

PAPER • OPEN ACCESS

Wirtinger gradient descent methods for low-dose Poisson phase retrieval




To cite this article: Benedikt Diederichs *et al* 2024 *Inverse Problems* **40** 125030

View the [article online](#) for updates and enhancements.

You may also like

- [Bayesian inverse Navier–Stokes problems: joint flow field reconstruction and parameter learning](#)
Alexandros Kontogiannis, Scott V Elgersma, Andrew J Sederman et al.
- [Time-harmonic optical flow with applications in elastography](#)
Oleh Melnyk, Michael Quellmalz, Gabriele Steidl et al.
- [Non-local structured adaptive dictionary learning method for seismic waveform inversion](#)
Hongyu Qi, Zhenwu Fu, Yang Li et al.

Wirtinger gradient descent methods for low-dose Poisson phase retrieval

Benedikt Diederichs^{1,2} , Frank Filbir^{1,3} 
and Patricia Römer^{1,3,*} 

¹ Mathematical Imaging and Data Analysis, Institute of Biological and Medical Imaging, Helmholtz Center Munich, Neuherberg, Germany

² Department of Chemistry and Center for NanoScience, Ludwig Maximilian University of Munich, Munich, Germany

³ Department of Mathematics, Technical University of Munich, Garching, Germany

E-mail: patricia.roemer@tum.de

Received 27 March 2024; revised 1 November 2024

Accepted for publication 27 November 2024

Published 13 December 2024



CrossMark

Abstract

The problem of phase retrieval has many applications in the field of optical imaging. Motivated by imaging experiments with biological specimens, we primarily consider the setting of low-dose illumination where Poisson noise plays the dominant role. In this paper, we discuss gradient descent algorithms based on different loss functions adapted to data affected by Poisson noise, in particular in the low-dose regime. Starting from the maximum log-likelihood function for the Poisson distribution, we investigate different regularizations and approximations of the problem to design an algorithm that meets the requirements that are faced in applications. In the course of this, we focus on low-count measurements. Based on an improved version of a variance-stabilizing transform for the Poisson distribution, we derive a decision rule for the regularization parameter in an averaged amplitude-based loss function. For all discussed loss functions, we study the convergence of the respective gradient descent algorithms to stationary points and find constant step sizes

* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

that guarantee descent of the loss in each iteration. Numerical experiments in the low-dose regime are performed to corroborate the theoretical observations.

Keywords: phase retrieval, Poisson noise, low-dose imaging, optimization, gradient descent

1. Introduction

Phase retrieval is a fundamental problem particularly in diffraction imaging techniques, where intensity measurements of several diffraction patterns of an object are recorded. In many instances, these measurements are taken in the far-field distance. It can be shown that such far-field intensity measurements are given by the squared absolute value of the Fourier transform of the object. To make this more concrete suppose a detector with n pixels is placed in the far-field distance. A discretized form of the signal at the detector plane is then given by the discrete Fourier transform

$$\sum_{j=0}^{n-1} x_j \exp\left(\frac{-2\pi i k j}{n}\right) = \langle x, u_k \rangle,$$

with vectors $u_k = (\exp(2\pi i k j/n))_{j=0}^{n-1}$ and $x = (x_j)_{j=0}^{n-1}$. The latter is a discrete representation of the object which we would like to reconstruct. According to this model, the vector u_k can be related to the k th frequency, or equivalently, with the k th pixel of the detector. As mentioned above, the detector itself measures only intensities, that is, it records only the squared modulus of $\langle x, u_k \rangle$. Consequently, we are faced with the challenging problem of reconstructing x from data of the form $|\langle x, u_k \rangle|^2$, $k = 0, \dots, n-1$.

It is known, however, that this data set does not contain enough information to make the inverse problem uniquely solvable [1], i.e. it is not sufficient to probe the object x only by the pure states u_k , $k = 0, \dots, n-1$. To avoid these difficulties, we have to insert a certain amount of redundancy into our data set. In order to introduce a sufficient portion of redundancy, one replaces the set of measurement vectors u_k by an over-complete system of vectors $a_k \in \mathbb{C}^n$, $k = 1, \dots, m$ with $m \gg n$, and measures $|\langle x, a_k \rangle|^2$, $k = 1, \dots, m$, instead of only $|\langle x, u_k \rangle|^2$, $k = 0, \dots, n-1$. A prominent realization for this approach is the so-called far-field ptychography. This method uses measurement vectors of the form $a_{k,\ell} = (w_{j,\ell} \exp(-2\pi i k j/n))_{j=0}^{n-1}$ with a mask $w_{j,\ell}$, which is given as translates $(w_{j-\ell \bmod n})_{j=0}^{n-1}$ of a vector $w = (w_j)_{j=0}^{n-1}$ with short support. The number of translates is arranged such that the object is scanned by the mask in a way that for adjacent scanning positions the supports of the mask overlap with each other by a certain fraction. We will not discuss ptychography further, instead we refer to [2–4].

Motivated by ptychography, we consider the following reconstruction problem. Given data of the form

$$\hat{y}_i = |\langle x, a_i \rangle|^2, \quad i = 1, \dots, m, \quad (1)$$

we have to reconstruct the vector $x \in \mathbb{C}^n$. This is the phase retrieval problem in its mathematical abstract form. The problem attracted a lot of attention during the last decades and a number of fundamental contributions were made [5–8]. The first fundamental problem is that of uniqueness. That is the question of how large m has to be in order to make the mapping $\mathbb{C}^n \rightarrow \mathbb{C}^m$, $x \mapsto (|\langle x, a_i \rangle|^2)_{i=1}^m$ injective. This problem was considered by several authors under

different assumptions on x and on the vectors a_i . We refer to the fundamental papers [6, 9] for further discussion. Our focus lies on the reconstruction of the vector x given \hat{y}_i , $i = 1, \dots, m$, and we will consider this problem under specific assumptions on the data set, which are again motivated by specific constraints in the experimental setup. Before going into the details of those, let us first recall the reconstruction techniques based on a variational approach. The variational method tries to recuperate x via a minimization

$$x = \arg \min_z \mathcal{L}(z), \quad (2)$$

using a suitable loss function $\mathcal{L} : \mathbb{C}^n \rightarrow \mathbb{R}$. For determining x , or at least a good approximation, usually gradient descent methods are applied to the problem (2), i.e.

$$z_{k+1} = z_k - \mu \nabla \mathcal{L}(z_k), \quad k \geq 0,$$

with learning rate $\mu > 0$ and some appropriate initial vector z_0 . For the phase retrieval problem, the loss function is of the form $\mathcal{L}(z) = \ell \circ \varphi(z)$, with a suitable function $\ell : \mathbb{R}^m \rightarrow \mathbb{R}$ and $\varphi : \mathbb{C}^n \rightarrow \mathbb{R}^m$, $z \mapsto (|\langle a_i, z \rangle|^2)_{i=1}^m$. Note also that the function φ usually makes the problem non-convex.

A frequently used loss function for reconstruction of x is the least squares loss

$$\mathcal{L}(z) = \sum_{i=1}^m (|\langle a_i, z \rangle|^2 - \hat{y}_i)^2,$$

or some regularized variants of it. Gradient descent algorithms for the corresponding minimization problem which apply Wirtinger calculus were investigated to some extent. These methods, now known as Wirtinger flow algorithms, were studied first by Candés, Li and Soltanolkotabi in [10]. Different variants were further discussed by other authors [11–19]. These works study the convergence of stochastic and non-stochastic gradient descent algorithms involving Wirtinger derivatives for different loss functions in case of random (Gaussian) measurement vectors a_i . Apart from convergence guarantees for random measurements, [20] and [21] analyze convergence of gradient descent algorithms to stationary points of an amplitude-based loss function for any type of measurement vectors, random or deterministic.

In all practical measurement scenarios, the data \hat{y}_i is corrupted by some sort of noise. Hence, we are given perturbed values y_i instead of \hat{y}_i . The perturbation may have many sources such as thermal noise, read-out noise, background noise, among others [22]. The problem of a certain type of background noise was recently tackled by one of the authors in [23]. Here, we concentrate on the perturbation which is caused by the operation mode of the detector. All modern detectors, such as CCD cameras, are ultimately counting devices. That means, the measurement process can be modeled as a counting process and can, therefore, mathematically be formulated in terms of a Poisson distributed discrete random variable, viz.

$$y_i \sim \text{Poisson} \left(|\langle a_i, x \rangle|^2 \right), \quad i = 1, \dots, m,$$

where $|\langle a_i, x \rangle|^2$ is the ground-truth. This means that the probability that y_i particles are counted at position i given the ground-truth $|\langle a_i, x \rangle|^2$ is $\frac{1}{y_i!} \exp(-|\langle a_i, x \rangle|^2) (|\langle a_i, x \rangle|^2)^{y_i}$. In order to incorporate the fact that the measurement process is a Poisson process into a variational reconstruction method, it is necessary to adapt the loss function accordingly. We follow a maximum

likelihood approach and replace the least squares loss function by the Poisson log-likelihood loss which reads as

$$\mathcal{L}_P(z) = \sum_{i=1}^m |\langle a_i, z \rangle|^2 - y_i \log(|\langle a_i, z \rangle|^2).$$

The minimization problem (2) for the Poisson likelihood loss \mathcal{L}_P was studied by many authors, see for example [13, 24–28] and references therein.

