
Linear Spectral Estimators and an Application to Phase Retrieval

Ramina Ghods¹ Andrew S. Lan² Tom Goldstein³ Christoph Studer¹

Abstract

Phase retrieval refers to the problem of recovering real- or complex-valued vectors from magnitude measurements. The best-known algorithms for this problem are iterative in nature and rely on so-called spectral initializers that provide accurate initialization vectors. We propose a novel class of estimators suitable for general nonlinear measurement systems, called linear spectral estimators (LSPEs), which can be used to compute accurate initialization vectors for phase retrieval problems. The proposed LSPEs not only provide accurate initialization vectors for noisy phase retrieval systems with structured or random measurement matrices, but also enable the derivation of sharp and nonasymptotic mean-squared error bounds. We demonstrate the efficacy of LSPEs on synthetic and real-world phase retrieval problems, and show that our estimators significantly outperform existing methods for structured measurement systems that arise in practice.

1. Introduction

Phase retrieval refers to the problem of recovering an unknown N -dimensional signal vector $\mathbf{x} \in \mathcal{H}^N$, with \mathcal{H} being the set of either real (\mathbb{R}) or complex (\mathbb{C}) numbers, from the following nonlinear measurement process:

$$\mathbf{y} = f(\mathbf{A}\mathbf{x} + \mathbf{e}^z) + \mathbf{e}^y. \quad (1)$$

Here, the measurement vector $\mathbf{y} \in \mathbb{R}^M$ contains M real-valued observations, for example measured through the nonlinear function $f(z) = |z|^2$ that operates element-wise on vectors, $\mathbf{A} \in \mathcal{H}^{M \times N}$ is a given measurement matrix, and the vectors $\mathbf{e}^z \in \mathcal{H}^M$ and $\mathbf{e}^y \in \mathbb{R}^M$ model signal and measurement noises, respectively. In contrast to the majority of

existing results on phase retrieval that assume randomness in the measurement matrix \mathbf{A} , we focus on the practical scenario in which the measurement matrix \mathbf{A} is deterministic, but the signal vector \mathbf{x} to be recovered as well as the two noise sources \mathbf{e}^z and \mathbf{e}^y are random.

1.1. Phase Retrieval

Phase retrieval has been studied extensively over the last decades (Gerchberg & Saxton, 1972; Fienup, 1982) and finds use in a range of applications, including imaging (Fogel et al., 2016; Yeh et al., 2015; Holloway et al., 2016), microscopy (Kou et al., 2010; Faulkner & Rodenburg, 2004), and X-ray crystallography (Harrison, 1993; Miao et al., 2008; Pfeiffer et al., 2006). Phase retrieval problems were solved traditionally using alternating projection methods, such as the Gerchberg-Saxton (Gerchberg & Saxton, 1972) and Fienup (Fienup, 1982) algorithms. More recent results have shown that semidefinite programming enables the design of algorithms with performance guarantees (Candès et al., 2013; Candès & Li, 2014; Candès et al., 2015a; Waldspurger et al., 2015). These methods lift the problem to a higher dimension, resulting in excessive complexity and memory requirements. To perform phase retrieval for high-dimensional problems with performance guarantees, a range of convex (Bahmani & Romberg, 2017; Goldstein & Studer, 2017; Hand & Voroninski, 2016; Dhifallah et al., 2017; Dhifallah & Lu, 2017; Yuan & Wang, 2017; Salehi et al., 2018) and nonconvex methods (Netrapalli et al., 2013; Schniter & Rangan, 2015; Candès et al., 2015b; Chen & Candès, 2015; Zhang & Liang, 2016; Wang et al., 2017a; Zhang et al., 2016; Wei, 2015; Sun et al., 2016; Zeng & So, 2017; Lu & Li, 2017; Ma et al., 2018) have been proposed recently.

1.2. Spectral Initializers

All of the above non-lifting-based phase retrieval methods rely on accurate initial estimates of the signal vector to be recovered. Such estimates are typically obtained by means of so-called *spectral initializers* put forward in (Netrapalli et al., 2013). Spectral initializers first compute a Hermitian matrix of the following form:

$$\mathbf{D}_\beta = \beta \sum_{m=1}^M \mathcal{T}(y_m) \mathbf{a}_m \mathbf{a}_m^H, \quad (2)$$

¹School of Electrical and Computer Engineering, Cornell University, Ithaca, NY ²Department of EE, Princeton University ³University of Maryland. Correspondence to: Ramina Ghods <rg548@cornell.edu>, Christoph Studer <studer@cornell.edu>.

where $\beta > 0$ is a suitably-chosen scaling factor, y_m denotes the m th measurement, \mathbf{a}_m^H corresponds to the m th row of the measurement matrix \mathbf{A} and $\mathcal{T} : \mathbb{R} \rightarrow \mathbb{R}$ is a (possibly nonlinear) *preprocessing function*. While the identity $\mathcal{T}(y) = y$ was used originally in (Netrapalli et al., 2013), recent results revealed that carefully crafted preprocessing functions yield more accurate estimates (Chen & Candès, 2015; Chen et al., 2015; Wang et al., 2017a;b; Lu & Li, 2017; Mondelli & Montanari, 2017). From the matrix \mathbf{D}_β in (2), one then extracts the (scaled) eigenvector $\hat{\mathbf{x}}$ associated with the largest eigenvalue, which serves as an initial estimate of the solution to the phase retrieval problem.

