

Gauss–Newton Optimization for Phase Recovery From the Bispectrum

James Lincoln Herring , James Nagy, and Lars Ruthotto

Abstract—Phase recovery from the bispectrum is a central problem in speckle interferometry which can be posed as an optimization problem minimizing a weighted nonlinear least-squares objective function. We look at two different formulations of the phase recovery problem from the literature, both of which can be minimized with respect to either the recovered phase or the recovered image. Previously, strategies for solving these formulations have been limited to gradient descent or quasi-Newton methods. This article explores Gauss–Newton optimization schemes for the problem of phase recovery from the bispectrum. We implement efficient Gauss–Newton optimization schemes for all the formulations. For the two of these formulations which optimize with respect to the recovered image, we also extend to projected Gauss–Newton to enforce element-wise lower and upper bounds on the pixel intensities of the recovered image. We show that our efficient Gauss–Newton schemes result in better image reconstructions with no or limited additional computational cost compared to previously implemented first-order optimization schemes for phase recovery from the bispectrum. MATLAB implementations of all methods and simulations are made publicly available in the BiBox repository on Github.

Index Terms—Phase recovery, bispectrum, bispectral imaging, Gauss–Newton method.

I. INTRODUCTION

IMAGE blurring due to turbulence poses a significant obstacle in many applications. One approach to obtaining high spatial frequency images of an object through turbulent optical systems is speckle interferometry, which is built upon Labeyrie’s observation that high spatial frequency information can be recovered from short-exposure images [1]. This work provided the means to obtain diffraction-limited reconstructions of an object’s Fourier modulus, or power spectrum, using short-exposure, photon-limited data.

While an object’s Fourier modulus may be sufficient in the case of some simple objects, many applications also require recovery of the object’s Fourier phase in order to produce high quality images. Thus, phase retrieval is often an essential subproblem when using speckle interferometry. Several methods

Manuscript received December 17, 2018; revised April 4, 2019 and June 25, 2019; accepted July 8, 2019. Date of publication October 21, 2019; date of current version January 13, 2020. This work was supported by the US National Science Foundation under Grants DMS-1819042 and 1522599. The associate editor coordinating the review of this manuscript and approving it for publication was K. J. Batenburg. (*Corresponding author: James Lincoln Herring.*)

J. L. Herring is with the Department of Mathematics, University of Houston, Houston, TX 77004 USA (e-mail: herring@math.uh.edu).

J. Nagy and L. Ruthotto are with the Department of Mathematics, Emory University, Atlanta, GA 30322 USA (e-mail: jnagy@emory.edu; lruthotto@emory.edu).

Digital Object Identifier 10.1109/TCI.2019.2948784

have been developed to solve the phase retrieval problem, most using relationships defined by high-order statistical correlation measures such as an object’s triple correlation or its Fourier counterpart, the bispectrum [2], [3]. These measures can be collected from the short-exposure data and used to reconstruct an object’s phase.

This work focuses on the problem of phase recovery from an object’s collected bispectrum. The broader problem of phase retrieval occurs in numerous engineering fields and the sciences including astronomy, electron microscopy, crystallography, and optical imaging. For a recent overview of applications and challenges, we recommend [4]. The problem of phase recovery from the bispectrum has its origins in astronomical imaging for low-light images in the visible spectrum [2], [3], [5]–[8]. More recently, the problem has been of interest in multiple applications including long-range horizontal and slant path imaging [9]–[11] and phase recovery in underwater imaging [12], [13]. The general problem of phase retrieval from Fourier measurements also continues to be an active area of research in the signal processing community [14]–[22]. Many phase retrieval applications considered in the literature are based on the relationship between an object’s phase and higher order statistical moments collected from the data (such as the bispectrum) and require solving constrained nonlinear least-squares problems; see, e.g., [16], [21], [22]. Thus while we consider only the problem of phase recovery from the bispectrum, the content of this paper is more broadly applicable to the phase retrieval problem in general.

For phase recovery from the bispectrum, strategies can be separated into two categories: recursive algorithms and weighted least-squares problems. Recursive strategies fix a small set of known phase values near the center of the Fourier domain and use these known phase values along with the collected bispectrum to recursively compute the remainder of the reconstructed phase. Such strategies are well explored in the literature [5], [6], [10], [23]–[28]. One limitation of recursive strategies is poor performance when attempting to reconstruct the phase values associated with high frequency information when using noisy data [7]. To improve on this, the phase recovery problem can be reformulated as a weighted least-squares problem minimizing the mismatch between the unknown object phase and the collected phase of the bispectrum. This least-squares problem is typically nonlinear because the bispectrum is collected modulo- 2π in the range $[-\pi, \pi]$, i.e., the bispectrum is “wrapped.” One can solve this nonlinear least-squares problem [7], [25], [26], [29]–[32]. Alternatively, techniques have been proposed to unwrap the collected bispectrum, e.g., [33], [34], resulting in

a linear least-squares problem, but these approaches have been shown to produce inferior results to solving the nonlinear formulation of the problem [29]. Previous approaches to solving both the linear and nonlinear least-squares formulations have focused on gradient-based first-order and quasi-Newton methods such as gradient descent and the limited memory Broyden–Fletcher–Goldfarb–Shanno method (L-BFGS) [7], [29], [30].

In this paper, we focus on the formulation of phase recovery from the bispectrum as a nonlinear weighted least-squares problem using the wrapped bispectrum. We present Gauss–Newton schemes as an alternative to the gradient-based optimization approaches previously used in the literature. Specifically, we make the following contributions:

- We implement efficient Gauss–Newton optimization methods for two nonlinear least-squares formulations of the phase recovery problem from the literature. Both of these formulations can be minimized with respect to the recovered phase or recovered image, resulting in four possible formulations. Our implementations exploit sparsity, matrix reorderings, incomplete factorization, and the speed of the fast Fourier transform (FFT) to reduce the cost associated with solving the linear system to calculate the Gauss–Newton step at each iteration of the optimization. The resulting schemes have per iteration costs on the same order of magnitude as gradient descent and quasi-Newton methods like NLCG and L-BFGS but benefit from faster convergence. This results in improved time-to-solution and lower overall computational cost compared with previous approaches.
- For the two formulations that are minimized with respect to the resulting image, we also explore the constrained problem with pixel-wise non-negativity constraints on the pixel intensities of the recovered image. We use a projected Gauss–Newton method to enforce these constraints. This strategy improves the quality of the recovered image. Additionally, it eliminates the need for a regularizer enforcing non-negativity and offers the option to use other regularization options while still enforcing non-negativity. We demonstrate this for two common regularizers, a discrete gradient regularizer, and a total variation regularizer. For comparison with constrained gradient-based optimization, we also test projected gradient descent with both regularizers for these two formulations.
- We show through numerical experiments that our proposed standard Gauss–Newton and projected Gauss–Newton implementations offer improvements on gradient descent, L-BFGS, and projected gradient descent in both time-to-solution and the quality of the resulting image for the four formulations of the phase recovery problem taken from the literature. We also show that these improvements in the quality of the resulting image are robust to problem parameters including atmospheric turbulence and noise level.
- We provide MATLAB implementations for all our methods and for all our numerical simulations in the BiBox repository on Github: <https://github.com/herrinj/BiBox>

The paper is organized as follows. Section II gives an overview of the collection of the bispectrum and the formulation of phase recovery from the bispectrum as a weighted nonlinear least-squares optimization problem. Section III details our proposed iterative Gauss–Newton optimization for solving the phase recovery problem from the previous section. Specifically, we introduce expressions for the gradient and Gauss–Newton approximations to the Hessians for four separate formulations of the phase recovery problem and discuss efficient strategies for solving the linear system associated with the Gauss–Newton step within each iteration of the optimization. Lastly, for two of the problem formulations, we extend our implementation to projected Gauss–Newton which allows for phase recovery while simultaneously imposing pixel-wise intensity bounds on the recovered image. Section IV presents numerical experiments. We compare our proposed Gauss–Newton schemes with common first-order optimization methods: gradient descent, projected gradient descent, and quasi-Newton L-BFGS method. We also demonstrate the robustness of our proposed Gauss–Newton schemes for a range of problem parameters. We end with concluding remarks in Section V.

II. PHASE RECOVERY PROBLEM

Most techniques for recovering an object’s phase from speckle image data rely on the object’s triple correlation and its Fourier transform, the bispectrum [2], [3]. An object’s triple correlation is a second-order statistical moment given by measuring an object against two independently shifted copies of itself. For a two dimensional object $o(\mathbf{x})$ with $\mathbf{x} \in \mathbb{R}^2$, this is expressed by

$$o^{TC}(\mathbf{x}_1, \mathbf{x}_2) = \iint_{-\infty}^{\infty} o^*(\mathbf{x}) o(\mathbf{x} + \mathbf{x}_1) o(\mathbf{x} + \mathbf{x}_2) d\mathbf{x}.$$

Taking the Fourier transform of this and using the convolution property, we get the object’s bispectrum,

$$O^{(3)}(\mathbf{u}, \mathbf{v}) = O(\mathbf{u}) O(\mathbf{v}) O^*(\mathbf{u} + \mathbf{v}).$$

Here, $O(\mathbf{u})$ is the Fourier transform of the object and $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$ are spatial frequencies. It follows that the phase of the object’s bispectrum and the object are related by

$$\beta(\mathbf{u}, \mathbf{v}) = \phi(\mathbf{u}) + \phi(\mathbf{v}) - \phi(\mathbf{u} + \mathbf{v}) \quad (1)$$

where $\phi(\mathbf{u})$ is the phase of $O(\mathbf{u})$ and $\beta(\mathbf{u}, \mathbf{v})$ is the phase of the object’s bispectrum corresponding to the triplet $(\mathbf{u}, \mathbf{v}, \mathbf{u} + \mathbf{v})$. This expression provides a deterministic relationship that provides the basis for various algorithms for phase recovery from the bispectrum.