Our motivation to reconsider the minimization problem (2) with Poisson log-likelihood \mathcal{L}_P originates in specific diffraction imaging setups. For imaging biological tissue like cells, viruses, etc it is generally not possible to work with a radiation beam of high intensity as such exposure would destroy the object instantaneously [29–31]. Therefore, it is essential to perform the measurement with a beam of suitably low intensity, which in turn leads to a weak signal at the detector. In realistic measurement scenarios, the counting rate can be in the range below 10 illumination particles per pixel, and at a number of pixels it can be even zero. This low-count scenario leads to serious problems with respect to the gradient descent reconstruction process as it causes a singularity in the gradient of \mathcal{L}_P . In order to take these problems into account, we have to introduce a regularization technique which deals with these singularities. Algorithmic approaches for solving phase retrieval problems with low-dose Poisson noisy data were also considered in [32–35].

With this paper, we contribute a decision rule for a regularization parameter based on a heuristic motivated by a new improved variance stabilization approach. Furthermore, we provide convergence guarantees which generalize the result of [20] and we extend this analysis to the loss using the Poisson maximum log-likelihood function.

The outline of the paper is as follows. In section 2, we summarize the theory on Wirtinger derivatives and some fundamental results on the convergence of the gradient descent algorithm with Wirtinger derivatives. Section 3 contains a discussion on the choice of loss function in the optimization problem formulated for solving the phase retrieval problem with Poisson noisy low-dose data. In section 4, we present a convergence analysis for the algorithms involving the different loss functions. We provide numerical justification of our theoretical results and corroborate our proposal of using the loss function derived via the improved variance stabilization approach in section 5. In section 6, we summarize our results and conclude with a brief outlook on further possible model adjustments.

2. Gradient descent with Wirtinger derivatives

We start with presenting some theory on the Wirtinger calculus. Consider a function

$$f(z) = u(x, y) + iv(x, y), \quad z = x + iy, \quad x, y \in \mathbb{R}^n,$$

with real-valued and differentiable functions u and v . Using conjugate variables $z = x + iy$ and $\bar{z} = x - iy$, the function f can be considered as a function of variables z and \bar{z} . Since u and v are differentiable, the function $f(z, \bar{z})$ is holomorphic w.r.t. z for fixed \bar{z} and vice versa. The Wirtinger calculus expresses the derivatives of f w.r.t. the real variables x and y in terms of the conjugate variables z and \bar{z} treating them as independent. The Wirtinger derivatives of f are given as

$$\partial_z f = \frac{1}{2} (\partial_x f - i \partial_y f), \quad \partial_{\bar{z}} f = \frac{1}{2} (\partial_x f + i \partial_y f).$$

This implies the relation

$$\overline{\partial_z f} = \partial_{\bar{z}} \bar{f}, \quad \text{and} \quad \overline{\partial_{\bar{z}} f} = \partial_z \bar{f}. \tag{3}$$

The Wirtinger derivatives $\partial_z f$ and $\partial_{\bar{z}} f$ can also be expressed as

$$\begin{aligned} \partial_z f &= \partial_z f(z, \bar{z}) \Big|_{\bar{z}=\text{const.}} = [\partial_{z_1} f(z, \bar{z}), \dots, \partial_{z_n} f(z, \bar{z})] \Big|_{\bar{z}=\text{const.}}, \\ \partial_{\bar{z}} f &= \partial_{\bar{z}} f(z, \bar{z}) \Big|_{z=\text{const.}} = [\partial_{\bar{z}_1} f(z, \bar{z}), \dots, \partial_{\bar{z}_n} f(z, \bar{z})] \Big|_{z=\text{const.}}. \end{aligned}$$

Consequently, the Wirtinger gradient and Wirtinger Hessian are given as

$$\nabla f(z) = \begin{pmatrix} (\partial_z f)^* \\ (\partial_{\bar{z}} f)^* \end{pmatrix}, \quad \nabla^2 f(z) = \begin{pmatrix} \partial_z (\partial_z f)^* & \partial_z (\partial_{\bar{z}} f)^* \\ \partial_{\bar{z}} (\partial_z f)^* & \partial_{\bar{z}} (\partial_{\bar{z}} f)^* \end{pmatrix}. \tag{4}$$

In case that f is a real-valued function, i.e. $f(z) = u(x, y)$, the relations in (3) provide

$$\partial_{\bar{z}} f = \overline{\partial_z f}, \quad \partial_{\bar{z}} (\partial_z f)^* = \overline{\partial_z (\partial_z f)^*}, \quad \partial_z (\partial_{\bar{z}} f)^* = \overline{\partial_z (\partial_{\bar{z}} f)^*}.$$

It is more convenient to use the following simplified notation

$$\nabla_z f := (\partial_z f)^*, \quad \nabla_{z, \bar{z}}^2 f := \partial_z (\partial_z f)^*, \quad \text{resp.} \quad \nabla_{\bar{z}, z}^2 f := \partial_{\bar{z}} (\partial_{\bar{z}} f)^*.$$

The second-order Taylor polynomial of f at a point z_0 is

$$P_f(v, z_0) = f(z_0) + (\nabla f(z_0))^* \begin{pmatrix} v \\ \bar{v} \end{pmatrix} + \begin{pmatrix} v \\ \bar{v} \end{pmatrix}^* \nabla^2 f(z) \begin{pmatrix} v \\ \bar{v} \end{pmatrix}.$$

In case of a real-valued function f , the quadratic term of the Taylor polynomial P_f can be expressed as

$$\begin{pmatrix} v \\ \bar{v} \end{pmatrix}^* \nabla^2 f(z) \begin{pmatrix} v \\ \bar{v} \end{pmatrix} = 2\text{Re}(v^* \nabla_{z, z}^2 f(z) v) + 2\text{Re}(v^* \nabla_{\bar{z}, z}^2 f(z) \bar{v}). \tag{5}$$

In this paper, we focus on real-valued functions f . For minimizing such a function f , we apply gradient descent

$$z_{k+1} = z_k - \mu_k \nabla_z f(z_k) \tag{6}$$

with some appropriate initial vector $z_0 \in \mathbb{C}^n$. The parameter $\mu_k > 0$ is called step size or learning rate. It can be chosen to be constant or adaptive, preferably such that descent in every iteration is guaranteed, i.e. $f(z_{k+1}) \leq f(z_k)$ for all $k \geq 0$. The proof of the following result can be found in [21].

Proposition 1. *Let $b \in \mathbb{R}, f: \mathbb{C}^n \rightarrow [b, \infty)$, be a twice Wirtinger differentiable function with a uniformly bounded Hessian, i.e.*

$$\begin{pmatrix} v \\ \bar{v} \end{pmatrix}^* \nabla^2 f(z) \begin{pmatrix} v \\ \bar{v} \end{pmatrix} \leq L \left\| \begin{pmatrix} v \\ \bar{v} \end{pmatrix} \right\|_2^2$$

for all $z, v \in \mathbb{C}^n$, with a constant $L > 0$ independent of z . Let the sequence $(z_k)_{k \geq 0}$ be generated by the update (6) with an arbitrary initialization $z_0 \in \mathbb{C}^n$. If $0 < \mu \leq L^{-1}$, then

$$f(z_k) - f(z_{k+1}) \geq \mu \|\nabla_z f(z_{k+1})\|_2^2$$

for all $k \geq 0$. Then, if f has compact sublevel sets $L_s(f) = \{z \in \mathbb{C}^n : f(z) \leq s\}$, the Wirtinger flow algorithm (6) is guaranteed to converge to a stationary point of f .

We now specify the requirements of this result to loss functions which are relevant in our context. As mentioned above, the loss functions we are going to consider are given as a composition $\ell \circ \varphi$ with a smooth function $\ell: \mathbb{R} \rightarrow \mathbb{R}$ and $\varphi: \mathbb{C}^n \rightarrow [0, \infty)$, $\varphi(z) = |\langle a, z \rangle|^2$. For such functions, we obtain the following bound on the Hessian.

Lemma 2. *Let $\ell: \mathbb{R} \rightarrow \mathbb{R}$ be twice differentiable and $\varphi: \mathbb{C}^n \rightarrow [0, \infty)$, $\varphi(z) = |\langle a, z \rangle|^2$ with $a \in \mathbb{C}^n$. Then*

$$\begin{aligned} \left(\frac{v}{\bar{v}} \right)^* \nabla^2 (\ell \circ \varphi)(z) \left(\frac{v}{\bar{v}} \right) &\leq \left(2\ell''(|\langle a, z \rangle|^2) |\langle a, z \rangle|^2 \right. \\ &\quad \left. + \ell'(|\langle a, z \rangle|^2) \right) \|a\|_2^2 \left\| \left(\frac{v}{\bar{v}} \right) \right\|_2^2 \end{aligned}$$

for all $z, v \in \mathbb{C}^n$.