As shown in (Netrapalli et al., 2013; Chen & Candès, 2015; Chen et al., 2015; Wang et al., 2017a;b; Lu & Li, 2017; Mondelli & Montanari, 2017), for i.i.d. Gaussian measurement matrices \mathbf{A} , sufficiently large measurement ratios $\delta = M/N$, and carefully crafted preprocessing functions \mathcal{T} , spectral initializers provide accurate initialization vectors. In fact, the results in (Mondelli & Montanari, 2017) for the large-system limit with δ fixed and $M \rightarrow \infty$ show that spectral initializers in combination with an optimal preprocessing function \mathcal{T} achieve the fundamental information-theoretic limits of phase retrieval. However, the assumption of having i.i.d. Gaussian measurement matrices \mathbf{A} is impractical—it is more natural to assume that the signal vector \mathbf{x} is random and the measurement matrix \mathbf{A} is deterministic and structured (Bendory & Eldar, 2017).

1.3. Contributions

We propose a novel class of estimators, called *linear spectral estimators (LSPEs)*, that provide accurate estimates for general nonlinear measurement systems of the form (1) and enable a nonasymptotic mean-squared error (MSE) analysis. We showcase the efficacy of LSPEs by applying them to phase retrieval problems, where we compute initialization vectors for real- and complex-valued systems with deterministic and finite-dimensional measurement matrices. For the proposed LSPEs, we derive nonasymptotic and sharp bounds on the MSE for signal estimation from phaseless measurements. We use synthetic and real-world phase retrieval problems to demonstrate that LSPEs are able to significantly outperform existing spectral initializers on systems that acquire structured measurements. We furthermore show that preprocessing the phaseless measurements enables LSPEs to generate improved initialization vectors for an even broader class of measurement systems.

1.4. Notation

Lowercase and uppercase boldface letters represent column vectors and matrices, respectively. For a matrix \mathbf{A} , its transpose and Hermitian conjugate is \mathbf{A}^T and \mathbf{A}^H , respectively, and the k th row and ℓ th column entry is $[\mathbf{A}]_{k,\ell} = A_{k,\ell}$. For

a vector \mathbf{a} , the k th entry is $[\mathbf{a}]_k = a_k$. The ℓ_2 -norm of \mathbf{a} is denoted by $\|\mathbf{a}\|_2$ and the Frobenius norm of \mathbf{A} by $\|\mathbf{A}\|_F$. The Kronecker product is \otimes , the Hadamard product is \odot , the Hadamard division is \oslash , and the trace operator is $\text{tr}(\cdot)$. The $N \times N$ identity matrix is denoted by \mathbf{I}_N ; the $M \times N$ all-zeros and all-ones matrices are denoted by $\mathbf{0}_{M \times N}$ and $\mathbf{1}_{M \times N}$, respectively. For a vector \mathbf{a} , $\text{diag}(\mathbf{a})$ is a square matrix with \mathbf{a} on the main diagonal; for a matrix \mathbf{A} , $\text{diag}(\mathbf{A})$ is a column vector containing the diagonal elements of \mathbf{A} .

2. Linear Spectral Estimators

We start by reviewing the essentials of spectral initializers and then, introduce linear spectral estimators (LSPEs) for measurement systems of the form (1) with general nonlinearities f . We furthermore provide nonasymptotic expressions for the associated estimation error, and we compare our analytical results to that of conventional spectral initializers in (2). In Section 3, we will apply LSPEs to phase retrieval.

2.1. Spectral Estimation and Initializers

One of the key issues of the phase retrieval problem is the fact that if \mathbf{x} is a solution to (1), then $e^{j\phi}\mathbf{x}$ for any $\phi \in [0, 2\pi)$ is also a valid solution (assuming $\mathcal{H} = \mathbb{C}$). Put simply, the solution is nonunique up to a global phase shift. One way of combating this issue is to directly recover the outer product $\mathbf{x}\mathbf{x}^H$ instead of \mathbf{x} , which is unaffected by phase shifts; this insight is the key underlying lifting-based phase retrieval methods (Candès et al., 2013; Candès & Li, 2014; Candès et al., 2015a; Waldspurger et al., 2015). With this in mind, one could envision the design of an estimator that directly minimizes the conditional MSE:

$$\hat{\mathbf{x}} = \arg \min_{\tilde{\mathbf{x}} \in \mathcal{H}^N} \mathbb{E} [\|\mathbf{x}\mathbf{x}^H - \tilde{\mathbf{x}}\tilde{\mathbf{x}}^H\|_F^2 \mid \mathbf{y}]. \quad (3)$$

Here, expectation is with respect to the signal vector \mathbf{x} and the two noise sources \mathbf{e}^z and \mathbf{e}^y . This optimization problem resembles that of a posterior mean estimator (PME) which is, in general, difficult to derive, even for simple observation models—for phase retrieval, we have two additional challenges: (i) nonlinear phaseless measurements as in (1) and (ii) the quantity $\tilde{\mathbf{x}}\tilde{\mathbf{x}}^H$ has rank-1.

Spectral initializers avoid the issues of the estimator in (3) by first replacing the true outer product $\mathbf{x}\mathbf{x}^H$ with a so-called spectral estimator matrix \mathbf{D}_β as in (2) that depends on the measurement vector \mathbf{y} . In a second step, one then computes the best rank-1 approximation as follows:

$$\hat{\mathbf{x}} = \arg \min_{\tilde{\mathbf{x}} \in \mathcal{H}^N} \|\mathbf{D}_\beta - \tilde{\mathbf{x}}\tilde{\mathbf{x}}^H\|_F^2 \quad (4)$$

from which the estimate $\hat{\mathbf{x}}$ can be extracted. By performing an eigenvalue decomposition $\mathbf{D}_\beta = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^H$ with $\mathbf{U}^H\mathbf{U} = \mathbf{I}_M$ and the eigenvalues in the diagonal matrix

$\Lambda = \text{diag}([\lambda_1, \dots, \lambda_M]^T)$ are sorted in descending order of their magnitudes, a spectral initializer is given by the *scaled leading eigenvector* $\hat{\mathbf{x}} = \sqrt{\lambda_1} \mathbf{u}_1$. In practice, one can use power iterations to efficiently compute $\hat{\mathbf{x}}$.