We highlight several characteristics of the relationship in (1). First, the object’s bispectrum is unknown in practice and is estimated using the average data bispectrum from a series of short-exposure images of the object, $i_k(\mathbf{x})$ for $k = 1, \dots, N$. To recover the object’s Fourier modulus, or power spectrum, an additional set of short-exposure images of an appropriate reference star, $s_k(\mathbf{x})$ for $k = 1, \dots, N$, is also necessary; see, e.g., [7]. Typically, the data bispectrum is collected modulo- 2π

or “wrapped.” This necessitates accounting for the modulus within the phase recovery problem or unwrapping the bispectrum before phase recovery; see, e.g., [33], [34]. We opt for the first strategy.

The size of the data bispectrum depends on the number of discrete $(\mathbf{u}, \mathbf{v}, \mathbf{u} + \mathbf{v})$ triplets for which the bispectrum is collected. Computational considerations and the optics of the problem make it inadvisable to range \mathbf{u} and \mathbf{v} over the entire Fourier plane. In practice, the number of discrete coordinates \mathbf{u} and \mathbf{v} used to determine these triplets is restricted by two radii: \mathbf{u} indices are restricted within a larger, “recovery radius,” and \mathbf{v} indices are restricted within a smaller radius [35]–[37]. The value of these radii vary depending on seeing conditions and parameters; we discuss the values we select in the discussion on data setup in Sec. IV. For computational purposes, the indices for the set of discrete $(\mathbf{u}, \mathbf{v}, \mathbf{u} + \mathbf{v})$ triplets are computed once and stored in an indexing structure which vectorizes the accumulation of the data bispectrum. This structure can be implemented to exploit the symmetries in the phase for real-valued images in Fourier space, which increases computational efficiency. We base our indexing structure on the work presented by [28].

A. Phase Recovery Schemes

Numerous algorithms for phase recovery from the modulo- 2π bispectrum have been proposed which build upon the relationship in (1). These can be separated into two categories: recursive algorithms; see, e.g., [5], [6], [10], [23]–[28], and nonlinear least-squares formulations; see, e.g., [7], [25], [26], [29]–[31]. We look at the second category. Specifically, we look at four nonlinear least-squares formulations proposed in the literature.

We first establish some notation. Let $\phi_{\text{true}} \in \mathbb{R}^n$ be the vector of true object phase values we aim to recover. Here, the dimension n is determined by seeing conditions and signal-to-noise (SNR) considerations. Let $\beta_{\text{true}} \in \mathbb{R}^m$ be the corresponding true, unwrapped phase of the data bispectrum collected for m distinct $(\mathbf{u}, \mathbf{v}, \mathbf{u} + \mathbf{v})$ triplets. These two vectors are related by

$$\beta_{\text{true}} = A\phi_{\text{true}} \quad (2)$$

where $A \in \mathbb{R}^{m \times n}$ is a sparse matrix with three non-zeros per row: two 1’s and one -1 corresponding to the signs of phase elements in (1). The indexing structure used to accumulate the data bispectrum also contains the information used to construct this matrix. The relationship in (2) provides the basis of the algorithms for recovering the unknown object phase, $\phi \in \mathbb{R}^n$, from the phase of the collected data bispectrum, $\beta \in \mathbb{R}^m$. The SNR for each entry in β varies, so we define a diagonal weighting matrix $W \in \mathbb{R}^{m \times m}$ with positive diagonal entries determined by the SNR of the bispectrum phase elements [7]. Let $W^{1/2}$ be the diagonal matrix with the square roots of the SNR weights on the diagonal.

Using these definitions, we introduce several optimization formulations for recovering an object’s phase from the phase of the data bispectrum from the literature. All of the formulations involve fitting the data through the minimization of an appropriate nonlinear least-squares objective function. Here, we introduce two such objective functions from the literature. The

first fits the object’s phase to the collected data bispectrum by solving the following minimization problem [29]:

$$\min_{\phi} \left\{ E_1(\phi) = \frac{1}{2} \|W^{1/2} \text{mod}_{2\pi}(\beta - A\phi)\|_2^2 \right\}. \quad (3)$$

Here, the modulus is introduced because the phase of the data bispectrum is collected modulo- 2π . This introduces several considerations. It causes $E_1(\phi)$ to be non-convex with periodic local minima every 2π . Additionally, $E_1(\phi)$ is periodically discontinuous with jumps every 2π where the misfit wraps from 0 to 2π .

An alternative formulation to (3) exists which avoids the issue of non-differentiability [29]. It exploits the identity $e^{i\theta} = \cos(\theta) + i\sin(\theta)$ to avoid the modulus. The resulting optimization problem is

$$\begin{aligned} \min_{\phi} \left\{ E_2(\phi) = \frac{1}{2} \|W^{1/2} (\cos \beta - \cos A\phi)\|_2^2 \right. \\ \left. + \frac{1}{2} \|W^{1/2} (\sin \beta - \sin A\phi)\|_2^2 \right\}, \end{aligned} \quad (4)$$

Both $E_1(\phi)$ and $E_2(\phi)$ are minimized with respect to the unknown object phase, ϕ . One shortcoming is that these objective functions solely minimize the data misfit between the recovered phase and the collected data bispectrum, i.e., they are blind to how the recovered phase impacts the recovered object when combined with the object’s power spectrum. An alternative is to optimize with respect to the resulting image, \mathbf{o} , i.e., to reformulate (3) and (4) as $E_1(\phi(\mathbf{o})) = E_1(\mathbf{o})$ and $E_2(\phi(\mathbf{o})) = E_2(\mathbf{o})$, respectively. This idea was originally proposed for $E_1(\mathbf{o})$ in [30]. Extending the idea to $E_2(\mathbf{o})$ is new to our knowledge.

Optimizing with respect to the resulting image is attractive for a number of reasons. In many applications, the image of the recovered object is the ultimate goal which necessitates solving the phase recovery problem. It then makes intuitive sense that the optimization should take into account this recovered image. Additionally, since the objective functions for both (3) and (4) are highly non-convex, optimizing with respect to the resulting image may allow us to converge to different minima, potentially resulting in improved recovery of the object. Lastly, optimizing with respect to the recovered image allows us to impose desirable characteristics on the solution image using appropriate regularization and bound-constraints. Such characteristics may include non-negativity or smoothness of the recovered object. With this in mind, we incorporate regularization and bound constraints into optimization formulations for $E_1(\mathbf{o})$ and $E_2(\mathbf{o})$ and consider the problems

$$\min_{\mathbf{o} \in \mathcal{C}} \{E_1(\mathbf{o}) + \alpha R(\mathbf{o})\} \quad (5)$$

and

$$\min_{\mathbf{o} \in \mathcal{C}} \{E_2(\mathbf{o}) + \alpha R(\mathbf{o})\}. \quad (6)$$

Here, \mathcal{C} denotes a closed, convex set enforcing element-wise bound constraints on the recovered object \mathbf{o} , for example, non-negativity. The operator $R(\mathbf{o})$ is a regularizer subject to the

weighting parameter $\alpha > 0$ for introducing desirable characteristics in the solution such as smoothness, sparsity, or non-negativity. In Sec. IV, we look at three potential regularizers. For the unconstrained problem ($\mathcal{C} = \mathbb{R}^n$), we explore a previously proposed penalty regularizer encouraging non-negative solutions [30]. We also solve the problem with non-negativity constraints, which we enforce using projected Gauss–Newton. For this formulation, we implement and test two regularizers: a quadratic regularizer with a discrete gradient operator, ∇_h , and a nonlinear total variation regularizer popular in many imaging applications [38]. Other potential options include the identity operator, p -norm based regularizers, or hybrid regularization techniques.

III. OPTIMIZATION

We now look at methods for minimizing the four optimization problems introduced in Sec. II. Previous work has relied on first-order and quasi-Newton optimization methods such as gradient descent, nonlinear conjugate gradient (NLCG), and the limited memory implementation of the Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) method [7], [29], [30]. For further references on these methods, see [39]. Our work implements efficient, Gauss–Newton optimization schemes. To begin our discussion, we provide a brief overview of the Gauss–Newton method. After this, we discuss extending the method to projected Gauss–Newton as a way to enforce element-wise bound constraints on the solution. Lastly, we narrow our focus and discuss specific approaches for solving the linear system to determine the Gauss–Newton step for each of the objective functions proposed in the previous section. Solving this linear system is the most computationally intensive step in Gauss–Newton optimization and represents the most significant contribution of our work. For (3) and (4), we introduce an efficient factorization based approach which uses sparse matrix reordering and factorization to solve the linear system directly. For (5) and (6), we solve for the Gauss–Newton step iteratively. We also extend the optimization for (5) and (6) to include regularization.

A. Gauss–Newton Method

The Gauss–Newton method is a common optimization scheme for solving unconstrained nonlinear least-squares problems such as (3) and (4) [39], [40]. Starting from an initial guess, the method iteratively updates the computed solution with a step towards minimizing the objective function. At each iteration, both the direction and length of this step must be computed.

The direction of this step at each iteration is determined by solving a linear system. For an arbitrary objective function $E(\mathbf{y})$ at the current iterate \mathbf{y} , this equation is

$$H_E(\mathbf{y})\mathbf{p} = -\nabla_{\mathbf{y}}E(\mathbf{y}). \quad (7)$$

Here, $\nabla_{\mathbf{y}}E(\mathbf{y})$ is the gradient, and $H_E(\mathbf{y}) \approx \nabla_{\mathbf{y}}^2E(\mathbf{y})$ is a symmetric positive semi-definite approximation to the Hessian; see, e.g., p. 246 in [39]. Both the gradient and the approximation to the Hessian must be updated at each iteration of the optimization. The system (7) can be solved using direct methods or an appropriate iterative solver, and solving this system is the most

significant cost in the Gauss–Newton method. Our work focuses on implementing efficient strategies for solving this system, and we use both direct and iterative solvers for (7) depending on which of the four formulations of the phase recovery problem we consider.

