

Multigrid Optimization for Large-Scale Ptychographic Phase Retrieval

Samy Wu Fung* Zichao (Wendy) Di †

February 20, 2019

Abstract

Ptychography is a popular imaging technique that combines diffractive imaging with scanning microscopy. The technique consists of a coherent beam that is scanned across an object in a series of overlapping positions, leading to reliable and improved reconstructions. Ptychographic microscopes allow for large fields to be imaged at high resolution at the cost of additional computational expense. In this work, we propose a multigrid-based optimization framework to reduce the computational burdens of large-scale ptychographic phase retrieval. Our proposed method exploits the inherent hierarchical structures in ptychography through tailored restriction and prolongation operators for the object and data domains. Our numerical results show that our proposed scheme accelerates the convergence of its underlying solver and outperforms the Ptychographic Iterative Engine (PIE), a workhorse in the optics community.

Keywords: phase retrieval, coherent diffractive imaging, ptychography, multigrid optimization, inverse problems.

1 Introduction

Ptychography is a coherent diffractive imaging (CDI) technique that arises in applications such as materials science [20, 21, 37], biology [31, 44], and x-ray crystallography [10]. The technique was originally proposed to improve the resolution in electron or x-ray microscopy by replacing single-element detectors with two-dimensional array detectors, combining diffractive imaging with scanning microscopy [18]. More precisely, a coherent beam is scanned across an object in a series of overlapping positions that couples information between successively collected diffraction patterns (see Fig. 1). As is well known

*Department of Mathematics, Emory University, Atlanta, GA. samy.wu@emory.edu,
<https://sites.google.com/site/samywufung/>

†Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL.
wendydi@anl.gov,
<http://www.mcs.anl.gov/person/zichao-wendy-di>

in single-pattern CDI, finite support constraints are crucial for the convergence of classical algorithms such as error reduction [14], hybrid input output (HIO) [13], gradient-based algorithms [7], relaxed averaged alternating reflections [25], and saddle point optimization [30, 47]. In ptychography, *a priori* knowledge of the scanning positions automatically delivers these constraints, leading to faster and more robust reconstructions than single-pattern CDI [5, 11, 22, 28].

The widespread application of ptychography has led to considerable research on methods for its numerical reconstruction. Among the most popular methods is the ptychographic iterative engine (PIE) [28, 40, 41]. The algorithm consists of alternating projections onto non-convex modulus constraint sets and is popular in the optics community [10, 15, 29, 31, 44]. Mathematically, PIE is equivalent to a projected steepest descent algorithm applied to a particular error metric (see Sec. 2). Other well-known approaches include standard gradient-based techniques [38, 52], and Wirtinger flow [7, 51], which uses careful initialization via a spectral method. Recently, a lifting approach (PhaseLift) was introduced, where the phase retrieval problem is reformulated as a convex optimization problem at the expense of solving for a quadratically increased number of unknown variables [6]; unfortunately, this quickly becomes intractable for large-scale problems. Despite serious efforts, large-scale ptychographic phase retrieval continues to be challenging as high resolution demands and small scanning beams generate large volumes of data.

In this work, we consider the multigrid-based optimization framework (MG/OPT) presented in [33] for solving the ptychographic phase retrieval. The MG/OPT scheme is a general-purpose framework designed to accelerate the large-scale nonlinear optimization problems. MG/OPT is inspired by the full approximation scheme (FAS) [3, 48], and has stronger convergence guarantees than traditional nonlinear multigrid methods [34]. Moreover, MG/OPT can be extended to optimization problems with equality and inequality constraints [24]. When applied to unconstrained optimization problems, MG/OPT is equivalent to applying FAS to the first-order optimality conditions. However, MG/OPT ensures a descent direction in the coarse-grid correction under mild assumptions; this leads to a globally convergent algorithm when a linesearch step is performed [34].

Our work is motivated by the successes of MG/OPT on different nonlinear problems [12, 24, 33], as well as a previously applied multilevel scheme [42]. The multilevel scheme proposed in [42] uses a *single slash* cycle on 1D single-pattern CDI phase retrieval problems; specifically, the coarse grid calculation only contributes to a better initial guess for the fine-grid problem. Our approach differs from [42] in that we use multiple V-cycles as well as a linesearch scheme that guarantees a descent direction during the coarse-grid correction. Our work is also motivated by the different hierarchical structures that ptychography exhibits, and which can be exploited by a multilevel scheme as we will show in Sec. 3.3.

The remainder of this paper is organized as follows. In Sec. 2, we review the mathematical formulation and some common algorithms for the ptychographic phase retrieval. In Sec. 3, we give an overview of the MG/OPT framework. In Sec. 4, we illustrate the poten-

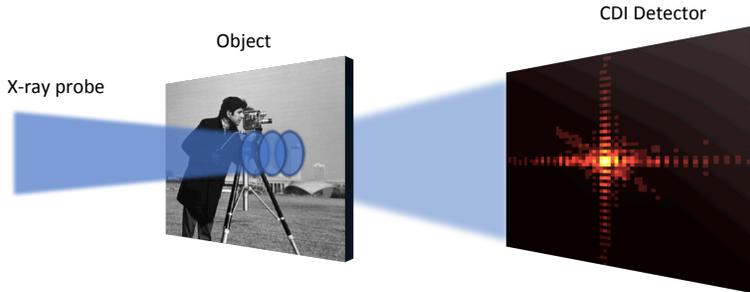


Fig. 1: Schematic illustration of the ptychography experiment with three successive overlapping scans.

tial of MG/OPT for solving ptychographic phase retrieval on a few synthetic experiments, and finally, we conclude with a discussion in Sec. 5.

2 Mathematical Background

In this section, we give an overview of the 2D ptychographic phase retrieval as well as some of the popular methods used for its reconstruction. Since we follow the discretize-then-optimize approach, we limit the discussion to the discrete setting - the continuous formulation of the general phase retrieval problem can be found in [26].

Let $\mathbf{z} = \mathbf{x} + \mathbf{y}i \in \mathbb{C}^{n^2}$ be the object of interest, and $\mathbf{d}_k \in \mathbb{R}^{m^2}$ be the observed data (or intensities) measured from the k^{th} probe, where n^2 and m^2 are the dimensions of the vectorized object and data resolution images, respectively. A ptychography experiment is modeled by

$$\mathbf{d}_k = |\mathcal{F}(\mathbf{Q}_k \mathbf{z})|^2 + \epsilon_k, \quad k = 1, \dots, N, \quad (2.1)$$

where N is the total number of probes (or scanning positions), $\mathcal{F}: \mathbb{C}^{n^2} \mapsto \mathbb{C}^{m^2}$ is the two-dimensional discrete Fourier operator, $\mathbf{Q}_k \in \mathbb{R}^{n^2 \times n^2}$ is the k^{th} probe (a diagonal *illumination* matrix), and $\epsilon_k \in \mathbb{R}^{m^2}$ is the noise corresponding to the k^{th} measurement error. The diagonal elements of \mathbf{Q}_k are nonzero in the columns corresponding to the pixels being illuminated in the object at scanning step k .

There are different ways for formulating the reconstruction problem. One such formulation is the intensity Gaussian error metric [38]

$$\min_{\mathbf{z}} \Phi_{\mathcal{IG}}(\mathbf{z}) = \frac{1}{2} \sum_{k=1}^N \|\ |\mathcal{F}(\mathbf{Q}_k \mathbf{z})|^2 - \mathbf{d}_k \|_2^2, \quad (2.2)$$

which involves minimizing the data misfit. Here, $\Phi_{\mathcal{IG}}: \mathbb{C}^{n^2} \mapsto \mathbb{R}$ is a real-valued cost function defined on the complex domain, and is therefore not complex-differentiable [39].