Proof. The Wirtinger derivative of $(\ell \circ \varphi)(z, \bar{z}) = \ell(\bar{z}^T a a^* z)$ is

$$\begin{aligned} \partial_z (\ell \circ \varphi)(z, \bar{z}) &= \partial_z (\ell \circ \varphi)(z, \bar{z})|_{\bar{z}=\text{const.}} = \ell'(\bar{z}^T a a^* z) \bar{z}^T a a^* \\ &= \ell'(|\langle a, z \rangle|^2) z^* a a^*. \end{aligned}$$

For the second derivatives we obtain

$$\begin{aligned} \partial_z (\partial_z (\ell \circ \varphi)(z, \bar{z}))^* &= \partial_z \left(\ell'(|\langle a, z \rangle|^2) a a^* z \right) = \ell''(|\langle a, z \rangle|^2) |\langle a, z \rangle|^2 a a^* \\ &\quad + \ell'(|\langle a, z \rangle|^2) a a^*, \end{aligned}$$

and analogously

$$\partial_{\bar{z}} (\partial_z (\ell \circ \varphi)(z, \bar{z}))^* = \partial_{\bar{z}} \left(\ell'(|\langle a, z \rangle|^2) a a^* z \right) = \ell''(|\langle a, z \rangle|^2) \langle a, z \rangle^2 a a^T.$$

In view of (5) we get

$$\begin{aligned} \left(\frac{v}{\bar{v}} \right)^* \nabla^2 (\ell \circ \varphi)(z) \left(\frac{v}{\bar{v}} \right) &= 2 \left(\ell''(|\langle a, z \rangle|^2) |\langle a, z \rangle|^2 + \ell'(|\langle a, z \rangle|^2) \right) |\langle a, v \rangle|^2 \\ &\quad + 2\text{Re} \left(\ell''(|\langle a, z \rangle|^2) \langle a, z \rangle^2 \langle a, v \rangle^2 \right) \\ &\leq 2 \left(2\ell''(|\langle a, z \rangle|^2) |\langle a, z \rangle|^2 + \ell'(|\langle a, z \rangle|^2) \right) |\langle a, v \rangle|^2 \\ &\leq \left(2\ell''(|\langle a, z \rangle|^2) |\langle a, z \rangle|^2 + \ell'(|\langle a, z \rangle|^2) \right) \\ &\quad \times \|a\|_2^2 \left\| \left(\frac{v}{\bar{v}} \right) \right\|_2^2, \end{aligned}$$

where we used $\text{Re}(\alpha^2 \beta^2) \leq |\alpha|^2 |\beta|^2$ for $\alpha, \beta \in \mathbb{C}$. □

In order to apply the convergence result of proposition 1, we need to bound $2\ell''(x)x + \ell'(x)$ by a constant independent of $x \in [0, \infty)$ and we have to show that the level sets of $\ell \circ \varphi$ are compact. The compactness of the level sets will be addressed in the following lemma.

Lemma 3. *Let $f: \mathbb{C}^n \rightarrow \mathbb{R}$, $z \mapsto \sum_{i=1}^m |\langle a_i, z \rangle|^2$ with $a_i \in \mathbb{C}^n$, $i = 1, \dots, m$. Denote with A the matrix with rows a_i^* , $i = 1, \dots, m$. For any $z_0 \in \mathbb{C}^n$, f restricted to $z_0 + \text{range}(A^*)$ has compact level sets.*

Proof. Consider any $z \in z_0 + \text{range}(A^*)$. We use that $\sum_{i=1}^m |\langle a_i, z \rangle|^2 = \|Az\|_2^2$. It is $z_0 = \hat{z}_0 + \tilde{z}_0$ with $\hat{z}_0 \in \ker(A)$ and $\tilde{z}_0 \in \text{range}(A^*)$. The vector $\tilde{z} = z - \hat{z}_0$ is the orthogonal projection of z onto the range of A^* . As \tilde{z} is orthogonal to the kernel of A , it is $\|A\tilde{z}\|_2^2 \geq \sigma_{\min}^2 \|\tilde{z}\|_2^2$, where σ_{\min} denotes the smallest non-zero singular value of A . Then, as $\|Az\|_2^2 = \|A\tilde{z}\|_2^2$, we can continue to bound

$$\|Az\|_2^2 \geq \sigma_{\min}^2 \|\tilde{z}\|_2^2 \geq \sigma_{\min}^2 (\|z\|_2^2 - \|\hat{z}_0\|_2^2).$$

Since σ_{\min} and $\|\hat{z}_0\|_2^2$ are constant for all $z \in z_0 + \text{range}(A^*)$, this shows that $\|Az\|_2^2$ is bounded from below by a scaled and shifted version of $\|z\|_2^2$ for all $z \in z_0 + \text{range}(A^*)$. As $z \mapsto \|z\|_2^2$ has compact level sets, f has compact level sets on $z_0 + \text{range}(A^*)$. \square

Remark 4. With lemma 3 we obtain that $\sum_{i=1}^m \ell_i(|\langle a_i, z \rangle|^2)$ has compact level sets on $z_0 + \text{range}(A^*)$ if the functions $\ell_i: [0, \infty) \rightarrow \mathbb{R}$, $i = 1, \dots, m$, are continuous, bounded from below, and satisfy $\ell_i(t) \rightarrow \infty$ for $t \rightarrow \infty$.

3. Phase retrieval as an optimization problem

As discussed above, our aim is to reconstruct an (approximate) solution to the phase retrieval problem (1) under certain assumptions on the random nature of the measurement process. Our focus lies on those cases where the measurement process can be modeled as a Poisson distributed random process. In order to associate this assumption with a variational reconstruction method, we use the maximum (log-) likelihood estimation.

3.1. Poisson log-likelihood loss

In order to maximize the (log-) likelihood function for the Poisson distributed random variable y_i with ground-truth $|\langle a_i, x \rangle|^2$, we have to determine z such that $\sum_{i=1}^m (y_i \log(|\langle a_i, z \rangle|^2) - |\langle a_i, z \rangle|^2)$ is maximal. This is equivalent to determining $\arg \min_z \mathcal{L}_P(z)$ with

$$\mathcal{L}_P(z) = \sum_{i=1}^m |\langle a_i, z \rangle|^2 - y_i \log(|\langle a_i, z \rangle|^2). \quad (7)$$

As mentioned earlier, we are interested in measurement scenarios where we can have no counts at certain pixels. In those cases we need to consider the value $|\langle a_i, z \rangle|^2 = 0$, which leads to singularities in the loss function \mathcal{L}_P . A simple strategy to deal with them is shifting the logarithmic term by a positive constant $\varepsilon_P > 0$, i.e.

$$\mathcal{L}_{P, \varepsilon_P}(z) = \sum_{i=1}^m |\langle a_i, z \rangle|^2 - y_i \log(|\langle a_i, z \rangle|^2 + \varepsilon_P). \quad (8)$$

The corresponding gradient descent update rule with constant step size reads as

$$z_{k+1} = z_k - \mu \sum_{i=1}^m \left(1 - \frac{y_i}{|\langle a_i, z_k \rangle|^2 + \varepsilon_P} \right) \langle a_i, z_k \rangle a_i.$$

A suitable choice of the parameter is, however, not obvious. In this regard, one could for example study a discrepancy principle similar to [36, 37]. We, instead, propose to consider an approximation of the Poisson loss (8). An appropriate approximation is discussed in the next section.

3.2. Gaussian log-likelihood and variance stabilization

Recall that the maximum log-likelihood loss function for a Gaussian distributed random variable y_i with distribution $\mathcal{N}(|\langle a_i, x \rangle|^2, \sigma_i^2)$, $i = 1, \dots, m$, is given as

$$\mathcal{L}(z) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \left(|\langle a_i, z \rangle|^2 - y_i \right)^2. \quad (9)$$

Often, this loss is used for image reconstruction even if the underlying process is a counting process, i.e. the random variable is Poisson distributed. This can be justified in cases for which the variance parameter λ in the Poisson distribution $Poisson(\lambda)$ is sufficiently large. In this situation, the central limit theorem shows that $Poisson(\lambda)$ can be approximated by a Gaussian distribution $\mathcal{N}(\lambda, \lambda)$. However, in a low-dose scenario this approximation is no longer suitable.