After choosing the step direction, we choose the step length via a backtracking Armijo line search [39]. The Armijo line search selects a step length $0 < \eta \leq 1$ by iteratively backtracking from the full Newton step $\eta = 1$ until a step length is found which guarantees a sufficient reduction of the objective function. Sufficient reduction is determined by the Armijo condition,

$$E(\mathbf{y} + \eta\mathbf{p}) \leq E(\mathbf{y}) + c\eta\nabla_{\mathbf{y}}E(\mathbf{y})^\top\mathbf{p}$$

where $c \in (0, 1)$ is a reduction constant. We use $c = 1 \times 10^{-4}$ as recommended in [39]. Note that if the data residual is small in the neighborhood of the solution, the Gauss–Newton approximation to the Hessian is more accurate, and we expect to take the full step $\eta = 1$. This reduces the number of backtracking line search iterations required per outer Gauss–Newton iteration and lowers the cost of the optimization.

The convergence of the Gauss–Newton method depends on the nonlinearity of the objective function and the norm of the data residual. In the best case, it can be shown to have near-quadratic convergence in the neighborhood of a minimizer. This makes it attractive compared to the linear convergence of gradient-based methods and the super-linear convergence of quasi-Newton methods like L-BFGS. We note that the backtracking Armijo line search we use does not satisfy the Wolfe conditions necessary for rigorous convergence results using Gauss–Newton, but this does not pose a problem in practice as observed in Sec. IV. For a more in-depth discussion of the convergence properties of Gauss–Newton and the other optimization methods; see, e.g., [39].

B. Projected Gauss–Newton Method

The standard Gauss–Newton method is sufficient for the unconstrained optimization problems in (3) and (4). However, additional considerations are required to solve constrained problems. For the formulations (5) and (6), we want to enforce element-wise bound constraints on the solution object. To do this, we use the projected Gauss–Newton method [41]. Projected Gauss–Newton incorporates constraints by combining Gauss–Newton with projected gradient descent. The goal of this combination is to benefit from the strengths of both methods: fast convergence in the neighborhood of the solution for Gauss–Newton and straightforward implementation of bound constraints for projected gradient descent.

We implement projected Gauss–Newton as follows. Given a feasible initial guess, the method separates the set of optimization variables into two sets at each iteration. The active set, \mathcal{A} , contains the indices of the variables for which the bound constraints are active, i.e., variables on the boundary of the feasible region. The inactive set, \mathcal{I} , contains the indices of the variables for which the bound constraints are inactive, i.e., variables on the interior of the feasible region. Let the vectors

$\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{I}}$ be vectors containing the entries of \mathbf{x} at the indices defined by the active and inactive sets, respectively.

On the inactive set, we take a Gauss–Newton step $\mathbf{p}_{\mathcal{I}}$ computed by solving a projected version of (7). This can be written

$$(I_{\mathcal{I}} H_E(\mathbf{y}) I_{\mathcal{I}}) \mathbf{p}_{\mathcal{I}} = -I_{\mathcal{I}} \nabla_{\mathbf{y}} E(\mathbf{y}),$$

where $I_{\mathcal{I}}$ is a projection operator. It is computed by modifying an identity matrix by setting diagonal entries in the columns corresponding to indices in \mathcal{A} to 0. This projected linear system can be solved using either a direct method or an iterative method with an appropriate preconditioner.

On the active set, we take a projected gradient descent step given by

$$\mathbf{p}_{\mathcal{A}} = -I_{\mathcal{A}} \nabla_{\mathbf{y}}^P E(\mathbf{y}).$$

Here, $I_{\mathcal{A}} = I - I_{\mathcal{I}}$ and $\nabla_{\mathbf{y}}^P E(\mathbf{y})$ is the projected gradient where the projection sets to zero entries of the gradient that would cause the updated solution to violate the bound constraints.

We then combine the steps on the inactive and active sets to get the full projected Gauss–Newton step, \mathbf{p} , given by

$$\mathbf{p} = \mathbf{p}_{\mathcal{I}} + \gamma \mathbf{p}_{\mathcal{A}}.$$

Here, the parameter $\gamma > 0$ is introduced to reconcile the difference in scale between the two steps. Following the recommendation in [41], we set $\gamma = \frac{\|\mathbf{p}_{\mathcal{I}}\|_{\infty}}{\|\mathbf{p}_{\mathcal{A}}\|_{\infty}}$.

Lastly, we determine the length of the step using a projected Armijo line search. It follows the same backtracking strategy presented previously for standard Gauss–Newton but uses a modified Armijo condition,

$$E(Q(\mathbf{y} + \eta \mathbf{p})) \leq E(\mathbf{y}) + c\eta \nabla_{\mathbf{y}}^P E(\mathbf{y})^\top \mathbf{p},$$

where Q is a projection onto the feasible region and $\nabla_{\mathbf{y}}^P E(\mathbf{y})$ is the projected gradient. Note that due to the projected gradient descent step on the active set, we expect projected Gauss–Newton to require more line search iterations than standard Gauss–Newton. This is particularly relevant when many bound-constraints are active, and we often observe this in practice. Also, variables may enter and leave the active and inactive sets with each update to the solution, so the active and inactive sets must be updated at each iteration of the optimization.

The convergence behavior of projected Gauss–Newton depends on the objective function and the number of variables in the active and inactive sets. For the best case where all variables lie in the inactive set, the method reverts to standard Gauss–Newton and has the corresponding convergence properties. In the worst case scenario where all variables are in the active set, convergence behavior is that of gradient descent, i.e., linear convergence. Our numerical experiments in Sec. IV show that for the phase recovery problems in this paper, we observe more of the former.

C. Efficient Approaches for Computing the Gauss–Newton Step

We now consider solving the linear system (7) associated with the Gauss–Newton step for the four formulations of the problem. We begin by considering the optimization problem associated

with $E_1(\phi)$ in equation (3). The gradient and the Gauss–Newton approximation to the Hessian are given by

$$\begin{aligned} \nabla_{\phi} E_1(\phi) &= -A^{\top} W \text{mod}_{2\pi}(\boldsymbol{\beta} - A\phi) \\ H_{E_1}(\phi) &= A^{\top} W A. \end{aligned} \quad (8)$$

For both derivatives, we ignore the modulus during differentiation. This strategy is used in the literature and proves effective in practice, but it may impact the optimization [29]. We also note that the approximation to the Hessian, $H_{E_1}(\phi)$, is independent of the variable ϕ and remains constant for each Gauss–Newton iteration. Furthermore, the matrix is large, sparse, and structured. It has a large number of zero rows and columns due to SNR considerations and symmetries in the phase for real-valued images. The object’s phase can only be recovered at points for which data, i.e., the phase of the bispectrum, has been collected. This is restricted with a problem specific ‘recovery radius’ determined by seeing conditions and the optics system [35]–[37]. Thus, columns and rows corresponding to entries outside this radius are set to zero. Furthermore, the phase for real-valued images displays a symmetry through D.C. That is, if we let D.C. be the origin, $(0, 0)$, then the phase at point (i, j) is the negative of the phase at point $(-i, -j)$, i.e., $-\phi(i, j) = \phi(-i, -j)$. Exploiting this allows us to further increase sparsity, and makes sparse, direct methods an attractive option for solving the linear system. Importantly, any factorization can be computed once offline and then reused for each iteration of the optimization.

With that in mind, we implement the following strategy. Prior to the Gauss–Newton optimization, we compute and factor the approximate Hessian offline. We perform a symmetric approximate minimum degree permutation to shift the non-zero rows and columns to the upper left-hand corner of the matrix [42]. This permutation is a heuristic designed to minimize the fill-in when the permuted matrix is factored. Then, we extract the $\tilde{n} \times \tilde{n}$ sub-matrix corresponding to the non-zero columns and rows of the permuted matrix. Here, the size \tilde{n} corresponds to the number of phase-values to be recovered and is small due to the SNR considerations and symmetries previously mentioned, $\tilde{n} < n$. The extracted sub-matrix is also symmetric positive definite. On this sub-matrix, we perform a zero fill-in incomplete Cholesky factorization and store the factors [43], [44]. The symmetric approximate minimum degree permutation helps to minimize the amount of information lost with the zero fill-in strategy. A full Cholesky factorization is also possible, but it may lead to some loss of sparsity due to fill-in and does not improve results for the problems presented in this paper. The stored incomplete Cholesky factors can then be used to solve for the Gauss–Newton step at each iteration of the standard Gauss–Newton method. Computing the factorization once offline reduces the cost of solving (7) at each iteration. The cost of solving the linear system associated with the Gauss–Newton step is thereby reduced to solving two sparse, triangular systems for a cost of $\mathcal{O}(\tilde{n}^2)$. Note that in practice, $\mathcal{O}(\tilde{n}^2)$ is a pessimistic bound on the cost of inverting the triangular factors because the incomplete Cholesky factorization preserves their sparsity. Fig. 1 gives a visual step-by-step of the strategy.

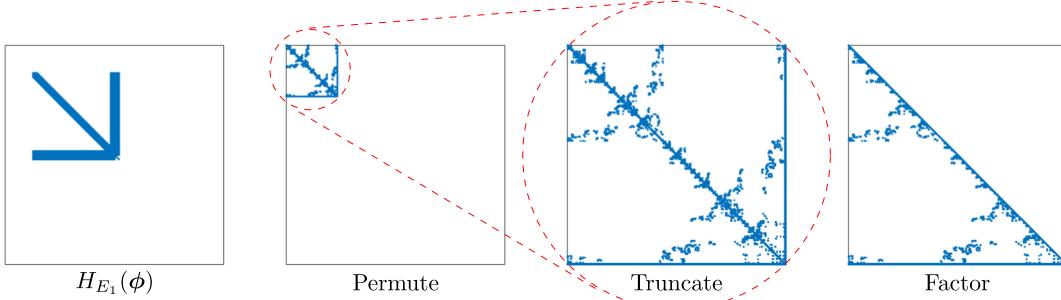


Fig. 1. This figure shows the permutation and factorization strategy used for the approximate Hessians of $E_1(\phi)$ and $E_2(\phi)$. From left to right, we start with the full Hessian, permute using a symmetric approximate minimum degree factorization, truncate out the zero rows and columns, and perform an incomplete Cholesky factorization on the truncated matrix. We use the resulting factors and the permutation to solve for the (7) at each iteration of the Gauss–Newton optimization.