To overcome the lack of complex-differentiability, it is common to employ the notion of $\mathbb{C}\mathbb{R}$ (Wirtinger) Calculus, where the derivatives of the real and imaginary parts of \mathbf{z} are computed independently [39, 43]. For these real-valued functions, the mere existence of these Wirtinger derivatives is necessary and sufficient for the existence of a stationary point [4, 39, 43]. Using Wirtinger calculus, the partial gradients for (2.2) can be computed as

$$\nabla_{\mathbf{x}}\Phi_{IG} = \sum_{k=1}^N \mathbf{Q}_k^\top \mathcal{F}(\overline{\mathcal{F}\mathbf{Q}_k\mathbf{z}} \odot \mathbf{r}_k), \quad \nabla_{\mathbf{y}}\Phi_{IG} = \sum_{k=1}^N i\mathbf{Q}_k^\top \mathcal{F}(\overline{\mathcal{F}\mathbf{Q}_k\mathbf{z}} \odot \mathbf{r}_k), \quad (2.3)$$

where $\mathbf{r}_k = |\mathcal{F}\mathbf{Q}_k\mathbf{z}|^2 - \mathbf{d}_k$ is the residual of the k^{th} probe, and the resulting Wirtinger derivative is

$$\nabla_{\mathbf{z}}\Phi_{IG} = [\nabla_{\mathbf{x}}\Phi_{IG}^\top \quad \nabla_{\mathbf{y}}\Phi_{IG}^\top]^\top. \quad (2.4)$$

Popular methods that solve this problem include classical gradient-based algorithms [38, 52] and Wirtinger flow [7, 51]. Variants of the intensity Gaussian error metric include the amplitude Gaussian metric [9, 38, 50], intensity Poisson metric [8, 9, 46], and the weighted intensity Gaussian metric [38], all of which measure some variants of the misfit between the forward model and the observed data.

Another formulation of the inverse problem involves solving a feasibility problem. In particular, let the k^{th} measurement constraint set and its corresponding projection operator be denoted by

$$\mathcal{M}_k = \{\mathbf{z} \in \mathbb{C}^{n^2} : |\mathcal{F}(\mathbf{Q}_k\mathbf{z})| = \mathbf{d}_k\}, \quad \text{and} \quad \mathcal{P}_{\mathcal{M}_k}(\mathbf{z}) = \mathcal{F}^{-1} \left[\sqrt{\mathbf{d}_k} \odot \exp(i\theta(\mathcal{F}(\mathbf{Q}_k\mathbf{z}))) \right], \quad (2.5)$$

respectively, where \odot is the Hadamard (or element-wise) product, the square root in $\mathcal{P}_{\mathcal{M}_k}$ is applied element-wise, and $\theta: \mathbb{C} \rightarrow [-\pi, \pi)$ is the element-wise principal argument function [39]. The feasibility problem is then formulated as

$$\text{find } \mathbf{z} \in \mathbb{C}^{n^2} \quad \text{such that} \quad \mathbf{z} \in \bigcap_{k=1}^N \mathcal{M}_k. \quad (2.6)$$

Perhaps the most well-known numerical technique for solving (2.6) is the alternating projection algorithm, PIE, and its variants [23, 28]. PIE consists of performing projections onto the sets \mathcal{M}_k that are calculated one at a time in a sequential manner, where $\gamma \in \mathbb{R}$ is a relaxation scalar that helps suppress the noise effect in the data (see Alg 1). These methods have enjoyed great success in the optics community [10, 15, 23, 31, 40]; however, convergence of these methods is not always guaranteed because the sets \mathcal{M}_k are non-convex [26].

Indeed, problem (2.6) can only be solved in the ideal, noise-free case. As soon as there is noise, we are left to minimize some distance metric based on (2.6). In fact, when $\gamma = 0$

Algorithm 1 Ptychographic Iterative Engine (PIE)

- initialize $N =$ number of probes
 - for $j = 1, 2, \dots$
 - for $k = 1, 2, \dots, N$
$$\mathbf{z}^{(j)} = \mathbf{z}^{(j)} + \left(\frac{|\mathbf{Q}_k|}{\max(\mathbf{Q}_k)} \frac{\mathbf{Q}_k^T}{|\mathbf{Q}_k|^2 + \gamma} \right) (\mathcal{P}_{\mathcal{M}_k}(\mathbf{z}^{(j)}) - \mathbf{Q}_k \mathbf{z}^{(j)})$$
 - end
 - check convergence criteria
-

and the probes are binary in Alg. 1, each update in PIE corresponds to a projected steepest descent iteration that solves

$$\min_{\mathbf{z}} \Phi_{\mathcal{M}}(\mathbf{z}) = \frac{1}{2} \sum_{k=1}^N \|\mathcal{P}_{\mathcal{M}_k}(\mathbf{z}) - \mathbf{Q}_k \mathbf{z}\|_2^2, \quad (2.7)$$

where $\Phi_{\mathcal{M}}: \mathbb{C}^{n^2} \mapsto \mathbb{R}$ describes the distance from the current point to the set $\bigcap_{k=1}^N \mathcal{M}_k$. Its corresponding Wirtinger derivative can be shown to be

$$\nabla_{\mathbf{z}} \Phi_{\mathcal{M}}(\mathbf{z}) = \sum_{k=1}^N \nabla_{\mathbf{z}} \Phi_{\mathcal{M}_k}(\mathbf{z}) = \sum_{k=1}^N \mathbf{Q}_k (\mathcal{P}_{\mathcal{M}_k}(\mathbf{z}) - \mathbf{Q}_k \mathbf{z}) \quad (2.8)$$

[2, 26, 27]. In particular, each inner iteration of PIE in Alg. 1 is equivalent to

$$\mathbf{z}^{(j)} = \mathbf{z}^{(j)} - \nabla_{\mathbf{z}} \Phi_{\mathcal{M}_k}(\mathbf{z}^{(j)}). \quad (2.9)$$

To the best of our knowledge, gradient-based algorithms for ptychography have been mostly used to solve the intensity Gaussian metric (2.2) [7, 38, 51, 52]. However, it has been observed that optimizing over $\Phi_{\mathcal{IG}}$ (as shown in Sec. 4.2) is prone to getting stuck in poor-quality local minima, i.e., artifacts in the reconstructions [38, 53]. This has resulted in different variants of $\Phi_{\mathcal{IG}}$ [38, 52], as well as careful spectral initialization using Wirtinger Flow [7] to be considered in the general phase retrieval community. In this work, we instead focus on the distance metric $\Phi_{\mathcal{M}}$ as the objective function for its stability.

Despite the aforementioned efforts, the ptychographic phase retrieval continues to be computationally demanding, especially when extended to 3D applications [16, 18], where the image sizes grow dramatically. The need to solve these large-scale problems quickly and accurately as a result of the increasing capacities of ptychographic microscopes thus motivates our pursuit of multigrid techniques for the efficient computations of their solutions.

3 Multigrid Optimization

MG/OPT is a multigrid-based optimization framework designed to solve large-scale non-linear optimization problems [33]. Its goal is to accelerate the convergence of traditional iterative algorithms by exploiting the hierarchy of the original optimization problems. The scheme has shown success in a broad class of problems such as PDE-constrained optimal control problems [24] as well as in the generation of centroidal Voronoi tessellations [12]. In this section, we review the MG/OPT framework by introducing some notation, followed by a description of an MG/OPT cycle, and finally with a discussion on its extension to ptychography.

3.1 Notation

We employ the standard notation used in the multigrid community, where the subscript h denotes the fine grid, and H denotes the coarse grid. In order to transfer information between grid levels, we denote the restriction operator by $\mathbf{I}_h^H \in \mathbb{R}^{n_H^2 \times n_h^2}$, which projects the variable and the gradient from the fine grid to the coarse grid, and the prolongation operator by $\mathbf{I}_H^h \in \mathbb{R}^{n_h^2 \times n_H^2}$, which interpolates the search direction from the coarse grid to the fine grid. We also denote the underlying optimization algorithm by "OPT". It is assumed to be convergent in the sense that, under appropriate assumptions on the objective function Φ ,

$$\lim_{j \rightarrow \infty} \left| \nabla \Phi(\mathbf{z}^{(j)}) \right| = 0, \quad (3.10)$$

where $\{\mathbf{z}^{(j)}\}$ are the iterates computed by OPT.

