

Non-Convex Phase Retrieval From STFT Measurements

Tamir Bendory, Yonina C. Eldar, *Fellow, IEEE*, and Nicolas Boumal

Abstract—The problem of recovering a one-dimensional signal from its Fourier transform magnitude, called Fourier phase retrieval, is ill-posed in most cases. We consider the closely-related problem of recovering a signal from its phaseless short-time Fourier transform (STFT) measurements. This problem arises naturally in several applications, such as ultra-short laser pulse characterization and ptychography. The redundancy offered by the STFT enables unique recovery under mild conditions. We show that in some cases the unique solution can be obtained by the principal eigenvector of a matrix, constructed as the solution of a simple least-squares problem. When these conditions are not met, we suggest using the principal eigenvector of this matrix to initialize non-convex local optimization algorithms and propose two such methods. The first is based on minimizing the empirical risk loss function, while the second maximizes a quadratic function on the manifold of phases. We prove that under appropriate conditions, the proposed initialization is close to the underlying signal. We then analyze the geometry of the empirical risk loss function and show numerically that both gradient algorithms converge to the underlying signal even with small redundancy in the measurements. In addition, the algorithms are robust to noise.

Index Terms—Phase retrieval, short-time Fourier transform, non-convex optimization, spectral initialization, least-squares, ptychography, ultra-short laser pulse characterization, optimization on manifolds.

I. INTRODUCTION

THE problem of recovering a signal from its Fourier transform magnitude arises in many areas in engineering and science, such as optics, X-ray crystallography, speech recognition, blind channel estimation, alignment and astronomy [5], [9], [27], [33], [41], [46], [66], [72]. This problem is called *Fourier phase retrieval* and can be viewed as a special case of a quadratic system of equations. The latter area received considerable attention recently, partially due to its strong connections with the fields of compressed sensing and

Manuscript received July 27, 2016; revised May 14, 2017; accepted July 5, 2017. Date of publication August 29, 2017; date of current version December 20, 2017. This work was supported in part by the European Union's Horizon 2020 Research and Innovation Program under Grant 646804-ERCCOG-BNYQ, in part by the Israel Science Foundation under Grant 335/14, and in part by the Research Grant from the Ollendorff Fund. T. Bendory was supported in part by Andrew and Erna Finci Viterbi Fellowship. N. Boumal was supported by NSF under Grant DMS-1719558.

T. Bendory and Y. C. Eldar are with the Department of Electrical Engineering, Technion Israel Institute of Technology, Haifa 3200003, Israel (e-mail: tamir.bendory@gmail.com).

N. Boumal is with the Department of Mathematics, Princeton University, Princeton, NJ 08544 USA.

Communicated by H. Rauhut, Associate Editor for Signal Processing.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2017.2745623

matrix completion; see for instance [15]–[17], [19], [25], [60], [64], [65]. Contemporary surveys of the phase retrieval problem from a signal processing point of view can be found in [8], [37], and [59]. Phase retrieval for one-dimensional (1D) signals is an ill-posed problem unless the signal has the minimum phase property [34], [56]. In this special case, the signal can be recovered by several tractable algorithms (see for instance of [21, Sec. 2.6]). Particularly, in [34] it was shown that a semidefinite program (SDP) relaxation achieves the optimal solution in the least-squares (LS) sense. For general signals, two main approaches are typically suggested. The first builds upon prior knowledge on the signal's support, such as sparsity or a portion of the underlying signal [28], [39], [53], [58], [60], [68]. An alternative strategy makes use of additional measurements. Such measurements can be obtained by structured illuminations and masks [15], [17], [32] or by measuring the magnitude of the short-time Fourier transform (STFT) [26], [38]. In [26], it was demonstrated that for the same number of measurements, the STFT magnitude leads to better performance than an over-sampled discrete Fourier transform (DFT).

This paper deals with the problem of recovering a 1D signal from its STFT magnitude. The STFT of a 1D signal $\mathbf{x} \in \mathbb{C}^N$ can be interpreted as the Fourier transform of the signal multiplied by a real sliding window $\mathbf{g} \in \mathbb{R}^N$ with support size W and is defined as

$$\mathbf{X}[m, k] := \sum_{n=0}^{N-1} \mathbf{x}[n] \mathbf{g}[mL - n] e^{-2\pi jkn/N}, \quad (1.1)$$

where $k = 0, \dots, N-1$, $m = 0, \dots, \lceil \frac{N}{L} \rceil - 1$ and L determines the separation in time between adjacent sections. The pseudo-inverse of the STFT is given by

$$\mathbf{x}[n] = \frac{\sum_{m=0}^{\lceil \frac{N}{L} \rceil - 1} \tilde{\mathbf{x}}[m, n] \overline{\mathbf{g}[mL - n]}}{\sum_{m=0}^{\lceil \frac{N}{L} \rceil - 1} |\mathbf{g}[mL - n]|^2}, \quad (1.2)$$

where $\tilde{\mathbf{x}}[m, n]$ is the inverse DFT of $\mathbf{X}[m, k]$ for fixed m with respect to the second variable [26]. In the sequel, all indices should be considered as modulo the signal's length N . We assume that \mathbf{x} and \mathbf{g} are periodically extended over the boundaries in (1.1).

The problem of recovering a signal from its STFT magnitude $|\mathbf{X}[m, k]|^2$, frequently called spectrogram, arises in several applications in optics and speech processing [31], [47]. Particularly, it serves as the model for a popular variant of an ultra-short laser pulse characterization technique

called Frequency-Resolved Optical Gating (referred to as X-FROG) [6], [10], [63]. Another important application is ptychography in which a moving probe is used to sense multiple diffraction measurements [44], [45], [54], [70].

Several algorithms were suggested to recover a signal from the magnitude of its STFT. The classical method, proposed by Griffin and Lim [31], is a modification of the alternating projection (or reduction error) algorithms of Gerchberg and Saxton [29] and Fienup [28]. The properties of this method are not well understood (for analysis of alternating projection algorithms in phase retrieval, see [45], [50], [65]). In [38], the authors prove that a non-vanishing signal can be recovered by an SDP with maximal overlap between adjacent windows ($L = 1$). They also demonstrate empirically that the algorithm works well with less restrictive requirements on the window and is robust to noise. Despite the appealing numerical performance, solving an SDP requires high computational resources. Recently, an interesting recovery approach was proposed in [52]. This paper suggests a multi-stage method, based on spectral clustering and phase synchronization. It is shown that the algorithm achieves stable estimation (and exact in the noise-free setting) with only $\mathcal{O}(N \log N)$ phaseless STFT measurements. However, this technique requires a random window of length $W = N$, while in most applications it is common to work with shorter windows. Another line of works [35], [36] suggest applying a phase synchronization framework [3], [4], [14], [21], [51], [61]. It was shown that even for short windows, the sought signal can be recovered exactly and efficiently by spectral and greedy techniques. These methods are accompanied by stability guarantees. Their main drawback is that they rely on reliable estimates of the temporal magnitude, which do not always exist.

Here, we take a different approach and propose a data-driven initialization technique, followed by non-convex gradient algorithms. We begin by taking the 1D DFT of the acquired data with respect to the frequency variable (the second variable of the STFT). This transformation reveals the underlying structure of the data and greatly simplifies the analysis. As a direct consequence, we show that for $L = 1$ and sufficiently long windows $W \geq \lceil \frac{N+1}{2} \rceil$ (and some mild additional conditions), one can recover the signal by extracting the principal eigenvector of a designed matrix, constructed as the solution of a simple linear LS problem. We refer to this matrix as the *approximation matrix* since it approximates the correlation matrix $\mathbf{X} := \mathbf{x}\mathbf{x}^*$.

When the conditions for a closed-form solution are not met, we propose using the principal eigenvector of the approximation matrix to initialize two non-convex algorithms. The first is based on minimizing a standard quadratic loss function, frequently called the empirical risk (ER). Inspired by the phasecut method [30], [64], we also propose a new phase retrieval algorithm, called Non-Convex PhaseCut (NCPC), that maximizes a quadratic function over the set of phases. Each step of the algorithm follows the component of the gradient which agrees with the phase constraints. As will be shown, the ER technique is more stable in the low signal-to-noise ratio (SNR) regimes, while NCPC is superior in high SNR environments and for short windows. Our approach

deviates in two important aspects from the recent line of work in non-convex phase retrieval [18], [22], [48], [57], [65], [67], [71]. First, all these papers focus their attention on the setup of phase retrieval with random sensing vectors and rely heavily on probabilistic considerations. In this case, efficient algorithms were designed to estimate the signal from $\mathcal{O}(N)$ measurements. In contrast, we consider a deterministic framework. Second, we construct our approximation matrix by the solution of a LS problem, whereas the aforementioned papers take a superposition of the measurements to approximate \mathbf{X} .

The properties of non-convex algorithms depend heavily on the initialization method and the geometry of the loss functions. For $L = 1$, we estimate the distance between the proposed initialization and the target signal, which decays to zero as W tends to $\frac{N+1}{2}$. If the signal has unit modulus entries, then a slight modification of our initialization recovers the signal exactly for $W \geq 2$. In the later case, we also prove the existence of a *basin of attraction* around the global minimum of the ER loss function and estimate its size. In the basin of attraction, the algorithm is guaranteed to converge to a global minimum at a geometric rate. We note that while the theoretical guarantees of the algorithms are limited, their experimental performance is significantly better. Particularly, the algorithms perform well with small redundancy in the measurements and are robust in the presence of noise.

The paper is organized as follows. We begin in Section II by formulating mathematically the problem of phase retrieval from STFT magnitude measurements. In Section III we discuss the uniqueness of the solution and present conditions under which it has a closed-form LS expression. Additionally, we present a method that recovers signals with unit modulus entries under mild conditions. Section IV presents the two non-convex algorithms with the proposed initialization. Section V shows numerical results and Section VI presents our theoretical findings regarding the proposed initialization and the ER loss function. Proofs are provided in Section VII. Section VIII concludes the paper, discusses its main implications and draws potential future research directions.

Throughout the paper we use the following notation. Boldface small and capital letters denote vectors and matrices, respectively. We use \mathbf{Z}^T and \mathbf{Z}^* for the transpose and Hermitian of a matrix \mathbf{Z} ; similar notation is used for vectors. We further use \mathbf{Z}^\dagger and $\text{tr}(\mathbf{Z})$ for the Moore–Penrose pseudo-inverse and the trace of the matrix \mathbf{Z} , respectively. The ℓ th circular diagonal of a matrix \mathbf{Z} is denoted by $\text{diag}(\mathbf{Z}, \ell)$. Namely, $\text{diag}(\mathbf{Z}, \ell)$ is a column vector with entries $\mathbf{Z}[i, (i + \ell) \bmod N]$ for $i = 0, \dots, N - 1$. We define the sign of a complex number a as $\text{phase}(a) := \frac{a}{|a|}$ for $a \neq 0$ and zero otherwise. We also use \odot , \circ' and \ast' for the Hadamard (point-wise) product, composition of functions and convolution, respectively. The set of all complex (real) signals of length N whose entries have modulus $a > 0$ are denoted by \mathbb{C}_a^N (\mathbb{R}_a^N). Namely, $\mathbf{z} \in \mathbb{C}_a^N$ means that $|\mathbf{z}[n]| = a$ for all n .

II. PROBLEM FORMULATION

We aim at recovering an underlying signal $\mathbf{x} \in \mathbb{C}^N$ from the magnitude of its STFT, i.e., from measurements

$$\mathbf{Z}[m, k] = |\mathbf{X}[m, k]|^2. \quad (\text{II.1})$$

Note that the signals \mathbf{x} and $\mathbf{x}e^{j\phi}$ yield the same measurements for any *global phase* $\phi \in \mathbb{R}$ and therefore the phase ϕ cannot be recovered by any method. This *global phase ambiguity* leads naturally to the following definition:

Definition 1: The distance between two vectors is defined as

$$d(\mathbf{z}, \mathbf{x}) = \min_{\phi \in [0, 2\pi)} \left\| \mathbf{z} - \mathbf{x}e^{j\phi} \right\|_2.$$

If $d(\mathbf{z}, \mathbf{x}) = 0$ then we say that \mathbf{x} and \mathbf{z} are equal *up to global phase*. The phase $\phi \in [0, 2\pi)$ attaining the minimum is denoted by $\phi(\mathbf{z})$, i.e.,

$$\phi(\mathbf{z}) = \arg \min_{\phi \in [0, 2\pi)} \left\| \mathbf{z} - \mathbf{x}e^{j\phi} \right\|_2.$$

In the sequel, we make use of the notion of *non-vanishing signals*, defined as follows:

Definition 2: A vector $\mathbf{z} \in \mathbb{C}^N$ is called non-vanishing if $\mathbf{z}[n] \neq 0$ for all $n = 0, \dots, N-1$.

Instead of treating the measurements (II.1) directly, we often consider the acquired data in a transformed domain by taking its 1D DFT with respect to the frequency variable (normalized by $1/N$). Then, our measurement model reads

$$\begin{aligned} \mathbf{Y}[m, \ell] &= \frac{1}{N} \sum_{k=0}^{N-1} \mathbf{Z}[m, k] e^{-2\pi j k \ell / N} \\ &= \sum_{n=0}^{N-1} \mathbf{x}[n] \overline{\mathbf{x}[n + \ell]} \mathbf{g}[mL - n] \overline{\mathbf{g}[mL - n - \ell]}. \end{aligned} \quad (\text{II.2})$$

When $W \leq \ell \leq N - W$, we have $\mathbf{Y}[m, \ell] = 0$ for all m . In this sense, $\mathbf{Y}[m, \ell]$ can be interpreted as a “ W – bandlimited” function. Observe that for fixed m , $\mathbf{Y}[m, \ell]$ is simply the auto-correlation of $\mathbf{x} \odot \mathbf{g}_{mL}$, where $\mathbf{g}_{mL} := \{\mathbf{g}[mL - n]\}_{n=0}^{N-1}$.

We will make repetitive use of several representations of the data. The first is based on a matrix formulation. Let $\mathbf{D}_{mL} \in \mathbb{R}^{N \times N}$ be a diagonal matrix composed of the entries of \mathbf{g}_{mL} . Let \mathbf{P}_ℓ be a matrix that shifts (circularly) the entries of a vector by ℓ locations, namely, $(\mathbf{P}_\ell \mathbf{x})[n] = \mathbf{x}[n + \ell]$. Then, the correlation matrix $\mathbf{X} := \mathbf{x}\mathbf{x}^*$ is mapped linearly to $\mathbf{Y}[m, \ell]$ as follows:

$$\begin{aligned} \mathbf{Y}[m, \ell] &= (\mathbf{D}_{mL-\ell} \mathbf{D}_{mL} \mathbf{P}_\ell \mathbf{x})^* \mathbf{x} \\ &= \mathbf{x}^* \mathbf{H}_{m, \ell} \mathbf{x} \\ &= \text{tr}(\mathbf{X} \mathbf{H}_{m, \ell}), \end{aligned} \quad (\text{II.3})$$

where

$$\mathbf{H}_{m, \ell} := \mathbf{P}_{-\ell} \mathbf{D}_{mL} \mathbf{D}_{mL-\ell}. \quad (\text{II.4})$$

Observe that $\mathbf{P}_\ell^T = \mathbf{P}_{-\ell}$ and $\mathbf{H}_{m, \ell} = 0$ for $W \leq \ell \leq N - W$. Similarly, the STFT magnitude in (II.1) (before the 1D DFT) can be written as

$$\mathbf{Z}[m, k] = \mathbf{x}^* \tilde{\mathbf{H}}_{m, k} \mathbf{x}, \quad (\text{II.5})$$

where

$$\tilde{\mathbf{H}}_{m, k} := \mathbf{D}_{mL} \mathbf{f}_k \mathbf{f}_k^* \mathbf{D}_{mL}, \quad (\text{II.6})$$

and \mathbf{f}_k^* is the k th row of the DFT matrix.