In the case of $E_2(\phi)$, the gradient and the Gauss–Newton approximation to the Hessian are

$$\begin{aligned}\nabla_\phi E_2(\phi) &= A^\top W D_1(\phi) \\ H_{E_2}(\phi) &= A^\top W D_2(\phi) A.\end{aligned}\quad (9)$$

The matrices $D_1(\phi)$ and $D_2(\phi)$ are diagonal matrices with diagonal entries given by

$$\begin{aligned}\text{diag}(D_1(\phi)) &= \cos \beta \odot \sin A\phi - \sin \beta \odot \cos A\phi \\ \text{diag}(D_2(\phi)) &= \cos \beta \odot \cos A\phi + \sin \beta \odot \sin A\phi,\end{aligned}$$

where \odot denotes the Hadamard product or component-wise multiplication of two vectors. Unlike (8), the Gauss–Newton approximation to the Hessian here is not independent of the phase due to the diagonal matrix $D_2(\phi)$, and makes solving the system in (9) more computationally expensive. However, we note that if $\text{mod}_{2\pi}(\beta - A\phi) = 0$, i.e., we perfectly match the recovered phase to the phase of the bispectrum, then the diagonal entries of $D_2(\phi)$ are 1, i.e., $D_2(\phi)$ becomes the identity. In practice, this is not true for noisy data and large residuals, but it serves as motivation for omitting the diagonal matrix $D_2(\phi)$ from the approximate Hessian. This makes $H_{E_1}(\phi) = H_{E_2}(\phi)$, and we can use the same permutation and factorization strategy as we use for (3) when computing the Gauss–Newton step for (4). This is effective in practice, as shown in the numerical experiments. We also remark that although we use the same approximate Hessian for both formulations, the objective functions and gradients differ resulting in different solutions.

For the optimization problems based on $E_1(\phi)$ and $E_2(\phi)$, that is solving (5) and (6) respectively, it is necessary to differentiate with respect to the resulting image. Expressions for the gradients and the Gauss–Newton approximations to the Hessians follow from (8) and (9) via the chain rule with additional gradients and derivatives for the regularization term. For the objective function in (5), we get

$$\begin{aligned}\nabla_\phi E_1(\phi) &= -\frac{\partial \phi^*}{\partial \phi} A^\top W \text{mod}_{2\pi}(\beta - A\phi) + \alpha \nabla_\phi R(\phi) \\ H_{E_1}(\phi) &= \frac{\partial \phi^*}{\partial \phi} A^\top W A \frac{\partial \phi}{\partial \phi} + \alpha \nabla_\phi^2 R(\phi)\end{aligned}\quad (10)$$

where $\frac{\partial \phi}{\partial \phi}$ is the derivative of the phase with respect to the resulting object and $\frac{\partial \phi^*}{\partial \phi}$ is its adjoint. Similarly, the derivatives for the objective function in (6) are given by

$$\begin{aligned}\nabla_\phi E_2(\phi) &= \frac{\partial \phi^*}{\partial \phi} A^\top W D_1(\phi) + \alpha \nabla_\phi R(\phi) \\ H_{E_2}(\phi) &= \frac{\partial \phi^*}{\partial \phi} A^\top W D_2(\phi) A \frac{\partial \phi}{\partial \phi} + \alpha \nabla_\phi^2 R(\phi).\end{aligned}\quad (11)$$

To derive expressions for $\frac{\partial \phi}{\partial \phi}$ and its adjoint, we observe that the object's phase can be expressed as a function of the object by

$$\phi(\phi) = \arctan \left(\frac{\text{Im}(\mathcal{F}\phi)}{\text{Re}(\mathcal{F}\phi)} \right),$$

where \mathcal{F} is a 2D Fourier transform matrix. Differentiating this with respect to ϕ , we get an expression for the action of the adjoint operator in the direction q ,

$$\begin{aligned}\frac{\partial \phi^*}{\partial \phi}(q) &= \frac{\text{Re}(\mathcal{F}\phi) \odot \text{Im}(\mathcal{F}q) - \text{Im}(\mathcal{F}\phi) \odot \text{Re}(\mathcal{F}q)}{|\mathcal{F}\phi|^2} \\ &= \frac{\text{Im}(\overline{\mathcal{F}\phi} \odot \mathcal{F}q)}{\mathcal{F}\phi \odot \mathcal{F}\phi} \\ &= \text{Im} \left(\frac{\mathcal{F}q}{\mathcal{F}\phi} \right).\end{aligned}$$

Here, the Hadamard product, division, square, and conjugation are all taken component-wise. This expression is also given in [30].

To evaluate the Gauss–Newton approximations to the Hessians, it is necessary to compute the forward operator, $\frac{\partial \phi}{\partial \phi}$. This operator is not included in previous work using first-order optimization methods. It can be computed and applied using a matrix-free approach. In the direction q , it is given by

$$\begin{aligned}\frac{\partial \phi}{\partial \phi}(q) &= \text{Im} \left(\mathcal{F} \left(\frac{\text{Re}(\mathcal{F}\phi) \odot q}{|\mathcal{F}\phi|^2} \right) \right) \\ &\quad - \text{Re} \left(\mathcal{F} \left(\frac{\text{Im}(\mathcal{F}\phi) \odot q}{|\mathcal{F}\phi|^2} \right) \right) \\ &= \text{Im} \left(\mathcal{F} \left(\frac{\overline{\mathcal{F}\phi} \odot q}{\mathcal{F}\phi \odot \mathcal{F}\phi} \right) \right) \\ &= \text{Im} \left(\mathcal{F} \left(\frac{q}{\mathcal{F}\phi} \right) \right).\end{aligned}$$

Again, all operations are taken component-wise. In practice, some additional considerations are necessary for both the forward and adjoint operators. Storing the 2D FFT matrices is inefficient and infeasible, so the operators are passed as function handles and evaluated as matrix-vector products in a specific direction, \mathbf{q} . To avoid division by zero, we set the gradient equal to zero when the denominator for indices where $\mathcal{F}\mathbf{o}$ is equal to zero. It is also important to scale the Fourier transformations appropriately. Lastly, we note that our code includes an adjoint test to verify the agreement of the forward and adjoint operators.

Sparse, direct methods are not a viable option when solving for the Gauss–Newton step using (10) and (11) because the 2D Fourier transform matrices in $\frac{\partial\phi}{\partial\mathbf{o}}$ and its adjoint are dense, complex matrices and should not be formed explicitly. Instead, we pass the approximate Hessians for $E_1(\mathbf{o})$ and $E_2(\mathbf{o})$ as function handles that evaluate the matrices’ product with a vector. Both matrices are symmetric, so separate implementations of their transposes are unnecessary. Many iterative methods that require only matrix-vector products can then be used to solve (7). We use the conjugate gradient method [45], [46]. We minimize the cost by solving the system to low accuracy, typically 1×10^{-1} , following the example of [39]. This makes solving (7) quite efficient. Potentially, we could increase efficiency further by introducing an appropriate preconditioner for the conjugate gradient method. This represents future work.

IV. NUMERICAL EXPERIMENTS

We now compare our proposed methods with methods found in the literature for phase recovery from the bispectrum in several experiments. First, we solve (3) and (4) to recover an object’s phase and compare results for our proposed Gauss–Newton strategy with gradient descent and L-BFGS. For these formulations of the problem, no regularization is included. We then run experiments with (5) and (6), optimizing with respect to the recovered object. Here, we compare gradient descent, projected gradient descent, L-BFGS, standard Gauss–Newton, and projected Gauss–Newton. For gradient descent, L-BFGS, and standard Gauss–Newton, we use the penalty term regularizer defined by Eq. 12 in [30] to encourage non-negativity in the solution. For the constrained problem with strict pixel-wise non-negativity in the solution image, we use projected Gauss–Newton to enforce the bound constraints and compare this to projected gradient descent, a common gradient-based approach for solving the constrained problem. For both projected methods, we test two regularization options for $R(\mathbf{o})$: a quadratic regularizer using a discrete gradient operator and a total variation regularizer. One could also consider a bound constrained approach to L-BFGS [47], but we do not explore that in this work.

We organize the experiments in the following way. We begin by discussing the setup of the phase recovery problem including simulation of the speckle data and an initial guess. This is followed by a discussion of method parameters for the optimization including line search parameters, stopping criteria, and regularization parameter selection. Then, we show the results for phase recovery when minimizing (3), (4), (5), and (6). These results

are split into two sets of experiments. The first set compares all the various optimization schemes for each of the four problem formulations for a problem with fixed parameters. This shows the utility of the proposed Gauss–Newton approach. The second set shows the robustness of the proposed Gauss–Newton approach by testing it across a range of problem parameters.

We compare the results of the various experiments by looking at several values. When comparing optimization methods, we compare the relative change in the objective function with respect to the initial guess (ROF). For an objective function $E(\cdot)$ at the iterate $\mathbf{y}^{(k)}$ with initial guess $\mathbf{y}^{(0)}$, this is given by $\frac{E(\mathbf{y}^{(k)})}{E(\mathbf{y}^{(0)})}$. We also compare the relative error of the recovered object (RE), number of optimization iterations (Its.), total optimization time (Time), and CPU time per optimization iteration (Time/It.). Lastly, we track the average number of line search iterations (inner iterations) per optimization iterate (outer iterations) during the optimization (LS/It.). The abbreviations listed in parentheses are used to reference these values in tables and figures.

We use the following notation to refer to each method in tables and figures: GD for gradient descent, PGD for projected gradient descent, L-BFGS for the limited memory Broyden–Fletcher–Goldfarb–Shanno method, GN for standard Gauss–Newton, and PGN for projected Gauss–Newton. When regularization is used for (5) and (6), we add a suffix of $+$, ∇_h , or TV in tables and figures to denote the penalty, discrete gradient, and total variation regularizers, respectively. For example, GD+ denotes gradient descent with the penalty regularizer.

