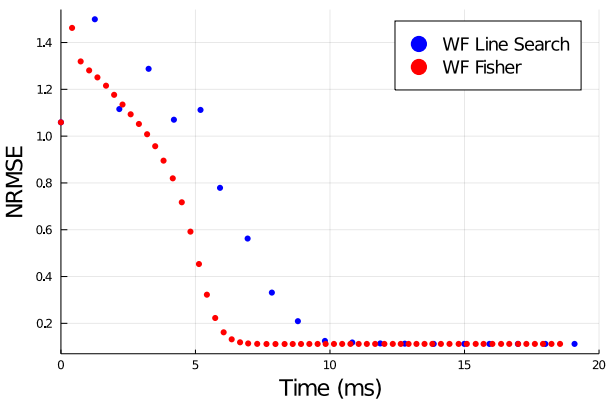

Fig. S.5: NRMSE vs. time (ms) using different step size strategy for WF. Results were averaged from 50 independent experiments. Here $\boldsymbol{A}$ is random Gaussian with the average of $\boldsymbol{a}_i'\boldsymbol{x}$ equals to 2, $M = 5000$ and $b_i = 0.1$. $\boldsymbol{x}_{\text{true}}$ was used as the true signal.

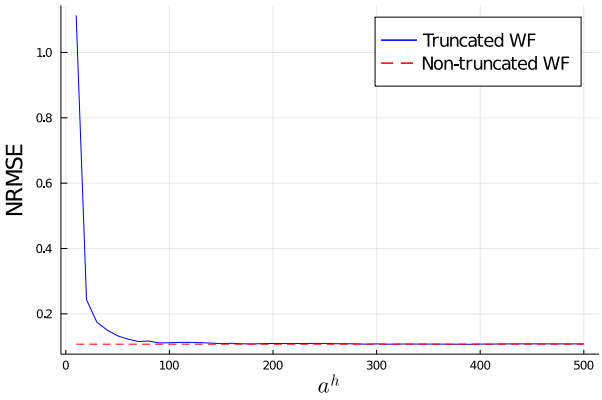

(a)

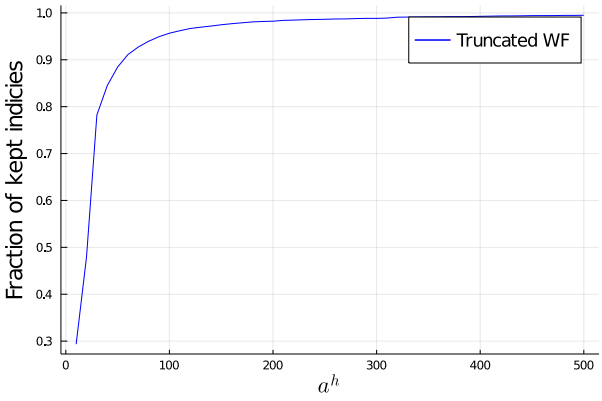

(b)

Fig. S.6: Comparison between truncated WF and non-truncated WF. a^h was set from 10 to 500 with interval 10. System matrix $\boldsymbol{A}$ was modeled as a random Gaussian. The average of $\boldsymbol{a}_i'\boldsymbol{x}$ is 2, with $b_i = 0.1$ and $M = 5000$.

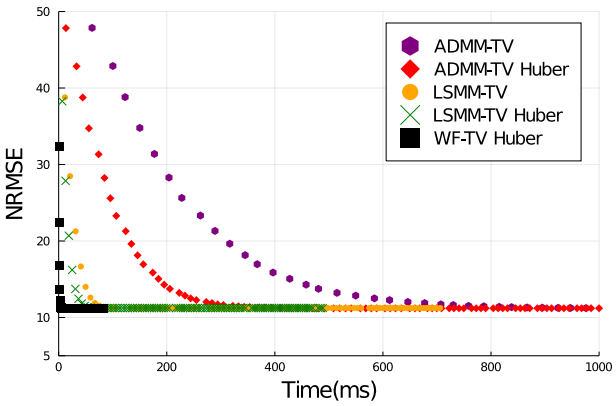

Fig. S.7: Speed comparison of WF, LSMM and ADMM using Huber function and alternating minimization for regularization. LSMM-TV and ADMM-TV denote the alternating minimization approach.

algorithms based on Gaussian ML model. Algorithms exploiting the assumed properties of true images can further lower the reconstruction error.