2.2. Linear Spectral Estimators

We now propose a novel class of estimators, which we call *linear spectral estimators (LSPEs)*, that provide accurate estimates for general nonlinear measurement systems of the form (1). To this end, we borrow ideas from the spectral initializer, the PME in (3), and the linear phase retrieval algorithm put forward in (Ghods et al., 2018). In the first step, LSPEs apply a linear estimator to the nonlinear observations in $\mathcal{T}(\mathbf{y})$ to construct a *spectral estimator matrix* \mathbf{D}_y for which the *spectral MSE* (or matrix MSE) defined as

$$S\text{-MSE} = \mathbb{E} \left[\left\| \mathbf{D}_y - \mathbf{x}\mathbf{x}^H \right\|_F^2 \right] \quad (5)$$

is minimal. We restrict ourselves to spectral estimator matrices \mathbf{D}_y that are *affine* in $\mathcal{T}(\mathbf{y})$, i.e., are of the form

$$\mathbf{D}_y = \mathbf{W}_0 + \sum_{m=1}^M \mathcal{T}(y_m) \mathbf{W}_m \quad (6)$$

with $\mathbf{W}_m \in \mathcal{H}^{N \times N}$, $m = 0, \dots, M$. In the second step, we use the spectral estimator matrix \mathbf{D}_y to extract a (scaled) leading eigenvector as in (3), which is the *linear spectral estimate* of the signal vector \mathbf{x} . Intuitively, if we can construct a matrix \mathbf{D}_y from the preprocessed measurements in $\mathcal{T}(\mathbf{y})$ for which the S-MSE in (5) is minimal, then we expect that computing its best rank-1 approximation would yield an accurate estimate of the signal vector \mathbf{x} up to a global phase shift. We will justify this claim in Section 2.3.

Mathematically, we wish to compute a matrix \mathbf{D}_y of the form (6) that is the solution to the following problem:

$$\underset{\substack{\mathbf{W}_m \in \mathcal{H}^{N \times N} \\ m=0, \dots, M}}{\text{minimize}} \mathbb{E} \left[\left\| \widetilde{\mathbf{W}}_0 + \sum_{m=1}^M \mathcal{T}(y_m) \widetilde{\mathbf{W}}_m - \mathbf{x}\mathbf{x}^H \right\|_F^2 \right]. \quad (7)$$

Clearly, the spectral estimator matrix \mathbf{D}_y will depend on the measurement matrix \mathbf{A} , the statistics of the signal to be estimated \mathbf{x} and the two noise sources \mathbf{e}^z and \mathbf{e}^y , the nonlinearity f , as well as the preprocessing function \mathcal{T} . For this setting, we have the following general result which summarizes the LSPE; the proof is given in Appendix A.

Theorem 1 (Linear Spectral Estimator). *Let the measurement vector \mathbf{y} be a result of the general measurement model in (1) and select a preprocessing function \mathcal{T} . Define the vector $\overline{\mathcal{T}}(\mathbf{y}) = \mathbb{E}[\mathcal{T}(\mathbf{y})]$ and assume the matrix*

$$\mathbf{T} = \mathbb{E}[(\mathcal{T}(\mathbf{y}) - \overline{\mathcal{T}}(\mathbf{y}))(\mathcal{T}(\mathbf{y}) - \overline{\mathcal{T}}(\mathbf{y}))^T]$$

is full rank. Let $\mathbf{t} \in \mathbb{R}^M$ satisfy $\mathbf{T}\mathbf{t} = \mathcal{T}(\mathbf{y}) - \overline{\mathcal{T}}(\mathbf{y})$ and

$$\mathbf{V}_m = \mathbb{E}[(\mathcal{T}(y_m) - \overline{\mathcal{T}}(y_m))(\mathbf{x}\mathbf{x}^H - \mathbf{K}_x)]$$

for $m = 1, \dots, M$ with $\mathbf{K}_x = \mathbb{E}[\mathbf{x}\mathbf{x}^H]$. Then, the LSPE matrix that minimizes the S-MSE in (5) is given by

$$\mathbf{D}_y = \mathbf{K}_x + \sum_{m=1}^M t_m \mathbf{V}_m. \quad (8)$$

The linear spectral estimate $\hat{\mathbf{x}}$ is then given by the scaled leading eigenvector of the matrix \mathbf{D}_y in (8).

The vector \mathbf{t} is the only quantity in Theorem 1 that depends on the *actual* (nonlinear) observations contained in the measurement vector \mathbf{y} . All other quantities depend only on the first two moments of $\mathbf{x}\mathbf{x}^H$ as well as the considered signal, noise, and measurement models. The key features of the LSPE are as follows: (i) the involved quantities can often be computed in closed form (see Section 3 for two applications to phase retrieval) and (ii) LSPEs enable a nonasymptotic and sharp analysis of the associated estimation error.

Remark 1. *Theorem 1 requires the matrix \mathbf{T} to be invertible. This condition is satisfied in most practical situations with nondegenerate measurement matrices \mathbf{A} or in situations with nonzero measurement noise.*

2.3. Estimation Error Analysis of LSPEs

The remaining piece of the proposed LSPE is to show that the result of this two-step estimation procedure indeed yields a vector that is close to the signal vector \mathbf{x} . We start with the following result; the proof is given in Appendix B.