A. Data Setup

To setup test problems, we simulate speckle imaging data for a known true object, \mathbf{o}_{true} . We use a 256×256 image of a satellite as the true object, which can be obtained from [48] (this test image is widely in the literature to evaluate algorithms for image restoration problems; see, e.g., [49]). We then generate 100 frames of short exposure data of the object and the reference star for a chosen Fried parameter, D/r_0 . Higher values of this parameter correspond to more atmospheric turbulence and blurrier images. The approach we use, which is implemented in the MATLAB package `IR Tools` [50] (software can be obtained from <https://github.com/jnagy1/IRtools>), is described in detail in [51]. Each data set is generated using a different seed for the random number generator to guarantee independence of the randomness of the two data sets. The object data is scaled to include 3×10^6 photo-events per data frame, and zero-mean Gaussian noise with a chosen standard deviation σ_{rn} is added. The reference star data is scaled to 5000 photo-events per frame. From this data, we recover the object’s power spectrum and accumulate the data bispectrum at a set of $(\mathbf{u}, \mathbf{v}, \mathbf{u} + \mathbf{v})$ triplets. The set of triplets is collected using selected ‘recovery radius,’ denoted by R , and a smaller radius of 5, i.e., the indices satisfy $|\mathbf{u}| \leq R$, $|\mathbf{u} + \mathbf{v}| \leq R$, and $|\mathbf{v}| \leq 5$. These triplets are stored in the indexing structure described in Sec. II and are used to generate the matrix A . The choices for the Fried parameter D/r_0 , recovery radius R , and standard deviation of the Gaussian noise σ_{rn} vary for different numerical experiments, so values for these parameters are listed in the subsections corresponding to the

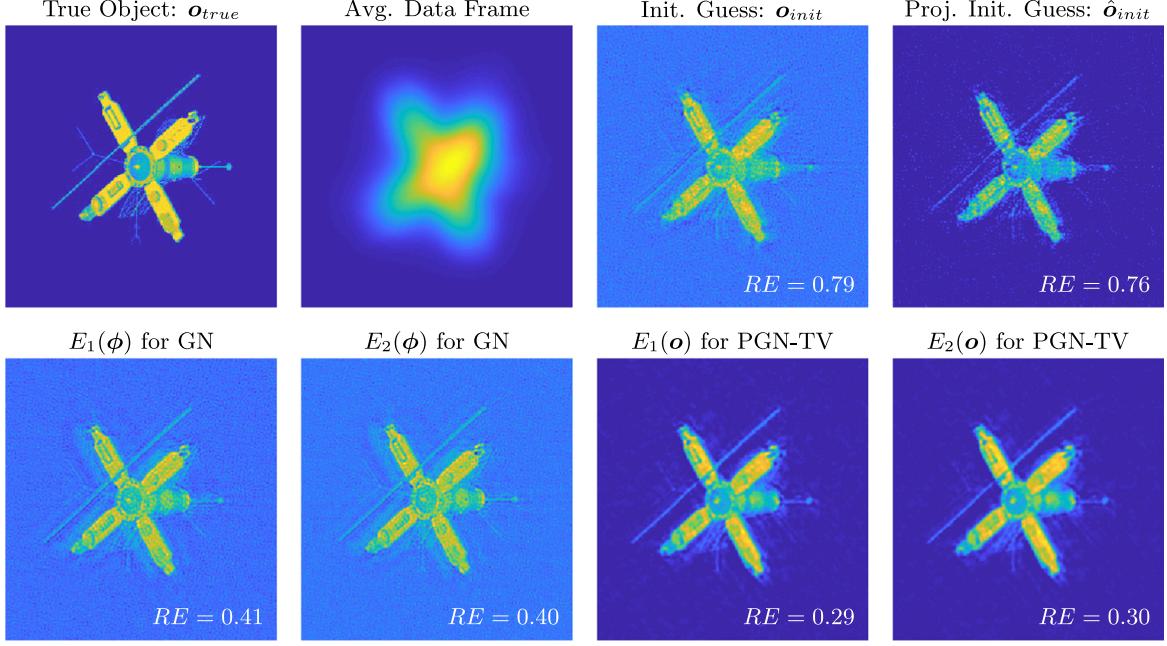


Fig. 2. The top row shows the true object, data, and initial guesses used to set up the phase recovery problem. The bottom row shows the best recovered images of the satellite. For $E_1(\phi)$ and $E_2(\phi)$, the results are shown for standard Gauss–Newton with no regularization on the solution phase. For $E_1(o)$ and $E_2(o)$, the results are shown for projected Gauss–Newton using a total variation regularizer. The relative error of the initial guesses and solution images are printed in each image’s bottom right-hand corner.

relevant set of experiments. All other parameters are fixed as listed here for all numerical experiments.

For all four formulations of the phase recovery problem, (3), (4), (5), and (6), the optimization methods are sensitive to the choice of initial guess for the object phase or image. To generate an appropriate initial guess for (3) and (4), we use the phase recovered by a single iteration of the recursive algorithm from [28]. We denote this by ϕ_{init} . For (5) and (6) in the unconstrained case, this is combined with the recovered power spectrum of the data to produce an initial guess for the recovered object, o_{init} . The relative error for this initial guess varies due to the random generation of the speckle data but is typically 0.75–0.85 depending on the problem parameters. When using projected Gauss–Newton, the initial guess for the recovered object must be within the feasible region, i.e., it must have strictly non-negative pixel intensities. To obtain a non-negative guess for the object, we project o_{init} using a method which finds the nearest non-negative image such that the sum of the pixel intensities is identical to o_{init} ; this is referred to as an energy preserving constraint, and the implementation we use can be found in the `IR Tools` package [50]. We denote this projected initial guess by \hat{o}_{init} . In practice, we also ‘bump’ this projected guess off the bound constraint by adding $\epsilon = 1 \times 10^{-4}$ to ensure the active set is empty for the first iteration of the projected Gauss–Newton method, i.e., no pixels intensities are exactly zero to start the optimization. We note that this projection changes both the power spectrum and phase of the initial guess. Another option for the projection is a single iteration of the error-reduction algorithm [52], [53]. Images for the true object, average speckle data frame, initial guess from the recursive algorithm, and projected initial guess can be seen in Fig. 2.

B. Parameter Selection

The gradient descent, projected gradient descent, L-BFGS, Gauss–Newton, and projected Gauss–Newton methods require setting several parameters. These include stopping criteria for the methods and parameters for the Armijo line search. Additionally, formulations (5) and (6) depend on the regularization parameter used for the objective functions. We discuss selections for these parameters here.

For stopping criteria, we monitor three values. The first two are the change in the objective function value and the 2-norm of the difference between successive iterates. The method stops when these two values fall below tolerances of 1×10^{-4} and 1×10^{-4} , respectively. Additionally for the Gauss–Newton and projected Gauss–Newton methods, we use an approximation to the Newton decrement as a stopping criteria [54]. The true Newton decrement monitors the reduction in the curvature of the Hessian near a minimizer and is defined by

$$(\nabla_y E(\mathbf{y})^\top \nabla_y^2 E(\mathbf{y})^{-1} \nabla_y E(\mathbf{y}))^{1/2}.$$

We approximate substituting our Gauss–Newton approximation to the Hessian, $H(\mathbf{y})$, instead of $\nabla_y^2 E(\mathbf{y})$. Note that Gauss–Newton step update is given by $\mathbf{p} = -H(\mathbf{y})^{-1} \nabla_y E(\mathbf{y})$, so we can cheaply compute the approximate Newton decrement by $(-\nabla_y E(\mathbf{y})^\top \mathbf{p})^{1/2}$. Both Gauss–Newton and projected Gauss–Newton stop if either the Newton decrement falls below a tolerance of 1×10^{-3} or if the above criteria for the objective function and difference between successive iterates are met.

The line search strategies differ for gradient descent, projected gradient descent, and L-BFGS compared to Gauss–Newton and projected Gauss–Newton. The full Newton step is perfectly

scaled, and it follows that $\eta = 1$ is a logical initial choice for the step length for Gauss–Newton and projected Gauss–Newton at each iteration when the line search is called. In contrast, the step updates for gradient descent, projected gradient descent, and L-BFGS do not have natural scaling. To account for this, we use an adaptive line search strategy. It works as follows. If the line search succeeds in its first iteration, i.e., it takes the longest step possible, then the line search doubles its initial guess for the step length at the next optimization iterate, i.e. the next time the line search is called. If the line search backtracks, i.e., it takes a shorter step than its initial guess, then it uses the accepted, shorter step as its initial guess for the step length the next time it is called. This adaptive strategy prevents the gradient-based optimization strategies from repeatedly taking initial line search steps which are too large or too small. This is important because each line search iteration requires reevaluating the objective function, which is costly. The adaptive strategy used for gradient descent, projected gradient descent, and L-BFGS helps to reduce the number of line search iterations per optimization iteration (and therefore reduce cost). This provides a fair grounds for comparison with Gauss–Newton and projected Gauss–Newton methods.

For (5) and (6), we include regularization on the recovered object. This requires selecting the regularization parameter α . Choosing this parameter is a challenging subproblem and an active area of research. Popular methods for selecting it include the unbiased predictive risk method, L-curve, and generalized cross-validation (GCV) [55], [56]. For our simulated problem where the true object is known, we solve both (5) and (6) for a wide range of parameters and choose the parameter which minimizes the relative error of the recovered image. For the unconstrained problem where the penalty-term regularizer is used, this results in $\alpha = 1 \times 10^3$ for both (5) and (6). For the constrained problem with the discrete gradient regularizer, $\alpha = 1 \times 10^{-2}$ provides the best solutions for both objective functions, and for the constrained problem with the total variation regularizer, we use $\alpha = 1 \times 10^4$. We note that these regularization parameters are sensitive to the data parameters and problem set up. Also, for real-world data where the true solution is unknown, this method for regularization parameter selection is infeasible. A more sophisticated strategy for regularization parameter selection represents future work, but by using the best case scenarios for choosing regularization parameters, we are able to provide a fair comparison between the computational efficiency of the various methods.