An alternative useful representation of (II.2) is as multiple systems of linear equations. For fixed $\ell \in \{-(W-1), \dots, W-1\}$ we have

$$\mathbf{y}_\ell = \mathbf{G}_\ell \mathbf{x}_\ell, \quad (\text{II.7})$$

where $\mathbf{y}_\ell := \{\mathbf{Y}[m, \ell]\}_{m=0}^{N-1}$ and $\mathbf{x}_\ell := \text{diag}(\mathbf{X}, \ell)$. The (m, n) th entry of the matrix $\mathbf{G}_\ell \in \mathbb{R}^{\lfloor \frac{N}{L} \rfloor \times N}$ is given by $\mathbf{g}[mL - n] \mathbf{g}[mL - n - \ell]$. For $L = 1$, \mathbf{G}_ℓ is a circulant matrix. We recall that a circulant matrix is diagonalized by the DFT matrix, namely, it can be factored as $\mathbf{G}_\ell = \mathbf{F}^{-1} \Sigma_\ell \mathbf{F}$, where \mathbf{F} is the DFT matrix and Σ_ℓ is a diagonal matrix, whose entries are given by the DFT of the first column of \mathbf{G}_ℓ . In this case, the first column is given by $\mathbf{g} \odot (\mathbf{P}_{-\ell} \mathbf{g})$. Therefore the matrix \mathbf{G}_ℓ is invertible if and only if the DFT of $\mathbf{g} \odot (\mathbf{P}_{-\ell} \mathbf{g})$ is non-vanishing.

Our problem of recovering \mathbf{x} from the measurements (II.1) can therefore be posed as a constrained LS problem:

$$\begin{aligned} \min_{\tilde{\mathbf{X}} \in \mathcal{H}^N} \sum_{\ell=-(W-1)}^{W-1} \left\| \mathbf{y}_\ell - \mathbf{G}_\ell \text{diag}(\tilde{\mathbf{X}}, \ell) \right\|_2^2 \\ \text{subject to } \tilde{\mathbf{X}} \succeq 0, \quad \text{rank}(\tilde{\mathbf{X}}) = 1, \end{aligned} \quad (\text{II.8})$$

where \mathcal{H}^N is the set of all Hermitian matrices of size $N \times N$. In the spirit of [15], [30], [60], [64], STFT phase retrieval may then be relaxed to a tractable SDP by dropping the rank constraint. In the noiseless case, this SDP relaxation is equivalent to the one suggested in [38] since the conditions on $\tilde{\mathbf{X}}$ to achieve zero objective function are the same, up to a Fourier transformation. While the SDP relaxation technique has shown good numerical performance for the recovery from phaseless STFT measurements, it requires solving the problem in a lifted domain with N^2 variables. We take a different route to reduce the computational load. In the next section, we show that (II.8) admits a unique solution under moderate conditions. We further show that it has a closed-form LS solution when the window \mathbf{g} is sufficiently long. If the conditions for the LS solution are not met, then we suggest two non-convex algorithms. To initialize the algorithms, we approximate (II.8) in two stages by first solving the LS objective function and then extracting its principal eigenvector.

III. UNIQUENESS AND BASIC ALGORITHMS

A fundamental question in phase retrieval problems is whether the quadratic measurement operator of (II.1), or equivalently the non-convex problem (II.8), determines the underlying signal \mathbf{x} uniquely (up to global phase, see Definition 1). In other words, one wants to know the conditions on the window \mathbf{g} and the signal \mathbf{x} such that the non-linear transformation that maps \mathbf{x} to \mathbf{Z} is injective. Before treating this question, we introduce some basic window definitions:

Definition 3: A window \mathbf{g} is called a *rectangular window of length W* if $\mathbf{g}[n] = 1$ for all $n = 0, \dots, W-1$ and zero elsewhere. It is a *non-vanishing window of length W* if $\mathbf{g}[n] \neq 0$ for all $n = 0, \dots, W-1$ and zero elsewhere.

According to (II.8), the injectivity of the measurement operator is related to the window’s length W and the invertibility of

TABLE I

THE MEAN AND MINIMAL VALUES OF $|\lambda_{\min}|$ FOR WINDOWS OF LENGTH W WITH I.I.D. NORMAL ENTRIES AS DEFINED IN (III.1) OVER 100 EXPERIMENTS FOR DIFFERENT WINDOW LENGTHS AND $N = 25$

Window's length \ $ \lambda_{\min} $	$W = 5$	$W = 10$	$W = 15$	$W = 20$
Mean	0.0463	0.0367	0.0426	0.0549
Min	0.0008	0.0021	0.0019	0.0031

the matrices \mathbf{G}_ℓ for $|\ell| < W$. For that reason, we give special attention to windows for which the associated matrices are invertible.

Definition 4: A window \mathbf{g} is called an *admissible window of length W* if for all $\ell = -(W-1), \dots, W-1$ the following two equivalent properties hold:

- 1) The DFT of the vector $\mathbf{g} \odot (\mathbf{P}_{-\ell}\mathbf{g})$ is non-vanishing.
- 2) The associated circulant matrices \mathbf{G}_ℓ as given in (II.7) are invertible.

An important example for an admissible window is a rectangular window. Specifically, we have the following lemma:

Lemma 5: A rectangular window \mathbf{g} of length $2 \leq W \leq N/2$ is an admissible window of length W if α and N are co-prime numbers for all $\alpha = 2, \dots, W$. This holds trivially when N is a prime number.

Proof: Observe that $\mathbf{g} \odot (\mathbf{P}_{-\ell}\mathbf{g})$ is a rectangular window of length $W - |\ell|$ for $\ell = -(W-1), \dots, W-1$. The DFT of a rectangular window of size $W - |\ell|$ is a Dirichlet kernel which is non-vanishing if $W - |\ell|$ and N are co-prime. ■

The family of admissible windows contains more examples. To demonstrate this, we consider a non-vanishing window of length W whose entries are i.i.d. normal variables. We then compute the minimal absolute value of the DFT of $\mathbf{g} \odot (\mathbf{P}_{-\ell}\mathbf{g})$ for all $\ell = -(W-1), \dots, W-1$, namely,

$$|\lambda_{\min}| = \min_{k, |\ell| \leq W} |(\mathbf{F}(\mathbf{g} \odot (\mathbf{P}_{-\ell}\mathbf{g}))) [k]|. \quad (\text{III.1})$$

We repeated this process 100 times for several values of W . As can be seen in Table I, $|\lambda_{\min}|$ is bounded away from zero, implying that the windows are indeed admissible.

We now analyze the uniqueness of the measurement operator for the case $L = 1$. Uniqueness results for $L > 1$ are discussed in [38], [47]. Our results are constructive in the sense that their proofs provide an explicit scheme to recover the signal.

Our first result concerns non-vanishing signals. In this case, the magnitude of the STFT determines the underlying signal uniquely under mild conditions. This conclusion was already derived in [11] based on different considerations. Nevertheless, the following proposition comes with an explicit recovery scheme as presented in Appendix A.

Proposition 6: Let $L = 1$. Suppose that \mathbf{x} is non-vanishing and that the DFT of $\mathbf{g} \odot (\mathbf{P}_{-\ell}\mathbf{g})$ is non-vanishing for $\ell = 0, 1$. Then, $|\mathbf{X}[m, k]|^2$ determines \mathbf{x} uniquely (up to global phase).

Proof: See Appendix A. ■

A similar uniqueness result was derived in [26]. There, it is required that the DFT of $|\mathbf{g}[n]|^2$ is non-vanishing, $N \geq 2W - 1$ and N and $W - 1$ are co-prime numbers.

In the special case in which the signal is known to have unit modulus entries, the signal can be recovered as the principal eigenvector of a matrix designed as follows:

Proposition 7: Let $L = 1$. Suppose that $\mathbf{x} \in \mathbb{C}^N_{1/\sqrt{N}}$ and that \mathbf{g} is an admissible window of length $W \geq 2$. Fix $M \in \{1, \dots, W-1\}$ and let \mathbf{X}_0 be a matrix defined by

$$\text{diag}(\mathbf{X}_0, \ell) = \begin{cases} \mathbf{G}_\ell^{-1} \mathbf{y}_\ell, & \ell = 0, M, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{III.2})$$

where \mathbf{G}_ℓ and \mathbf{y}_ℓ are defined in (II.7). Then, \mathbf{x} (up to global phase) is a principal eigenvector of \mathbf{X}_0 .

Proof: See Appendix B. ■

For general signals (not necessarily non-vanishing) and $L = 1$, we next derive a LS algorithm that stably recovers any complex signal if the window is sufficiently long. In the absence of noise, the recovery is exact (up to global phase). The method, summarized in Algorithm 1, is based on constructing a matrix \mathbf{X}_0 that approximates the correlation matrix $\mathbf{X} := \mathbf{x}\mathbf{x}^*$. The ℓ th diagonal of \mathbf{X}_0 is chosen as the solution of the LS problem $\min_{\tilde{\mathbf{x}} \in \mathbb{C}^N} \|\mathbf{y}_\ell - \mathbf{G}_\ell \tilde{\mathbf{x}}\|_2$ (see (II.7)). If the matrix \mathbf{G}_ℓ is invertible, then

$$\text{diag}(\mathbf{X}_0, \ell) = \mathbf{G}_\ell^{-1} \mathbf{y}_\ell = \text{diag}(\mathbf{X}, \ell).$$

Therefore, when all matrices \mathbf{G}_ℓ are invertible, $\mathbf{X}_0 = \mathbf{X}$. In order to estimate \mathbf{x} , the (unit-norm) principal eigenvector of \mathbf{X}_0 is normalized by

$$\alpha = \frac{1}{\sqrt{\sum_{n \in P} (\mathbf{G}_0^\dagger \mathbf{y}_0)[n]}}, \quad (\text{III.3})$$

where $P := \{n : (\mathbf{G}_0^\dagger \mathbf{y}_0)[n] > 0\}$. If \mathbf{G}_0 is invertible then

$$\sum_{n=0}^{N-1} (\mathbf{G}_0^{-1} \mathbf{y}_0)[n] = \sum_{n=0}^{N-1} (\text{diag}(\mathbf{X}, 0))[n] = \|\mathbf{x}\|_2^2 = \lambda_0,$$

where λ_0 is the top eigenvalue of \mathbf{X} . If \mathbf{G}_0 is not invertible or in the presence of noise, some terms of the vector $\mathbf{G}_0^\dagger \mathbf{y}_0$ might be negative. In this case, we estimate $\|\mathbf{x}\|_2$ by summing only the positive terms (the set P in (III.3)). Note that all matrix inversions can be performed efficiently using the FFT due to the circulant structure of \mathbf{G}_ℓ .

The following proposition shows that Algorithm 1 recovers the underlying signal for $L = 1$ if the window is sufficiently long and satisfies some additional technical conditions. In [11], an equivalent uniqueness result was derived but without providing an algorithm. Algorithm 1 is equivalent to the discretized version of Wigner deconvolution that was suggested previously without theoretical analysis in [55] and [69].

Proposition 8: Let $L = 1$ and suppose that \mathbf{g} is an admissible window of length $W \geq \lceil \frac{N+1}{2} \rceil$ (see Definition 4). Then, Algorithm 1 recovers any complex signal uniquely up to global phase.

Proof: See Appendix C. ■

In many cases, the window is shorter than $\lceil \frac{N+1}{2} \rceil$ so that (II.8) may not admit a closed-form LS solution. In these cases, we propose two non-convex recovery algorithms. The first is a standard ER minimization that seems to work well in low SNR regimes. The second maximizes a quadratic function

Algorithm 1 Least-Squares Algorithm for $L = 1$ **Input:** The measurements $\mathbf{Z}[m, k]$ as given in (II.1).**Output:** \mathbf{x}_0 : estimation of \mathbf{x} .

- 1) Compute $\mathbf{Y}[m, \ell]$, the 1D DFT with respect to the second variable of $\mathbf{Z}[m, k]$ as given in (II.2).
- 2) Construct a matrix \mathbf{X}_0 such that

$$\text{diag}(\mathbf{X}_0, \ell) = \begin{cases} \mathbf{G}_\ell^\dagger \mathbf{y}_\ell & \ell = -(W-1), \dots, (W-1), \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathbf{G}_\ell \in \mathbb{R}^{N \times N}$ are defined in (II.7).

- 3) Let \mathbf{x}_p be the principal (unit-norm) eigenvector of \mathbf{X}_0 . Then,

$$\mathbf{x}_0 = \sqrt{\sum_{n \in P} (\mathbf{G}_0^\dagger \mathbf{y}_0)[n] \mathbf{x}_p},$$

where $P := \{n : (\mathbf{G}_0^\dagger \mathbf{y}_0)[n] > 0\}$.

over the manifold of phases. This approach shows superior performance for short windows and high SNR. In order to initialize these algorithms, we use the same LS-based method of Algorithm 1. However, for short windows we cannot estimate $\text{diag}(\mathbf{X}, \ell)$ for $\ell = W, \dots, N - W$ as the matrices \mathbf{G}_ℓ are simply zero. Nonetheless, we will show by both theoretical results and numerical experiments that under appropriate conditions, the principal eigenvector of the approximation matrix \mathbf{X}_0 , with appropriate normalization, is a good initial estimator of \mathbf{x} .

IV. LOCAL NON-CONVEX ALGORITHMS

In this section we present our main algorithmic approach to recover a signal from its STFT magnitude (II.1). First, we propose two non-convex gradient algorithms to estimate the signal. As the problem is inherently non-convex, we then suggest a systematic, data-driven, technique for initialization. This non-convex approach for STFT phase retrieval is summarized in Algorithm 2. The code for all algorithms is publicly available at <http://webee.technion.ac.il/Sites/People/YoninaEldar>.

Algorithm 2 Non-Convex Approach for STFT Phase Retrieval**Input:** The measurements $\mathbf{Z}[m, k]$ as given in (II.1).**Output:** $\hat{\mathbf{x}}$: estimation of \mathbf{x} .

- 1) **Initialization:** Apply Algorithm 1 (for $L = 1$) or Algorithm 5 (for $L > 1$).
- 2) **Refinement:** Use the output of stage 1 to initialize a gradient algorithm that minimizes the empirical risk (Section IV-A) or the Non-Convex Phase-Cut (NCPC) of Algorithm 3.

A. Empirical Risk Minimization

Recall that the STFT magnitude can be written as $\mathbf{Z}[m, k] = \mathbf{x}^* \tilde{\mathbf{H}}_{m,k} \mathbf{x}$, where $\tilde{\mathbf{H}}$ is given in (II.6). Alternatively, by taking

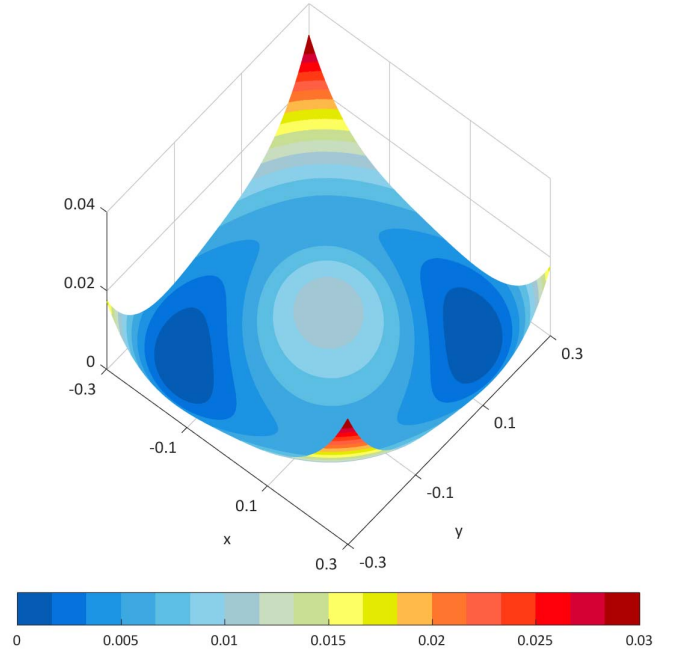


Fig. 1. The two-dimensional (first two variables) plane of the loss function (IV.2) of the signal $\mathbf{x} = [0.2, 0.2, 0, 0, 0]$ (i.e., $N = 5$) with $L = 1$ and a rectangular window of length $W = 2$.

the 1D DFT with respect to the frequency variable, the measurement model becomes $\mathbf{Y}[m, \ell] = \mathbf{x}^* \mathbf{H}_{m,\ell} \mathbf{x}$, where $\mathbf{H}_{m,\ell}$ is defined in (II.4). It is therefore natural to minimize the empirical risk (ER) loss function:

$$f(\mathbf{u}) = \frac{1}{2} \sum_{m=0}^{\lceil N/L \rceil - 1} \sum_{k=0}^{N-1} |\mathbf{u}^* \tilde{\mathbf{H}}_{m,k} \mathbf{u} - \mathbf{Z}[m, k]|^2 \quad (\text{IV.1})$$

$$= \frac{1}{2} \sum_{m=0}^{\lceil N/L \rceil - 1} \sum_{\ell=-(W-1)}^{W-1} |\mathbf{u}^* \mathbf{H}_{m,\ell} \mathbf{u} - \mathbf{Y}[m, \ell]|^2. \quad (\text{IV.2})$$

The equality between the two loss functions is proven in Appendix D. In the sequel, we use both formulations.