Theorem 2 (S-MSE of the LSPE). *Let the assumptions of Theorem 1 hold. Then, the S-MSE in (5) for the LSPE matrix in (8) is given by*

$$S\text{-MSE}_{\text{LSPE}} = C_{\mathbf{x}\mathbf{x}^H} - \sum_{m=1}^M \sum_{m'=1}^M [\mathbf{T}^{-1}]_{m,m'} \text{tr}(\mathbf{V}_m^H \mathbf{V}_{m'}) \quad (9)$$

with $C_{\mathbf{x}\mathbf{x}^H} = \mathbb{E}[\|\mathbf{x}\mathbf{x}^H - \mathbf{K}_x\|_F^2]$.

With this result, we are ready to establish a bound on the estimation error of the LSPE. The proof of the following result follows from Theorem 2 and is given in Appendix C.

Corollary 1 (LSPE Estimation Error). *Let the assumptions of Theorem 1 hold. Then, the estimation error (EER) of the LSPE satisfies the following inequality:*

$$EER_{\text{LSPE}} = \mathbb{E}[\|\hat{\mathbf{x}}\hat{\mathbf{x}}^H - \mathbf{x}\mathbf{x}^H\|_F^2] \leq 4 S\text{-MSE}_{\text{LSPE}}. \quad (10)$$

This result implies that by minimizing the S-MSE in (5) via (7), we are also reducing the EER of the LSPE. In other words, if the spectral error $\mathbf{E} = \mathbf{D}_y - \hat{\mathbf{x}}\hat{\mathbf{x}}^H$ is small, then the EER of the LSPE (10) will be small.

Remark 2. *Corollary 1 is nonasymptotic and depends on the instance of measurement matrix \mathbf{A} . This result is in stark contrast to existing performance bounds for spectral initializers (Netrapalli et al., 2013; Chen & Candès, 2015; Chen et al., 2015; Wang et al., 2017a;b) that strongly rely on randomness in the measurement matrix. In addition to randomness, the sharp performance guarantees in (Lu & Li, 2017; Mondelli & Montanari, 2017) focus on the asymptotic regime for which $\delta = M/N$ is fixed and $M \rightarrow \infty$.*

2.4. S-MSE of Spectral Initializers

We can also derive an exact expression for the S-MSE of the conventional spectral initializer in (2). We assume optimal scaling, i.e., the parameter β is set to minimize the S-MSE. The following result characterizes the S-MSE of such a scaled spectral initializer; the proof is given in Appendix D.

Proposition 1 (S-MSE of the Spectral Initializer). *Let \mathbf{D}_β be the conventional spectral initializer matrix in (2). Then, the optimally-scaled S-MSE defined as*

$$S\text{-MSE}_{SI} = \min_{\beta \in \mathcal{H}} \mathbb{E} [\|\mathbf{D}_\beta - \mathbf{x}\mathbf{x}^H\|_F^2] \quad (11)$$

is given by

$$S\text{-MSE}_{SI} = R_{\mathbf{x}\mathbf{x}^H} - \frac{\left| \sum_{m=1}^M \mathbf{a}_m^H \tilde{\mathbf{V}}_m \mathbf{a}_m \right|^2}{\sum_{m=1}^M \sum_{m'=1}^M \tilde{T}_{m,m'} |\mathbf{a}_m^H \mathbf{a}_{m'}|^2}, \quad (12)$$

where $R_{\mathbf{x}\mathbf{x}^H} = \mathbb{E} [\|\mathbf{x}\mathbf{x}^H\|_F^2]$, $\tilde{\mathbf{V}}_m = \mathbb{E} [\mathcal{T}(y_m)\mathbf{x}\mathbf{x}^H]$, $m = 1, \dots, M$, and $\tilde{\mathbf{T}} = \mathbb{E} [\mathcal{T}(\mathbf{y})\mathcal{T}(\mathbf{y})^T]$.

Since the matrix in (2) is a special case of the LSPE matrix in (6), we have the following simple yet important property:

$$S\text{-MSE}_{\text{LSPE}} \leq S\text{-MSE}_{SI}.$$

In words, the spectral MSE of the LSPE cannot be worse than that of a spectral initializer. As we will show in Section 4, LSPEs are able to outperform spectral initializers on both synthetic and real-world phase retrieval problems given that the same preprocessing function \mathcal{T} is used.

3. LSPEs for Phase Retrieval Problems

The LSPE provides a framework for estimating signal vectors from the general observation model in (1). To make the concept of LSPEs explicit and to demonstrate their efficacy in practice, we now show two application examples to phase retrieval in complex-valued systems. The LSPE for real-valued phase retrieval can be found in Appendix E.

3.1. Phase Retrieval without Preprocessing

We first focus on the case where the signal vector \mathbf{x} to be estimated and the measurement matrix \mathbf{A} are both complex-

valued. The phaseless measurements \mathbf{y} , however, remain real-valued. We need the following assumptions.