C. Comparison of Gauss–Newton and Gradient-Based Optimization

To compare our proposed Gauss–Newton schemes with previous gradient-based optimization approaches, we solved 50 different phase recovery problems with independently generated speckle data using all four formulations of the phase recovery problem. For these problems, we used the data parameters and problem setup outlined previously and set the Fried parameter $D/r_0 = 30$, recovery radius $R = 96$, and standard deviation of the Gaussian noise $\sigma_{rn} = 5$. Table I provides results detailing

TABLE I
OPTIMIZATION RESULTS FOR THE FOUR FORMULATIONS OF THE PHASE RECOVERY PROBLEM

$E_1(\phi)$						
Method	ROF	RE	Its.	Time	Time/It.	LS/It.
GD	0.95	0.44	100.0	9.8e0	9.8e-2	1.7
L-BFGS	0.78	0.44	100.0	2.9e1	2.9e-1	1.7
GN	0.74	0.43	24.1	3.5e0	1.5e-1	1.0
$E_2(\phi)$						
Method	ROF	RE	Its.	Time	Time/It.	LS/It.
GD	0.96	0.44	98.0	1.8e1	1.9e-1	1.7
L-BFGS	0.81	0.44	100.0	2.0e1	2.0e-1	1.7
GN	0.73	0.43	60.4	1.1e1	1.8e-1	1.0
$E_1(o)$						
Method	ROF	RE	Its.	Time	Time/It.	LS/It.
GD +	0.37	0.53	54.2	1.4e1	2.7e-1	1.8
PGD-TV	0.60	0.43	100.0	4.0e1	4.0e-1	1.7
PGD- ∇_h	0.62	0.43	100.0	2.9e1	2.8e-1	1.7
L-BFGS+	0.42	0.79	64.2	1.8e1	2.8e-1	1.9
GN+	0.39	0.39	11.0	9.5e0	8.7e-1	1.0
PGN-TV	0.68	0.28	13.7	2.5e1	1.8e0	1.1
PGN- ∇_h	0.65	0.31	29.1	5.8e1	2.0e0	1.2
$E_2(o)$						
Method	ROF	RE	Its.	Time	Time/It.	LS/It.
GD +	0.29	0.61	83.9	1.9e1	2.2e-1	1.7
PGD-TV	0.60	0.34	100.0	3.5e1	3.5e-1	1.7
PGD- ∇_h	0.62	0.38	100.0	2.1e1	2.1e-1	1.7
L-BFGS+	0.32	0.82	87.1	2.0e1	2.3e-1	1.8
GN+	0.32	0.38	12.3	1.1e1	9.1e-1	1.0
PGN-TV	0.67	0.28	14.1	3.3e1	2.3e0	1.0
PGN- ∇_h	0.66	0.30	23.1	3.9e1	1.7e0	1.0

From *left to right*, the columns show the optimization method used, minimum relative objective function (ROF), minimum relative error (RE), number of iterations (Its.), CPU time (Time), CPU time per iteration (Time/It.), and line search iterations per optimization iteration (LS/It.). All values were averaged over 50 separate simulated problems.

the cost of the optimization and quality of the resulting solutions averaged over all 50 problems. Fig. 3 compares the convergence of the various optimization schemes using the relative reduction in the objective function for a single problem. Also, the best recovered images for the single problem are displayed in Fig. 2. To comment on the results, we separate the discussion for the formulations of the problem which solve for the phase, (3) and (4), and those which solve for the resulting object, (5) and (6).

Looking at the results for $E_1(\phi)$ and $E_2(\phi)$ in Table I, we see that the Gauss–Newton method outperforms both gradient descent and L-BFGS in terms of minimizing the objective function. This results in marginally lower relative error values for the solution object for the Gauss–Newton method, and all methods offer marginal improvements on the initial guess provided by the recursive algorithm. This is true for both formulations of the problem. In terms of cost, the gradient-based methods require less time per iteration due to the additional computational cost of solving for the step direction at each iteration of the Gauss–Newton method. However, this additional cost is offset by the faster convergence of the Gauss–Newton method

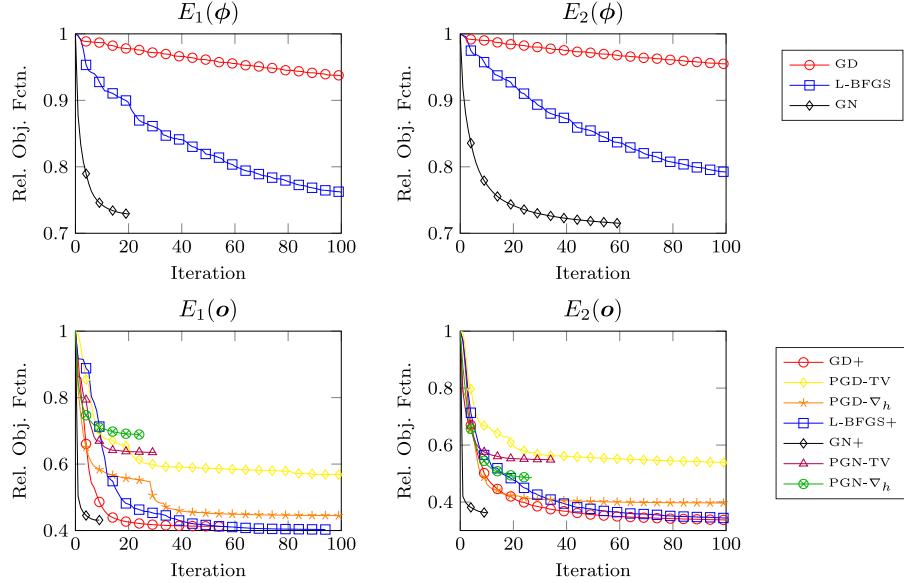


Fig. 3. This figure shows convergence plots of the relative objective function for all four formulations of the problem. Note that the Gauss–Newton approaches converge to their minima in fewer iterations than the gradient-based approaches. Also, note that when minimizing $E_1(o)$ and $E_2(o)$, projected gradient descent and projected Gauss–Newton converge to different minima than the other approaches due to the different regularizers, different regularization parameters, and constraints.

(it requires significantly fewer iterations for convergence than either gradient-based optimization scheme). This is observable in Fig. 3. Overall, the time-to-solution for (3) and (4) using Gauss–Newton is reduced for both formulations, and as mentioned above, this speed-up is coupled with better solutions in terms of relative error. This suggests that Gauss–Newton optimization is a preferable option to the gradient-based methods for these formulations of the problem.

The results for (5) and (6) differ from the two formulations above which consider only the recovered phase. Both standard Gauss–Newton with the penalty term regularizer and projected Gauss–Newton with the total variation and discrete gradient regularizers offer significant improvements in the quality of the recovered object. This can be seen from the minimum relative error values in Table I, where both Gauss–Newton approaches offer significant improvements on the relative error of the initial guess. Projected gradient descent also reduces the relative error of the recovered object for both regularization options, although not as much as the projected Gauss–Newton approach. However, the table shows that gradient descent and L-BFGS with the penalty regularizer actually worsen the relative error of the recovered image compared to the initial guess despite reducing the objective function. This emphasizes the importance of appropriate regularization and bound constraints in driving the solution to an appropriate minimum. The recovered images also show the advantages of the Gauss–Newton approaches, where the regularization terms in (5) and (6) reduce ringing and graininess in the solution images in Fig. 2 compared to the initial guess.

When comparing the different optimization methods for solving (5) and (6), we note that the different regularizers and regularization parameters used make the relative objective function a suboptimal way of comparing the methods, see Fig. 3. Instead,

we note that projected Gauss–Newton with non-negativity constraints and the total variation regularizer achieves the best solution images in terms of relative error out of all the methods. However, this improvement comes at a cost as the projected Gauss–Newton approach is the slowest of the four methods. This represents a trade-off, but the method may be preferable if quality of the recovered object is more important than time-to-solution. Also, the projected Gauss–Newton approach allows for the use of various regularizers, so an appropriate regularizer can be chosen for a specific application to optimize solution quality.

Using standard Gauss–Newton to solve (5) and (6) with the penalty term regularizer offers the best compromise of speed and quality for our problem. It improves on solution quality and time-to-solution compared to all the gradient-based methods. It is also faster than the projected Gauss–Newton approach, but the resulting images have slightly higher relative error.

Overall, we observe that for all four formulations of the problem, the standard and projected Gauss–Newton optimization schemes offer improvements in recovered solution images over the gradient descent, L-BFGS, and projected gradient descent methods. Comparing the formulations of the problem, (5) and (6) are preferable to (3) and (4) due to additional regularization and constraints on the solution image that these formulations offer. This results in improvements in the relative error of the recovered image, and the solution images in Fig. 2 are less grainy and have less ringing. In terms of cost, the faster convergence for standard Gauss–Newton schemes results in lower time-to-solution for all formulations. This is not the case when solving (5) and (6) with constraints using projected Gauss–Newton. However, the introduction of constraints and appropriate regularization using this approach gives the best-quality images in terms of relative error. This may make this strategy attractive despite the increased cost.

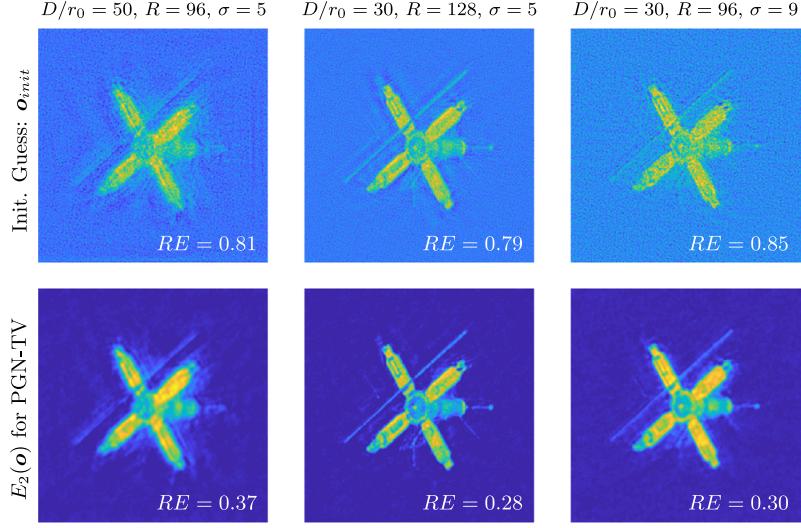


Fig. 4. The rows from *top* to *bottom* show the initial guesses given by the recursive algorithm and optimization solutions for $E_2(\mathbf{o})$ using projected Gauss–Newton with the total variation regularizer (PGN-TV). The columns from *left* to *right* indicate the method parameters used for each problem. The relative error of each image is printed in its bottom right-hand corner.