Figure 1 presents the two-dimensional (first two variables) plane of the loss function (IV.1) for the signal $\mathbf{x} = [0.2, 0.2, 0, 0, 0]$ (i.e., $N = 5$) with $L = 1$ and a rectangular window of length $W = 2$. The function has no sharp transitions and contains two saddle points and two global minima (as a result of the global phase ambiguity). Accordingly, in this specific case and bearing in mind that our view is restricted to two of the five dimensions only, it seems that a gradient descent algorithm will converge to a global minimum from almost any initialization (see also [43]). While this phenomenon does not occur for any arbitrary parameter selection, this example motivates applying a gradient algorithm directly on the non-convex loss function (for a similar demonstration of the loss function with random sensing vectors, see [62]).

One way to minimize the ER loss function (IV.1) or (IV.2) is by employing a gradient algorithm, where the k th iteration takes on the form

$$\mathbf{x}_k = \mathbf{x}_{k-1} - \mu \nabla f(\mathbf{x}_{k-1}),$$

for step size μ . For real signals, direct computation of the gradient in (IV.1) gives

$$\nabla f(\mathbf{u}) = \sum_{m=0}^{\lfloor N/L \rfloor - 1} \sum_{k=0}^{N-1} (h(\mathbf{u}) - \mathbf{Z}[m, k]) \nabla h(\mathbf{u}), \quad (\text{IV.3})$$

where

$$h(\mathbf{u}) := \mathbf{u}^T \tilde{\mathbf{H}}_{m,k} \mathbf{u}, \quad \nabla h(\mathbf{u}) = 2\tilde{\mathbf{H}}_{m,k} \mathbf{u}.$$

Similar computations can be performed for (IV.2). If the signal is complex, then one can use the elegant formulation of Wirtinger derivatives, see [18], [42], [62]. The loss functions (IV.1) or (IV.2) may be minimized by many other methods. For instance, in Section V we employ a trust-region algorithm.

B. Non-Convex PhaseCut (NCPC)

1) *The Algorithm:* When minimizing the empirical risk (IV.1) or (IV.2), the unknown signal itself is the optimization variable. Alternatively, we may take the point of view that the unknowns are the phases of the STFT measurements. Indeed, if these phases were known, then one could recover the signal by applying (I.2). We may therefore rework the problem into one where only the phases are variables [64].

Thus, we aim to estimate $\mathbf{U} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N}$ with unit-modulus entries such that $\mathbf{X} \approx \mathbf{Z}^{1/2} \odot \mathbf{U}$, that is, we wish to recover the missing phases. One may propose to estimate these and the signal \mathbf{x} simultaneously by minimizing $\|\mathbf{Z}^{1/2} \odot \mathbf{U} - \text{STFT}(\mathbf{x})\|_{\mathbb{F}}^2$ over both \mathbf{x} and \mathbf{U} , where $\text{STFT}(\mathbf{x})$ maps \mathbf{x} to its STFT following (I.1). Assuming \mathbf{U} is fixed, the solution for \mathbf{x} is $\mathbf{x} = \text{STFT}^\dagger(\mathbf{Z}^{1/2} \odot \mathbf{U})$, where the operator STFT^\dagger is given by (I.2). By substitution, we obtain an optimization problem in terms of \mathbf{U} only:

$$\begin{aligned} \min_{\mathbf{U} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N}} \quad & \|(\mathbf{I} - \text{STFT} \circ \text{STFT}^\dagger)(\mathbf{Z}^{1/2} \odot \mathbf{U})\|_{\mathbb{F}}^2 \\ \text{subject to} \quad & |\mathbf{U}[m, k]| = 1, \quad \forall m, k. \end{aligned}$$

Since $\mathbf{I} - \text{STFT} \circ \text{STFT}^\dagger$ is an orthogonal projector, this further simplifies into the following non-convex optimization problem over complex phases:

$$\begin{aligned} \min_{\mathbf{U} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N}} \quad & \left\langle \mathbf{Z}^{1/2} \odot \mathbf{U}, (\mathbf{I} - \text{STFT} \circ \text{STFT}^\dagger)(\mathbf{Z}^{1/2} \odot \mathbf{U}) \right\rangle \\ \text{subject to} \quad & |\mathbf{U}[m, k]| = 1, \quad \forall m, k, \end{aligned}$$

where we use the Frobenius inner product

$$\langle \mathbf{A}, \mathbf{B} \rangle = \Re \{ \text{Trace}(\mathbf{A}^* \mathbf{B}) \}. \quad (\text{IV.4})$$

The term involving the identity operator \mathbf{I} is constant under the constraints, so that the problem is equivalent to the following maximization problem:

$$\begin{aligned} \max_{\mathbf{U} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N}} \quad & \left\langle \mathbf{Z}^{1/2} \odot \mathbf{U}, \text{STFT} \circ \text{STFT}^\dagger(\mathbf{Z}^{1/2} \odot \mathbf{U}) \right\rangle \\ \text{subject to} \quad & |\mathbf{U}[m, k]| = 1, \quad \forall m, k. \end{aligned} \quad (\text{IV.5})$$

Notice that $\text{STFT} \circ \text{STFT}^\dagger$ is the orthogonal projector onto the subspace of matrices which are the STFT of some signal.

As a result, applying $\text{STFT} \circ \text{STFT}^\dagger$ to the matrix $\mathbf{Z}^{1/2} \odot \mathbf{U}$ produces the matrix which, in the LS sense, is closest to being the STFT of a signal. Thus, the cost function in (IV.5) favors phases \mathbf{U} such that $\mathbf{Z}^{1/2} \odot \mathbf{U}$ is as close as possible to an STFT. We recall that this projection operator can be computed efficiently by applying (I.1) and (I.2) using FFT.

Problem (IV.5) resembles the phase synchronization problem [3], [14], [61]. In [64], the authors pursue a convex relaxation of (IV.5) named *phasecut*. Here, following [14], we use the Manopt toolbox to run local optimization of (IV.5) over the manifold of phases [13]. In its simplest form, the algorithm follows the gradient's component which is consistent with the feasible set of solutions (see details below). To initialize the local optimization algorithm, we set \mathbf{U}_0 to be the phases of $\text{STFT}(\mathbf{x}_0)$, where \mathbf{x}_0 is the initialization used by Algorithm 2. This approach is summarized in Algorithm 3.

Algorithm 3 Non-Convex PhaseCut (NCPC)

Input: The measurements $\mathbf{Z} \approx |\text{STFT}(\mathbf{x})|^2$ as given in (II.1).

Output: $\hat{\mathbf{x}}$: estimation of \mathbf{x} .

- 1) Compute the initialization \mathbf{x}_0 with Algorithm 2 to obtain $\mathbf{U}_0 = \text{phase}(\text{STFT}(\mathbf{x}_0))$.
- 2) Using \mathbf{U}_0 as initialization, use a local optimization algorithm to try to compute a solution $\hat{\mathbf{U}}$ to

$$\begin{aligned} \max_{\mathbf{U} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N}} \quad & \left\langle \mathbf{Z}^{1/2} \odot \mathbf{U}, \text{STFT} \circ \text{STFT}^\dagger(\mathbf{Z}^{1/2} \odot \mathbf{U}) \right\rangle \\ \text{subject to} \quad & |\mathbf{U}[m, k]| = 1, \quad \forall m, k. \end{aligned}$$

See Algorithm 4 for a simple Riemannian gradient method; see [1], [13] for Riemannian trust regions.

- 3) Return $\hat{\mathbf{x}} = \text{STFT}^\dagger(\mathbf{Z}^{1/2} \odot \hat{\mathbf{U}})$.
-

For completeness, we provide a brief overview of step 2 of Algorithm 3, that is, optimization of the phases. We restrict attention to a simple Riemannian optimization algorithm, namely, the gradient ascent algorithm. See [1] for details about the more sophisticated Riemannian trust-region method (RTR), which we use in practice.

The variable \mathbf{U} lives on a smooth manifold, namely, the set of phases

$$\mathcal{M} = \{ \mathbf{U} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N} : |\mathbf{U}[m, k]| = 1 \text{ for all } m, k \},$$

which is a Cartesian product of unit circles in the complex plane (a torus). This smooth nonlinear space can be linearized about every point \mathbf{U} by differentiating the constraints. This yields a linear subspace known as the tangent space to \mathcal{M} at \mathbf{U} :

$$\mathbf{T}_{\mathbf{U}}\mathcal{M} = \{ \dot{\mathbf{U}} \in \mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N} : \Re\{\overline{\mathbf{U}} \odot \dot{\mathbf{U}}\} = 0 \}.$$

Each tangent space of \mathcal{M} can be endowed with the inner product (IV.4) (simply by restricting it to each particular subspace), which turns \mathcal{M} into a Riemannian submanifold of $\mathbb{C}^{\lfloor \frac{N}{L} \rfloor \times N}$. This makes it particularly easy to compute the gradient of the objective function $f: \mathcal{M} \rightarrow \mathbb{R}$,

$$f(\mathbf{U}) = \left\langle \mathbf{Z}^{1/2} \odot \mathbf{U}, \text{STFT} \circ \text{STFT}^\dagger(\mathbf{Z}^{1/2} \odot \mathbf{U}) \right\rangle. \quad (\text{IV.6})$$

Indeed, following [2, eq. (3.37)], the gradient of f at \mathbf{U} restricted to \mathcal{M} —known as the Riemannian gradient $\text{grad } f(\mathbf{U})$ —is the orthogonal projection of the classical (unconstrained) gradient of f , denoted by $\nabla f(\mathbf{U})$, to the tangent space $\text{T}_{\mathbf{U}}\mathcal{M}$:

$$\begin{aligned}\nabla f(\mathbf{U}) &= 2\mathbf{Z}^{1/2} \odot \left(\text{STFT} \circ \text{STFT}^\dagger (\mathbf{Z}^{1/2} \odot \mathbf{U}) \right), \\ \text{grad } f(\mathbf{U}) &= \text{Proj}_{\mathbf{U}}(\nabla f(\mathbf{U})).\end{aligned}\quad (\text{IV.7})$$

The orthogonal projector $\text{Proj}_{\mathbf{U}}: \mathbb{C}^{\lceil \frac{N}{L} \rceil \times N} \rightarrow \text{T}_{\mathbf{U}}\mathcal{M}$ is given by

$$\text{Proj}_{\mathbf{U}}(\mathbf{V}) = \mathbf{V} - \Re\{\overline{\mathbf{U}} \odot \mathbf{V}\} \odot \mathbf{U}.$$

That is, it subtracts from each entry $\mathbf{V}[m, k]$ its component aligned with $\mathbf{U}[m, k]$. Explicitly, the Riemannian gradient is then given by

$$\text{grad } f(\mathbf{U}) = \nabla f(\mathbf{U}) - \Re\{\overline{\mathbf{U}} \odot \nabla f(\mathbf{U})\} \odot \mathbf{U}.$$

Now that we are equipped with a notion of gradient on the manifold, the only missing ingredient to implement a gradient ascent optimization algorithm is a means of moving away from a point (a current iterate) along a chosen tangent direction (here, the gradient vector), while remaining on the manifold \mathcal{M} . The standard tool to achieve this is known as a *retraction* [2, Sec.4.1]. An obvious retraction for \mathcal{M} is

$$\text{Retr}_{\mathbf{U}}(\dot{\mathbf{U}}) = \text{phase}(\mathbf{U} + \dot{\mathbf{U}}).$$

Indeed, for $\mathbf{U} \in \mathcal{M}$ and $\dot{\mathbf{U}} \in \text{T}_{\mathbf{U}}\mathcal{M}$, the result of this operation is always on \mathcal{M} and locally (that is, for small $\dot{\mathbf{U}}$) the change is along the prescribed tangent direction $\dot{\mathbf{U}}$.

The gradient ascent algorithm takes the form

$$\mathbf{U}_{k+1} = \text{Retr}_{\mathbf{U}_k}(\eta_k \text{grad } f(\mathbf{U}_k)), \quad (\text{IV.8})$$

where $\eta_k > 0$ is an appropriately chosen step size (typically using a form of line-search [2, Sec. 4]) and $\mathbf{U}_0 \in \mathcal{M}$ is a given initial guess. Owing to \mathcal{M} being a compact submanifold of $\mathbb{C}^{\lceil \frac{N}{L} \rceil \times N}$ and to f being smooth, both Riemannian gradient ascent (with appropriate line-search) and RTR are guaranteed to converge to points which satisfy first-order necessary optimality conditions, that is, $\|\text{grad } f(\mathbf{U})\| = 0$ (and even second-order conditions for RTR) regardless of initialization, with known worst-case bounds on iteration counts [12]. Explicitly, at a critical point \mathbf{U} the algorithm satisfies:

$$\mathbf{U} = \text{phase} \left(\text{STFT} \circ \text{STFT}^\dagger (\mathbf{Z}^{1/2} \odot \mathbf{U}) \right). \quad (\text{IV.9})$$

As will be shown next, this is also the stagnation point of Fienup's algorithm. This approach is summarized in Algorithm 4.

We stress that this approach is different from a projected gradient method. Indeed, in a projected gradient method, one would alternate between following the classical gradient $\nabla f(\mathbf{U})$ and projecting to \mathcal{M} with the phase operator. That is, each iteration resembles (IV.8) with ∇f instead of $\text{grad } f$. In contrast, the Riemannian gradient method follows the tangent part of the gradient, $\text{grad } f$ (IV.7) and then projects onto \mathcal{M} . One advantage is that, close to convergence, the Riemannian gradient has small norm (as expected), whereas the classical gradient may still be large.

Algorithm 4 Riemannian Gradient Method for NCPC

Input: The measurements $\mathbf{Z} \approx |\text{STFT}(\mathbf{x})|^2$ as given in (II.1), initial guess $\mathbf{U}_0 \in \mathcal{M}$ and tolerance $\varepsilon > 0$.

Output: $\hat{\mathbf{U}} \in \mathcal{M}$ satisfying $\|\text{grad } f(\hat{\mathbf{U}})\|_{\text{F}} \leq \varepsilon$.

For $k = 0, 1, \dots$

a) Compute:

$$\text{grad } f(\mathbf{U}_k) = \nabla f(\mathbf{U}_k) - \Re\{\overline{\mathbf{U}_k} \odot \nabla f(\mathbf{U}_k)\} \odot \mathbf{U}_k,$$

where

$$\nabla f(\mathbf{U}_k) = 2\mathbf{Z}^{1/2} \odot \left(\text{STFT} \circ \text{STFT}^\dagger (\mathbf{Z}^{1/2} \odot \mathbf{U}_k) \right).$$

b) If $\|\text{grad } f(\mathbf{U}_k)\|_{\text{F}} \leq \varepsilon$, return $\hat{\mathbf{U}} = \mathbf{U}_k$.

c) Compute a step size η_k with a classical line-search algorithm, e.g., [2, Sec.4].

d) Set $\mathbf{U}_{k+1} = \text{phase}(\mathbf{U}_k + \eta_k \text{grad } f(\mathbf{U}_k))$.