Assumptions 1. *Let $\mathcal{H} = \mathbb{C}$. Assume square absolute measurements $f(z) = |z|^2$ and the identity preprocessing function $\mathcal{T}(y) = y$. Assume that the signal vector $\mathbf{x} \in \mathbb{C}^N$ is i.i.d. circularly-symmetric complex Gaussian with covariance matrix $\mathbf{C}_\mathbf{x} = \sigma_x^2 \mathbf{I}_N$, i.e., $\mathbf{x} \sim \mathcal{CN}(\mathbf{0}_{N \times 1}, \sigma_x^2 \mathbf{I}_N)$. Assume that the signal noise vector \mathbf{e}^z is circularly-symmetric complex Gaussian with covariance matrix $\mathbf{C}_{\mathbf{e}^z}$, i.e., $\mathbf{e}^z \sim \mathcal{CN}(\mathbf{0}_{M \times 1}, \mathbf{C}_{\mathbf{e}^z})$, and the measurement noise vector \mathbf{e}^y is a real-valued Gaussian vector with mean $\bar{\mathbf{e}}^y$ and covariance matrix $\mathbf{C}_{\mathbf{e}^y}$, i.e., $\mathbf{e}^y \sim \mathcal{N}(\bar{\mathbf{e}}^y, \mathbf{C}_{\mathbf{e}^y})$. Furthermore assume that \mathbf{x} , \mathbf{e}^z , and \mathbf{e}^y are independent.*

Under these assumptions, we can derive the following LSPE which we call LSPE- \mathbb{C} ; the detailed derivations of this spectral estimator are given in Appendix G.

Estimator 1 (LSPE- \mathbb{C}). *Let Assumptions 1 hold. Then, the spectral estimation matrix is given by*

$$\mathbf{D}_\mathbf{y}^{\mathbb{C}} = \mathbf{K}_\mathbf{x} + \sum_{m=1}^M t_m \mathbf{V}_m, \quad (13)$$

where $\mathbf{K}_\mathbf{x} = \sigma_x^2 \mathbf{I}_N$, the vector $\mathbf{t} \in \mathbb{R}^M$ is given by the solution to the linear system $\mathbf{T}\mathbf{t} = \mathbf{y} - \bar{\mathbf{y}}$ with

$$\begin{aligned} \bar{\mathbf{y}} &= \text{diag}(\mathbf{C}_\mathbf{z}) + \bar{\mathbf{e}}^y \\ \mathbf{C}_\mathbf{z} &= \sigma_x^2 \mathbf{A}\mathbf{A}^H + \mathbf{C}_{\mathbf{e}^z} \\ \mathbf{T} &= \mathbf{C}_\mathbf{z} \odot \mathbf{C}_\mathbf{z}^* + \mathbf{C}_{\mathbf{e}^y} \end{aligned}$$

and $\mathbf{V}_m = \sigma_x^4 \mathbf{a}_m \mathbf{a}_m^H$, $m = 1, \dots, M$. The spectral estimate $\hat{\mathbf{x}}$ is given by the (scaled) leading eigenvector of $\mathbf{D}_\mathbf{y}^{\mathbb{C}}$ in (13). Furthermore, the S-MSE is given by Theorem 2.

We emphasize that the spectral estimator matrix in (13) resembles that of the conventional spectral initializer matrix (2) with the following key differences. First and foremost, each outer product contained in $\mathbf{V}_m = \sigma_x^4 \mathbf{a}_m \mathbf{a}_m^H$ in Estimator 1 is weighted by t_m , which is a function of all phaseless measurements in \mathbf{y} and of the covariance matrix $\mathbf{C}_\mathbf{x}$. In contrast, each outer product in the conventional spectral initializer matrix in (2) is only weighted by the associated measurement y_m . This difference enables the LSPE to weight each outer product depending on correlations in the phaseless measurements caused by structure in the matrix \mathbf{A} . Second, the spectral estimator matrix includes a mean term $\mathbf{K}_\mathbf{x}$, which is absent in the spectral initializer matrix. As we will show in Section 4, for the same preprocessing function \mathcal{T} , Estimator 1 is able to outperform spectral initializers for systems with structured measurement matrices \mathbf{A} . For large i.i.d. Gaussian measurement matrices, there is no particular correlation structure to exploit and LSPEs perform on par with spectral initializers.

3.2. Phase Retrieval with Exponential Preprocessing

To demonstrate the flexibility and generality of our framework, we now design an LSPE with an exponential preprocessing function for complex-valued phase retrieval. We derive the LSPE under the following assumptions.

Assumptions 2. Let $\mathcal{H} = \mathbb{C}$. Assume square absolute measurements $f(z) = |z|^2$ and the exponential preprocessing function $\mathcal{T}(y) = \exp(-\gamma y)$ with $\gamma > 0$, i.e., we consider

$$\mathcal{T}(y) = \exp(-\gamma(|z|^2 + \mathbf{e}^y)) \quad \text{and} \quad \mathbf{z} = \mathbf{A}\mathbf{x} + \mathbf{e}^z,$$

where the exponential function is applied element-wise to vectors. The remaining assumptions are the same as in Assumptions 1.

We now derive the following LSPE called LSPE-Exp; the derivation of this spectral estimator is given in Appendix H.

Estimator 2 (LSPE-Exp). Let Assumptions 2 hold. Then, the spectral estimation matrix is given by

$$\mathbf{D}_y^{\text{Exp}} = \mathbf{K}_x + \sum_{m=1}^M t_m \mathbf{V}_m, \quad (14)$$

where $\mathbf{K}_x = \sigma_x^2 \mathbf{I}_N$, the vector $\mathbf{t} \in \mathbb{R}^M$ is given by the solution to the linear system $\mathbf{T}\mathbf{t} = \mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y})$ with

$$\begin{aligned} \bar{\mathcal{T}}(\mathbf{y}) &= \mathbf{p}_\gamma \odot \mathbf{q}_\gamma \\ \mathbf{T} &= (\mathbf{p}_\gamma \mathbf{p}_\gamma^T) \odot \exp(\gamma^2 \mathbf{C}_{\mathbf{e}^y}) \odot (\mathbf{q}_\gamma \mathbf{q}_\gamma^T - \gamma^2 \mathbf{C}_z \odot \mathbf{C}_z^*) \\ &\quad - (\mathbf{p}_\gamma \mathbf{p}_\gamma^T) \odot (\mathbf{q}_\gamma \mathbf{q}_\gamma^T) \\ \mathbf{V}_m &= -\frac{\gamma \sigma_x^4 [\mathbf{p}_\gamma]_m}{(\gamma [\mathbf{C}_z]_{m,m} + 1)^2} \mathbf{a}_m \mathbf{a}_m^H, \quad m = 1, \dots, M, \end{aligned}$$