2) Convergence speed comparison: Fig. S.8 shows how fast algorithms decrease NRMSE for the binary 'UM1817' image. Among unregularized algorithms, GS converged the fastest, perhaps due to the simple inverse operation of the diagonal matrix $\boldsymbol{A}'\boldsymbol{A}$. Perhaps for the same reason, ADMM converged faster than LSMM, for both unregularized and regularized cases. Thus, based on the convergence rate of algorithms shown in Fig. S.8, an efficient strategy would be to first run few iterations of GS then switch to the ADMM-ODWT method.

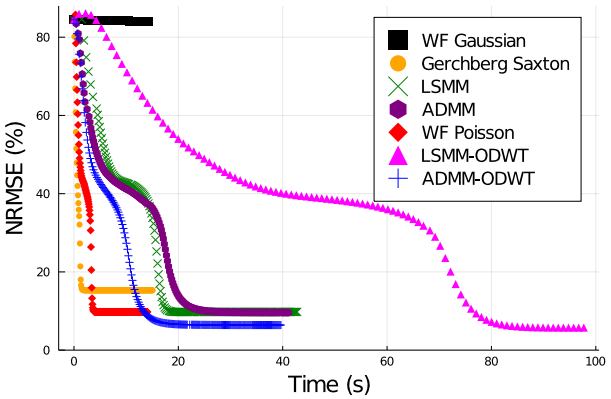

Fig. S.8: NRMSE vs. time (s) when $\boldsymbol{A}$ is a DFT matrix. The true image is shown at the upper-left corner of Fig. S.9.

H. Experiments on an Empirical Transmission Matrix

As an interesting experiment (for completeness), we also tested our algorithms using an empirical transmission matrix provided by [41]. As discussed in [41, 42], a transmission matrix $\boldsymbol{A}$ can be learned empirically through physical experiments. We chose the empirical transmission matrix (of size 65536×256) that was learned using the prVAMP algorithm [42] and uniformly sampled 10000 rows of it so that $\boldsymbol{A} \in \mathbb{C}^{10000\times 256}$. Unlike in the previous sections where we scaled the system matrix $\boldsymbol{A}$, here we keep $\boldsymbol{A}$ fixed and scaled pixel values in the true image (Fig. S.10) to have pixel range from 0 to 0.25. With this scaling, the average of $|\boldsymbol{a}_i'\boldsymbol{x}_{\text{true}}|^2$ was 1.14. The true image was selected from the dataset provided in [42], where all images have pixel values ranging from 0 to 255 before scaling. The background counts $\boldsymbol{b}$ were set to 0.1. Again, we chose $\boldsymbol{T}$ to be an ODWT matrix to exploit sparsity of the true image's discrete wavelet transform. We chose the regularization parameter to be $\beta = 32$ for both LSMM-ODWT and ADMM-ODWT and we ran 3 inner iterations of POGM for each outer iteration.

Fig. S.10 shows the images reconstructed using different algorithms. Again, as expected, the Poisson methods (WF Poisson, LSMM, ADMM) all demonstrated noticeably improved image quality (reduced NRMSE) compared to Gaussian methods (WF Gaussian, GS). The regularized Poisson methods (LSMM-ODWT, ADMM-ODWT) that exploit the assumed properties of the true image further improved the reconstruction quality.

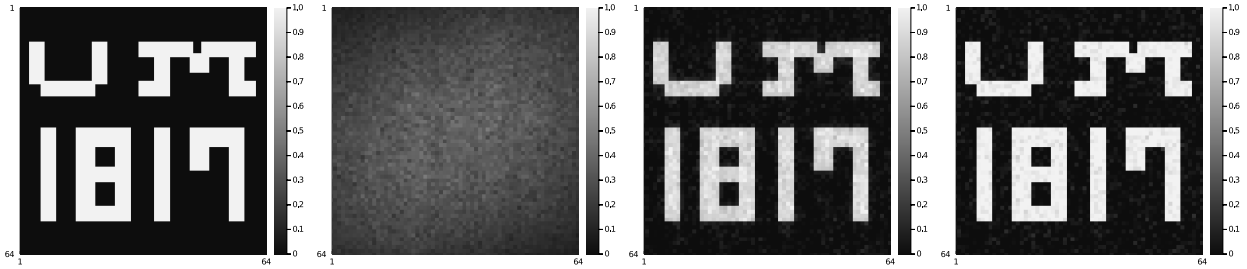

(a) True image

(b) WF Gaussian (NRMSE = 84.0%)