Another approach to approximate a Poisson random variable by a Gaussian random variable is the so-called variance stabilization method [38–41]. This method transforms the Poissonian data such that the resulting random variable has approximately constant variance. The rationale of this method is as follows. Let $\lambda > 0$ and $X \sim Poisson(\lambda)$. Assume $f: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a sufficiently smooth function. Then, its first order Taylor approximation around the variance λ is $f(t) \approx f(\lambda) + (t - \lambda)f'(\lambda)$, and, hence, $\mathbb{V}(f(X)) \approx \mathbb{V}(X) (f'(\lambda))^2$. In order to obtain an approximate constant variance, set the right-hand side of the latter relation equal to σ^2 . This leads to $f(t) = 2\sigma\sqrt{t}$. The approach goes back to [38, 39] and was modified by Anscombe in [41] using a fifth order Taylor approximation in order to handle cases where the variance λ is rather small. Anscombe considered the shifted square-root transform $f(t) = \sqrt{t+c}$. Following the same arguments as above, one arrives at

$$\begin{aligned} \mathbb{V}(\sqrt{X+c}) &\approx (\lambda+c) \cdot \mathbb{V} \left[1 + \frac{1}{2} \cdot \frac{X-\lambda}{\sqrt{\lambda+c}} - \frac{1}{8} \cdot \left(\frac{X-\lambda}{\sqrt{\lambda+c}} \right)^2 + \frac{1}{16} \cdot \left(\frac{X-\lambda}{\sqrt{\lambda+c}} \right)^3 \right. \\ &\quad \left. - \frac{5}{128} \cdot \left(\frac{X-\lambda}{\sqrt{\lambda+c}} \right)^4 + \frac{7}{256} \cdot \left(\frac{X-\lambda}{\sqrt{\lambda+c}} \right)^5 \right] \\ &\approx \frac{1}{4} \cdot \left(1 + \frac{\frac{3}{8}-c}{\lambda} + \frac{32c^2-52c+17}{32\lambda^2} \right), \end{aligned}$$

which suggests to choose $c = 3/8$ to achieve $\mathbb{V}(X) \approx 1/4$.

However, neither the simple square-root transform nor the Anscombe transform performs well when $\lambda \in [0, 2]$, hence, we aim for a better suited variance-stabilizing transform of square-root type $f(t) = \sqrt{t+c}$. We determine the parameter $c > 0$ such that $\mathbb{V}(\sqrt{X+c}) = 1/4$ by considering

$$\begin{aligned} \mathbb{V}(\sqrt{X+c}) &= \mathbb{E}(X) + c - \left[\mathbb{E}(\sqrt{X+c}) \right]^2 \\ &= \lambda + c - \left(\sum_{k=0}^{\infty} \sqrt{k+c} \cdot \frac{\exp(-\lambda)\lambda^k}{k!} \right)^2, \end{aligned}$$

and obtain an approximately optimal value for c by setting this equal to $1/4$. By this method, we obtain, e.g. for $\lambda = 1$ the optimal value $c \approx 0.12$ and for $\lambda = 2$ the value $c \approx 0.27$.

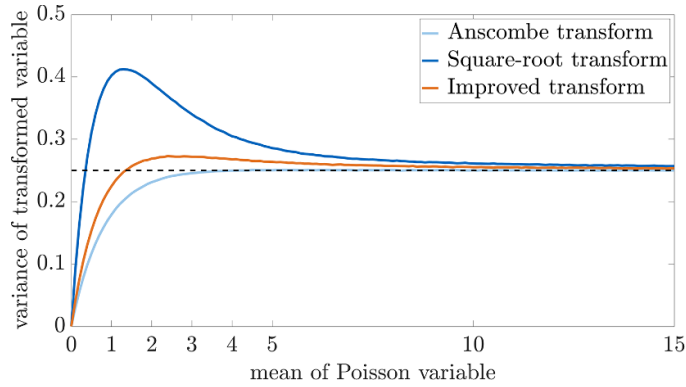


Figure 1. Variance-stabilizing transforms.

To cover not only one optimal choice of c for one specific value of λ , we propose to consider an averaging transform

$$f(t) = \frac{1}{2} (\sqrt{t+c_1} + \sqrt{t+c_2}) \tag{10}$$

with $c_1, c_2 \geq 0$. An example for such a transform is the Tukey–Freeman transform [42]

$$f(t) = \frac{1}{2} (\sqrt{t} + \sqrt{t+1}) \tag{11}$$

that is known to perform well for small λ . We advance this idea with using $c_1 = 0.12$ and $c_2 = 0.27$ in an experiment dominated by 1 and 2 counts. This is justified as for measurements of this size it is very likely that the underlying ground-truth is close to 1 or 2. The performance of this transform compared to the square-root and the Anscombe transform can be studied in figure 1.

By means of figure 1 we note that for λ close to 0 the square-root transform $f(t) = \sqrt{t}$ might be preferred over the other variance-stabilizing transforms. Also by the method described above we find that for $\lambda \rightarrow 0$ the optimal choice is c close to 0.

Returning back to the discussion on suitable loss functions, we can use any of the presented variance-stabilizing transforms f and consider the loss

$$\mathcal{L}(z) = \sum_{i=1}^m 2 \left(g(|\langle a_i, z \rangle|^2) - f(y_i) \right)^2,$$

which represents the log-likelihood function for a Gaussian distribution $\mathcal{N}(g(|\langle a_i, z \rangle|^2), \frac{1}{4})$, with function g satisfying $g(X) = \mathbb{E}(f(X))$.

If we consider a transform $f(t) = \sqrt{t+c}$ with $c \geq 0$ that stabilizes the variance of a Poisson random variable around $1/4$, the mean of the transformed random variable is approximately

$$\mathbb{E}(\sqrt{X+c}) = \sqrt{\mathbb{E}(X+c) - \mathbb{V}(\sqrt{X+c})} \approx \sqrt{\mathbb{E}(X) + c - \frac{1}{4}}. \tag{12}$$

Hence, we would work with a loss function

$$\mathcal{L}(z) = \sum_{i=1}^m 2 \left(\sqrt{|\langle a_i, z \rangle|^2 + c - \frac{1}{4}} - \sqrt{y_i + c} \right)^2. \tag{13}$$

A variant of this type of loss function is the amplitude loss

$$\mathcal{L}(z) = \sum_{i=1}^m 2 \left(\sqrt{|\langle a_i, z \rangle|^2 + \varepsilon_A} - \sqrt{y_i} \right)^2, \quad (14)$$

with $\varepsilon_A > 0$, which has attracted interest recently in the phase retrieval community, see [20, 21]. This loss function is motivated by the Gaussian model involving a square-root transform. However, this formulation disregards the relation (12) but uses the coarse approximation $\mathbb{E}(\sqrt{X}) \approx \sqrt{\mathbb{E}(X)}$.

A loss function based on the maximum likelihood model for a Gaussian distribution using the Anscombe transform, i.e. with $c = 3/8$, was considered before, e.g. in [43], and studied for mixed Poisson-Gaussian noise in [44]. These works also used the approximation $\mathbb{E}(\sqrt{X + 3/8}) \approx \sqrt{\mathbb{E}(X + 3/8)}$ instead of (12).

While a loss function (13) fits very well the problem setting we want to consider, applying a gradient descent method to this loss function is problematic if $c \leq 1/4$. The considered function is not well-defined for z with $|\langle a_i, z \rangle|^2 \in [0, 1/4 - c)$, and not differentiable at $|\langle a_i, z \rangle|^2 = 1/4 - c$. Hence, we need to regularize again to counteract this issue. We decide to also utilize the coarse approximation $\mathbb{E}(\sqrt{X + c}) \approx \sqrt{\mathbb{E}(X + c)}$ and consider

$$\mathcal{L}(z) = \sum_{i=1}^m 2 \left(\sqrt{|\langle a_i, z \rangle|^2 + c} - \sqrt{y_i + c} \right)^2. \quad (15)$$

The idea of adapting the expectation of the distribution accordingly is also not reflected in works as, e.g. [16]. The intention of those approaches is to smoothen the loss function. We, additionally, aim at formulating the loss such that it is close to the Poissonian model. However, in an experiment with too low illumination dose, resulting in ground-truth values much smaller than 1, the variance after variance-stabilizing transform is rather close to 0. Hence, the approximation (12) is not good for an extreme low-dose experiment and, thus, we abstain from subtracting $1/4$ in the mean approximation and prefer using (15).

More generally, we can also use an averaging transform (10), with values $c_1, c_2 \geq 0$, for variance stabilization. For the same reasons as before, we use the coarse approximation

$$\mathbb{E}(\sqrt{X + c_1} + \sqrt{X + c_2}) \approx \sqrt{\mathbb{E}(X) + c_1} + \sqrt{\mathbb{E}(X) + c_2},$$

and consider loss functions of the form

$$\mathcal{L}(z) = \sum_{i=1}^m \frac{1}{2} \left(\sqrt{|\langle a_i, z \rangle|^2 + c_1} + \sqrt{|\langle a_i, z \rangle|^2 + c_2} - \sqrt{y_i + c_1} - \sqrt{y_i + c_2} \right)^2.$$

All discussed variance-stabilizing transforms result in a variance approximately equal to $1/4$ only for distribution parameters $\lambda > 0$. Hence, it might not be advisable to use the loss functions resulting from the variance-stabilizing transforms in case of ground-truth values equal or close to 0. In an experiment, the ground-truth values are not known. However, if $y_i = 0$, it is likely that the corresponding ground-truth is close to 0 and we wish to avoid variance-stabilizing transforms for these instances. Note that for $y_i = 0$ no regularization is required in the exact Poisson log-likelihood model as the loss function reduces to the term $|\langle a_i, z \rangle|^2$. Therefore, we propose to consider

$$\begin{aligned} \mathcal{L}_0(z) &= \sum_{i=1}^m \mathbb{1}_{y_i=0} \cdot |\langle a_i, z \rangle|^2 \\ &\quad + \mathbb{1}_{y_i>0} \cdot \frac{1}{2} \left(\sqrt{|\langle a_i, z \rangle|^2 + c_1} + \sqrt{|\langle a_i, z \rangle|^2 + c_2} - C_i \right)^2, \end{aligned} \quad (16)$$

with $c_2 \geq c_1 > 0, C_i \geq 0, i = 1, \dots, m$, if one is working with low-dose data dominated by zero measurements.