where we use the following definitions:

$$\begin{aligned} \mathbf{q}_\gamma &= \gamma \text{diag}(\mathbf{C}_z) + \mathbf{1}_{M \times 1} \\ \mathbf{p}_\gamma &= \exp(-\gamma \bar{\mathbf{e}}^y + \gamma^2 \frac{1}{2} \text{diag}(\mathbf{C}_{\mathbf{e}^y})) \\ \mathbf{C}_z &= \sigma_x^2 \mathbf{A} \mathbf{A}^H + \mathbf{C}_{\mathbf{e}^z}. \end{aligned}$$

The spectral estimate $\hat{\mathbf{x}}$ is given by the (scaled) leading eigenvector of $\mathbf{D}_y^{\text{Exp}}$ in (14). Furthermore, the S-MSE of this estimator is given by Theorem 2.

At first sight, the choice of the exponential preprocessing function used in Estimator 2 seems to be arbitrary. We emphasize, however, that this particular function is inspired by the asymptotically-optimal preprocessing function for properly-normalized Gaussian measurement ensembles proposed in (Mondelli & Montanari, 2017) which is given by

$$\mathcal{T}_{\text{opt}}(y) = \frac{y-1}{y + \sqrt{\delta} - 1}. \quad (15)$$

As it turns out, we can scale, negate, and shift the exponential preprocessing function $\mathcal{T}(y) = \exp(-\gamma y)$ to make it

take a similar shape as the function in (15). More concretely, exponential preprocessing as well as $\mathcal{T}_{\text{opt}}(y)$ enables one to attenuate the effect of measurements with large magnitude, which is also the idea underlying the class of orthogonal spectral initializers, as proposed in (Chen et al., 2015; Wang et al., 2017a;b), that perform well in practice.

4. Numerical Results

We now compare the performance of our LSPEs against existing spectral initializers proposed for phase retrieval on synthetic and real image data. All our results use the spectral initializers and experimental setups provided by PhasePack (Chandra et al., 2017).

4.1. Impact of Measurement Ensemble

We start by comparing the normalized MSE (N-MSE) defined as (Chandra et al., 2017)

$$N\text{-MSE} = \frac{\min_{\alpha \in \mathcal{H}} \|\mathbf{x} - \alpha \hat{\mathbf{x}}\|^2}{\|\mathbf{x}\|^2}$$

for a range of spectral initializers on different measurement ensembles. Specifically, we focus on the complex-valued case and consider (i) an i.i.d. Gaussian measurement matrix with signal dimension $N = 16$, (ii) an i.i.d. Gaussian measurement matrix with $N = 256$, and (iii) the structured ‘‘transmission matrix’’ used for image recovery through multiple scattering media as detailed in (Metzler et al., 2017). We vary the oversampling ratio $\delta = M/N$ and compare the N-MSE of the proposed complex-valued LSPEs, LSPE-C (Estimator 1) and LSPE-Exp (Estimator 2 with $\gamma = 0.001$), to the following spectral initializers: the original spectral initializer (Netrapalli et al., 2013; Candès et al., 2015a) called ‘‘spectral,’’ truncated spectral initializer (Chen & Candès, 2015) called ‘‘truncated,’’ weighted spectral initializer (Wang et al., 2017b) called ‘‘weighted,’’ amplitude spectral initializer (Wang et al., 2017a) called ‘‘amplitude,’’ orthogonal spectral initializer (Chen et al., 2015) called ‘‘orthogonal,’’ and the asymptotically-optimal spectral initializer (Mondelli & Montanari, 2017) called ‘‘optimal.’’ For the following synthetic experiments, we generate the signals to be recovered according to Assumptions 1 and Assumptions 2 for LSPE-C and LSPE-Exp, respectively.

Figure 1a shows that the proposed LSPEs significantly outperform all existing spectral initializers for small problem dimensions with Gaussian measurements; this improvement is even more pronounced for large oversampling ratios. The reason is that since we randomly generate a low-dimensional sensing matrix, the system will exhibit strong correlations among the measurements that can be exploited by LSPEs. For larger dimensions with Gaussian measurements, we see in Figure 1b that the proposed LSPEs do not provide an advantage over other methods. In fact, only LSPE-Exp is




Figure 1: Comparison of normalized MSE (N-MSE) as a function of the oversampling ratio $\delta = M/N$ for complex-valued phase retrieval with different spectral initializers and with different measurement matrices. The proposed LSPEs perform well on low-dimensional problems, for structured measurement ensembles, or at high oversampling ratios δ .

able to perform as well as the orthogonal spectral initializer, which achieves the best performance in this scenario. This behavior can be attributed to the facts that (i) for large random matrices there is no particular correlation structure among the measurements to exploit and (ii) ignoring measurements associated to large values in y_m is increasingly important. For structured measurements, as it is the case for the transmission matrix from (Metzler et al., 2017), we see in Figure 1c that LSPEs significantly outperform existing methods that are designed for random measurement ensembles. In this scenario, exponential preprocessing does not improve performance since correlations in the transmission matrix are dominating the performance.