2) *Relation to Fienup's Algorithm:* Our method can be compared with the classical Fienup algorithm for the STFT case, also called Griffin–Lim algorithm [31], as follows. One approach to optimize (IV.5), instead of the Riemannian gradient iterations that we describe in Algorithm 4, is an iterative technique called projected power method (PPM), or generalized power method [14], [40]. This algorithm iterates as the power method, with the difference that, at each iteration, it keeps only the phases of the current iterate. Specifically, the k th iteration is of the form

$$\mathbf{U}_k = \text{phase} \left(\text{STFT} \circ \text{STFT}^\dagger \left(\mathbf{U}_{k-1} \odot \mathbf{Z}^{1/2} \right) \right). \quad (\text{IV.10})$$

Similarly to NCPC, the algorithm stops when (IV.9) is satisfied. On the other hand, each iteration of Fienup's algorithm takes on the form

$$\mathbf{x}_k = \text{STFT}^\dagger \left(\text{phase}(\text{STFT}(\mathbf{x}_{k-1})) \odot \mathbf{Z}^{1/2} \right). \quad (\text{IV.11})$$

Applying the operator $\text{phase} \circ \text{STFT}$ on the iterations of (IV.11) shows that it is equivalent to PPM through the mapping $\mathbf{U}_k = \text{phase} \circ \text{STFT}(\mathbf{x}_k)$. In this sense, one can understand Fienup's algorithm as a particular iterative method to solve the optimization problem (IV.5).

According to [14, Lemma 15], all fixed points of (IV.10)–and hence of (IV.11)–map to critical points of the optimization problem (IV.5), that is, they map to points \mathbf{U}_k where the Riemannian gradient is zero. These are only the first-order necessary optimality conditions. Numerical experiments (not displayed here) show that some of the stable fixed points of (IV.11) map to critical points which do not satisfy the second-order necessary optimality conditions (their Riemannian Hessian admits a positive eigenvalue) and are therefore suboptimal. In contrast, such points would be unstable fixed points for any reasonable Riemannian optimization algorithm as confirmed in the same experiments. This distinction at least partially explains why the empirical performance of the NCPC algorithm is superior to that of Fienup's algorithm, as demonstrated in Section V.

C. Initialization

1) *Initialization for $L = 1$* : Since the phase retrieval problem is inherently non-convex, it is not clear whether the proposed refinement algorithms will converge to a global minimum from an arbitrary initialization. When $L = 1$, we propose initializing the iterations by using Algorithm 1. As explained in Section III, for $W \geq \lceil \frac{N+1}{2} \rceil$ the algorithm returns \mathbf{x} exactly. However, when $W < \lceil \frac{N+1}{2} \rceil$, $\mathbf{G}_\ell = \mathbf{0}$ for $\ell = W, \dots, N - W$ so that the output is not necessarily \mathbf{x} . Nevertheless, in Section VI we provide theoretical guarantees establishing that under appropriate conditions, this initialization results in a good approximation.

In practical applications, a variety of approaches are used to initialize the refinement techniques. While the specific initialization method is application-dependent, these approaches can be broadly classified into two categories. The first is based on the structure of the expected signal. For instance, in some applications it is common to use a Gaussian pulse with random phases as an initial point [23]. This, however, may lead to a phenomenon called *model bias* in which the estimate tends to capture characteristics of the model rather than the true signal. An alternative strategy, also used by commercial software, is based on random initialization. This is very different from our initialization which exploits the acquired data.

2) *Initialization for $L > 1$* : Until now we focused on maximal overlap between adjacent windows, namely, $L = 1$. When $L > 1$, (II.7) results in an underdetermined system of equations since $\mathbf{y}_\ell \in \mathbb{R}^{\lceil \frac{N}{L} \rceil}$, $\mathbf{G}_\ell \in \mathbb{R}^{\lceil \frac{N}{L} \rceil \times N}$ and $\mathbf{x}_\ell \in \mathbb{R}^N$. In this case, the LS solution $\mathbf{G}_\ell^\dagger \mathbf{y}_\ell$ is the vector with minimal ℓ_2 norm among the set of feasible solutions. This approximation is quite poor in general.

We notice that the measurements \mathbf{y}_ℓ are a downsampled version by a factor L of the case of maximal overlap ($L = 1$). Therefore, we suggest upsampling \mathbf{y}_ℓ to approximate the maximal overlap setting based on the averaging nature of the window \mathbf{g} . In order to motivate our approach, we start by considering an ideal situation. Suppose that for some ℓ , the DFT of the first column of \mathbf{G}_ℓ , denoted by $\hat{\mathbf{g}}_\ell$, is an ideal low-pass with integer bandwidth N/L_{BW} . Namely,

$$\hat{\mathbf{g}}_\ell[k] = \begin{cases} 1, & k = 0, \dots, N/L_{BW} - 1, \\ 0, & \text{otherwise.} \end{cases}$$

The following lemma states that in this case, no information is lost by choosing $L = L_{BW}$ compared to taking maximal overlap $L = 1$. Moreover, it suggests to upsample the measurement vector by expansion and low-pass interpolation. Our technique resembles standard upsampling arguments in digital signal processing (DSP) (see for instance Section 4.6 of [49]).

*Lemma 9: Let $\tilde{\mathbf{g}} := \{\mathbf{g}[(-n) \bmod N]\}_{n=0}^{N-1}$. Suppose that $\tilde{\mathbf{g}} \in \mathbb{R}^N$ is an ideal low-pass with integer bandwidth N/L and $\mathbf{y} = \mathbf{g} * \mathbf{x}$ for some $\mathbf{x} \in \mathbb{C}^N$ (or equivalently, $\mathbf{y} = \mathbf{G}\mathbf{x}$, where \mathbf{G} is a circulant matrix whose first column is $\tilde{\mathbf{g}}$). Let $\mathbf{y}_L \in \mathbb{C}^{\frac{N}{L}}$ be its L -downsampled version, i.e.,*

$$\mathbf{y}_L[n] = \mathbf{y}[nL], \quad n = 0, \dots, N/L - 1.$$

Then, $\mathbf{y} = (\mathbf{F}_p^* \mathbf{F}_p) \tilde{\mathbf{y}}_L$, where

$$\tilde{\mathbf{y}}_L[n] = \begin{cases} \mathbf{y}_L[m], & n = mL, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{IV.12})$$

and \mathbf{F}_p is a partial Fourier matrix consisting of the first N/L rows of the DFT matrix \mathbf{F} .

Proof: See Appendix E. ■

While Lemma 9 shows that no information is lost using an ideal low-pass window with integer bandwidth N/L , in practice we do not use these windows. Instead, we approximate the low-pass interpolation of $\mathbf{F}_p^* \mathbf{F}_p$ as suggested in Lemma 9 by a simple smooth interpolation. This leads to better numerical results and reduces the computational complexity. In Section V we show simulations with both linear and cubic interpolations.

Following the upsampling stage, the algorithm proceeds as for $L = 1$ by extracting the principal eigenvector (with the appropriate normalization) of an approximation matrix. This initialization is summarized in Algorithm 5.

Algorithm 5 Least-Squares Initialization for $L > 1$

Input: The measurements $\mathbf{Z}[m, k]$ as given in (II.1) and a smooth interpolation filter $\mathbf{h}_L \in \mathbb{R}^N$ that approximates a low-pass filter with bandwidth $\lceil N/L \rceil$.

Output: \mathbf{x}_0 : Estimation of \mathbf{x} .

- 1) Compute $\mathbf{Y}[m, \ell]$, the 1D DFT with respect to the second variable of $\mathbf{Z}[m, k]$ as given in (II.2).
- 2) Upsampling: For each $\ell \in [-(W-1), \dots, (W-1)]$:

- a) Let $\mathbf{y}_\ell[m] := \{\mathbf{Y}[m, \ell]\}_{m=0}^{\lceil \frac{N}{L} \rceil - 1}$.
- b) Expansion:

$$\tilde{\mathbf{y}}_\ell[n] := \begin{cases} \mathbf{y}_\ell[m], & n = mL, \\ 0, & \text{otherwise.} \end{cases}$$

- c) Interpolation:

$$\bar{\mathbf{y}}_\ell = \tilde{\mathbf{y}}_\ell * \mathbf{h}_L.$$

- 3) Construct a matrix \mathbf{X}_0 such that

$$\text{diag}(\mathbf{X}_0, \ell) = \begin{cases} \mathbf{G}_\ell^\dagger \bar{\mathbf{y}}_\ell, & \ell = -(W-1), \dots, (W-1), \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathbf{G}_\ell \in \mathbb{R}^{N \times N}$ are defined as in (II.7) for $L = 1$.

- 4) Let \mathbf{x}_p be the principal (unit-norm) eigenvector of \mathbf{X}_0 . Then,

$$\mathbf{x}_0 = \sqrt{\sum_{n \in P} (\mathbf{G}_0^\dagger \mathbf{y}_0)[n]} \mathbf{x}_p,$$

where $P := \{n : (\mathbf{G}_0^\dagger \mathbf{y}_0)[n] > 0\}$.

V. NUMERICAL RESULTS

This section is devoted to numerical experiments examining the proposed non-convex algorithms. In all experiments, the underlying signal was drawn from $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I})$, where \mathbf{I} is the identity matrix. The measurements \mathbf{Z} (II.1) were contaminated with either i.i.d. additive Gaussian noise or

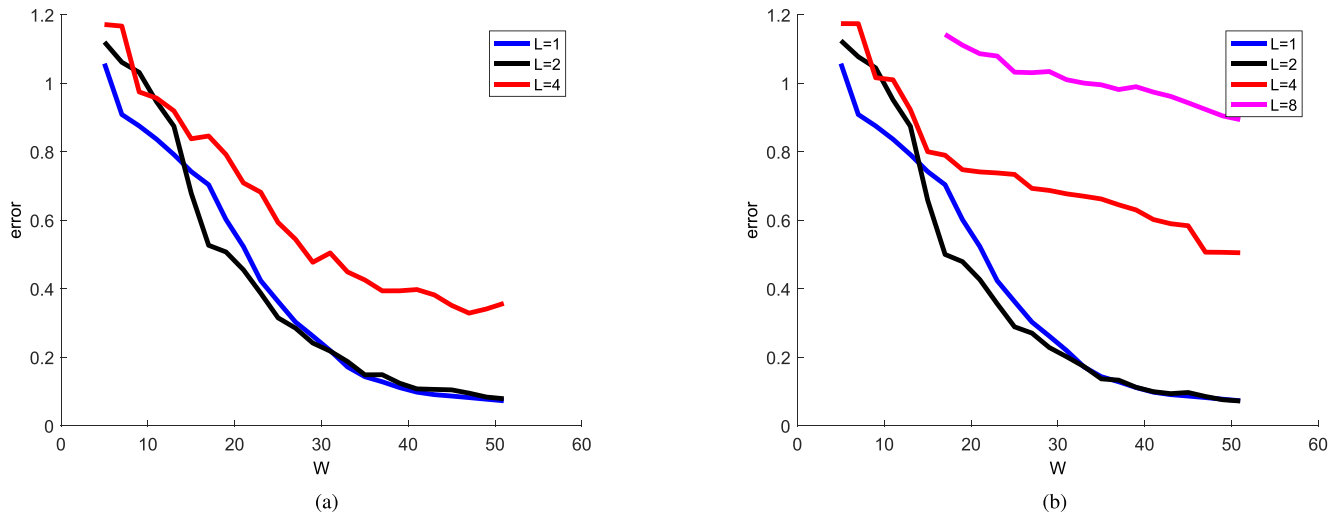


Fig. 2. Average error (over 50 experiments) of the initialization of Algorithm 5 as a function of W and L . The experiments were conducted on a signal of length $N = 101$ with a Gaussian window $e^{-\frac{n^2}{2\sigma^2}}$ and linear or cubic interpolation. The window length was set to be $W = 3\sigma$. (a) Initialization with linear interpolation. (b) Initialization with cubic interpolation.

Poisson noise. The recovery error is computed by $\frac{d(\mathbf{x}, \hat{\mathbf{x}})}{\|\mathbf{x}\|_2}$, where $\hat{\mathbf{x}}$ is the estimated signal and the distance function $d(\cdot, \cdot)$ is defined in Definition 1. We optimize both the empirical risk loss function (IV.1) and the non-convex phase-cut (NCPC) objective function by a trust-region algorithm using the Manopt toolbox [13].

The first experiment examines the estimation quality of the initialization method described in Algorithm 5. Figure 2 presents the initialization error as a function of the window's length. We considered a Gaussian window defined by $\mathbf{g}[n] = e^{-\frac{n^2}{2\sigma^2}}$ and cubic and linear interpolations. For $n > 3\sigma$, we set the entries of the window to be zero so that $W = 3\sigma$. The results demonstrate the effectiveness of the smooth interpolation technique. For low values of L , it seems that the two interpolations achieve similar performance. For larger L , namely, fewer measurements, cubic interpolation outperforms linear interpolation. In the following experiments we use cubic interpolation.

The next experiment aims to estimate the *basin of attraction* of the loss function (IV.1) or (IV.2). That is to say, the area in which a local optimization method will converge to a global minimum. To do that, we set the initialization vector to be $\mathbf{x}_0 = \mathbf{x} + \mathbf{z}$, where $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I})$ is the underlying signal. The perturbation vector \mathbf{z} takes on the values $\pm\sigma$ (with random signs) for some $\sigma > 0$ so that $d(\mathbf{x}_0, \mathbf{x}) \leq \sqrt{N}\sigma$. Then, we applied the trust-region algorithm and checked whether the algorithm converges to \mathbf{x} . As can be seen in Figure 3, the algorithm converges to the global minimum as long as $\sigma \leq 0.3$ for $L = 1, 2$ (the case of $L = 1$ is not presented in the figure) and $\sigma \leq 0.25$ for $L = 4$. These experimental results indicate that the actual basin of attraction is larger than our theoretical estimation in Section VI and Theorem 11.

Figure 4 shows a representative example of the performance of Algorithm 2 where we minimized the empirical risk loss function (IV.1) or (IV.2). The experiment was conducted on

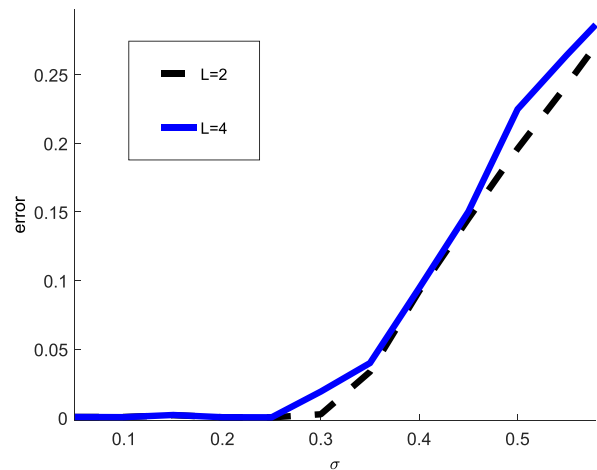


Fig. 3. Average recovery error (over 100 experiments) of minimizing the ER loss function (IV.1) or (IV.2) for signals of length $N = 43$ and a rectangular window of length $W = 11$. The initialization was set as $\mathbf{x}_0 = \mathbf{x} + \mathbf{z}$, where \mathbf{x} is the underlying signal and the perturbation vector \mathbf{z} takes the values of $\pm\sigma$ for some $\sigma > 0$ where the sign is drawn randomly.

a signal of length $N = 23$ with a rectangular window in a noisy environment of SNR= 20 dB.