4. Convergence analysis of gradient descent algorithms for phase retrieval

In the following, we analyze the convergence of gradient descent algorithms for optimization problems using the loss functions discussed in sections 3.1 and 3.2.

Firstly, we state a convergence guarantee for a gradient descent algorithm according to the Poisson log-likelihood loss.

Theorem 5. *Let $\varepsilon_P > 0$ and $\mathcal{L}_{P,\varepsilon_P}(z) = \sum_{i=1}^m |\langle a_i, z \rangle|^2 - y_i \log(|\langle a_i, z \rangle|^2 + \varepsilon_P)$. Let M be the matrix with rows $\sqrt{1 + \frac{y_i}{8\varepsilon_P}} \cdot a_i^*, i = 1, \dots, m$, and choose $\mu \leq \|M\|^{-2}$. Then, the sequence $(z_k)_{k \geq 0} \subset \mathbb{C}^n$ defined by*

$$z_{k+1} = z_k - \mu \cdot \sum_{i=1}^m \left(1 - \frac{y_i}{|\langle a_i, z_k \rangle|^2 + \varepsilon_P} \right) \langle a_i, z_k \rangle a_i$$

converges to a stationary point of $\mathcal{L}_{P,\varepsilon_P}$.

Proof. We shall apply proposition 1. Since $\mathcal{L}_{P,\varepsilon_P}(z) = \sum_{i=1}^m \ell_i \circ \varphi(z)$ with $\ell_i : [0, \infty) \rightarrow \mathbb{R}, t \mapsto t - y_i \log(t + \varepsilon_P)$ and $\varphi(z) = |\langle a_i, z \rangle|^2$, we compute

$$\begin{aligned} 2 \cdot \ell_i''(t) \cdot t + \ell_i'(t) &= 2 \cdot \frac{y_i}{(t + \varepsilon_P)^2} \cdot t + 1 - \frac{y_i}{t + \varepsilon_P} \\ &= 1 + \frac{y_i(t - \varepsilon_P)}{(t + \varepsilon_P)^2} \leq 1 + \frac{y_i}{8\varepsilon_P}, \end{aligned}$$

where we used that $\frac{s - \varepsilon_P}{(s + \varepsilon_P)^2} \leq \frac{1}{8\varepsilon_P}$ for every $s \in \mathbb{R}$ and any fixed $\varepsilon_P > 0$. With lemma 2 and this bound, we obtain

$$\begin{aligned} \left(\begin{array}{c} v \\ \bar{v} \end{array} \right)^* \nabla^2 \mathcal{L}_{P,\varepsilon_P}(z) \left(\begin{array}{c} v \\ \bar{v} \end{array} \right) &\leq 2 \cdot \sum_{i=1}^m \left| \left\langle \sqrt{1 + \frac{y_i}{8\varepsilon_P}} \cdot a_i, v \right\rangle \right|^2 \\ &\leq 2 \|M\|^2 \|v\|_2^2 = \|M\|^2 \left\| \left(\begin{array}{c} v \\ \bar{v} \end{array} \right) \right\|_2^2 \end{aligned}$$

for all $z, v \in \mathbb{C}^n$. It remains to show that $\mathcal{L}_{P,\varepsilon_P}$ has compact sublevel sets on the subspace $z_0 + \text{range}(A^*)$, which contains all possibly attainable iterates of the algorithm. This follows from lemma 3 and remark 4. Clearly, all functions ℓ_i are continuous and bounded from below. Further, we can find $k_1 > 0$ and $k_2 \in \mathbb{R}$ with $\ell_i(t) \geq k_1 \cdot t + k_2$ for all $t \in [0, \infty)$, hence $\ell_i(t) \rightarrow \infty$ for $t \rightarrow \infty$. Then, using proposition 1, we conclude that the considered gradient descent algorithm converges to a stationary point of the loss $\mathcal{L}_{P,\varepsilon_P}$. \square

Our main conclusion of theorem 5 is that the step size which guarantees convergence must be of the order of the regularization parameter ε_P . If we are interested in affecting the problem with only a small regularization parameter, this involves choosing a small step size, but using a constant small step size in all iterations results in slow convergence.

In [13], a convergence analysis for a similar algorithm is presented. In contrast to our theory on convergence to stationary points, the authors of [13] state a guarantee for convergence to the ground-truth solution. However, while this guarantee only holds in case of a good initialization and is only applicable for measurement vectors a_i being Gaussian random vectors, our result holds independently of the quality of the initialization and for arbitrary measurement systems and, hence, is more relevant for applications.

We proceed with a convergence analysis for the gradient descent algorithms based on the class of loss functions considered in section 3.2.

Theorem 6. *Let*

$$\mathcal{L}_{\text{avg}}(z) = \sum_{i=1}^m \frac{1}{2} \left(\sqrt{|\langle a_i, z \rangle|^2 + c_1} + \sqrt{|\langle a_i, z \rangle|^2 + c_2} - C_i \right)^2, \tag{17}$$

with constants $c_2 \geq c_1 > 0, C_i \geq 0, i = 1, \dots, m$. For arbitrary $z_0 \in \mathbb{C}^n$ define a sequence $(z_k)_{k \geq 1}$ by

$$z_{k+1} = z_k - \mu \cdot \frac{1}{2} \sum_{i=1}^m \left(\sqrt{|\langle a_i, z_k \rangle|^2 + c_1} + \sqrt{|\langle a_i, z_k \rangle|^2 + c_2} - C_i \right) \cdot \left(\frac{1}{\sqrt{|\langle a_i, z_k \rangle|^2 + c_1}} + \frac{1}{\sqrt{|\langle a_i, z_k \rangle|^2 + c_2}} \right) \langle a_i, z_k \rangle a_i.$$

Then, $(z_k)_{k \geq 0} \subset \mathbb{C}^n$ converges to a stationary point of the loss function \mathcal{L}_{avg} provided $\mu \leq 2 \left((3 + \sqrt{c_2/c_1}) \|A\|^2 \right)^{-1}$, where A denotes the matrix with rows $a_i^*, i = 1, \dots, m$.

Proof. We proceed as in the proof of theorem 5. A bound for the Hessian of \mathcal{L}_{avg} is found by using lemma 2 with $\ell_i : [0, \infty) \rightarrow [0, \infty), t \mapsto \frac{1}{2} (\sqrt{t+c_1} + \sqrt{t+c_2} - C_i)^2$. We have

$$\begin{aligned} \ell'(t) &= \frac{1}{2} (\sqrt{t+c_1} + \sqrt{t+c_2} - C_i) \cdot \left(\frac{1}{\sqrt{t+c_1}} + \frac{1}{\sqrt{t+c_2}} \right) \\ &= \frac{1}{2} \left[2 + \frac{\sqrt{t+c_1}}{\sqrt{t+c_2}} + \frac{\sqrt{t+c_2}}{\sqrt{t+c_1}} - C_i \cdot \left(\frac{1}{\sqrt{t+c_1}} + \frac{1}{\sqrt{t+c_2}} \right) \right], \end{aligned}$$

which leads to

$$\begin{aligned} \ell''(t) &= \frac{1}{4} \left(\frac{1}{\sqrt{t+c_1}} + \frac{1}{\sqrt{t+c_2}} \right)^2 - \frac{1}{4} (\sqrt{t+c_1} + \sqrt{t+c_2} - C_i) \cdot \left(\frac{1}{(t+c_1)^{\frac{3}{2}}} + \frac{1}{(t+c_2)^{\frac{3}{2}}} \right) \\ &= \frac{1}{4} \left[\frac{2}{\sqrt{t+c_1}\sqrt{t+c_2}} - \frac{\sqrt{t+c_1}}{(t+c_2)^{\frac{3}{2}}} - \frac{\sqrt{t+c_2}}{(t+c_1)^{\frac{3}{2}}} \right. \\ &\quad \left. + C_i \cdot \left(\frac{1}{(t+c_1)^{\frac{3}{2}}} + \frac{1}{(t+c_2)^{\frac{3}{2}}} \right) \right] \leq \frac{1}{4} C_i \cdot \left(\frac{1}{(t+c_1)^{\frac{3}{2}}} + \frac{1}{(t+c_2)^{\frac{3}{2}}} \right). \end{aligned}$$