Moreover, we define OPT as a function of the form $\mathbf{z}^+ \leftarrow \text{OPT}(\Phi, \mathbf{v}, \bar{\mathbf{z}}, K)$ which applies K iterations of OPT to the problem

$$\min_{\mathbf{z}_H} \Phi(\mathbf{z}_H) - \mathbf{v}^T \mathbf{z}_H \quad (3.11)$$

with initial guess $\bar{\mathbf{z}}$ to obtain \mathbf{z}^+ . If OPT is required to proceed until convergence, we rewrite this as $\mathbf{z}^+ \leftarrow \text{OPT}(\Phi, v, \bar{\mathbf{z}})$.

3.2 MG/OPT Cycle

Given an initial estimate of the solution $\mathbf{z}_h^{(0)}$ on the fine grid, set $\mathbf{v}_h = 0$. Select non-negative integers k_1 and k_2 satisfying $k_1 + k_2 > 0$. Then for $j = 0, 1, \dots$, set

$$\mathbf{z}_h^{(j+1)} \leftarrow \text{MG/OPT}(\Phi_h, \mathbf{v}_h, \mathbf{z}_h^{(j)}), \quad (3.12)$$

where the function MG/OPT is defined as follows.

- *Coarse-grid solve:* If on the coarsest grid, then solve the optimization problem:

$$\mathbf{z}_h^{(j+1)} \leftarrow \text{OPT}(\Phi_h, \mathbf{v}_h, \mathbf{z}_h^{(j)}).$$

Otherwise,

- *Pre-smoothing:*

$$\bar{\mathbf{z}}_h \leftarrow \text{OPT}(\Phi_h, \mathbf{v}_h, \mathbf{z}_h^{(j)}, k_1)$$

- *Coarse-grid correction:*

- Compute

$$\begin{aligned} \bar{\mathbf{z}}_H &= \mathbf{I}_h^H \bar{\mathbf{z}}_h \\ \bar{\mathbf{v}} &= \mathbf{I}_h^H \mathbf{v}_h + \nabla \Phi_H(\bar{\mathbf{z}}_H) - \mathbf{I}_h^H \nabla \Phi_h(\bar{\mathbf{z}}_h) \end{aligned}$$

- Apply MG/OPT recursively to the surrogate model:

$$\mathbf{z}_H^+ \leftarrow \text{MG/OPT}(\Phi_H, \bar{\mathbf{v}}, \bar{\mathbf{z}}_H)$$

- Compute the search directions

$$\mathbf{e}_H = \mathbf{z}_H^+ - \bar{\mathbf{z}}_H$$

and

$$\mathbf{e}_h = \mathbf{I}_H^h \mathbf{e}_H.$$

- Use a linesearch to determine

$$\mathbf{z}_h^+ = \bar{\mathbf{z}}_h + \alpha \mathbf{e}_h$$

satisfying $\Phi_h(\mathbf{z}_h^+) \leq \Phi_h(\bar{\mathbf{z}}_h)$.

- *Post-smoothing:*

$$\mathbf{z}_h^{(j+1)} \leftarrow \text{OPT}(\Phi_h, \mathbf{v}_h, \mathbf{z}_h^+, k_2)$$

3.3 Extension to Ptychography

We now describe MG/OPT in the context of ptychography for the distance error metric (2.7). For brevity, we consider a two-grid hierarchy, however, the discussion presented below can naturally be extended to multiple levels of grids as well as the 3D case [16].

In ptychography, the fine-grid problem in MG/OPT is given by

$$\min_{\mathbf{z}_h} \Phi_{\mathcal{M},h}(\mathbf{z}_h) = \frac{1}{2} \sum_{k=1}^N \left\| \mathcal{P}_{\mathcal{M}_{k,h}}(\mathbf{z}_h) - \mathbf{Q}_{k,h} \mathbf{z}_h \right\|_2^2, \quad (3.13)$$

where

$$\mathcal{P}_{\mathcal{M}_{k,h}}(\mathbf{z}_h) = \mathcal{F}^{-1} \left(\sqrt{\mathbf{d}_{k,h}} \odot \exp \left(i \theta \left(\mathcal{F} \left(\mathbf{Q}_{k,h} \mathbf{z}_h \right) \right) \right) \right), \quad (3.14)$$

and the additional subscript h represents a certain fine grid level. The coarse-grid surrogate problem can then be written as

$$\begin{aligned} \min_{\mathbf{z}_H} \quad & \Phi_{\mathcal{M},H}(\mathbf{z}_H) - \mathbf{v}^T \mathbf{z}_H \\ = & \frac{1}{2} \sum_{k=1}^N \left\| \mathcal{P}_{\mathcal{M}_{k,H}}(\mathbf{z}_H) - \mathbf{Q}_{k,H} \mathbf{z}_H \right\|_2^2 - \left(\nabla_{\mathbf{z}_H} \Phi_{\mathcal{M},H}(\mathbf{I}_h^H \bar{\mathbf{z}}_h) - \mathbf{I}_h^H \nabla_{\mathbf{z}_h} \Phi_{\mathcal{M},h}(\bar{\mathbf{z}}_h) \right)^T \mathbf{z}_H, \end{aligned} \quad (3.15)$$

where

$$\mathcal{P}_{\mathcal{M}_{k,H}}(\mathbf{z}_H) = \mathcal{F}^{-1} \left(\sqrt{\tilde{\mathbf{I}}_h^H \mathbf{d}_{k,h}} \exp \left(i \theta \left(\mathcal{F} \left(\mathbf{I}_h^H \mathbf{Q}_{k,h} \mathbf{z}_H \right) \right) \right) \right). \quad (3.16)$$

Here, we also add a subscript h or H to the grid-independent objective function $\Phi_{\mathcal{M}}$ to distinguish the objectives across grid levels.

Unlike the standard MG/OPT where only the variable and gradients are transversed across levels, we also consider transvering the data via the data-restriction operator $\tilde{\mathbf{I}}_h^H \in \mathbb{R}^{m_H^2 \times m_h^2}$. To maintain the data fidelity across different grids, we employ the *low-pass filtering* concept [1, 17], where we only sample the low-frequency components of the data for the coarse-grid problem. More precisely, the data-restriction operator $\tilde{\mathbf{I}}_h^H \in \mathbb{R}^{m_H^2 \times n_h^2}$ is a binary diagonal matrix

$$[\tilde{\mathbf{I}}_h^H]_{j,j} = \begin{cases} 1 & \text{if } j \text{ corresponds to the chosen low-frequency pixel} \\ 0 & \text{otherwise,} \end{cases} \quad (3.17)$$

which crops out the high-frequency components and keeps the low-frequency components; see Fig. 2a. This choice of restriction is motivated by the fact that in the frequency domain, the high frequency modes correspond to the detailed features of the object (e.g., the sharp edges) and can only be captured by fine resolutions on the object domain. We illustrate this phenomenon in Fig. 2, where Fig. 2a corresponds to the Fourier transform of the original object, Fig. 2b corresponds to applying a direct inverse Fourier transform which recovers the original object, and Fig. 2c corresponds to applying an inverse Fourier transform to the cropped 64×64 part shown in Fig. 2a, which recovers a coarsened version of the object.