D. Robustness of Gauss–Newton Schemes

We also tested the robustness of the our proposed Gauss–Newton schemes by solving several phase recovery problems over a range of parameter values for the Fried parameter D/r_0 , recovery radius R , and standard deviation of the Gaussian noise σ_{rn} . We set up the problems in the following way. To test the performance with respect to atmospheric blur, we fixed the recovery radius $R = 96$ and standard deviation $\sigma_{rn} = 5$ and solved the problem for Fried parameters $D/r_0 = 10, 20, 30, 40$, and 50 . Here, increasing values for D/r_0 indicate more atmospheric turbulence and blurrier data. For each value of D/r_0 , we averaged results over 10 problems. We then tested the methods for a range of recovery radii, $R = 64, 80, 96, 112$, and 128 while fixing $D/r_0 = 30$ and $\sigma_{rn} = 5$. Increasing the recovery radius introduces values of the data bispectrum and recovered phase corresponding to higher spatial frequencies. This presents a tradeoff as data for higher spatial frequencies can potentially enable better recovery of the object being imaged but is more negatively affected by data noise. As with the Fried parameter, we solved 10 problems for each different value of R . Lastly, we solved the phase recovery over a range of values for the standard deviation of the Gaussian noise, $\sigma_{rn} = 1, 3, 5, 7$, and 9 for fixed $D/r_0 = 30$ and $R = 96$. This corresponds to increased noise in the problem data. As with the previous setups, we solved 10 problems for each value of σ_{rn} .

We solved the problems for each of the four objective functions using the respective Gauss–Newton approach that performed best in terms of relative error in the previous set of experiments. For (3) and (4), this was the Gauss–Newton approach using the direct permutation and factorization approach described in III, denoted GN in Table I. For (5) and (6), we used projected Gauss–Newton with the total variation regularizer, denoted PGN-TV.

We compare the results using the average relative error of the recovered object from solving each objective function across the

TABLE II
OPTIMIZATION RESULTS SHOWING THE AVERAGE RELATIVE ERROR OF THE RECOVERED OBJECT USING THE PROPOSED GAUSS–NEWTON SCHEMES FOR VARIOUS VALUES OF FRIED PARAMETER D/r_0 , RECOVERY RADIUS R , AND STANDARD DEVIATION OF THE GAUSSIAN NOISE σ_{rn}

D/r_0	\mathbf{o}_{init}	$\hat{\mathbf{o}}_{init}$	$E_1(\phi)$	$E_2(\phi)$	$E_1(\mathbf{o})$	$E_2(\mathbf{o})$
10	0.82	0.79	0.38	0.37	0.20	0.20
20	0.83	0.78	0.47	0.46	0.27	0.26
30	0.83	0.77	0.50	0.50	0.31	0.31
40	0.82	0.75	0.56	0.55	0.35	0.35
50	0.83	0.76	0.62	0.62	0.45	0.46

R	\mathbf{o}_{init}	$\hat{\mathbf{o}}_{init}$	$E_1(\phi)$	$E_2(\phi)$	$E_1(\mathbf{o})$	$E_2(\mathbf{o})$
64	0.77	0.76	0.35	0.35	0.28	0.28
80	0.76	0.75	0.39	0.38	0.27	0.27
96	0.79	0.76	0.44	0.43	0.30	0.29
112	0.80	0.75	0.47	0.47	0.30	0.30
128	0.82	0.77	0.51	0.52	0.32	0.31

σ_{rn}	\mathbf{o}_{init}	$\hat{\mathbf{o}}_{init}$	$E_1(\phi)$	$E_2(\phi)$	$E_1(\mathbf{o})$	$E_2(\mathbf{o})$
1	0.75	0.74	0.29	0.29	0.29	0.28
3	0.75	0.73	0.34	0.33	0.27	0.27
5	0.79	0.76	0.43	0.43	0.29	0.28
7	0.81	0.77	0.51	0.50	0.30	0.30
9	0.85	0.78	0.58	0.57	0.31	0.31

From *left* to *right*, the columns show the value of the problem parameter being tested and the relative errors for the initial guess guess, projected initial guess, solution to $E_1(\phi)$ using GN, solution to $E_2(\phi)$ using GN, solution to $E_1(\mathbf{o})$ minimized using PGN-TV, and the solution to $E_2(\mathbf{o})$ minimized using PGN-TV. All relative error values are averaged over 10 separate problems.

range problem parameters. This is contrasted with the average relative error of the initial guess, \mathbf{o}_{init} , from the recursive algorithm and its projection, $\hat{\mathbf{o}}_{init}$. These relative error values are displayed in Table II with the best relative error results highlighted. The performance of the optimization (reduction of the relative objective function, iterations, and time-to-solution, etc.) is similar to the previous set of experiments, so we omit

these values. We display the initial guesses and solution images minimizing (6) using projected Gauss–Newton with the total variation regularizer for three separate problems in Fig. 4. As expected, the results show that the average relative error of the recovered objects increases as the problems become more difficult, i.e., blurrier or noisier data. Also, phase recovery using the nonlinear least-squares formulations results in better solutions than the initial guesses based on the recursive algorithm. This is likely due to the recursive algorithm’s limited ability to recover high spatial frequency information and is evident over the range of problem parameters. The results also reinforce the conclusion that the non-negativity constraints and regularization options available for (5) and (6) make these formulations of the phase recovery problem preferable to (3) and (4), especially if the quality of the recovered object is paramount. This is well illustrated by comparing the initial guesses and solution images using projected Gauss–Newton with the total variation regularizer in Fig. 4.

V. CONCLUSION

This paper revisits multiple formulations for phase recovery from the bispectrum, a central problem in speckle interferometry. The formulations considered lead to weighted nonlinear least-squares problems which can be solved for either the phase or the object itself. Previous approaches in the literature focused on gradient-based optimization schemes for solving these least-squares problem, including gradient descent and L-BFGS.

In this work, we implement efficient Gauss–Newton schemes to solve the phase recovery problem. We implement these schemes for four formulations of the problem taken from the literature. To reduce the computational cost of solving the linear system associated with the Gauss–Newton step, we develop tailored approaches for each formulation which exploit the structure and sparsity of the problem. For two formulations of the problem which solve for the solution image, we also extend standard Gauss–Newton to projected Gauss–Newton to allow for element-wise non-negativity constraints on the solution object, a desirable characteristic in many imaging applications. Additionally, the non-negativity constraints within the optimization allows flexibility to introduce additional regularization on the problem. We show this by running numerical experiments with both total variation and discrete gradient regularizers, but the formulation allows for other options.

Our numerical experiments show that our implementations offer improvements in terms of time-to-solution and quality of solution compared to previously used gradient-based approaches. Our Gauss–Newton schemes produce recovered objects with lower relative error and reduce image artifacts in the resulting images. Furthermore, our implementations achieve these results in less time than the previously implemented first-order methods in most cases. In the cases where Gauss–Newton optimization is more expensive than gradient-based approaches, the increased cost is offset by significantly better solutions. We also show that the improvements in solution quality are robust across a range of problem parameters, particularly for formulations (5) and (6) where the introduction of bound constraints and appropriate

regularization helps to introduce desirable qualities into the solutions of phase retrieval problems.

Several directions for future work also present themselves. Many recent phase retrieval applications from the signal processing literature involve solving large nonlinear least-squares problems like the ones considered in this paper. It would be interesting to extend the ideas presented in this work to those applications. Also, any adaptation of this work for real-world applications would benefit from a more in-depth discussion of regularization parameter selection and appropriate problem-specific regularizers.

Lastly, the codes and methods in this paper are available publicly on Github in the BiBox repository for other researchers looking to use bispectral imaging as part of their work.

ACKNOWLEDGMENT

The authors would like to thank Dr. Brandoch Calef at the Maui High Performance Computing Center for helpful advice and guidance.

REFERENCES

- [1] A. Labeyrie, “Attainment of diffraction limited resolution in large telescopes by Fourier analyzing speckle patterns in star images,” *Astron. Astrophys.*, vol. 6, pp. 85–87, 1970.
- [2] K. T. Knox and B. J. Thompson, “Recovery of images from atmospherically degraded short-exposure photographs,” *Astrophys. J.*, vol. 193, pp. L45–L48, 1974.
- [3] G. Weigelt, “Modified speckle interferometry: Speckle masking,” *Opt. Commun.*, vol. 21, pp. 55–59, 1977.
- [4] Y. Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, “Phase retrieval with application to optical imaging: A contemporary overview,” *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 87–109, May 2015.
- [5] A. W. Lohmann, G. Weigelt, and B. Wirnitzer, “Speckle masking in astronomy: Triple correlation theory and applications,” *Appl. Opt.*, vol. 22, no. 24, pp. 4028–4037, 1983.
- [6] A. W. Lohmann and B. Wirnitzer, “Triple correlations,” *Proc. IEEE*, vol. 72, no. 7, pp. 889–901, Jul. 1984.
- [7] P. Negrete-Regagnon, “Practical aspects of image recovery by means of the bispectrum,” *JOSA A*, vol. 13, no. 7, pp. 1557–1576, 1996.
- [8] B. Wirnitzer, “Bispectral analysis at low light levels and astronomical speckle masking,” *JOSA A*, vol. 2, no. 1, pp. 14–21, 1985.
- [9] G. E. Archer, J. P. Bos, and M. C. Roggemann, “Comparison of bispectrum, multiframe blind deconvolution and hybrid bispectrum-multiframe blind deconvolution image reconstruction techniques for anisoplanatic, long horizontal-path imaging,” *Opt. Eng.*, vol. 53, no. 4, 2014, Art. no. 043109.
- [10] J. P. Bos and M. C. Roggemann, “Mean squared error performance of speckle-imaging using the bispectrum in horizontal imaging applications,” in *Proc. SPIE*, vol. 8056, 2011, Art. no. 805603.
- [11] C. J. Carrano, “Speckle imaging over horizontal paths,” in *Proc. SPIE*, in *High-Resolution Wavefront Control: Methods, Devices, and Applications IV*, vol. 48252002, pp. 109–121.
- [12] Z. Y. Wen, D. Fraser, A. Lambert, and H. D. Li, “Reconstruction of underwater image by bispectrum,” in *Proc. IEEE Image Process., Int. Conf.*, vol. 3, 2007, pp. III–545.
- [13] Z. Y. Wen, A. Lambert, D. Fraser, and H. D. Li, “Bispectral analysis and recovery of images distorted by a moving water surface,” *Appl. Opt.*, vol. 49, no. 33, pp. 6376–6384, 2010.
- [14] T. Bendory, R. R. Beinert, and Y. C. Eldar, “Fourier phase retrieval: Uniqueness and algorithms,” in *Compressed Sensing and Its Applications*. Berlin, Germany: Springer, 2017, pp. 55–91.
- [15] T. Bendory, P. Sodorenko, and Y. C. Eldar, “On the uniqueness of frog methods,” *IEEE Signal Process. Lett.*, vol. 24, no. 5, pp. 722–726, May 2017.
- [16] T. Bendory, Y. C. Eldar, and N. Boumal, “Non-convex phase retrieval from STFT measurements,” *IEEE Trans. Inf. Theory*, vol. 64, no. 1, pp. 467–484, Jan. 2018.