For the bound of the second derivative we have used that $\frac{2}{\alpha\beta} \leq \frac{\alpha}{\beta^3} + \frac{\beta}{\alpha^3}$ for $\alpha, \beta \in \mathbb{R}$. Thus, we arrive at

$$\begin{aligned} 2t\ell''(t) + \ell'(t) &\leq \frac{1}{2}C_i \cdot \left(\frac{t}{(t+c_1)^{\frac{3}{2}}} + \frac{t}{(t+c_2)^{\frac{3}{2}}} \right) \\ &\quad + \frac{1}{2} \left[2 + \frac{\sqrt{t+c_1}}{\sqrt{t+c_2}} + \frac{\sqrt{t+c_2}}{\sqrt{t+c_1}} - C_i \cdot \left(\frac{1}{\sqrt{t+c_1}} + \frac{1}{\sqrt{t+c_2}} \right) \right] \\ &= \frac{1}{2} \left[-C_i \cdot \left(\frac{c_1}{\sqrt{t+c_1}} + \frac{c_2}{\sqrt{t+c_2}} \right) + 2 + 1 + \sqrt{\frac{c_2}{c_1}} \right] \\ &\leq \frac{1}{2} \left(3 + \sqrt{\frac{c_2}{c_1}} \right), \end{aligned}$$

and, consequently,

$$\begin{aligned} \begin{pmatrix} v \\ \bar{v} \end{pmatrix}^* \nabla^2 \mathcal{L}_{\text{avg}}(z) \begin{pmatrix} v \\ \bar{v} \end{pmatrix} &\leq \sum_{i=1}^m \left(3 + \sqrt{\frac{c_2}{c_1}} \right) |\langle a_i, v \rangle|^2 \\ &\leq \frac{1}{2} \left(3 + \sqrt{\frac{c_2}{c_1}} \right) \|A\|^2 \left\| \begin{pmatrix} v \\ \bar{v} \end{pmatrix} \right\|_2^2 \end{aligned}$$

for all $z, v \in \mathbb{C}^n$. This suggests choosing $\mu = 2 \left(\left(3 + \sqrt{c_2/c_1} \right) \|A\|^2 \right)^{-1}$.

The function \mathcal{L}_{avg} restricted to $z_0 + \text{range}(A^*)$ has compact level sets by lemma 3 and remark 4, as all functions ℓ_i are continuous, bounded from below, and satisfy $\ell_i(t) \rightarrow \infty$ for $t \rightarrow \infty$. By proposition 1, the gradient descent algorithm converges to a stationary point of the loss function \mathcal{L}_{avg} . □

By this result, we also obtain a rule for the step size when using a gradient descent algorithm for a loss (13) or (14).

Corollary 7. *Convergence to a stationary point of a loss function (14) for some regularization parameter $\varepsilon_A > 0$ can be achieved by the corresponding gradient descent algorithm using a step size $\mu \leq \left(2 \|A\|^2 \right)^{-1}$.*

This is the same result as theorem A.1 in [20], which means that we generalized the convergence guarantee of [20] for the amplitude loss to the class of loss functions (17).

The important difference of this result to the convergence guarantee for the Poisson log-likelihood loss, as stated in theorem 5, is that the bound on the step size necessary for convergence is independent of a potentially small regularization parameter.

Corollary 8. *A loss function using the Tukey–Freeman transform (11) for variance stabilization would be of the form*

$$\mathcal{L}(z) = \sum_{i=1}^m \frac{1}{2} \left(\sqrt{|\langle a_i, z \rangle|^2 + \varepsilon} + \sqrt{|\langle a_i, z \rangle|^2 + 1} - \sqrt{y_i} - \sqrt{y_i + 1} \right)^2,$$

with $\varepsilon > 0$ to avoid the singularity in the derivative. By theorem 6, we conclude that a gradient descent algorithm using this loss function requires a step size $\mu \leq 2 \left(\left(3 + \sqrt{1/\varepsilon} \right) \|A\|^2 \right)^{-1}$ for guaranteed convergence.




Figure 2. Distribution of measurements for a low-dose and a higher-dose experiment.

This means that for this type of variance-stabilizing transform we face the same problem as for the Poisson log-likelihood loss.

Also for the loss \mathcal{L}_0 finally proposed by us we can use the same fixed step size and can guarantee convergence to a stationary point of the loss function.

Corollary 9. *Convergence to a stationary point of loss \mathcal{L}_0 defined in (16) can be achieved by the corresponding gradient descent algorithm with a step size $\mu \leq 2 \left((3 + \sqrt{c_2/c_1}) \|A\|^2 \right)^{-1}$.*

5. Numerical experiments

We corroborate our theoretical consideration with numerical experiments on the reconstruction of an object from simulated low-dose Poissonian phase retrieval measurements. Our simulations are based on a test object $x \in \mathbb{C}^n$ and complex Gaussian measurement vectors $a_i \in \mathbb{C}^n$, $i = 1, \dots, m$. In the initial subsections, we work with $n = 256$ and $m = 10n$, whereas in section 5.4, we give a more detailed analysis for different object dimensions n and oversampling ratios $\frac{m}{n}$.

We normalize the measurements such that $\sum_{i=1}^m |\langle a_i, x \rangle|^2 = 1$ and understand a noiseless value $|\langle a_i, x \rangle|^2$ as the probability that photons arrive at the i th detector pixel. We choose a dose $d \in \mathbb{N}$ and work with measurements

$$y_i \sim \text{Poisson} \left(d \cdot |\langle a_i, x \rangle|^2 \right), \quad i = 1, \dots, m.$$

Unless reported otherwise, we experiment with doses $d \in \{500, 1000, 1500, \dots, 3000\}$. The histograms in figure 2 show exemplary realizations of measurement distributions for $d = 500$ and $d = 3000$. The lowest dose corresponds to a signal-to-noise ratio

$$SNR := \left\| (d \cdot |\langle a_i, x \rangle|^2)_{i=1}^m \right\|_2 / \left\| (y_i - d \cdot |\langle a_i, x \rangle|^2)_{i=1}^m \right\|_2$$

of approximately 0.64, the highest dose corresponds to $SNR \approx 1.44$.

All considered algorithms are initialized by the result of a few iterations of the power method proposed in [10]. Further, the algorithms are stopped at the latest after 1000 iterations or as soon as a stopping criterion is met, which requires the respective loss function to drop by less than 10^{-6} within one iteration. Our trials showed that this enabled the algorithms to saturate at a local minimum in the majority of cases.




Figure 3. Performance of the Poisson flow with regularization parameter $\varepsilon_P \in \{10^{-3}, 0.1, 0.25, 0.5, 1\}$. For each dose, the experiments were repeated 20 times and averages of the reconstruction results are shown.

The performance of the various investigated algorithms is compared in terms of the relative reconstruction error $\min_{\theta \in (0, 2\pi]} \|x - e^{i\theta} x_T\|_2 / \|x\|_2$, where x is the ground-truth and x_T is the reconstruction result obtained after T iterations of the respective algorithm.

5.1. Gradient descent with Poisson log-likelihood

First, we compare the performance of the gradient descent algorithm using the Poisson log-likelihood loss (8) for different regularization parameters ε_P . The respective algorithm is named ‘Poisson flow’ (PF).

Figure 3 depicts the performance of the discussed algorithms for this experiment in terms of the relative reconstruction error. This experiment shows that the performance of the algorithm clearly depends on the choice of the regularization parameter ε_P . While in a higher-dose experiment the algorithm is not very sensible to the parameter selection, a low-dose experiment requires knowledge about a good choice for ε_P .

We would like to comment here that, alternatively to (8), we could also work with the loss

$$\mathcal{L}(z) = \sum_{i=1}^m |\langle a_i, z \rangle|^2 - (y_i + \varepsilon) \log \left(|\langle a_i, z \rangle|^2 + \varepsilon \right).$$

The corresponding method can be understood as unbiased since this loss is minimized by z satisfying $|\langle a_i, z \rangle|^2 = y_i$, $i = 1, \dots, m$. Hence, this method would be less sensible to the selection of ε . However, we found in our numerical analysis that the corresponding algorithm performs worse than the algorithm using loss (8) with a good parameter ε_P . Moreover, the problem that the recommended step size is parameter dependent remains.

5.2. Gradient descent with Gaussian log-likelihood after improved variance stabilization

Further, we compare the Poisson flow with the gradient descent algorithm using the Gaussian log-likelihood loss with the proposed improved variance stabilization transform. The algorithm we label here as ‘Flow with improved variance stabilization’ (FIVS) works with the loss function (16) with $c_1 = 0.12$, $c_2 = 0.27$, and $C_i = \sqrt{y_i + c_1} + \sqrt{y_i + c_2}$, $i = 1, \dots, m$, and algorithm




Figure 4. Performance of the gradient descent algorithms for the proposed loss functions based on the improved variance stabilization approach. The gray lines correspond to the Poisson flow as in figure 3. The plot shows averages over 20 trials.