For the object \mathbf{z} and the gradient $\nabla_{\mathbf{z}} \Phi_{\mathcal{M}}$, we choose the restriction operator

$$\mathbf{I}_h^H = \frac{1}{4} \begin{pmatrix} 1 & 1 & 0 & \dots & 0 & 1 & 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 1 & 0 & \dots & 0 & 1 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & & \ddots & \ddots & \ddots & & \vdots \\ 0 & \dots & 0 & 1 & 1 & 0 & 0 & \dots & 0 & 1 & 1 \end{pmatrix} \in \mathbb{R}^{n_H^2 \times n_h^2} \quad (3.18)$$

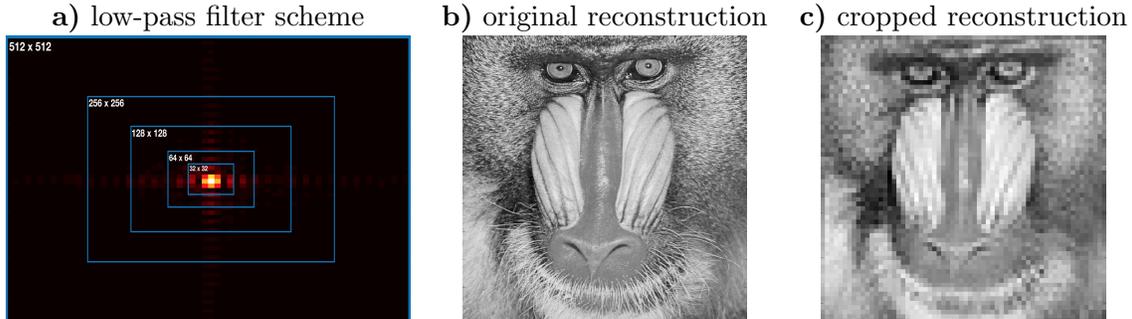


Fig. 2: *a) Data coarsening at different grids. b) Reconstruction of linear problem from original (512×512) diffraction pattern. c) Reconstruction of linear problem from cropped high-frequencies (64×64) diffraction pattern.*

as described in [48, pg. 69], which consists of a four-point average that maps from the vertices to the cell centers. The prolongation operator is chosen in the standard fashion as

$$\mathbf{I}_H^h = c (\mathbf{I}_h^H)^\top, \quad (3.19)$$

for some constant $c > 0$. As a standard choice for the cell-centered prolongation operator, we follow [48, pg. 61] and choose $c = 4$ in our numerical experiments.

3.4 Discussion

MG/OPT is not an algorithm; instead, it is a framework that allows for different choices of discretization, restriction/prolongation operators, and optimization algorithms as its underlying solver. This flexibility is of particular importance in the context of ptychography as it allows us to exploit the natural hierarchy embedded in both the object domain and the data domain. Furthermore, since each MG/OPT cycle contains at least one iteration of the fine-grid optimization algorithm (OPT with $k_1 + k_2 > 0$ in Sec. 3.2), the linesearch in the coarse-grid correction guarantees the convergence of MG/OPT in the same fashion as OPT [33].

4 Numerical Results

In this section, we compare the performance of derivative-based algorithms for the two objective functions (2.2) and (2.7) described in Sec 2. We then demonstrate the potential of the MG/OPT scheme applied to a 2D ptychographic phase retrieval problem.

4.1 Ptychography Setup

For our experiments, we consider a 2D ptychographic phase retrieval modeled by (2.1). We simulate our data using two $n \times n$ images, baboon and cameraman from MATLAB's demo

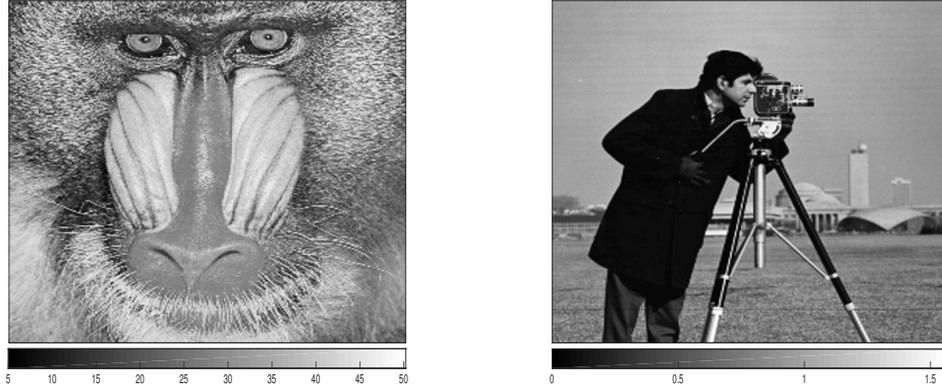


Fig. 3: *Ground truth used to simulate data in numerical experiments. The baboon image is used as the magnitude and the camera man image is used as the phase of the object of interest.*

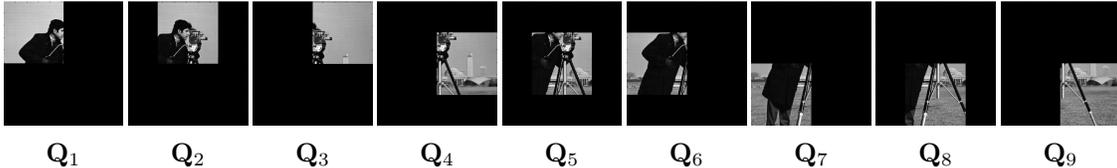


Fig. 4: *Scanning positions (probes) described in Sec. 4.1. Here, each probe illuminates a window of size 256×256 pixels and shifts 128 pixels at a time, leading to 50% overlap between consecutive probes.*

images. We assume $m = n$, which is a practical assumption since the object resolution is typically determined by the data resolution. We set the baboon image to be the magnitude and the cameraman image to be the phase of our ground truth (see Fig. 3). In order to avoid the ambiguity of a global constant phase offset [27] (common in phase retrieval), we set the range of the true phase to be $[0, \frac{\pi}{2}]$. Consequently, we compute the phase errors in our experiments by first mapping the current iterates to this range before comparing with the true phase. For simplicity, we use binary probes, where each probe illuminates $\frac{n}{2} \times \frac{n}{2}$ pixels, that is, $\mathbf{Q}_k \in \mathbb{R}^{n^2 \times n^2}$ is a diagonal matrix that satisfies

$$[\mathbf{Q}_k]_{j,j} = \begin{cases} 1 & \text{if pixel } j \text{ is illuminated} \\ 0 & \text{otherwise.} \end{cases} \quad (4.20)$$

The illumination window is shifted $\frac{n}{4}$ pixels at a time starting from the top-left corner to the right until it reaches the top-right corner. The probe is then shifted downwards and then to the left until it reaches the left edge again. This scanning procedure is continued until we have covered the entire image, and results in a total of 9 probes with 50% overlap between adjacent probes. We illustrate this in Fig. 4 for all the different scanning positions.

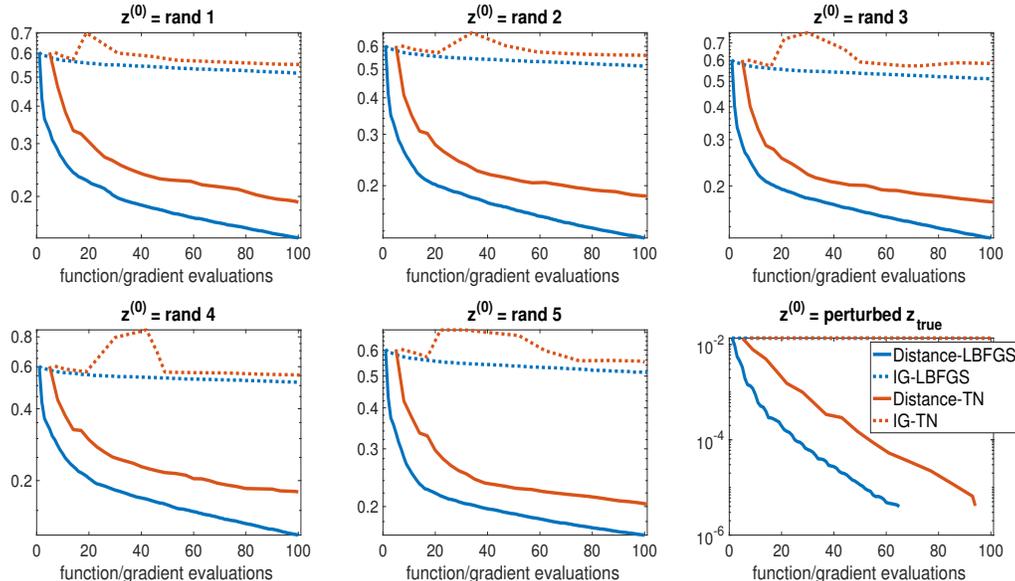


Fig. 5: Relative errors for LBFGS and TN for two error metrics: the distance metric $\Phi_{\mathcal{M}}$ and Intensity Gaussian Φ_{IG} .

We note that the scanning patterns and overlap percentage in our setup are consistent with the recommended experimental setups for ptychography [5, 22].