4.2. S-MSE Expressions and Approximation Error

We now validate our theoretical S-MSE expressions in Theorem 2 and Proposition 1, and confirm the accuracy of the EER bound given in Corollary 1. In the following experiment, we set $M = 8N$ and vary the dimension N from 8 to 64. For each pair (M, N) , we randomly generate one instance of an i.i.d. circularly symmetric complex Gaussian measurement matrix and average the different errors (S-MSE and EER) over 10,000 Monte-Carlo trials. We consider a noiseless setting and assume identity preprocessing, i.e., $\mathcal{T}(y) = y$. The signal vectors are generated according to an i.i.d. circularly complex Gaussian random vector. From Figure 2, we see that our analytical S-MSE expressions for the LSPE-C and spectral initializers match their empirical values. We furthermore see that the empirical EER is only about 6 dB to 10 dB lower than our non-asymptotic upper bound given in Corollary 1.

4.3. Real-World Image Recovery

We finally illustrate the efficacy of LSPEs in a more realistic scenario. In particular, we show results for a real image reconstruction task by using LSPEs and spectral initializers




Figure 2: Comparison of the analytical and empirical spectral MSE (S-MSE) and estimation error (EER) for LSPEs and spectral initializers (SI) at oversampling ratio $\delta = 8$. Our analytical expressions in Theorem 2 and Proposition 1 match the empirical S-MSE; the upper bound in Corollary 1 accurately characterizes the empirical EER.

only, i.e., we are *not* using any additional phase retrieval algorithm. Our goal is to recover a 16×16 -pixel and a 40×40 -pixel image that was captured through a multiple scattering media using the deterministic and highly-structured transmission matrix as detailed in (Metzler et al., 2017). We compare the proposed LSPEs to the same set of spectral initializers as in Section 4.1. The signal priors are as in Assumptions 1 (LSPE-C) and Assumptions 2 (LSPE-Exp).

Figures 3 and 4 show the recovered images along with the N-MSE values. The proposed LSPEs (often significantly) outperform all spectral initializers in terms of visual quality as well as the N-MSE. This result confirms the observations made in Figure 1c that LSPEs outperform existing spectral initializers for structured measurement matrices. We note that exponential preprocessing for LSPEs does not noticeably improve the N-MSE (over LSPE-C) in this setting since correlations in the transmission measurement matrix are dominating the recovery performance.

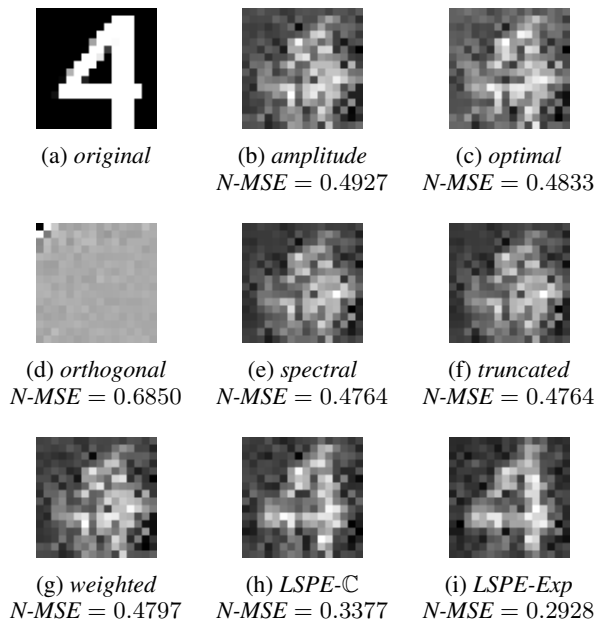


Figure 3: Recovery of a 16×16 image from with $M = 5N$ measurements captured through a scattering medium without the use of a phase retrieval algorithm. LSPEs outperform all initializers for structured measurements.

5. Conclusions

We have proposed a novel class of estimators, called *linear spectral estimators* (LSPEs), which are suitable for the recovery of signals from general nonlinear measurement systems. We have developed nonasymptotic and deterministic performance guarantees for LSPEs that provide accurate bounds on the estimation error, especially for structured or low-dimensional measurement systems. To demonstrate the efficacy of LSPEs in practice, we have applied them to complex-valued phase retrieval problems, in which LSPEs can be used to compute accurate signal estimates or initialization vectors for other convex or nonconvex phase retrieval algorithms. We have shown that properly preprocessing the nonlinear measurements can further improve the performance of LSPEs in practical scenarios. Our simulations with synthetic and real data have shown that LSPEs are able to significantly outperform existing spectral initializers, especially for low-dimensional problems, for structured measurement matrices, or for large oversampling ratios.

There are many avenues for future work. First, one could derive LSPEs for the asymptotically-optimal preprocessing function in (15) or for other commonly used functions, which may lead to further performance improvements. Second, the proposed error analysis could be used to generate improved measurement matrices. Third, an exploration of LSPEs for other nonlinearities that arise in machine learning and signal processing applications is left for future work.




Figure 4: Recovery of a 40×40 image from with $M = 10N$ measurements captured through a scattering medium without the use of a phase retrieval algorithm. LSPEs outperform all initializers for structured measurements.

Acknowledgments

R. Ghods and C. Studer were supported in part by Xilinx, Inc. and by the US National Science Foundation (NSF) under grants ECCS-1408006, EECS-1740286, CCF-1535897, CCF-1652065, and CNS-1717559. T. Goldstein was supported by the US NSF under grant CCF-1535902, the US ONR under grant N00014-15-1-2676, the DARPA Lifelong Learning Machines program, and the Sloan Foundation.

A. Proof of Theorem 1

The proof proceeds in two steps detailed as follows.