Figure 5 presents the success rate of the algorithms as a function of the window's length in a noise-free environment. As can be seen, NCPC achieves the highest success rate, implying that it requires less redundancy in the data. Figure 6 presents the recovery error for different noise models. Figure 6a shows the error when the measurements are contaminated with normal noise as a function of the SNR level. The proposed algorithms are compared with Fienup's method [31] that iterates according to (IV.11). In the low SNR regime, minimizing the ER loss function seems to be better. Figure 6b shows the error with Poisson noise as a function of W . For short windows, NCPC works best. The performance for longer windows is comparable for all algorithms.

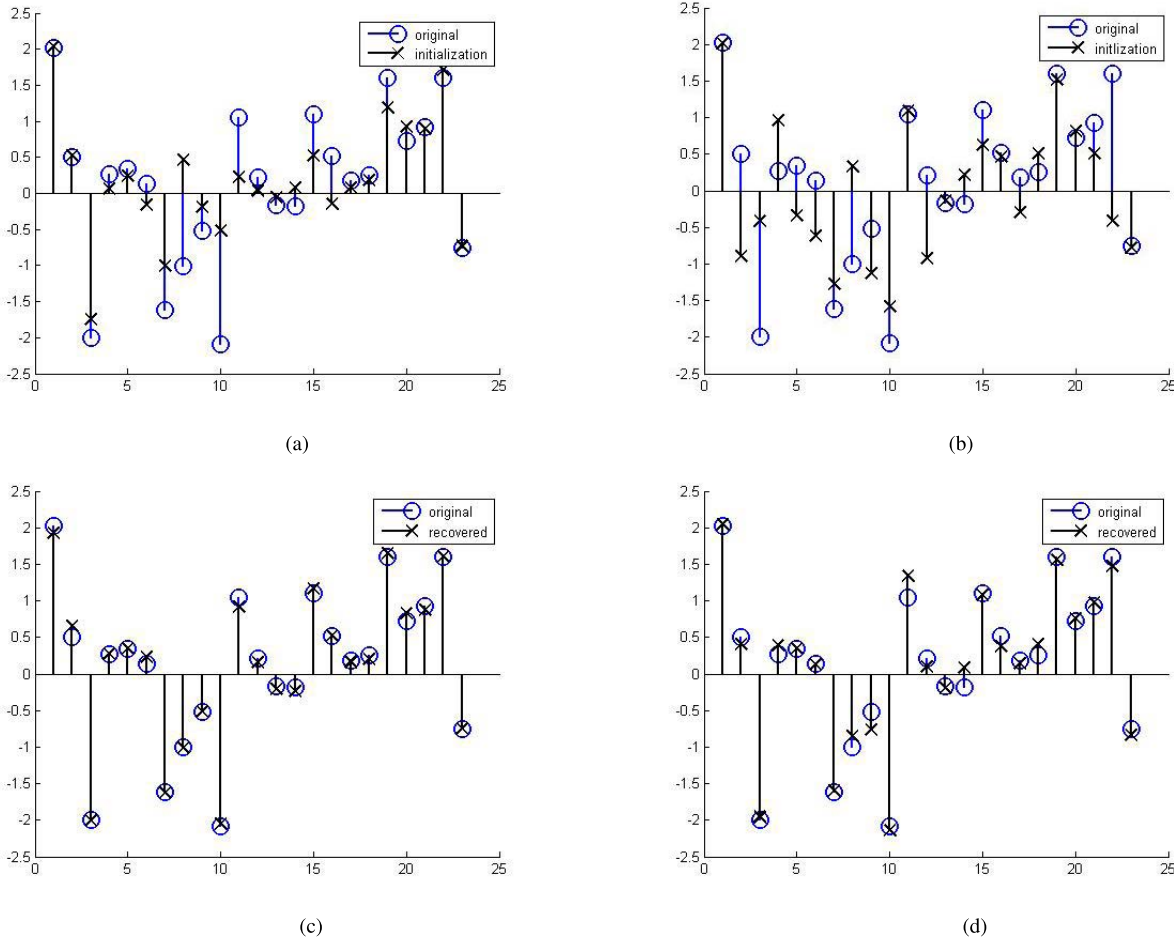


Fig. 4. Recovery of a signal of length $N = 23$ with a rectangular window in a noisy environment of $\text{SNR} = 20$ dB. We used a trust-region algorithm to minimize the ER loss function (IV.1). The experiments were conducted with $W = 7$ and $L = 1$ and $W = 11$ and $L = 3$ in the left and the right columns, respectively. (a) Initialization with $W = 7$ and $L = 1$. (b) Initialization with $W = 11$ and $L = 3$. (c) Recovery with $W = 7$ and $L = 1$. (d) Recovery with $W = 11$ and $L = 3$.

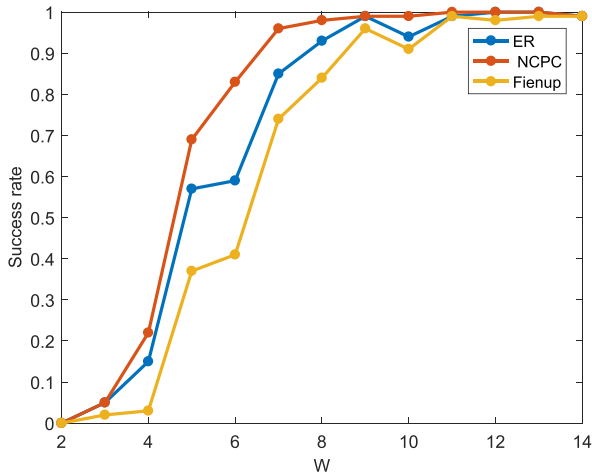


Fig. 5. Success rate as a function of W over 100 experiments conducted with $N = 31$, $L = 2$ and a rectangular window. We compared three algorithms: minimizing the ER loss function (IV.1), NCPC and Fienup. A success was declared for recovery error less than 10^{-3} .

Figure 7 presents the same experiments with low-pass data. This reflects a phenomenon that typically occurs in optical applications in which the fine details of the data are blurred

by the measurement process. Estimating a signal from its low-resolution measurements, when the phases are available, has been investigated thoroughly in the last years, see for instance [7], [20]. Accordingly, we assume that we can acquire the data $\mathbf{Z}[m, k]$ for all m but only for $k = -K_{\max}, \dots, K_{\max}$ for some cut-off frequency K_{\max} . Particularly, in Figure 7 we consider $N = 53$ and $K_{\max} = 18$ (i.e., 70% of the spectral content) for the two proposed algorithms. In this case, if the SNR is not too bad, then NCPC works significantly better than ER in both cases. As in Figure 6a, in the low SNR regime, minimizing the ER loss function achieves superior performance for Gaussian noise.

VI. THEORY

This section presents the theoretical contribution of this work, focusing on the case of maximum overlap between adjacent windows ($L = 1$). As explained and demonstrated numerically, the non-convex approaches also tend to work well for $L > 1$ and when the high-frequencies of the data are suppressed.

In our first theoretical result, Theorem 10, we analyze the initialization algorithm presented in Algorithm 1 and estimate the distance between the initialization vector and

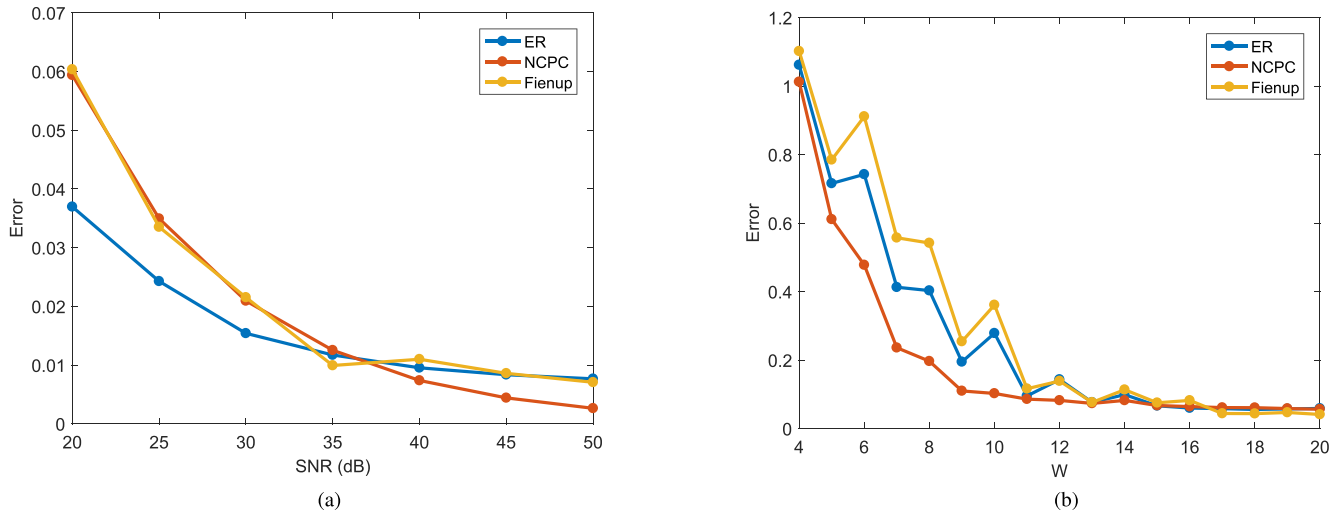


Fig. 6. Comparison of the average recovery error (over 100 experiments) of three algorithms: minimizing the ER loss function (IV.1), NCP and Fienup [31]. The experiments were conducted on signals of length $N = 53$ with a rectangular window and $L = 2$. (a) Recovery error with Gaussian i.i.d. noise as a function of the SNR with $W = 15$. (b) Recovery error with Poisson noise as a function of the window's length W .

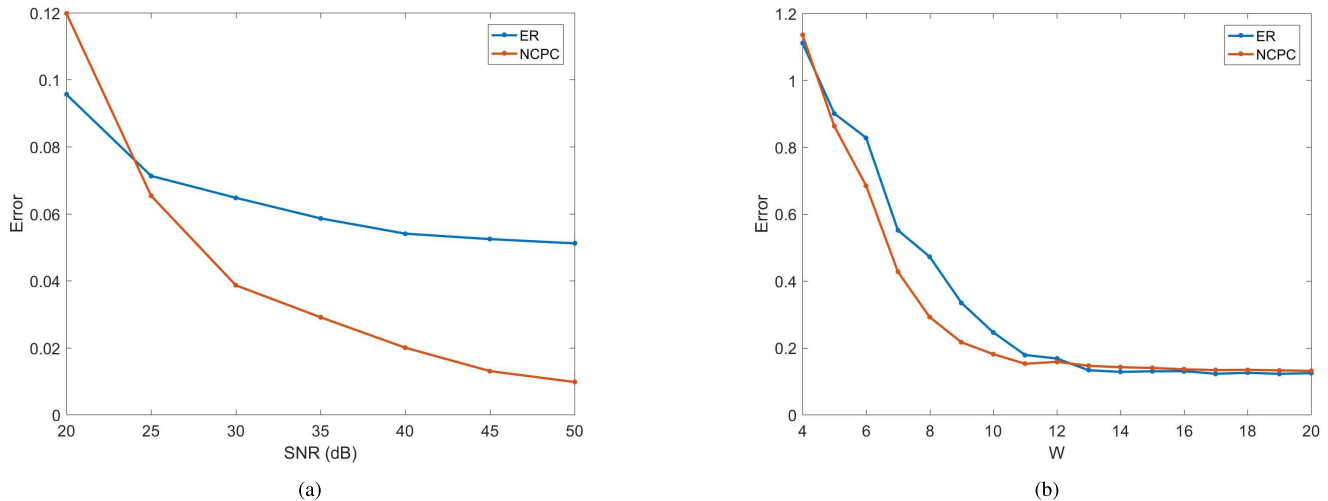


Fig. 7. Average recovery error (over 100 experiments with signals of length $N = 53$, a rectangular window and $L = 2$) of minimizing the ER loss function (IV.1) and NCP with low-passed data. Particularly, we used the measured data $\mathbf{Z}[m, k]$ for all m and $k = -K_{\max}, \dots, K_{\max}$ with $K_{\max} = 18$. (a) Recovery error with Gaussian i.i.d. noise as a function of the SNR with $W = 15$. (b) Recovery error with Poisson noise as a function of the window's length W .

the ground truth. Next, we study the geometry of the loss function (IV.2), which controls the behavior of our ER minimization algorithm. To this end, suppose we minimize the ER loss function (IV.2) using gradient descent followed by a thresholding step that can be used if the signal is bounded. This scheme is presented in Algorithm 6. In Theorem 11 we establish the existence of a *basin of attraction* of size $\frac{1}{8\sqrt{N}W^2}$ around the global minimum for signals with unit modulus entries. In the basin of attraction, a gradient algorithm is guaranteed to converge to a global minimum at a geometric rate. This result is true for any gradient scheme with a thresholding step as in Algorithm 6. We stress that the theoretical contribution of this result is limited. As presented in Corollary 12, the estimated basin of attraction is small so that theoretically Algorithm 6 converges in the same area in

which the problem has a closed linear LS solution. To the best of our knowledge, this is the first result quantifying the size of the basin of attraction of a gradient algorithm in a deterministic phase retrieval setup. This is in contrast to the basin of attraction of random phase retrieval setups which is quite well-understood.

A crucial condition for the success of gradient algorithms is that its initialization is sufficiently close to the global minimum. The following result quantifies the estimation error of the proposed initialization presented in Algorithm 1 for bounded signals and $L = 1$. The error reduces to zero as W approaches $\frac{N+1}{2}$. The case of $L > 1$ is discussed briefly in Section IV. The result is stated for a normalized signal. The norm of the signal can be estimated easily from the main diagonal of \mathbf{xx}^* as explained in Section III.

Algorithm 6 Gradient Descent Algorithm to Minimize the ER Loss Functions (IV.1) or (IV.2)

Input: The measurements $\mathbf{Z}[m, k]$ as given in (II.1) and (optional) thresholding parameter $B > 0$.

Output: Estimation of \mathbf{x} .

- 1) Initialization by Algorithm 1 (for $L = 1$) or Algorithm 5 (for $L > 1$).
- 2) Apply the update rule until convergence:
 - a) Gradient step:

$$\tilde{\mathbf{x}}_k = \mathbf{x}_{k-1} - \mu \nabla f(\mathbf{x}_{k-1}),$$

for step size μ and ∇f given in (IV.3).

- b) Optional thresholding:

$$\mathbf{x}_k[n] = \begin{cases} \tilde{\mathbf{x}}_k[n] & \text{if } |\tilde{\mathbf{x}}_k[n]| \leq B, \\ B \cdot \text{phase}(\tilde{\mathbf{x}}_k[n]) & \text{if } |\tilde{\mathbf{x}}_k[n]| > B. \end{cases}$$

Theorem 10: Suppose that $L = 1$, $\|\mathbf{x}\|_2 = 1$, \mathbf{g} is an admissible window of length $W \geq 2$ and that $\|\mathbf{x}\|_\infty \leq \sqrt{\frac{B}{N}}$ for some $0 < B \leq \frac{N}{2(N-2W+1)}$. Then under the measurement model of (II.1), the initialization vector given in Algorithm 1 satisfies

$$d^2(\mathbf{x}_0, \mathbf{x}) \leq 2 \left(1 - \sqrt{1 - 2B \frac{N-2W+1}{N}} \right).$$

Proof: See Section VII-A. ■

The properties of the gradient algorithm minimizing the ER rely on the geometry of the loss function (IV.2) near the global minimum. The following result quantifies the size of the basin of attraction of the loss function (IV.2), namely, the area in which a gradient algorithm is guaranteed to converge to a global minimum at a geometric rate. As demonstrated in Figure 3, in practice the basin of attraction is quite large for a broad family of signals. The proof relies on a geometric analysis of the loss function as presented in Lemmas 15 and 16.