4.2 Error Metric Comparison

We begin by comparing the performance of derivative-based optimization algorithms for the two error metrics: Φ_{IG} and $\Phi_{\mathcal{M}}$. To this end, we employ the truncated Newton (TN) algorithm described in [32], which uses a preconditioned conjugate gradient to solve the Newton system and a finite difference scheme to approximate the action of the Hessian on a vector. We also employ an LBFGS algorithm that stores the 25 most recent vectors to approximate the inverse of the Hessian [36]. For both solvers, we use a basic linesearch scheme that takes a step once a descent direction is computed, and set a maximum of 50 linesearch iterations. We run the inversions for 6 different initial guesses: 5 random guesses, and a good initial guess that consists of a perturbation of the true solution. In Fig. 5, we show performance of LBFGS and TN applied to Φ_{IG} and $\Phi_{\mathcal{M}}$, in terms the relative reconstruction error

$$\frac{\|\mathbf{z}^{(j)} - \mathbf{z}_{\text{true}}\|}{\|\mathbf{z}_{\text{true}}\|}, \quad (4.21)$$

where \mathbf{z}_{true} is the ground truth. For a fair comparison, the performance is reported based on the number of function/gradient evaluations.

For all initial conditions, we observe the following. First, we obtain much higher accuracies when optimizing over $\Phi_{\mathcal{M}}$, since gradient-based algorithms applied to $\Phi_{\mathcal{I}G}$ tend to get stuck at potentially poor-quality local minima. This observation is consistent with the success of PIE and other alternating projection algorithms, since they optimize the more stable objective function $\Phi_{\mathcal{M}}$ for ptychography. This observation is also seen in other ptychography experiments [38, 53]. Second, in terms of computational work, LBFGS outperforms TN for this particular application because each inner CG iteration in TN requires an additional gradient computation, which leads to an overall higher computational load; we discuss this in Sec 4.3.1. Consequently, in the remainder of our experiments, we consider only the LBFGS as OPT to the distance metric $\Phi_{\mathcal{M}}$.

4.3 Multilevel Ptychography

We illustrate the potential of MG/OPT for the ptychographic phase retrieval problem described in Sec. 4.1.

4.3.1 Computational Costs

We measure the computational cost of different algorithms by counting the number of fine-grid function/gradient evaluations. Each evaluation of $\Phi_{\mathcal{M}}$ requires N projections $\mathcal{P}_{\mathcal{M}_k}$, and the residuals from $\Phi_{\mathcal{M}}$ can be used to compute the gradients at negligible costs (see (2.8)). Moreover, each projection consists of one Fourier transform and one inverse Fourier transform, which have complexity $\mathcal{O}(n \log(n))$. Therefore, one function/gradient evaluation has a total complexity of $\mathcal{O}(Nn \log(n))$. Similarly, each outer PIE iteration also contains N projections $\mathcal{P}_{\mathcal{M}_k}$ and therefore has the same complexity as one function/gradient evaluation of $\Phi_{\mathcal{M}}$.

To account for the computational cost of MG/OPT (with LBFGS as its underlying solver), we determine (for each level) the relative cost of a function/gradient evaluation compared to an evaluation on the finest grid. That is, given our restriction approach, the computational cost of one function/gradient evaluation on the coarser grid is roughly one fourth of the cost of one function/gradient evaluation on the previous finer grid. Following this relationship, the computational work shown in Table 1 amounts to 100 function/gradient evaluations for each MG/Opt scheme.

4.3.2 Inversion Setup

We solve the phase retrieval problem using four approaches: PIE, LBFGS, a 2-level, and 5-level MG/OPT with LBFGS as its underlying solver. In the latter three approaches, we use the same linesearch scheme described in Sec. 4.2 with a maximum of 50 iterations. In LBFGS, we store the 25 most recent vectors used to approximate the inverse of the Hessian at each iteration. We run the inversions for 0%, 5%, and 10% added Gaussian

function/gradient evaluations of MG/OPT across grids						
	noise level	$n = 512$	$n = 256$	$n = 128$	$n = 64$	$n = 32$
2-grid MG/OPT	0%	68	137			
	5%	68	130			
	10%	72	114			
5-grid MG/OPT	0%	88	32	24	6	409
	5%	82	55	46	6	337
	10%	78	66	62	9	486

Table 1: Allocation of computational work for MG/OPT inversions across multiple grids for different noise levels. We show the different number of linesearch iterations performed on the different grids.

n	convergence factor		number of cycles	
	LBFGS	MG/OPT	LBFGS	MG/OPT
64	0.49	0.26	35	9
128	0.74	0.40	33	11
256	0.87	0.63	64	19
512	0.90	0.69	86	22

Table 2: Comparison of convergence factors between LBFGS and MG/OPT. We run LBFGS and MG/OPT until convergence. In MG/OPT, the grid is always projected onto a mesh of size 32×32 .

noise using a random initial guess, and stop after 100 function/gradient evaluations have been performed.

In the MG/OPT scheme, we set a maximum of 3 and 100 iterations at the coarsest grids of the 2-level and 5-level schemes per cycle, respectively. We do this since an iteration on the coarsest grid of a 5-level V-cycle is much cheaper than an iteration on the coarsest grid of a 2-level V-cycle (see Sec 4.3.1). As for the remaining finer grids, we set a maximum of 1 iteration, i.e., $k_1 = k_2 = 1$. For the inner stopping criteria, we stop when the relative gradient norm is less than 10^{-4} at the coarsest grid, and 10^{-3} at the remaining finer grids. In the 2-level MG/OPT, we project once from $\mathbb{C}^{512 \times 512}$ onto $\mathbb{C}^{256 \times 256}$. In the 5-level MG/OPT, we project four times: $\mathbb{C}^{512 \times 512} \mapsto \mathbb{C}^{256 \times 256} \mapsto \mathbb{C}^{128 \times 128} \mapsto \mathbb{C}^{64 \times 64} \mapsto \mathbb{C}^{32 \times 32}$.

4.3.3 Results

In Fig. 6, we show the convergence plots of the objective function values and relative errors of the three algorithms for 0%, 5%, and 10% noise levels. As expected, MG/OPT accelerates the convergence of LBFGS as a result of a substantial allocation of work to the coarser grids (see Table 1 for work allocation). In our experiments, we observe that the coarse grids contribute mainly in the earlier iterations, when there are coarse feature errors. After the first few cycles, the stopping criteria in the coarser grids are satisfied,

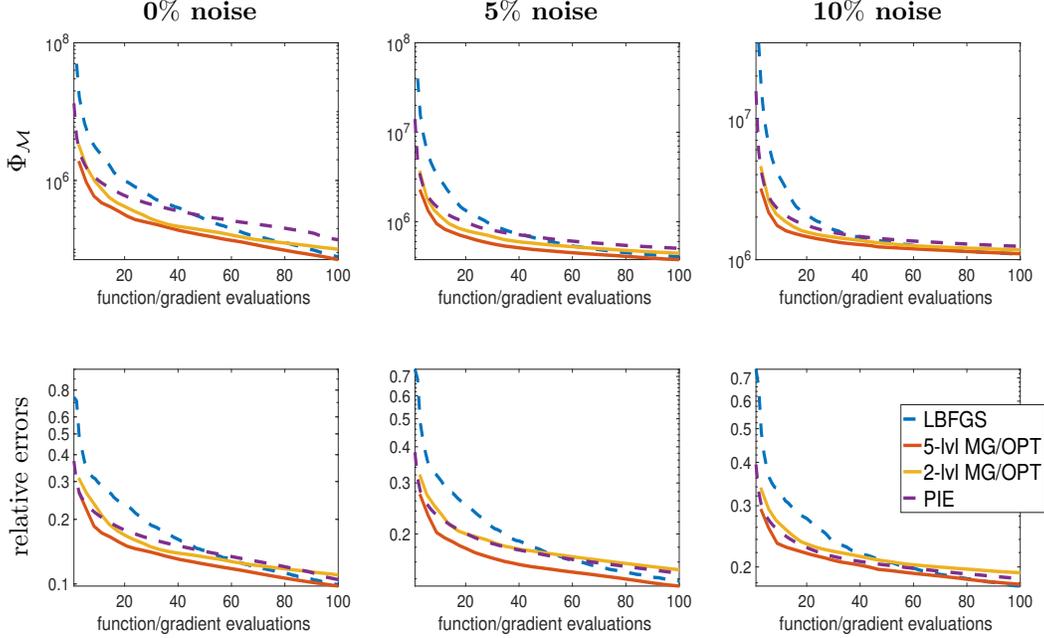


Fig. 6: Objective function and relative error histories for 0%, 5%, and 10% noise levels.