‘Flow with improved variance stabilization without adaption for zero measurements’ (FIVS-WA) uses loss (17) with the same parameters.

The plot of the relative reconstruction errors for different doses in figure 4 indicates that the FIVS algorithm, using the suggested loss function with the improved variance-stabilizing transform including the Poisson log-likelihood adaption to zero measurements, performs in all instances comparably to the Poisson flow with the respective close to optimal choice of regularization parameter. If one is aware of an optimal parameter selection, it is reasonable to apply the Poisson flow. However, if no decision rule is available, FIVS is an advisable alternative.

Note that the proposed adaption (16) of the loss for zero measurements is essential for a low-dose experiment dominated by zero measurements. Consequently, we focus solely on FIVS (with the adaption for zero measurements) in the following.

5.3. Wirtinger flow and truncated Wirtinger flow

Next, we compare the performance of FIVS with two approaches suggested in the literature that do not use a variance stabilization and do not require a choice of a regularization parameter as in the Poisson flow.

We denote with ‘Wirtinger flow’ (WF), as common in the literature, the algorithm using the loss (9) with constant variance $\sigma_i^2 = 1/4$. This is the first variant of a Wirtinger gradient descent algorithm and was proposed in [10]. Further, we study the ‘truncated Wirtinger flow’ (TWF) as described in [13], based on a truncation principle for the Poisson maximum log-likelihood loss. For this method, we did not tune the required parameters but used the numbers suggested by the authors.

From the results shown in figure 5 we conclude that neither of these two methods can keep up with the proposed method (FIVS) in a low-dose setting.

5.4. Amplitude flow

Interpreting the reconstruction results reported up to here, we have to note that we are in the low-dose regime and cannot expect very small reconstruction errors without much redundancy




Figure 5. Performance of the gradient descent algorithm for the proposed loss function based on the improved variance stabilization approach (FIVS) versus the Wirtinger flow algorithm [10] and the truncated Wirtinger flow [13]. The plot shows averages over 20 trials.

in the data, i.e. a large amount of measurements. Furthermore, we have to mind that for drawing more measurements, the dose per measurement needs to be reduced accordingly in a low-dose experiment because more measurements cause more damage. To explore the trade-off between the oversampling ratio (m/n) and the illumination dose used per object pixel (d/n) in more detail, we extend the analysis to a large range of oversampling-dose-pairings.

In these experiments, we compare the proposed algorithm, FIVS, with the gradient descent algorithm for the loss (14) for various choices of regularization parameters ε_A . We refer to this algorithm as ‘amplitude flow’ (AF). While our proposed method is a variant of amplitude flow with a parameter choice based on the heuristics of the improved variance stabilization, we aim to decide whether this parameter rule is indeed reasonable in a range of applications. The experiments are summarized in the plots shown in figure 6.

It can be recognized that for a very low-dose experiment or in case of a small oversampling ratio the AF with a different parameter choice than in FIVS can be favorable. The experiments suggest that in extreme cases with a small oversampling ratio the parameter ε_A should be chosen rather large. However, no definite rule on the choice of parameter can be deduced from this empirical study. This argues in favor of the proposed FIVS method, i.e. the algorithm using the loss function built on the heuristics motivated by the improved variance stabilization transform.

5.5. Discussion of the step size selection

For the Wirtinger flow algorithm using the Gaussian log-likelihood loss we do not find a constant step size that guarantees a decrease of the loss function in each iteration as for the other loss functions discussed in this paper. We can choose an iteratively adapted step size, based on an analysis similar to the proof ideas used over all this paper, where it is not possible to eliminate the dependence on the iterates. However, this implementation is computationally expensive and makes the choice of this loss function less attractive. For the experiments, we used the computationally more efficient step size proposed in [45]. For all other algorithms, we use the step sizes proposed by the respective theoretical results, i.e. theorems 5 and 6.

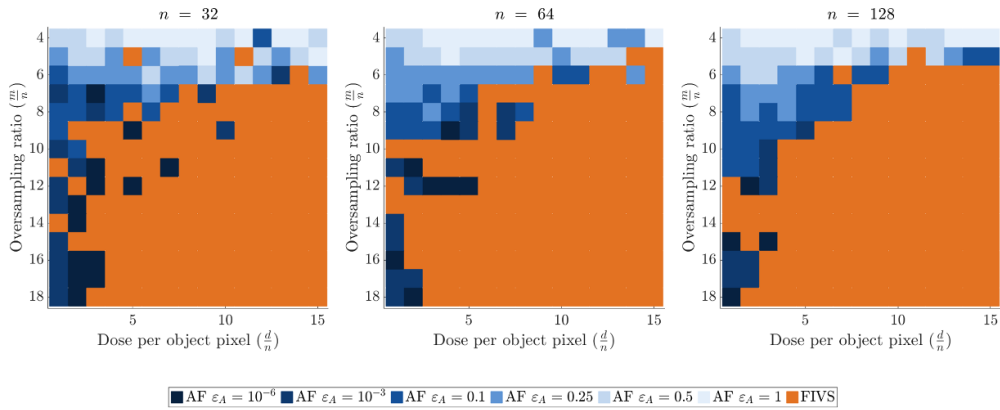


Figure 6. The plot displays the best performing method per oversampling-dose-pairing, represented by the respective assigned color. Performance is measured in terms of relative reconstruction error after the described stopping criterion on the respective loss is satisfied. The results are averages over five trials. For all three experiments for different object dimensions n , the proposed method FIVS showed the best performance in a majority of combinations of oversampling ratios $\frac{m}{n}$ and doses ratios $\frac{d}{n}$. In the cases where AF with a certain parameter choice ε_A obtained better reconstruction results, we observe the tendency that a smaller oversampling ratio requires a larger value for ε_A . Nevertheless, no clear decision rule can be derived for the regularization parameter of the amplitude flow, whereas FIVS, a method based on a heuristic determining the parameter choice, demonstrates good performance in a wide range of cases.

Li et al[34] suggest to use a step size based on an approximation of the Hessian using observed Fisher information for the Poisson log-likelihood loss. This step size is claimed to accelerate the convergence as it is larger than the step size involving the inverse of the factor $1 + y_i/(8\varepsilon_P)$. However, in comparison to the step size derived in theorem 5, the Fisher information based step size solves a line search in each iteration, which results in computational expense. Here, we face a trade-off between having a constant step size, independent of the approximate obtained in each iteration, that guarantees convergence or having a step size rule that can be made independent of ε_P by performing a step size optimization in each iteration. Due to the computational advantage, our interest lies rather in the constant step size. In our numerical experiments we found that for larger regularization parameters ε_P the step size proposed by [34], as contrasted with our step size rule, does not guarantee descent of the loss in each iteration. While for $\varepsilon_P \rightarrow 0$ our step size rule becomes impractical, the iteration dependent step size becomes more useful as it does not decrease on the order of ε_P . On the other hand, our experiments showed that it is not reasonable to work with an extremely small regularization parameter ε_P .

We conclude that using the constant step size in combination with ε_P large enough seems to be reasonable. Other than that, we would like to emphasize again that this trade-off problem can be avoided when using the Gaussian log-likelihood with a good variance stabilization as suggested here.

6. Conclusion

In this paper, we studied gradient descent algorithms for suitable regularizations and approximations of the Poisson log-likelihood problem. Motivated by the nature of low-dose imaging

experiments, we designed a method that allows for improved treatment of zero measurements and other small counts. For this and all other discussed algorithms, we provided a convergence analysis including a step size rule.

In terms of applying such methods to real-world data, a next step is to incorporate a suitable regularization term in the optimization problem, for example of the type discussed in [46] or [47]. The optimal choice for the regularization term depends highly on the type of object under consideration. It is future work to find a reasonable problem formulation for objects such as, for example, viruses that shall be imaged using low-dose illumination. A possible follow-up work on this paper would be to extend the convergence analysis to algorithms involving such regularization attempts.

During the preparation of this manuscript, one of the authors continued the research on low-dose Poisson phase retrieval, especially in the very extreme regime of only 0 and 1 counts. In this extreme setting, the newly developed algorithms presented in [48] outperform the algorithms discussed in this work. However, in a moderately low-dose scenario as studied in the present paper, the optimization approach using the maximum likelihood loss and approximations thereof are our method of choice.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://github.com/patricia-roemer/LowDosePoissonPR>.


Acknowledgments

The authors acknowledge support by the Helmholtz Association under the Contracts Nos. ZT-I-0025 (Ptychography 4.0), ZT-I-PF-4-018 (AsoftXm), ZT-I-PF-5-28 (EDARTI), ZT-I-PF-4-024 (BRLEMMM). Patricia Römer was further financially supported by the German Research Foundation (DFG) Grants KR 4512/1-1 and KR 4512/2-2.