(c) GS (NRMSE = 15.3%)

(d) WF Poisson (NRMSE = 9.8%)

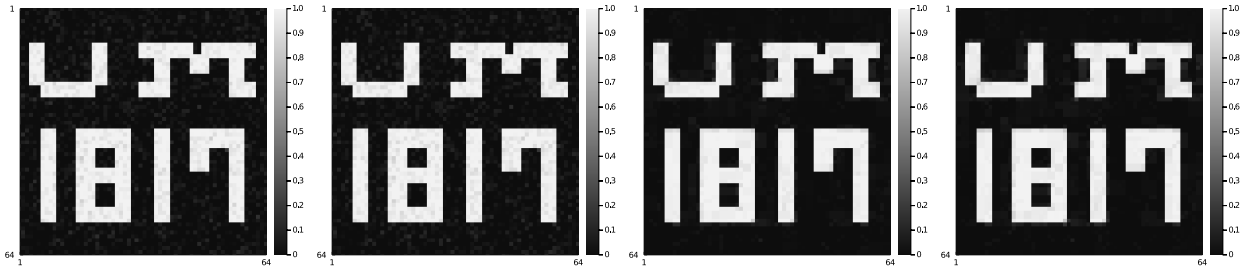

(e) LSMM (NRMSE = 9.8%)

(f) ADMM (NRMSE = 9.6%)

(g) LSMM-ODWT (NRMSE = 5.8%)

(h) ADMM-ODWT (NRMSE = 6.4%)

Fig. S.9: Reconstructed image and corresponding NRMSE compared to the true image of size 64×64, for a sensing system with $L = 21$ masked DFT matrices. The average of $|\boldsymbol{a}_i'\boldsymbol{x}|^2$ and b_i were set to 1 and 0.1, respectively.

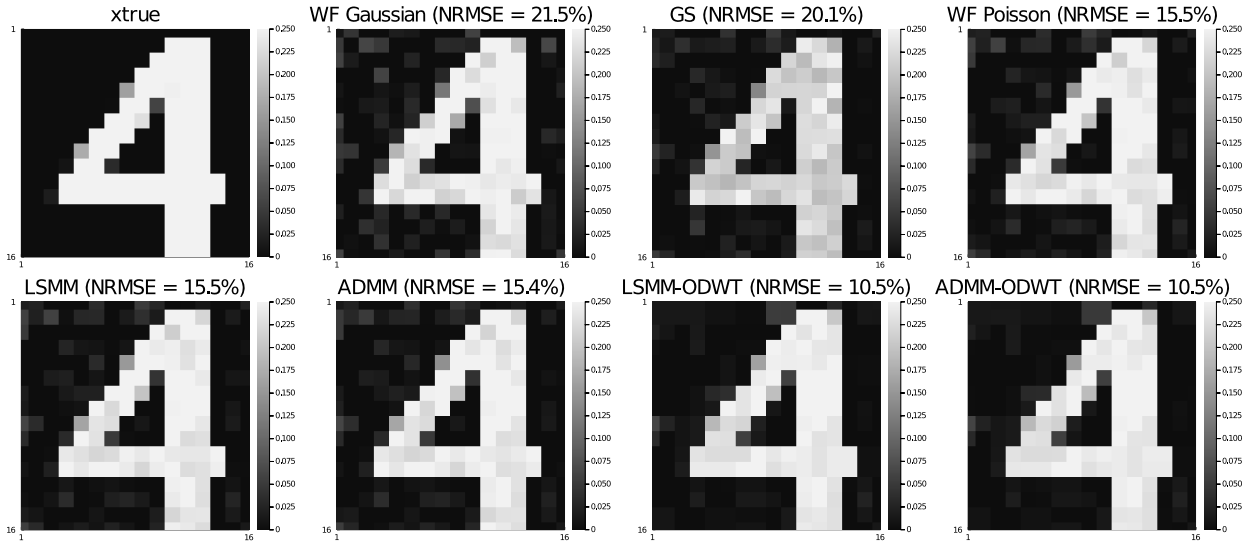

Fig. S.10: Reconstructed images and their corresponding NRMSE compared to the true image (of size 16×16), using an empirical system matrix from [42] with $M = 10000$.