and most of the work is then shifted to the finer grids. This behavior also leads to a lack of mesh-independence, where the convergence factor defined as

$$c = \left(\frac{\Phi_{\mathcal{M}}(\mathbf{z}^{(j)}) - \Phi_{\mathcal{M}}(\mathbf{z}_{true})}{\Phi_{\mathcal{M}}(\mathbf{z}^{(0)}) - \Phi_{\mathcal{M}}(\mathbf{z}_{true})} \right)^{1/(j+1)} \quad (4.22)$$

is not fixed as the size of the problem grows. This can be seen in Table 2, where we show these estimates as the size of the problem increases for a stopping criteria of the relative gradient norm ($\|\nabla\Phi_{\mathcal{M}}(\mathbf{z}^j)\|/\|\nabla\Phi_{\mathcal{M}}(\mathbf{z}^0)\| \leq 10^{-3}$). Nonetheless, the contribution from the coarser grids causes the convergence factors of MG/OPT to be less sensitive to the problem size than the single-level LBGS.

To complement Fig. 6, we show the reconstruction history of the magnitude and phase, separately, in Fig. 7. We use the same relative error metric shown in (4.21) for the magnitudes. For the phase, however, we employ the structural similarity index metric (SSIM) [49] as its error metric to further emphasize the perception change. In particular, we measure the phase error by first performing a linear map of the current phase iterate to $[0, \frac{\pi}{2}]$ (the domain of the true phase), and then perform the SSIM with the true phase as the reference image. An SSIM value of 1 means the images are identical, and an SSIM value of 0 means that there are no structural similarities. The SSIM is an appropriate choice since it measures the image degradation as perceived change in structural information; this accounts for any potential global phase offset in our reconstructions.

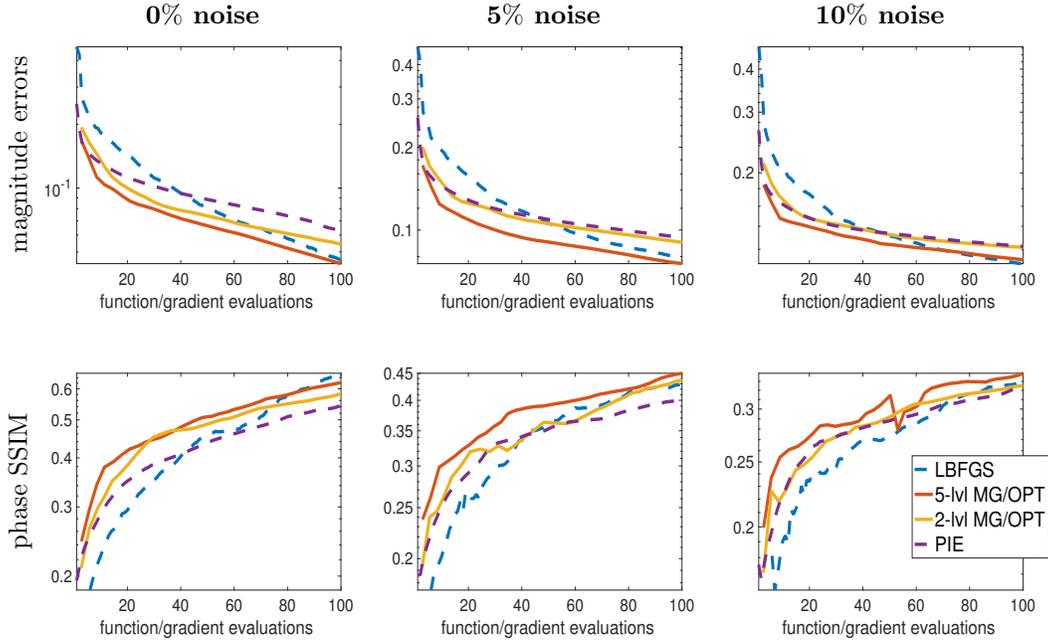


Fig. 7: *Relative magnitude error and phase SSIM histories for 0%, 5%, and 10% noise levels.*

As expected, the 5-level MG/OPT scheme outperforms the 2-level MG/OPT schemes, as we obtain higher computational savings coming from a more distributed allocation of work across grids (see Table 1). Fig. 6 and Fig. 7 also show that MG/OPT with LBFGS as its underlying solver outperforms PIE, a workhorse in the optics community.

In Fig. 8, we demonstrate the robustness of MG/OPT by plotting the relative errors vs. the noise levels after 38, 78, and 100 function/gradient evaluations. As can be seen in the plot, the 2-level and 5-level MG/OPT schemes are more robust to noise. We note that we are only able to record the values of the relative errors in the MG/OPT schemes after every V-cycle. Moreover, since the number of function/gradient evaluations performed after each V-cycle in MG/OPT is arbitrary (considering the relationship of the computational costs across grids described in Sec. 4.3.1), we are unable to plot the relative errors after exactly 38, 78, or 100 function/gradient evaluations. Instead, we find an approximate function/gradient iteration value where all the MG/OPT schemes best agree (in our case, after 38, 78, and 100 function/gradient evaluations). For instance, the relative error shown in Fig. 8b for the 5-level MG/OPT at 10% noise is recorded after 28 V-cycles, which in this particular case is equivalent to having performed 77.4 function/gradient evaluations.

In Fig. 9, we show the ptychographic phase retrieval reconstructions for PIE, LBFGS, the 2-level MG/OPT, and the 5-level MG/OPT schemes for 0%, 5% and 10% noise levels. To allow for a fair comparison, the color axis is chosen identically to the ones shown in

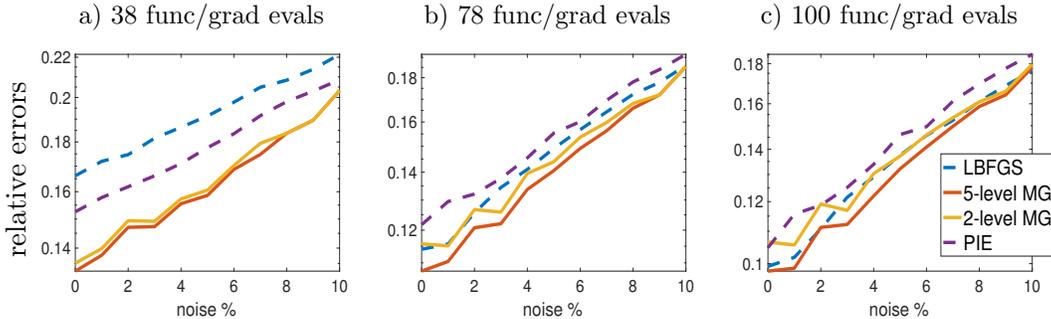


Fig. 8: Relative errors vs. noise after 38, 78, and 100 function/gradient evaluations.

Fig. 3. The robustness of MG/OPT with respect to noise can be seen clearly in the phase reconstructions.

5 Discussion

We present a multigrid approach for solving large-scale ptychographic phase retrieval problems. To this end, we consider the Multigrid Optimization (MG/OPT) scheme in order to improve the performance of gradient-based optimization algorithms for the ptychographic phase retrieval. The MG/OPT framework provides plenty of flexibility as it allows us to exploit the different hierarchical structures that ptychography exhibits. Our approach is motivated by the Full Approximation Scheme, however, rather than attempting to solve a nonlinear system of equations, MG/OPT is a multigrid optimization framework where the coarse-grid subproblem is a first-order approximation to the fine grid problem; this guarantees a descent direction in the coarse-grid correction, and ultimately, convergence when a linesearch is performed (see Sec. 3) [34].