ORCID iDs

Benedikt Diederichs  <https://orcid.org/0009-0002-2301-9583>

Frank Filbir  <https://orcid.org/0000-0002-9700-8505>

Patricia Römer  <https://orcid.org/0009-0005-3362-4957>

References

- [1] Beinert R and Plonka G 2015 Ambiguities in one-dimensional discrete phase retrieval from Fourier magnitudes *J. Fourier Anal. Appl.* **21** 1169–98
- [2] Rodenburg J and Maiden A 2019 Ptychography *Springer Handbook Microscopy* (Springer) pp 819–904
- [3] Pfeiffer F 2018 x-ray ptychography *Nat. Photon.* **12** 9–17
- [4] Thibault P, Dierolf M, Bunk O, Menzel A and Pfeiffer F 2009 Probe retrieval in ptychographic coherent diffractive imaging *Ultramicroscopy* **109** 338–43
- [5] Fienup J R 1982 Phase retrieval algorithms: a comparison *Appl. Opt.* **21** 2758–69
- [6] Balan R, Casazza P and Edidin D 2006 On signal reconstruction without phase *Appl. Comput. Harmon. Anal.* **20** 345–56
- [7] Candès E J, Eldar Y C, Strohmer T and Voroninski V 2015 Phase retrieval via matrix completion *SIAM Rev.* **57** 225–51

- [8] Sun J, Qu Q and Wright J 2018 A geometric analysis of phase retrieval *Found. Comput. Math.* **18** 1131–98
- [9] Conca A, Eddidin D, Hering M and Vinzant C 2015 An algebraic characterization of injectivity in phase retrieval *Appl. Comput. Harmon. Anal.* **38** 346–56
- [10] Candès E J, Li X and Soltanolkotabi M 2015 Phase retrieval via Wirtinger flow: theory and algorithms *IEEE Trans. Inf. Theory* **61** 1985–2007
- [11] Zhang H, Liang Y and Chi Y 2017 A nonconvex approach for phase retrieval: reshaped Wirtinger flow and incremental algorithms *J. Mach. Learn. Res.* **18** 1–35 (available at: <https://www.jmlr.org/papers/volume18/16-572/16-572.pdf>)
- [12] Wang G, Giannakis G B and Eldar Y C 2017 Solving systems of random quadratic equations via truncated amplitude flow *IEEE Trans. Inf. Theory* **64** 773–94
- [13] Chen Y and Candès E J 2017 Solving random quadratic systems of equations is nearly as easy as solving linear systems *Commun. Pure Appl. Math.* **70** 822–83
- [14] Kolte R and Özgür A 2016 Phase retrieval via incremental truncated Wirtinger flow (arXiv:1606.03196)
- [15] Wang G, Giannakis G B and Chen J 2017 Scalable solvers of random quadratic equations via stochastic truncated amplitude flow *IEEE Trans. Signal Process.* **65** 1961–74
- [16] Gao B, Sun X, Wang Y and Xu Z 2020 Perturbed amplitude flow for phase retrieval *IEEE Trans. Signal Process.* **68** 5427–40
- [17] Tan Y S and Vershynin R 2019 Phase retrieval via randomized Kaczmarz: theoretical guarantees *Inf. Inference: A J. IMA* **8** 97–123
- [18] Huang M and Wang Y 2022 Linear convergence of randomized Kaczmarz method for solving complex-valued phaseless equations *SIAM J. Imaging Sci.* **15** 989–1016
- [19] Römer P, Filbir F and Kraher F 2021 On the randomized Kaczmarz algorithm for phase retrieval *2021 55th Asilomar Conf. on Signals, Systems and Computers* (IEEE) pp 847–51
- [20] Xu R, Soltanolkotabi M, Haldar J P, Unglaub W, Zusman J, Levi A F and Leahy R M 2018 Accelerated Wirtinger flow: a fast algorithm for ptychography (arXiv:1806.05546)
- [21] Filbir F and Melnyk O 2023 Image recovery for blind polychromatic ptychography *SIAM J. Imaging Sci.* **16** 1308–37
- [22] Chang H, Enfedaque P, Zhang J, Reinhardt J, Enders B, Yu Y S, Shapiro D, Schroer C G, Zeng T and Marchesini S 2019 Advanced denoising for x-ray ptychography *Opt. Express* **27** 10395–418
- [23] Melnyk O and Römer P 2024 Background removal for Ptychography via Wigner distribution deconvolution *SIAM J. Imaging Sci.* **17** 1978–2014
- [24] Thibault P and Guizar-Sicairos M 2012 Maximum-likelihood refinement for coherent diffractive imaging *New J. Phys.* **14** 063004
- [25] Yeh L H, Dong J, Zhong J, Tian L, Chen M, Tang G, Soltanolkotabi M and Waller L 2015 Experimental robustness of Fourier ptychography phase retrieval algorithms *Opt. Express* **23** 33214–40
- [26] Hohage T and Werner F 2016 Inverse problems with Poisson data: statistical regularization theory, applications and algorithms *Inverse Problems* **32** 093001
- [27] Chang H, Lou Y, Duan Y and Marchesini S 2018 Total variation-based phase retrieval for Poisson noise removal *SIAM J. Imaging Sci.* **11** 24–55
- [28] Fatima G, Li Z, Arora A and PDMM B P 2022 A novel primal-dual majorization-minimization algorithm for poisson phase-retrieval problem *IEEE Trans. Signal Process.* **70** 1241–55
- [29] Glaeser R M 1971 Limitations to significant information in biological electron microscopy as a result of radiation damage *J. Ultrastructure Res.* **36** 466–82
- [30] Wang J, Morin C, Li L, Hitchcock A P, Scholl A and Doran A 2009 Radiation damage in soft x-ray microscopy *J. Electron Spectrosc. Relat. Phenom.* **170** 25–36
- [31] Küçükoglu B et al 2024 Low-dose cryo-electron ptychography of proteins at sub-nanometer resolution *Nat. Commun.* **15** 8062
- [32] Pelz P M, Qiu W X, Bücken R, Kassier G and Miller R D 2017 Low-dose cryo electron ptychography via non-convex Bayesian optimization *Sci. Rep.* **7** 9883
- [33] Godard P, Allain M, Chamard V and Rodenburg J 2012 Noise models for low counting rate coherent diffraction imaging *Opt. Express* **20** 25914–34
- [34] Li Z, Lange K and Fessler J 2022 Poisson phase retrieval in very low-count regimes *IEEE Trans. Comput. Imaging* **8** 838–50
- [35] Leidl M L, Diederichs B, Sachse C and Müller-Caspary K 2024 Influence of loss function and electron dose on ptychography of 2D materials using the Wirtinger flow *Micron* **185** 103688

- [36] Bertero M, Boccacci P, Talenti G, Zanella R and Zanni L 2010 A discrepancy principle for Poisson data *Inverse Problems* **26** 105004
- [37] Bevilacqua F, Lanza A, Pragliola M and Sgallari F 2021 Nearly exact discrepancy principle for low-count Poisson image restoration *J. Imaging* **8** 1
- [38] Tippett L 1935 2-statistical methods in textile research. Part 2-Uses of the binomial and poisson distributions *J. Text. Inst. Trans.* **26** T13–T50
- [39] Bartlett M S 1936 The square root transformation in analysis of variance *Suppl. J. R. Stat. Soc.* **3** 68–78
- [40] Curtiss J H 1943 On transformations used in the analysis of variance *Ann. Math. Stat.* **14** 107–22
- [41] Anscombe F J 1948 The transformation of Poisson, binomial and negative-binomial data *Biometrika* **35** 246–54
- [42] Freeman M F and Tukey J W 1950 Transformations related to the angular and the square root *Ann. Math. Stat.* **21** 607–11
- [43] Konijnenberg A, Coene W and Urbach H 2017 Study of cost functionals for ptychographic phase retrieval to improve the robustness against noise, and a proposal for another noise-robust ptychographic phase retrieval scheme (arXiv:1703.10892)
- [44] Zhang Y, Song P and Dai Q 2017 Fourier ptychographic microscopy using a generalized Anscombe transform approximation of the mixed Poisson-Gaussian likelihood *Opt. Express* **25** 168–79
- [45] Chen Z, Zhou Y, Liang Y and Lu Z 2023 Generalized-smooth nonconvex optimization is as efficient as smooth nonconvex optimization *Proc. 40th Int. Conf. on Machine Learning (Proc. of Machine Learning Research* vol 202 (PMLR) pp 5396–427 (available at: <https://proceedings.mlr.press/v202/chen23ar/chen23ar.pdf>)
- [46] Le T, Chartrand R and Asaki T J 2007 A variational approach to reconstructing images corrupted by Poisson noise *J. Math. Imaging Vis.* **27** 257–63
- [47] Loock S and Plonka G 2014 Phase retrieval for Fresnel measurements using a shearlet sparsity constraint *Inverse Problems* **30** 055005
- [48] Römer P and Kraemer F 2024 A one-bit quantization approach for low-dose poisson phase retrieval 2024 *Int. Workshop on the Theory of Computational Sensing and its Applications to Radar, Multimodal Sensing and Imaging (CoSeRa)* (IEEE) pp 42–46