Theorem 11: Let $L = 1$ and suppose that $\mathbf{x} \in \mathbb{R}_{1/\sqrt{N}}^N$ and \mathbf{g} is a rectangular window of length W . Additionally, suppose that $d(\mathbf{x}_0, \mathbf{x}) \leq \frac{1}{8\sqrt{N}W^2}$, where \mathbf{x}_0 obeys $\|\mathbf{x}_0\|_\infty \leq \frac{1}{\sqrt{N}}$. Then, under the measurement model (II.1), Algorithm 6 with thresholding parameter $B = \frac{1}{\sqrt{N}}$ and step size $0 < \mu \leq 2/\beta$ achieves the following geometric convergence:

$$d^2(\mathbf{x}_k, \mathbf{x}) \leq \left(1 - \frac{2\mu}{\alpha} \right)^k d^2(\mathbf{x}_0, \mathbf{x}),$$

where $\alpha \geq \frac{4N}{W}$ and $\beta \geq 256N^2W^3$.

Proof: See Section VII-B. ■

Combining Theorems 10 and 11 leads to the following corollary:

Corollary 12: Suppose that $L = 1$, $\mathbf{x} \in \mathbb{R}_{1/\sqrt{N}}^N$, N is a prime number and \mathbf{g} is a rectangular window of length W that satisfies:

$$2W - 1 + \frac{1}{128W^4} \geq N.$$

Then, under the measurement model of (II.1), Algorithm 6, initialized by Algorithm 1, with thresholding parameter $B = \frac{1}{\sqrt{N}}$ and step size $0 < \mu \leq 2/\beta$ achieves the following geometric convergence:

$$d^2(\mathbf{x}_k, \mathbf{x}) \leq \left(1 - \frac{2\mu}{\alpha} \right)^k d^2(\mathbf{x}_0, \mathbf{x}),$$

where $\alpha \geq \frac{4N}{W}$ and $\beta \geq 256N^2W^3$.

Proof: See Section VII-C. ■

We mention that the result of Corollary 12 is good merely for long windows. However, in practice we observe that the algorithm works well also for short windows. As we discuss in Section VIII, bridging this theoretical gap is an important direction for future research.

VII. PROOFS

A. Proof of Theorem 10

The initialization is based on extracting the principal eigenvector of the matrix \mathbf{X}_0 defined in Algorithm 1. By assumption, \mathbf{G}_ℓ are invertible matrices for $\ell = -(W-1), \dots, W-1$ for some $W \geq 2$ and hence we can compute (see (II.7))

$$\text{diag}(\mathbf{X}_0, \ell) = \mathbf{G}_\ell^{-1} \mathbf{y}_\ell = \text{diag}(\mathbf{X}, \ell).$$

For $\ell = W, \dots, N-W$ we have $\text{diag}(\mathbf{X}_0, \ell) = 0$. Let us take a look at the matrix $\mathbf{E} := \mathbf{X} - \mathbf{X}_0$. Clearly, \mathbf{E} is not zero at most on $N - 2W + 1$ diagonals. In other words, in each row and column, there are at most $N - 2W + 1$ non-zero values. Let Ω_i be the set of non-zero values of the i th row of \mathbf{E} with cardinality $|\Omega_i| \leq N - 2W + 1$. Using the fact that $\|\mathbf{x}\|_\infty = \sqrt{\frac{B}{N}}$ we can estimate

$$\begin{aligned} \|\mathbf{E}\|_\infty &:= \max_i \sum_j |\mathbf{X}[i, j] - \mathbf{X}_0[i, j]| \\ &= \max_i \sum_{j \in \Omega_i} |\mathbf{X}[i, j]| \\ &= \max_i \sum_{j \in \Omega_i} |\mathbf{x}[i] \mathbf{x}[j]| \\ &\leq \frac{B(N - 2W + 1)}{N}. \end{aligned}$$

The same bound holds for $\|\mathbf{E}\|_1 := \max_j \sum_i |\mathbf{E}[i, j]|$ and therefore by Hölder's inequality we get

$$\|\mathbf{E}\|_2 \leq \sqrt{\|\mathbf{E}\|_\infty \|\mathbf{E}\|_1} = \frac{B(N - 2W + 1)}{N}.$$

In order to complete the proof, we still need to show that if $\|\mathbf{X} - \mathbf{X}_0\|_2$ is small, then $d(\mathbf{x}, \mathbf{x}_0)$ is small as well, where \mathbf{x}_0 is the principal eigenvector of \mathbf{X}_0 with appropriate normalization. To show that, we follow the outline of in [18, Sec. 7.8]. Observe that as \mathbf{G}_0 is invertible by assumption, the norm of \mathbf{x} is known by

$$\|\mathbf{x}\|_2^2 = \sum_{n=0}^{N-1} (\text{diag}(\mathbf{X}, 0)[n]) = \sum_{n=0}^{N-1} (\mathbf{G}_0^{-1} \mathbf{y}_0)[n].$$

Accordingly, we assume hereinafter without loss of generality that \mathbf{x} and \mathbf{x}_0 have unit norm. Let λ_0 be the top eigenvalue

of \mathbf{X}_0 , associated with \mathbf{x}_0 . We observe that

$$\begin{aligned} \left| \lambda_0 - |\mathbf{x}_0^* \mathbf{x}|^2 \right| &= \left| \mathbf{x}_0^* \mathbf{X}_0 \mathbf{x}_0 - \mathbf{x}_0^* \mathbf{x} \mathbf{x}^* \mathbf{x}_0 \right| \\ &\leq \left\| \mathbf{X}_0 - \mathbf{x} \mathbf{x}^* \right\|_2. \end{aligned}$$

Furthermore, as $\|\mathbf{x}\|_2 = 1$ we also have

$$\begin{aligned} \lambda_0 &\geq \mathbf{x}^* \mathbf{X}_0 \mathbf{x} = \mathbf{x}^* (\mathbf{X}_0 - \mathbf{x} \mathbf{x}^*) \mathbf{x} + 1 \\ &\geq 1 - \left\| \mathbf{X}_0 - \mathbf{x} \mathbf{x}^* \right\|_2. \end{aligned}$$

Combining the last two inequalities we get

$$\begin{aligned} |\mathbf{x}_0^* \mathbf{x}|^2 &\geq 1 - 2 \left\| \mathbf{X}_0 - \mathbf{x} \mathbf{x}^* \right\|_2 \\ &\geq 1 - 2B \frac{N - 2W + 1}{N}. \end{aligned}$$

It then follows that W

$$d^2(\mathbf{x}_0, \mathbf{x}) \leq 2 \left(1 - \sqrt{1 - 2B \frac{N - 2W + 1}{N}} \right),$$

where the term in the square root is positive by assumption.

B. Proof of Theorem 11

For fixed \mathbf{x} , let \mathcal{E} be the set of vectors in \mathbb{R}^N satisfying $\|\mathbf{z}\|_\infty \leq \frac{1}{\sqrt{N}}$ and $d(\mathbf{x}, \mathbf{z}) \leq \frac{1}{8\sqrt{N}W^2}$. We first need the following definition:

Definition 13: We say that a function f satisfies the *regularity condition* in \mathcal{E} if for all vectors $\mathbf{z} \in \mathcal{E}$ we have

$$\left\langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x} e^{j\phi(\mathbf{z})} \right\rangle \geq \frac{1}{\alpha} d^2(\mathbf{z}, \mathbf{x}) + \frac{1}{\beta} \|\nabla f(\mathbf{z})\|_2^2,$$

for some positive constants α, β .

The following lemma states that if the regularity condition is met, then the gradient step converges to a global minimum at a geometric rate.

Lemma 14: Assume that f satisfies the regularity condition for all $\mathbf{z} \in \mathcal{E}$. Consider the following update rule

$$\mathbf{z}_k = \mathbf{z}_{k-1} - \mu \nabla f(\mathbf{z}_{k-1}),$$

for $0 < \mu \leq 2/\beta$. Then,

$$d^2(\mathbf{z}_k, \mathbf{x}) \leq \left(1 - \frac{2\mu}{\alpha} \right) d^2(\mathbf{z}_{k-1}, \mathbf{x}).$$

Proof: See in [18, Sec. 7.4]. ■

In order to show that the regularity condition of Definition 13 is met, we present two lemmas for signals with unit modulus entries. The first result shows that the gradient of the loss function (IV.2), given explicitly in (IV.3), is bounded near its global minimum. This implies that the loss function is smooth. We consider here only the case of a rectangular window \mathbf{g} of length W . The extension to non-vanishing windows of length W is straightforward (see remark in Appendix F):

Lemma 15: Suppose that $\mathbf{x} \in \mathbb{R}_{1/\sqrt{N}}^N$, $\|\mathbf{z}\|_\infty \leq \frac{1}{\sqrt{N}}$ and $d(\mathbf{x}, \mathbf{z}) \leq \frac{1}{\sqrt{N}}$. Let \mathbf{g} be a rectangular window of length W . Then, $\nabla f(\mathbf{z})$ as given in (IV.3) satisfies

$$\|\nabla f(\mathbf{z})\|_2 \leq \frac{8}{L} W^2 \sqrt{N} d(\mathbf{x}, \mathbf{z}).$$

Proof: See Appendix F. ■

The second lemma shows that the inner product between the gradient and the vector $\mathbf{z} - \mathbf{x} e^{j\phi(\mathbf{z})}$ is positive if $d(\mathbf{x}, \mathbf{z}) \leq \frac{1}{8\sqrt{N}W^2}$. This result implies that $-\nabla f(\mathbf{z})$ points approximately towards \mathbf{x} . As in Lemma 15, we consider for simplicity rectangular windows of length W . Yet, the analysis can be extended to non-vanishing windows of length W . In this case, the bounds are dependent on the dynamic range of \mathbf{g} (for details, see remark in Appendix G).

Lemma 16: Suppose that $L = 1$ and \mathbf{g} is a rectangular window of length W . For any $\mathbf{x} \in \mathbb{R}_{1/\sqrt{N}}^N$ and $\|\mathbf{z}\|_\infty \leq \frac{1}{\sqrt{N}}$, if $d(\mathbf{x}, \mathbf{z}) \leq \frac{1}{8\sqrt{N}W^2}$, then

$$\left\langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x} e^{j\phi(\mathbf{z})} \right\rangle \geq \frac{W d^2(\mathbf{x}, \mathbf{z})}{2N},$$

where $\nabla f(\mathbf{z})$ is given in (IV.3).

Proof: See Appendix G. ■

We notice that the thresholding stage of Algorithm 6 cannot increase the error as the signal is assumed to be bounded. The proof of Theorem 11 is then completed by directly leveraging lemmas 15 and 16 and seeing that Definition 13 holds in our case with constants $\alpha \geq \frac{4N}{W}$ and $\beta \geq 256N^2W^3$.

C. Proof of Corollary 12

As N is a prime number, \mathbf{g} is an admissible window of length W (see Lemma 5). According to Theorem 11, we merely need to show that the initialization point is within the basin of attraction, namely, $d(\mathbf{x}, \mathbf{x}_0) \leq \frac{1}{8\sqrt{N}W^2}$. From Lemma 10, we know that the initialization obeys

$$d^2(\mathbf{x}_0, \mathbf{x}) \leq 2 \left(1 - \sqrt{1 - 2 \frac{N - 2W + 1}{N}} \right).$$

Using the fact that $a \leq \sqrt{a}$ for all $0 \leq a \leq 1$ and some standard algebraic calculations, we conclude that the initialization of Algorithm 1 is within the basin of attraction as long as

$$2W - 1 + \frac{1}{128W^4} \geq N,$$

which completes the proof.

VIII. DISCUSSION

This paper explores practical, efficient, non-convex phase retrieval algorithms with some deterministic theoretical guarantees. Particularly, we propose two local optimization methods based on minimizing the ER loss function and optimizing on the manifold of phases. The latter is a new phase retrieval algorithm that takes into account the special geometry of the phase retrieval problem.

Since the optimization problems are non-convex, we also propose an initialization method. The method is based on the insight that, for sufficiently long windows, the signal can be recovered as the solution of a linear LS problem. While this may not be true for shorter windows, we use the LS solution to construct a special matrix and initialize the local optimization algorithms with the principal eigenvector of this matrix. Similar initialization approaches were suggested recently for

phase retrieval problems. However, they are mainly focused on random setups and based on probabilistic considerations. For $L = 1$, we estimate the distance between the initialization point and the ground truth. The case of $L > 1$ raises some interesting questions. As a heuristic, we suggested to smoothly interpolate the missing entries. This practice works quite well since the window acts as an averaging operator. Clearly, the interpolation method depends on the window shape. A main challenge for future research is analyzing the setting of $L > 1$.

For signals with unit modulus entries, we prove in Theorem 11 that the ER loss function has a basin of attraction. We show numerically that the actual basin of attraction is larger than the theoretical bound and exists for a broader family of signals. The gap between the actual size of the basin of attraction and the theoretical result is the bottleneck that prevents a full theoretical understanding of the proposed algorithms. Specifically, improving Lemma 16 will lead directly to tighter estimation of the size of the basin of attraction. Ideally, this would lead to the conclusion that the proposed initial guess lies in the basin.

APPENDIX

A. Proof of Proposition 6

By assumption, the DFT of $\mathbf{g} \odot (\mathbf{P}_{-\ell} \mathbf{g})$ is non-vanishing for $\ell = 0, 1$, and the matrices \mathbf{G}_ℓ , $\ell = 0, 1$ as given in (II.7) are invertible. Then, we can compute

$$\mathbf{x}_\ell = \mathbf{G}_\ell^{-1} \mathbf{y}_\ell, \quad \ell = 0, 1,$$

where $\mathbf{X} = \mathbf{x}\mathbf{x}^*$, $\mathbf{x}_\ell = \text{diag}(\mathbf{X}, \ell)$ and $\mathbf{y}_\ell := \{\mathbf{Y}[m, \ell]\}_{m=0}^{N-1}$. Because of the fundamental ambiguity of phase retrieval, the first entry can be set arbitrarily to $\sqrt{\mathbf{x}_0[0]} = |\mathbf{x}[0]|$. Then, as we assume non-vanishing signals, the rest of the entries are determined recursively for $n = 1 \dots, N-1$ by

$$\frac{\mathbf{x}_1[n-1]}{\mathbf{x}[n-1]} = \frac{\mathbf{x}[n-1]\mathbf{x}[n]}{\mathbf{x}[n-1]} = \mathbf{x}[n].$$

This completes the proof.

B. Proof of Proposition 7

By assumption, \mathbf{G}_ℓ is an invertible matrix for $|\ell| \leq W-1$ for some $W \geq 2$ (see (II.7)). Hence, we can compute $\text{diag}(\mathbf{X}, \ell) = \mathbf{G}_\ell^{-1} \mathbf{y}_\ell$ for $\ell = 0, M$ for any $1 \leq M \leq W-1$. The proof is a direct corollary of the following lemma:

Lemma: Let $L = 1$. Suppose that $\mathbf{x} \in \mathbb{C}^N_{1/\sqrt{N}}$ and let $\mathbf{X} = \mathbf{x}\mathbf{x}^*$. Fix $M \in \{1, \dots, N-1\}$ and let \mathbf{X}_0 be a matrix obeying

$$\text{diag}(\mathbf{X}_0, \ell) = \begin{cases} \text{diag}(\mathbf{X}, \ell), & \ell = 0, M, \\ 0, & \text{otherwise.} \end{cases}$$

Then, \mathbf{x} is a principal eigenvector of \mathbf{X}_0 (up to global phase).