As expected, our numerical results show that MG/OPT improves the convergence of LBFGS for the ptychographic phase retrieval problem. Our numerical experiments also show that a 2-level and a 5-level MG/OPT with LBFGS as its underlying solver is competitive when compared to the Ptychographic Iterative Engine (PIE). In particular, the 5-level MG/OPT outperforms all other algorithms, which motivates more aggressive coarsening of the grids for large-scale problems. MG/OPT reduces the computational costs and accelerates the convergence of the ptychographic phase retrieval by shifting a substantial amount of work to the coarser grid. The acceleration of the convergence is particularly evident in the early iterations, when there is the most contribution coming from the coarser grids. Further benefits are to be expected for larger problems, where deeper hierarchical structures can be exploited. We intend to extend our work for blind ptychography [19, 28, 35, 45], where the probes must also be recovered, as well as for 3D ptycho-tomography [18, 16].

Acknowledgments

We thank Stefan Wild and Doga Gursoy for valuable discussions in the preparation of this paper. We also thank Meily Wu Fung for designing the schematic illustration in Fig. 1. This material is supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357.

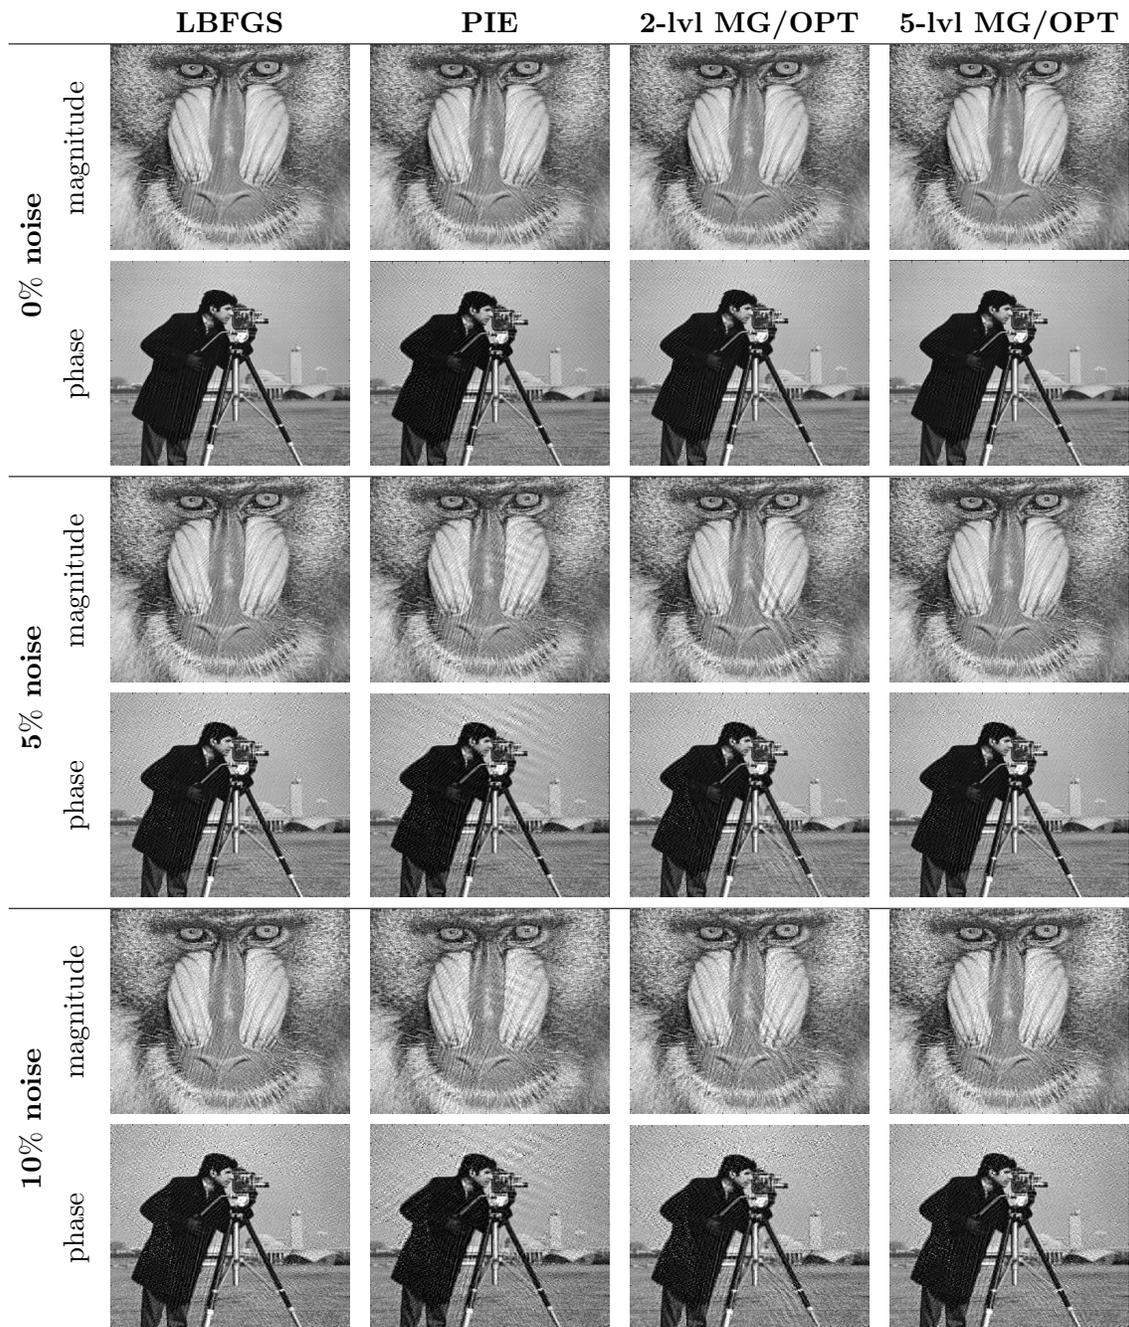


Fig. 9: *Reconstructions of the magnitude (baboon) and phase (cameraman) for different noise levels.*

References

- [1] A. Antoniou. *Digital Signal Processing*. McGraw-Hill, New York City, 2016.
- [2] R. Barakat and G. Newsam. Algorithms for reconstruction of partially known, band-limited Fourier-transform pairs from noisy data. *JOSA A*, 2(11):2027–2039, 1985.
- [3] A. Brandt and O. E. Livne. *Multigrid Techniques: 1984 Guide with Applications to Fluid Dynamics*, volume 67. SIAM, 2011.
- [4] D. Brandwood. A complex gradient operator and its application in adaptive array theory. In *IEE Proceedings H-Microwaves, Optics and Antennas*, volume 130, pages 11–16. IET, 1983.
- [5] O. Bunk, M. Dierolf, S. Kynde, I. Johnson, O. Marti, and F. Pfeiffer. Influence of the overlap parameter on the convergence of the ptychographical iterative engine. *Ultramicroscopy*, 108(5):481–487, 2008.
- [6] E. J. Candes, Y. C. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM Review*, 57(2):225–251, 2015.
- [7] E. J. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- [8] H. Chang, Y. Lou, Y. Duan, and S. Marchesini. Total variation–based phase retrieval for Poisson noise removal. *SIAM Journal on Imaging Sciences*, 11(1):24–55, 2018.
- [9] H. Chang, S. Marchesini, Y. Lou, and T. Zeng. Variational phase retrieval with globally convergent preconditioned proximal algorithm. *SIAM Journal on Imaging Sciences*, 11(1):56–93, 2018.
- [10] L. De Caro, D. Altamura, M. Arciniegas, D. Siliqi, M. R. Kim, T. Sibillano, L. Manna, and C. Giannini. Ptychographic imaging of branched colloidal nanocrystals embedded in free-standing thick polystyrene films. *Scientific Reports*, 6:19397, 2016.
- [11] J. Deng, Y. S. Nashed, S. Chen, N. W. Phillips, T. Peterka, R. Ross, S. Vogt, C. Jacobsen, and D. J. Vine. Continuous motion scan ptychography: characterization for increased speed in coherent x-ray imaging. *Optics Express*, 23(5):5438–5451, 2015.
- [12] Z. Di, M. Emelianenko, and S. Nash. Truncated Newton-based multigrid algorithm for centroidal Voronoi diagram calculation. *Numerical Mathematics: Theory, Methods and Applications*, 5(2):242–259, 2012.
- [13] J. R. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758–2769, 1982.