[17] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer, "Bispectrum inversion with application to multireference alignment," *IEEE Trans. Signal Process.*, vol. 66, no. 4, pp. 1037–1050, Feb. 2018.

[18] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Convex algorithms and limitations," in *Proc. IEEE Int. Symp. Inf. Theory*, 2013, pp. 1022–1026.

[19] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Uniqueness guarantees and recovery algorithms," *IEEE Trans. Signal Process.*, vol. 65, no. 9, pp. 2402–2410, May 2017.

[20] S. Pinilla, J. Bacca, and H. Arguello, "Phase retrieval algorithm via nonconvex minimization using a smoothing function," *IEEE Trans. Signal Process.*, vol. 66, no. 17, pp. 4574–4584, Sep. 2018.

[21] A. Repetti, E. Chouzenoux, and J. C. Pesquet, "A nonconvex regularized approach for phase retrieval," in *Proc. IEEE Int. Conf. Image Process.*, 2014, pp. 1753–1757.

[22] G. Wang, L. Zhang, G. B. Giannakis, M. Akçakaya, and J. Chen, "Sparse phase retrieval via truncated amplitude flow," *IEEE Trans. Signal Process.*, vol. 66, no. 2, pp. 479–491, Jan. 2018.

[23] D. J. Cho, E. A. Watson, and G. M. Morris, "Application of bispectral speckle imaging to near-diffraction-limited imaging in the presence of unknown aberrations," *Appl. Opt.*, vol. 34, no. 5, pp. 854–864, 1995.

[24] M. G. Kang, K. T. Lay, and A. K. Katsaggelos, "Image restoration algorithms based on the bispectrum," in *Proc. SPIE, in Visual Commun. and Image Process. '91: Image Process.*, vol. 1606, 1991, pp. 408–419.

[25] C. L. Matson, "Weighted-least-squares phase reconstruction from the bispectrum," *JOSA A*, vol. 8, no. 12, pp. 1905–1913, 1991.

[26] J. Meng, G. J. M. Aitken, E. K. Hege, and J. S. Morgan, "Triple-correlation subplane reconstruction of photon-address stellar images," *JOSA A*, vol. 7, no. 7, pp. 1243–1250, 1990.

[27] M. J. Northcott, G. R. Ayers, and J. C. Dainty, "Algorithms for image reconstruction from photon-limited data using the triple correlation," *JOSA A*, vol. 5, no. 7, pp. 986–992, 1988.

[28] D. Tyler and K. Schulze, "Fast phase spectrum estimation using the parallel part-bispectrum algorithm," *Pub. Astron. Soc. Pacific*, vol. 116, no. 815, pp. 65–76, 2004.

[29] C. A. Haniff, "Least-squares Fourier phase estimation from the modulo 2π bispectrum phase," *JOSA A*, vol. 8, no. 1, pp. 134–140, 1991.

[30] A. Glindemann and J. C. Dainty, "Object fitting to the bispectral phase by using least squares," *JOSA A*, vol. 10, no. 5, pp. 1056–1063, 1993.

[31] H. Takajo and T. Takahashi, "Least-squares phase estimation from the phase difference," *JOSA A*, vol. 5, no. 3, pp. 416–425, 1988.

[32] A. T. Erdem and M. I. Sezan, "Least-squares reconstruction of an image from its noisy observations using the bispectrum," in *Proc. IEEE Sixth SP Workshop Statistical Signal Array Process.*, 1992, pp. 156–159.

[33] D. M. Goodman, T. W. Lawrence, J. P. Fitch, and E. M. Johansson, "Bispectral-based optimization algorithms for speckle imaging," in *Proc. SPIE, in Digit. Image Synthesis and Inverse Opt.*, vol. 1351, 1990, pp. 546–561.

[34] J. C. Marron, P. P. Sanches, and R. C. Sullivan, "Unwrapping algorithm for least-squares phase recovery from the modulo 2π bispectrum phase," *JOSA A*, vol. 7, no. 1, pp. 14–20, 1990.

[35] D. F. Buscher and C. A. Haniff, "Diffraction-limited imaging with partially redundant masks: II. optical imaging of faint sources," *JOSA A*, vol. 10, no. 9, pp. 1882–1894, 1993.

[36] T. Nakajima, "Signal-to-noise ratio of the bispectral analysis of speckle interferometry," *JOSA A*, vol. 5, no. 9, pp. 1477–1491, 1988.

[37] M. C. Roggeman and C. L. Matson, "Power spectrum and fourier phase spectrum estimation by using fully and partially compensating adaptive optics and bispectrum postprocessing," *JOSA A*, vol. 9, no. 9, pp. 1525–1535, 1992.

[38] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1–4, pp. 259–268, 1992.

[39] J. Nocedal and S. Wright, *Numerical Optimization*. Berlin, Germany: Springer, 1999.

[40] A. Beck, *Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications With MATLAB*. Philadelphia, PA, USA: SIAM, 2014.

[41] E. Haber, *Computational Methods in Geophysical Electromagnetics*. Philadelphia, PA, USA: SIAM, 2014.

[42] T. A. Davis, J. R. Gilbert, S. I. Larimore, and E. G. Ng, "Algorithm 836: COLAMD, a column approximate minimum degree ordering algorithm," *ACM Trans. Math. Softw. (TOMS)*, vol. 30, no. 3, pp. 377–380, 2004.

[43] T. A. Manteuffel, "An incomplete factorization technique for positive definite linear systems," *Math. Comput.*, vol. 34, no. 150, pp. 473–497, 1980.

[44] Y. Saad, *Iterative Methods for Sparse Linear Systems*, vol. 82. Philadelphia, PA, USA: SIAM, 2003.

[45] Å. Björck, *Numerical Methods for Least Squares Problems*. Philadelphia, PA, USA: SIAM, 1996.

[46] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 4th ed. Baltimore, MD, USA: John Hopkins University Press, 2013.

[47] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM J. Scientific Comput.*, vol. 16, no. 5, pp. 1190–1208, 1995.

[48] J. G. Nagy, K. M. Palmer, and L. Perrone, "Iterative methods for image deblurring: A Matlab object oriented approach," *Numer. Algorithms*, vol. 36, pp. 73–93, 2004.

[49] M. C. Roggeman and B. M. Welsh, *Imaging Through Turbulence*. Boca Raton, FL, USA: CRC-Press, 1996.

[50] S. Gazzola, P. C. Hansen, and J. G. Nagy, "IR Tools: A MATLAB package of iterative regularization methods and large-scale test problems," *Numer. Algorithms* [Online]. Available: <https://github.com/jnagy1/IRtools>, doi: 10.1007/s11075-018-0570-7.

[51] J. D. Schmid, *Numerical Simulation of Optical Wave Propagation With Examples in MATLAB*. Bellingham, WA, USA: SPIE Press, 2010.

[52] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, 1982.

[53] P. Negrete-Regagnon, "Phase recovery from the bispectrum aided by the error-reduction algorithm," *Opt. Lett.*, vol. 21, no. 4, pp. 275–277, 1996.

[54] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[55] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems* (SIAM Monographs on Mathematical Modeling and Computation). Philadelphia, PA, USA: SIAM, 1998.

[56] C. R. Vogel, *Computational Methods for Inverse Problems*. Philadelphia, PA, USA: SIAM, 2002.



James Lincoln Herring received the undergraduate and doctorate degrees from Emory University, Atlanta, GA, USA, in 2010 and 2018, respectively. He is a Postdoctoral Researcher with the Department of Mathematics at the University of Houston in Houston, Texas, USA. His research focuses on applied mathematics and scientific computing applications in medical and astronomical imaging with interests in numerical solutions to inverse problems, numerical methods for optimization, optimal control, and machine learning.



James Nagy received the Ph.D. degree in applied mathematics from North Carolina State University, in 1991. He is a Professor of Mathematics and Computer Science with Emory University, Atlanta, GA, USA. His research interests include numerical linear algebra, structured matrix computations, numerical solution of inverse problems, and image processing. Dr. Nagy is a Fellow of the Society for Industrial and Applied Mathematics (SIAM). Before joining Emory University in 1999 he had Postdoctoral research fellowships with the Institute for Mathematics and Its Applications (IMA) at the University of Minnesota, Minneapolis, MN, USA, with the National Science Foundation (NSF) at the University of Maryland, College Park, MD, USA, and was on the faculty at Southern Methodist University, Dallas, TX, USA. He is on the editorial boards of *SIAM Journal on Scientific Computing (SISC)*, *SIAM Journal on Matrix Analysis and Applications (SIMAX)*, *Numerical Algorithms*, and the *Electronic Journal of Linear Algebra (ELA)*.



Lars Ruthotto received the diploma and Ph.D. degree in mathematics from the University of Mnster, in 2010 and 2012, respectively. He is an Assistant Professor of Mathematics and Computer Science with Emory University, Atlanta, GA, USA. His research interests include numerical analysis (particularly numerical methods for optimization, linear algebra, and partial differential equations) and scientific computing with applications in medical and geophysical imaging and machine learning.