Proof: Based on the special structure of \mathbf{X}_0 , the following calculation shows that \mathbf{x} is an eigenvector of \mathbf{X}_0 with $\frac{2}{N}$ as

the associated eigenvalue:

$$\begin{aligned} (\mathbf{X}_0 \mathbf{x})[i] &= \sum_{j=1}^N \mathbf{X}_0[i, j] \mathbf{x}[j] \\ &= \mathbf{X}_0[i, i] \mathbf{x}[i] + \mathbf{X}_0[i, i+M] \mathbf{x}[i+M] \\ &= \mathbf{x}[i] |\mathbf{x}[i]|^2 + \mathbf{x}[i] |\mathbf{x}[i+M]|^2 \\ &= \frac{2}{N} \mathbf{x}[i]. \end{aligned}$$

We still need to show that \mathbf{x} is a principal eigenvector of \mathbf{X}_0 . Since each column and row of \mathbf{X}_0 is composed of two non-zero values, it is evident that

$$\|\mathbf{X}_0\|_\infty := \max_i \sum_j |\mathbf{X}_0[i, j]| = \frac{2}{N}.$$

In the same manner

$$\|\mathbf{X}_0\|_1 := \max_j \sum_i |\mathbf{X}_0[i, j]| = \frac{2}{N}.$$

Hence by Hölder inequality we get

$$\|\mathbf{X}_0\|_2 \leq \sqrt{\|\mathbf{X}_0\|_1 \|\mathbf{X}_0\|_\infty} = \frac{2}{N}.$$

completing the proof. \blacksquare

C. Proof of Proposition 8

As the matrices \mathbf{G}_ℓ are invertible by assumption for all $\ell = -(W-1), \dots, (W-1)$, we can compute

$$\text{diag}(\mathbf{X}_0, \ell) = \mathbf{G}_\ell^{-1} \mathbf{y}_\ell = \text{diag}(\mathbf{X}, \ell).$$

The assumption $W \geq \lceil \frac{N+1}{2} \rceil$ implies that $\mathbf{X}_0 = \mathbf{X}$. Specifically, observe that it is sufficient to consider only $W = \lceil \frac{N+1}{2} \rceil$ since for any $|\ell_1| > \lceil \frac{N+1}{2} \rceil$, the window $\mathbf{g} \odot (\mathbf{P}_{-\ell_1})$ is equal to another window $\mathbf{g} \odot (\mathbf{P}_{-\ell_2})$ for some $|\ell_2| \leq \lceil \frac{N+1}{2} \rceil$.

Let $\tilde{\mathbf{x}} := \mathbf{x}/\|\mathbf{x}\|_2$. Then, $\tilde{\mathbf{x}}$ is the principal eigenvector of \mathbf{X} and the normalization stage of Algorithm 1 gives

$$\sqrt{\sum_{n=0}^{N-1} (\mathbf{G}_0^{-1} \mathbf{y}_0)[n]} = \|\mathbf{x}\|_2.$$

D. Proof of the Equality Between the Loss Functions (IV.2) and (IV.1)

Recall that

$$\begin{aligned} f(\mathbf{u}) &= \frac{1}{2} \sum_{m=0}^{\lceil \frac{N}{L} \rceil - 1} \sum_{k=0}^{N-1} (\mathbf{u}^* \tilde{\mathbf{H}}_{m,k} \mathbf{u} - \mathbf{Z}[m, k])^2 \\ &= \frac{1}{2} \sum_{m=0}^{\lceil \frac{N}{L} \rceil - 1} \|\tilde{\mathbf{H}}_m - \mathbf{Z}_m\|_2^2, \end{aligned}$$

where $\mathbf{Z}_m := \{\mathbf{Z}[m, k]\}_{k=0}^{N-1} \in \mathbb{R}^N$ and $\tilde{\mathbf{H}}_m := \{\mathbf{u}^* \tilde{\mathbf{H}}_{m,k} \mathbf{u}\}_{k=0}^{N-1} \in \mathbb{R}^N$.

Let \mathbf{U} be a unitary matrix. Since unitary matrices do not change the length of a vector, we have

$$\begin{aligned} f(\mathbf{u}) &= \frac{1}{2} \sum_{m=0}^{\lceil \frac{N}{L} \rceil - 1} \|\mathbf{U}(\tilde{\mathbf{H}}_m - \mathbf{Z}_m)\|_2^2 \\ &= \frac{1}{2} \sum_{m=0}^{\lceil \frac{N}{L} \rceil - 1} \|\mathbf{U}\tilde{\mathbf{H}}_m - \mathbf{U}\mathbf{Z}_m\|_2^2. \end{aligned}$$

By choosing \mathbf{U} to be the DFT matrix and normalize, we get exactly the loss function in (IV.2).

E. Proof of Lemma 9

We identify the convolution $\mathbf{g} * \mathbf{x}$ by the matrix-vector product $\mathbf{G}\mathbf{x}$, where $\mathbf{G} \in \mathbb{R}^{N \times N}$ is a circulant matrix whose first column is given by $\tilde{\mathbf{g}} := \{\mathbf{g}[(-n) \bmod N]\}_{n=0}^{N-1}$. For $L = 1$, we can then write

$$\mathbf{y} = \mathbf{G}\mathbf{x} = \mathbf{F}^* \Sigma \mathbf{F}\mathbf{x},$$

where \mathbf{F} is a DFT matrix and Σ is a diagonal matrix whose entries are the DFT of $\tilde{\mathbf{g}}$. By assumption, the first N/L entries of Σ are ones and the rest are zeros. Hence, we may write

$$\mathbf{y} = \mathbf{F}_p^* \mathbf{F}_p \mathbf{x}, \quad (\text{E.1})$$

where $\mathbf{F}_p \in \mathbb{C}^{N/L \times N}$ consists of the first N/L rows of \mathbf{F} .

Let $\mathbf{G}_L \in \mathbb{R}^{N/L \times N}$ be a matrix consists of the $\{jL : j = 0, \dots, N/L - 1\}$ rows of \mathbf{G} . For $L > 1$, we get the downsampled system of equations

$$\mathbf{y}_L = \mathbf{G}_L \mathbf{x} = \mathbf{F}_L^* \Sigma \mathbf{F}\mathbf{x},$$

where \mathbf{F}_L consists of the $\{jL : j = 0, \dots, N/L - 1\}$ columns of \mathbf{F} (notice the difference between \mathbf{F}_L and \mathbf{F}_p). We aim at showing that expanding and interpolating \mathbf{y}_L as explained in Lemma 9 results in \mathbf{y} . Direct computation shows that the expansion stage as described in (IV.12) is equivalent to multiplying both sides by $\mathbf{F}^* \mathbf{F}_L$:

$$\tilde{\mathbf{y}}_L = \mathbf{F}^* \mathbf{F}_L \mathbf{y}_L = \mathbf{F}^* (\mathbf{F}_L \mathbf{F}_L^*) \Sigma \mathbf{F}\mathbf{x}.$$

Let us denote $\mathbf{T} := \mathbf{F}_L \mathbf{F}_L^*$, which is a Toeplitz matrix with L on the $\frac{jN}{L}$ diagonals for $j = 0, \dots, N/L - 1$ and zero otherwise. Because of the structure of Σ we can then write

$$\tilde{\mathbf{y}}_L = \mathbf{F}^* \mathbf{T}_p \mathbf{F}_p \mathbf{x},$$

where $\mathbf{T}_p \in \mathbb{R}^{N \times \frac{N}{L}}$ consists of the first N/L columns of \mathbf{T} . Direct calculation shows that $\mathbf{F}_p \mathbf{F}_p^* \mathbf{T}_p = \mathbf{I}$, where \mathbf{I} is the identity matrix. Therefore we conclude that

$$\left(\mathbf{F}_p^* \mathbf{F}_p\right) \tilde{\mathbf{y}}_L = \mathbf{F}_p^* \mathbf{F}_p \mathbf{x}. \quad (\text{E.2})$$

Comparing (E.2) with (E.1) completes the proof.

F. Proof of Lemma 15

Recall that

$$\begin{aligned} \nabla f(\mathbf{z}) &= \sum_{m=0}^{\lceil N/L \rceil - 1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{Y}[m, \ell] \right) \\ &\quad \cdot \left(\mathbf{H}_{m,\ell} + \mathbf{H}_{m,\ell}^T \right) \mathbf{z}, \end{aligned}$$

where

$$\mathbf{H}_{m,\ell} := \mathbf{P}_{-\ell} \mathbf{D}_{mL} \mathbf{D}_{mL-\ell},$$

\mathbf{D}_{mL} is a diagonal matrix whose entries are $\{\mathbf{g}[mL - n]\}_{n=0}^{N-1}$ for fixed m and \mathbf{P}_ℓ is a matrix that shifts (circularly) the entries of an arbitrary vector by ℓ entries. We observe that for a rectangular window of length W and $\|\mathbf{z}\|_\infty \leq \frac{1}{\sqrt{N}}$, we have $\|\mathbf{z}\|_2 \leq 1$ and

$$\|\mathbf{H}_{m,\ell} \mathbf{z}\|_2 \leq \|\mathbf{H}_{m,\ell}\|_2 \|\mathbf{z}\|_2 \leq 1,$$

so that

$$\|\nabla f(\mathbf{z})\|_2 \leq 2 \sum_{m=0}^{\lceil N/L \rceil - 1} \sum_{\ell=-(W-1)}^{W-1} \left| \mathbf{Y}[m, \ell] - \mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} \right|. \quad (\text{F.1})$$

For convenience, let us denote $d(\mathbf{x}, \mathbf{z}) = \frac{\varepsilon}{\sqrt{N}}$ for some $\varepsilon \leq 1$ and therefore $|\mathbf{z}[n]| \geq \frac{1-\varepsilon}{\sqrt{N}}$ for all n . Accordingly, for any (n, k) ,

$$\frac{(1-\varepsilon)^2}{N} \leq |\mathbf{z}[n]\mathbf{z}[n+k]| \leq \frac{1}{N}.$$

Since $\mathbf{x}[n]$ and $\mathbf{z}[n]$ have the same sign pattern, we have

$$\begin{aligned} |\mathbf{x}[n]\mathbf{x}[n+k] - \mathbf{z}[n]\mathbf{z}[n+k]| &\leq \frac{1}{N} \left| 1 - (1-\varepsilon)^2 \right| \\ &\leq \frac{2\varepsilon}{N}, \end{aligned}$$

and for all $m, \ell \geq 0$,

$$\begin{aligned} &\left| \mathbf{Y}[m, \ell] - \mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} \right| \\ &\leq \sum_{k=m-(W-1)}^{m-\ell} |\mathbf{x}[n]\mathbf{x}[n+k] - \mathbf{z}[n]\mathbf{z}[n+k]| \\ &\leq \sum_{k=m-(W-1)}^{m-\ell} \frac{2\varepsilon}{N} \leq \frac{2W\varepsilon}{N} = \frac{2Wd(\mathbf{x}, \mathbf{z})}{\sqrt{N}}. \quad (\text{F.2}) \end{aligned}$$

The same bound holds for $\ell < 0$. Combining (F.1) and (F.2) we conclude that

$$\begin{aligned} \|\nabla f(\mathbf{z})\|_2 &\leq 2 \sum_{m=0}^{\lceil N/L \rceil - 1} \sum_{\ell=-(W-1)}^{W-1} \frac{2Wd(\mathbf{x}, \mathbf{z})}{\sqrt{N}} \\ &= \frac{8}{L} W^2 \sqrt{N} d(\mathbf{x}, \mathbf{z}). \end{aligned}$$

Remark: In case of a non-vanishing window of length W , one can easily bound the gradient using the same technique, while taking into account $\max_n |\mathbf{g}[n]|$ in the inequalities.

G. Proof of Lemma 16

Recall that (see (IV.3))

$$\begin{aligned} & \langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \rangle \\ &= \sum_{m=0}^{\lceil N/L \rceil - 1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right) \\ & \quad \cdot \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right)^T \left(\mathbf{H}_{m,\ell} + \mathbf{H}_{m,\ell}^T \right) \mathbf{z}. \end{aligned}$$

Since $\mathbf{x}^T \mathbf{H}_{m,\ell}^T \mathbf{z} = \mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{x}$ we have for fixed (m, ℓ) and $\phi(\mathbf{z}) \in \{0, \pi\}$:

$$\begin{aligned} & \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right)^T \left(\mathbf{H}_{m,\ell} + \mathbf{H}_{m,\ell}^T \right) \mathbf{z} \\ &= \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right)^T \mathbf{H}_{m,\ell} \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right) \\ & \quad + \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} & \langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \rangle \\ &= \sum_{m=0}^{N-1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right)^2 \\ & \quad + \sum_{m=0}^{N-1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right) \\ & \quad \cdot \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right)^T \mathbf{H}_{m,\ell} \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right). \quad (\text{G.1}) \end{aligned}$$

Clearly, if $\mathbf{z} = \mathbf{x}e^{j\phi(\mathbf{z})}$ then $\langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \rangle = 0$. Otherwise, the first term of (G.1) is strictly positive. Hence, in order to achieve a lower bound on (G.1), we first derive an upper bound on the second term and then bound the first term from below.

By assumption $d(\mathbf{x}, \mathbf{z}) \leq \frac{1}{\sqrt{N}}$. Denote $|\mathbf{x}[n]e^{j\phi(\mathbf{z})} - \mathbf{z}[n]| := \frac{\varepsilon_n}{\sqrt{N}}$ for some $\varepsilon_n \leq 1$. We observe that $\sum_n \left(\frac{\varepsilon_n}{\sqrt{N}}\right)^2 = d^2(\mathbf{x}, \mathbf{z})$. For fixed $\ell \geq 0$, we can use the Cauchy-Schwarz inequality to obtain:

$$\begin{aligned} & \sum_{m=0}^{N-1} \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right)^T \mathbf{H}_{m,\ell} \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right) \\ &= \sum_{m=0}^{N-1} \sum_{n=m-(W-1)}^{m-\ell} \left(\mathbf{z}[n] - \mathbf{x}[n]e^{j\phi(\mathbf{z})} \right) \\ & \quad \cdot \left(\mathbf{z}[n+\ell] - \mathbf{x}[n+\ell]e^{j\phi(\mathbf{z})} \right) \\ &\leq W \sqrt{\sum_{m=0}^{N-1} \frac{\varepsilon_m^2}{N}} \sqrt{\sum_{m=0}^{N-1} \frac{\varepsilon_{m+\ell}^2}{N}} = Wd^2(\mathbf{x}, \mathbf{z}). \end{aligned}$$

The same bound holds for $\ell < 0$. Combining the last result with (F.2) we get for the second term in (G.1) that

$$\begin{aligned} & \sum_{m=0}^{N-1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right) \\ & \quad \cdot \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right)^T \mathbf{H}_{m,\ell} \left(\mathbf{z} - \mathbf{x}e^{j\phi(\mathbf{z})} \right) \leq \frac{4}{\sqrt{N}} W^3 d^3(\mathbf{x}, \mathbf{z}). \quad (\text{G.2}) \end{aligned}$$

Next, we aim to bound the first term of (G.1) from below as follows:

$$\begin{aligned} & \sum_{m=0}^{N-1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right)^2 \\ &\geq \sum_{m=0}^{N-1} \left(\mathbf{z}^T \mathbf{H}_{m,0} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,0} \mathbf{x} \right)^2 \\ &= \sum_{m=0}^{N-1} \left(\sum_{n=m-(W-1)}^m \mathbf{z}^2[n] - \mathbf{x}^2[n] \right)^2 \\ &\geq \sum_{m=0}^{N-1} \sum_{n=m-(W-1)}^m \left(\mathbf{z}^2[n] - \mathbf{x}^2[n] \right)^2, \quad (\text{G.3}) \end{aligned}$$

where the last inequality is true since $\mathbf{x}^2[n] \geq \mathbf{z}^2[n]$ and for any positive (or negative) sequence $\{a_i\}$ we have $(\sum_i a_i)^2 \geq \sum_i a_i^2$. Furthermore, since $|\mathbf{z}[n]| = \frac{1-\varepsilon_n}{\sqrt{N}}$ we have

$$\begin{aligned} & \sum_{m=0}^{N-1} \sum_{n=m-(W-1)}^m \left(\mathbf{z}^2[n] - \mathbf{x}^2[n] \right)^2 \\ &= \frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{n=m-(W-1)}^m \left(1 - (1 - \varepsilon_n)^2 \right)^2 \\ &= \frac{W}{N^2} \sum_{n=0}^{N-1} \left(2\varepsilon_n - \varepsilon_n^2 \right)^2. \end{aligned}$$

Therefore, since $\varepsilon_n \leq 1$ for all n we conclude that

$$\begin{aligned} & \sum_{m=0}^{N-1} \sum_{\ell=-(W-1)}^{W-1} \left(\mathbf{z}^T \mathbf{H}_{m,\ell} \mathbf{z} - \mathbf{x}^T \mathbf{H}_{m,\ell} \mathbf{x} \right)^2 \\ &\geq \frac{W}{N^2} \sum_{n=0}^{N-1} \varepsilon_n^2 = \frac{Wd^2(\mathbf{x}, \mathbf{z})}{N}. \quad (\text{G.4}) \end{aligned}$$

Plugging (G.2) and (G.4) into (G.1) yields

$$\begin{aligned} \langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x} \rangle &\geq \frac{Wd^2(\mathbf{x}, \mathbf{z})}{N} \left(1 - 4\sqrt{N}W^2d(\mathbf{x}, \mathbf{z}) \right) \\ &\geq \frac{Wd^2(\mathbf{x}, \mathbf{z})}{2N}, \end{aligned}$$

where the last inequality holds for $d(\mathbf{x}, \mathbf{z}) \leq \frac{1}{8\sqrt{N}W^2}$.