- [14] R. W. Gerchberg. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.
- [15] K. Giewekemeyer, M. Beckers, T. Gorniak, M. Grunze, T. Salditt, and A. Rosenhahn. Ptychographic coherent x-ray diffractive imaging in the water window. *Optics Express*, 19(2):1037–1050, 2011.
- [16] M. Gilles, Y. Nashed, M. Du, C. Jacobsen, and S. Wild. 3d x-ray imaging of continuous objects beyond the depth of focus limit. *Optica*, 5(9):1078–1086, 2018.
- [17] R. C. Gonzalez, R. E. Woods, et al. *Digital Image Processing*. Prentice Hall, New Jersey, 2002.
- [18] D. Gürsoy. Direct coupling of tomography and ptychography. *Optics Letters*, 42(16):3169–3172, 2017.
- [19] R. Hesse, D. R. Luke, S. Sabach, and M. K. Tam. Proximal heterogeneous block implicit-explicit method and application to blind ptychographic diffraction imaging. *SIAM Journal on Imaging Sciences*, 8(1):426–457, 2015.
- [20] M. Holler, M. Guizar-Sicairos, E. H. Tsai, R. Dinapoli, E. Müller, O. Bunk, J. Raabe, and G. Aeppli. High-resolution non-destructive three-dimensional imaging of integrated circuits. *Nature*, 543(7645):402, 2017.
- [21] R. Hoppe, J. Reinhardt, G. Hofmann, J. Patommel, J.-D. Grunwaldt, C. D. Damsgaard, G. Wellenreuther, G. Falkenberg, and C. Schroer. High-resolution chemical imaging of gold nanoparticles using hard x-ray ptychography. *Applied Physics Letters*, 102(20):203104, 2013.
- [22] X. Huang, H. Yan, R. Harder, Y. Hwu, I. K. Robinson, and Y. S. Chu. Optimization of overlap uniformness for ptychography. *Optics Express*, 22(10):12634–12644, 2014.
- [23] A. Konijnenberg, W. Coene, S. Pereira, and H. Urbach. Combining ptychographical algorithms with the Hybrid Input-Output (HIO) algorithm. *Ultramicroscopy*, 171:43–54, 2016.
- [24] R. M. Lewis and S. G. Nash. Model problems for the multigrid optimization of systems governed by differential equations. *SIAM Journal on Scientific Computing*, 26(6):1811–1837, 2005.
- [25] D. R. Luke. Relaxed averaged alternating reflections for diffraction imaging. *Inverse Problems*, 21(1):37, 2004.
- [26] D. R. Luke, J. V. Burke, and R. G. Lyon. Optical wavefront reconstruction: Theory and numerical methods. *SIAM Review*, 44(2):169–224, 2002.

- [27] A. Maiden, D. Johnson, and P. Li. Further improvements to the ptychographical iterative engine. *Optica*, 4(7):736–745, 2017.
- [28] A. M. Maiden and J. M. Rodenburg. An improved ptychographical phase retrieval algorithm for diffractive imaging. *Ultramicroscopy*, 109(10):1256–1262, 2009.
- [29] S. Marchesini. A unified evaluation of iterative projection algorithms for phase retrieval. *Review of Scientific Instruments*, 78(1):011301, 2007.
- [30] S. Marchesini. Phase retrieval and saddle-point optimization. *JOSA A*, 24(10):3289–3296, 2007.
- [31] J. Marrison, L. Rätty, P. Marriott, and P. O’toole. Ptychography—a label free, high-contrast imaging technique for live cells using quantitative phase information. *Scientific Reports*, 3:2369, 2013.
- [32] S. G. Nash. Preconditioning of truncated-Newton methods. *SIAM Journal on Scientific and Statistical Computing*, 6(3):599–616, 1985.
- [33] S. G. Nash. A multigrid approach to discretized optimization problems. *Optimization Methods and Software*, 14(1-2):99–116, 2000.
- [34] S. G. Nash. Convergence and descent properties for a class of multilevel optimization algorithms. Technical report, George Mason University, Fairfax, VA, 2010.
- [35] Y. S. Nashed, D. J. Vine, T. Peterka, J. Deng, R. Ross, and C. Jacobsen. Parallel ptychographic reconstruction. *Optics Express*, 22(26):32082–32097, 2014.
- [36] J. Nocedal and S. J. Wright. *Numerical Optimization, 2nd edition*. Springer, New York City, 2006.
- [37] P. M. Pelz, M. Guizar-Sicairos, P. Thibault, I. Johnson, M. Holler, and A. Menzel. On-the-fly scans for X-ray ptychography. *Applied Physics Letters*, 105(25):251101, 2014.
- [38] J. Qian, C. Yang, A. Schirotzek, F. Maia, and S. Marchesini. Efficient algorithms for ptychographic phase retrieval. *Inverse Problems and Applications, Contemporary Mathematics*, 625:261–280, 2014.
- [39] R. Remmert. *Theory of Complex Functions*, volume 122. Springer Science & Business Media, 2012.
- [40] J. M. Rodenburg. Ptychography and related diffractive imaging methods. *Advances in Imaging and Electron Physics*, 150:87–184, 2008.

- [41] J. M. Rodenburg and H. M. Faulkner. A phase retrieval algorithm for shifting illumination. *Applied Physics Letters*, 85(20):4795–4797, 2004.
- [42] B. Seifert, H. Stolz, M. Donatelli, D. Langemann, and M. Tasche. Multilevel Gauss–Newton methods for phase retrieval problems. *Journal of Physics A: Mathematical and General*, 39(16):4191, 2006.
- [43] L. Sorber, M. V. Barel, and L. D. Lathauwer. Unconstrained optimization of real functions in complex variables. *SIAM Journal on Optimization*, 22(3):879–898, 2012.
- [44] A. Suzuki, K. Shimomura, M. Hirose, N. Burdet, and Y. Takahashi. Dark-field X-ray ptychography: Towards high-resolution imaging of thick and unstained biological specimens. *Scientific Reports*, 6:35060, 2016.
- [45] P. Thibault, M. Dierolf, O. Bunk, A. Menzel, and F. Pfeiffer. Probe retrieval in ptychographic coherent diffractive imaging. *Ultramicroscopy*, 109(4):338–343, 2009.
- [46] P. Thibault and M. Guizar-Sicairos. Maximum-likelihood refinement for coherent diffractive imaging. *New Journal of Physics*, 14(6):063004, 2012.
- [47] A. Tripathi, S. Leyffer, T. Munson, and S. M. Wild. Visualizing and improving the robustness of phase retrieval algorithms. *Procedia Computer Science*, 51:815–824, 2015.
- [48] U. Trottenberg, C. W. Oosterlee, and A. Schuller. *Multigrid*. Elsevier, 2000.
- [49] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [50] Z. Wen, C. Yang, X. Liu, and S. Marchesini. Alternating direction methods for classical and ptychographic phase retrieval. *Inverse Problems*, 28(11):115010, 2012.
- [51] R. Xu, M. Soltanolkotabi, J. P. Haldar, W. Unglaub, J. Zusman, A. F. Levi, and R. M. Leahy. Accelerated Wirtinger flow: A fast algorithm for ptychography. *arXiv preprint 1806.05546*, 2018.
- [52] C. Yang, J. Qian, A. Schirotzek, F. Maia, and S. Marchesini. Iterative algorithms for ptychographic phase retrieval. *arXiv preprint 1105.5628*, 2011.
- [53] L.-H. Yeh, J. Dong, J. Zhong, L. Tian, M. Chen, G. Tang, M. Soltanolkotabi, and L. Waller. Experimental robustness of fourier ptychography phase retrieval algorithms. *Optics Express*, 23(26):33214–33240, 2015.