Remark: Observe that the analysis for non-vanishing windows of length W requires only a small modification. In this case, one should use the maximal and the minimal values of the window in the above inequalities. For instance, one would need to take $g_{\min} := \min_{n=0, \dots, W-1} |g[n]|$ into account in (G.3).

ACKNOWLEDGMENT

We would like to thank Mahdi Soltanolkotabi, Irène Waldspurger, Pavel Sidorenko and Laura Waller for their remarks on an initial draft of this paper.

REFERENCES

- [1] P.-A. Absil, C. G. Baker, and K. A. Gallivan, "Trust-region methods on Riemannian manifolds," *Found. Comput. Math.*, vol. 7, no. 3, pp. 303–330, 2007.
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ, USA: Princeton Univ. Press, 2008.
- [3] A. S. Bandeira, N. Boumal, and A. Singer, "Tightness of the maximum likelihood semidefinite relaxation for angular synchronization," *Math. Program.*, vol. 163, nos. 1–2, pp. 145–167, 2016.
- [4] A. S. Bandeira, Y. Chen, and A. Singer. (2015). "Non-unique games over compact groups and orientation estimation in cryo-EM." [Online]. Available: <https://arxiv.org/abs/1505.03840>
- [5] B. Baykal, "Blind channel estimation via combining autocorrelation and blind phase estimation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 6, pp. 1125–1131, Jun. 2004.
- [6] T. Bendory, P. Sidorenko, and Y. C. Eldar, "On the uniqueness of FROG methods," *IEEE Signal Process. Lett.*, vol. 24, no. 5, pp. 722–726, May 2017.
- [7] T. Bendory, "Robust recovery of positive stream of pulses," *IEEE Trans. Signal Process.*, vol. 65, no. 8, pp. 2114–2122, Apr. 2017.
- [8] T. Bendory, R. Beinert, and Y. C. Eldar. (2017). "Fourier phase retrieval: Uniqueness and algorithms." [Online]. Available: <https://arxiv.org/abs/1705.09590>
- [9] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer. (2017). "Bispectrum inversion with application to multireference alignment." [Online]. Available: <https://arxiv.org/abs/1705.00641>
- [10] T. Bendory, D. Edidin, and Y. C. Eldar. (2017). "On signal reconstruction from frog measurements." [Online]. Available: <https://arxiv.org/abs/1706.08494>
- [11] I. Bojarovska and A. Flinth, "Phase retrieval from Gabor measurements," *J. Fourier Anal. Appl.*, vol. 22, no. 3, pp. 542–567, 2016.
- [12] N. Boumal, P.-A. Absil, and C. Cartis. (2016). "Global rates of convergence for nonconvex optimization on manifolds." [Online]. Available: <https://arxiv.org/abs/1605.08101>
- [13] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a Matlab toolbox for optimization on manifolds," *J. Mach. Learn. Res.*, vol. 15, pp. 1455–1459, Jan. 2014. [Online]. Available: <http://www.manopt.org>
- [14] N. Boumal, "Nonconvex phase synchronization," *SIAM J. Optim.*, vol. 26, no. 4, pp. 2355–2377, 2016.
- [15] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM Rev.*, vol. 57, no. 2, pp. 225–251, 2015.
- [16] E. J. Candès and X. Li, "Solving quadratic equations via Phaselift when there are about as many equations as unknowns," *Found. Comput. Math.*, vol. 14, no. 5, pp. 1017–1026, 2014.
- [17] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval from coded diffraction patterns," *Appl. Comput. Harmon. Anal.*, vol. 39, no. 2, pp. 277–299, 2015.
- [18] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via wirtinger flow: theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, Apr. 2015.
- [19] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [20] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Commun. Pure Appl. Math.*, vol. 67, no. 6, pp. 906–956, Jun. 2014.
- [21] Y. Chen and E. J. Candès. (2016). "The projected power method: An efficient algorithm for joint alignment from pairwise differences." [Online]. Available: <https://arxiv.org/abs/1609.05820>
- [22] Y. Chen and E. J. Candès, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," *Commun. Pure Appl. Math.*, vol. 70, no. 5, pp. 822–883, 2017.
- [23] K. W. DeLong, R. Trebino, J. Hunter, and W. E. White, "Frequency-resolved optical gating with the use of second-harmonic generation," *J. Opt. Soc. Amer. B*, vol. 11, no. 11, pp. 2206–2215, 1994.
- [24] B. Dumitrescu, *Positive Trigonometric Polynomials and Signal Processing Applications*. Berlin, Germany: Springer, 2007.
- [25] Y. C. Eldar and S. Mendelson, "Phase retrieval: Stability and recovery guarantees," *Appl. Comput. Harmon. Anal.*, vol. 36, no. 3, pp. 473–494, 2014.
- [26] Y. C. Eldar, P. Sidorenko, D. G. Mixon, S. Barel, and O. Cohen, "Sparse phase retrieval from short-time Fourier measurements," *IEEE Signal Process. Lett.*, vol. 22, no. 5, pp. 638–642, May 2015.
- [27] C. Fienup and J. Dainty, "Phase retrieval and image reconstruction for astronomy," in *Image Recovery: Theory Application*. Orlando, FL, USA: Academic, 1987, pp. 231–275.
- [28] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, Aug. 1982.
- [29] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik*, vol. 35, p. 237, Nov. 1972.
- [30] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *J. Assoc. Comput. Mach.*, vol. 42, no. 6, pp. 1115–1145, 1995.
- [31] D. Griffin and J. Lim, "Signal estimation from modified short-time Fourier transform," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-32, no. 2, pp. 236–243, Apr. 1984.
- [32] D. Gross, F. Kraemer, and R. Kueng, "Improved recovery guarantees for phase retrieval from coded diffraction patterns," *Appl. Comput. Harmon. Anal.*, vol. 42, no. 1, pp. 37–64, 2015.
- [33] R. Harrison, "Phase problem in crystallography," *J. Opt. Soc. Amer. A*, vol. 10, no. 5, pp. 1046–1055, 1993.
- [34] K. Huang, Y. C. Eldar, and N. D. Sidiropoulos, "Phase retrieval from 1 D Fourier measurements: Convexity, uniqueness, and algorithms," *IEEE Trans. Signal Process.*, vol. 64, no. 23, pp. 6105–6117, Dec. 2016.
- [35] M. A. Iwen, B. Preskitt, R. Saab, and A. Viswanathan. (2016). "Phase retrieval from local measurements: Improved robustness via eigenvector-based angular synchronization." [Online]. Available: <https://arxiv.org/abs/1612.01182>
- [36] M. A. Iwen, A. Viswanathan, and Y. Wang, "Fast phase retrieval from local correlation measurements," *SIAM J. Imag. Sci.*, vol. 9, no. 4, pp. 1655–1688, 2016.
- [37] K. Jaganathan, Y. C. Eldar, and B. Hassibi, *Phase Retrieval: An Overview of Recent Developments*, A. Stern, Ed. Boca Raton, FL, USA: CRC Press, 2016.
- [38] K. Jaganathan, Y. C. Eldar, and B. Hassibi, "STFT phase retrieval: Uniqueness guarantees and recovery algorithms," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 770–781, Apr. 2016.
- [39] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Convex algorithms and limitations," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2013, pp. 1022–1026.
- [40] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre, "Generalized power method for sparse principal component analysis," *J. Mach. Learn. Res.*, vol. 11, pp. 517–553, Feb. 2010.
- [41] B. Juang and L. Rabiner, "Fundamentals of speech recognition," in *Signal Processing Series*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [42] K. Kreutz-Delgado. (2009). "The complex gradient operator and the CR-calculus." [Online]. Available: <https://arxiv.org/abs/0906.4835>
- [43] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. (2016). "Gradient descent converges to minimizers." [Online]. Available: <https://arxiv.org/abs/1602.04915>
- [44] A. Maiden, M. J. Humphry, F. Zhang, and J. M. Rodenburg, "Superresolution imaging via ptychography," *J. Opt. Soc. Amer. A*, vol. 28, no. 4, pp. 604–612, 2011.
- [45] S. Marchesini, Y. Tu, and H. Wu, "Alternating projection, ptychographic imaging and phase synchronization," *Appl. Comput. Harmon. Anal.*, vol. 41, no. 3, pp. 815–851, 2015.
- [46] R. Millane, "Phase retrieval in crystallography and optics," *J. Opt. Soc. Amer. A*, vol. 7, no. 3, pp. 394–411, 1990.
- [47] S. Nawab, T. Quatieri, and J. Lim, "Signal reconstruction from short-time Fourier transform magnitude," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 31, no. 4, pp. 986–998, Aug. 1983.
- [48] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 18, pp. 4814–4826, Sep. 2015.
- [49] A. Oppenheim and R. Schaffer, *Discrete-Time Signal Processing*. London, U.K.: Pearson Ed., 2010.
- [50] E. Pauwels, A. Beck, Y. C. Eldar, and S. Sabach. (2017). "On Fienup methods for regularized phase retrieval." [Online]. Available: <https://arxiv.org/abs/1702.08339>
- [51] A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra. (2016). "Message-passing algorithms for synchronization problems over compact groups." [Online]. Available: <https://arxiv.org/abs/1610.04583>
- [52] G. E. Pfander and P. Salanevich. (2016). "Robust phase retrieval algorithm for time-frequency structured measurements." [Online]. Available: <https://arxiv.org/abs/1611.02540>

- [53] J. Ranieri, A. Chebira, Y. M. Lu, and M. Vetterli. (2013). "Phase retrieval for sparse signals: Uniqueness conditions." [Online]. Available: <https://arxiv.org/abs/1308.3058>
- [54] J. Rodenburg, "Ptychography and related diffractive imaging methods," *Adv. Imag. Electron Phys.*, vol. 150, no. 7, pp. 87–184, 2008.
- [55] J. Rodenburg and R. Bates, "The theory of super-resolution electron microscopy via Wigner-distribution deconvolution," *Phil. Trans. R. Soc. Lond. A, Math. Phys. Sci.*, vol. 339, no. 1655, pp. 521–553, 1992.
- [56] C. Rusu and J. Astola, "Extending a sequence into a minimum-phase sequence," in *Proc. Int. TICSP Workshop Spectral Methods Multirate Signal Process. (SMMSP)*, Moscow, Russia, Sep. 2007, pp. 1–2.
- [57] S. Sanghavi, R. Ward, and C. D. White, "The local convexity of solving systems of quadratic equations," *Results Math.*, vol. 71, pp. 1–40, Jun. 2017.
- [58] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: Efficient phase retrieval of sparse signals," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 928–938, Apr. 2014.
- [59] Y. Shechtman, Y. C. Eldar, O. Cohen, H. 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, Mar. 2015.
- [60] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," *Opt. Exp.*, vol. 19, no. 16, pp. 14807–14822, 2011.
- [61] A. Singer, "Angular synchronization by eigenvectors and semidefinite programming," *Appl. Comput. Harmon. Anal.*, vol. 30, no. 1, pp. 20–36, 2011.
- [62] J. Sun, Q. Qu, and J. Wright. (2016). "A geometric analysis of phase retrieval." [Online]. Available: <https://arxiv.org/abs/1602.06664>
- [63] R. Trebino, *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses*. New York, NY, USA: Springer, 2012.
- [64] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Math. Programm.*, vol. 149, nos. 1–2, pp. 47–81, 2015.
- [65] I. Waldspurger. (2016). "Phase retrieval with random Gaussian sensing vectors by alternating projections." [Online]. Available: <https://arxiv.org/abs/1609.03088>
- [66] A. Walther, "The question of phase retrieval in optics," *J. Modern Opt.*, vol. 10, no. 1, pp. 41–49, 1963.
- [67] G. Wang, G. Giannakis, and Y. C. Eldar. (2016). "Solving systems of random quadratic equations via truncated amplitude flow." [Online]. Available: <https://arxiv.org/abs/1605.08285>
- [68] Y. Wang and Z. Xu, "Phase retrieval for sparse signals," *Appl. Comput. Harmon. Anal.*, vol. 37, no. 3, pp. 531–544, 2014.
- [69] C. Yang, J. Qian, A. Schirrotzek, F. Maia, and S. Marchesini. (2011). "Iterative algorithms for ptychographic phase retrieval." [Online]. Available: <https://arxiv.org/abs/1105.5628>
- [70] L.-H. Yeh *et al.*, "Experimental robustness of Fourier ptychography phase retrieval algorithms," *Opt. Exp.*, vol. 23, no. 26, pp. 33214–33240, 2015.
- [71] H. Zhang and Y. Liang, "Reshaped Wirtinger flow for solving quadratic system of equations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2622–2630.
- [72] V. Elser and T.-Y. Lan. (2017). "Benchmark problems for phase retrieval." [Online]. Available: <https://arxiv.org/abs/1706.00399>

Tamir Bendory received his PhD degree from the Electrical Engineering department at the Technion - Israel Institute of Technology in 2015. Between September 2015 and September 2016 he was a postdoctoral researcher with the SAMPL lab in the Electrical Engineering department at the Technion - Israel Institute of Technology. Currently, he is a postdoctoral researcher in the Program in Applied and Computational Mathematics at Princeton University. His research focuses on the design and the mathematical understanding of efficient algorithms for data science and signal processing applications.

Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in Physics in 1995 and the B.Sc. degree in Electrical Engineering in 1996 both from Tel-Aviv University (TAU), Tel-Aviv, Israel, and the Ph.D. degree in Electrical Engineering and Computer Science in 2002 from the Massachusetts Institute of Technology (MIT), Cambridge.

She is currently a Professor in the Department of Electrical Engineering at the Technion - Israel Institute of Technology, Haifa, Israel, where she holds the Edwards Chair in Engineering. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT, an Adjunct Professor at Duke University, and was a Visiting Professor at Stanford University, Stanford, CA. She is a member of the Israel Academy of Sciences and Humanities (elected 2017), an IEEE Fellow and a EURASIP Fellow.

Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014), and the IEEE Kiyu Tomiyasu Award (2016). She was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technion's Award for Excellence in Teaching (two times). She received several best paper awards and best demo awards together with her research students and colleagues including the SIAM outstanding Paper Prize, the UFFC Outstanding Paper Award, the Signal Processing Society Best Paper Award and the IET Circuits, Devices and Systems Premium Award, and was selected as one of the 50 most influential women in Israel.

She was a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She is the Editor in Chief of Foundations and Trends in Signal Processing, a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees, and served as an associate editor for the IEEE Transactions on Signal Processing, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*. She was Co-Chair and Technical Co-Chair of several international conferences and workshops.

She is author of the book "Sampling Theory: Beyond Bandlimited Systems" and co-author of the books "Compressed Sensing" and "Convex Optimization Methods in Signal Processing and Communications", all published by Cambridge University Press.

Nicolas Boumal is an instructor with the mathematics department and the program in computational and applied mathematics at Princeton University since February 2016. He obtained his PhD in applied mathematics from the Université catholique de Louvain in 2014. In between, he was a postdoctoral researcher with Inria in the SIERRA team in Paris. He is interested in optimization in general, and optimization on manifolds in particular, for which he develops Manopt.