Mean Matrix We first compute the mean matrix \mathbf{W}_0 . Since (7) is a quadratic form, we can take the derivative in $\widetilde{\mathbf{W}}_0^H$ and set it to zero, i.e.,

$$\frac{d}{d\widetilde{\mathbf{W}}_0^H} \mathbb{E} \left[\left\| \widetilde{\mathbf{W}}_0 + \sum_{m=1}^M \mathcal{T}(y_m) \widetilde{\mathbf{W}}_m - \mathbf{x}\mathbf{x}^H \right\|_F^2 \right] = 0.$$

Basic matrix calculus yields

$$\widetilde{\mathbf{W}}_0 = \mathbf{K}_x - \sum_{m=1}^M \overline{\mathcal{T}}(y_m) \widetilde{\mathbf{W}}_m \quad (16)$$

with $\overline{\mathcal{T}}(y_m) = \mathbb{E}[\mathcal{T}(y_m)]$ and $\mathbf{K}_x = \mathbb{E}[\mathbf{x}\mathbf{x}^H]$.

Linear Estimation Matrix With (16) and the fact that (7) is a quadratic form in the matrices \mathbf{W}_m , $m = 1, \dots, M$,

we take the derivatives in \mathbf{W}_m^H and setting them to zero:

$$\frac{d}{d\mathbf{W}_m^H} \mathbb{E} \left[\left\| \sum_{m=1}^M (\mathcal{T}(y_m) - \bar{\mathcal{T}}(y_m)) \widetilde{\mathbf{W}}_m - (\mathbf{x}\mathbf{x}^H - \mathbf{K}_x) \right\|_F^2 \right] = 0.$$

By interchanging the derivative with expectation and with basic manipulations, we obtain the following set of optimality conditions for \mathbf{W}_m for $m = 1, \dots, M$:

$$\begin{aligned} \sum_{m'=1}^M \widetilde{\mathbf{W}}_{m'} \mathbb{E} [(\mathcal{T}(y_m) - \bar{\mathcal{T}}(y_m))(\mathcal{T}(y_{m'}) - \bar{\mathcal{T}}(y_{m'}))] \\ = \mathbb{E} [(\mathcal{T}(y_m) - \bar{\mathcal{T}}(y_m))(\mathbf{x}\mathbf{x}^H - \mathbf{K}_x)]. \end{aligned} \quad (17)$$

In compact matrix form, the above condition reads

$$(\mathbf{T} \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{W}} = \underline{\mathbf{V}}, \quad (18)$$

where we used the following shortcuts:

$$\begin{aligned} \mathbf{T} &= \mathbb{E} [(\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))(\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))^T] \\ \underline{\mathbf{W}} &= [\widetilde{\mathbf{W}}_1^T, \dots, \widetilde{\mathbf{W}}_m^T, \dots, \widetilde{\mathbf{W}}_M^T]^T \\ \mathbf{V}_m &= \mathbb{E} [(\mathcal{T}(y_m) - \bar{\mathcal{T}}(y_m))(\mathbf{x}\mathbf{x}^H - \mathbf{K}_x)], m = 1, \dots, M \\ \underline{\mathbf{V}} &= [\mathbf{V}_1^T, \dots, \mathbf{V}_m^T, \dots, \mathbf{V}_M^T]^T. \end{aligned}$$

The condition in (18) can be solved for the estimation matrices in $\underline{\mathbf{W}}$ leading to $\underline{\mathbf{W}} = (\mathbf{T}^{-1} \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{V}}$, where we require the matrix \mathbf{T} to be full rank. To obtain the linear spectral estimator matrix, we simplify as

$$\begin{aligned} \mathbf{D}_y &= \mathbf{K}_x + ((\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))^T \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{W}} \\ &= \mathbf{K}_x + \sum_{m=1}^M t_m \mathbf{V}_m, \end{aligned}$$

where we define the vector $\mathbf{t} = \mathbf{T}^{-1}(\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))$.

B. Proof of Theorem 2

To compute the spectral MSE in (5), we simplify

$$S\text{-MSE} = \mathbb{E} \left[\left\| \sum_{m=1}^M t_m \mathbf{V}_m - (\mathbf{x}\mathbf{x}^H - \mathbf{K}_x) \right\|_F^2 \right].$$

We expand this expression into four terms

$$\begin{aligned} &\mathbb{E} \left[\left\| \sum_{m=1}^M t_m \mathbf{V}_m - (\mathbf{x}\mathbf{x}^H - \mathbf{K}_x) \right\|_F^2 \right] \\ &= \mathbb{E} \left[\left\| \sum_{m=1}^M t_m \mathbf{V}_m \right\|_F^2 \right] \end{aligned} \quad (19)$$

$$+ \mathbb{E} \left[\left\| \mathbf{x}\mathbf{x}^H - \mathbf{K}_x \right\|_F^2 \right] - \mathbb{E} \left[\text{tr} \left((\mathbf{x}\mathbf{x}^H - \mathbf{K}_x)^H \left(\sum_{m=1}^M t_m \mathbf{V}_m \right) \right) \right] \quad (20)$$

$$- \mathbb{E} \left[\text{tr} \left(\left(\sum_{m=1}^M t_m \mathbf{V}_m \right)^H (\mathbf{x}\mathbf{x}^H - \mathbf{K}_x) \right) \right] \quad (21)$$

and simplify each expression individually. We start with (19) and use the fact that

$$\sum_{m=1}^M t_m \mathbf{V}_m = ((\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))^T \mathbf{T}^{-1} \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{V}}$$

and rewrite the quantity within expectation as follows:

$$\begin{aligned} &\underline{\mathbf{V}}^H (\mathbf{T}^{-1} (\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y})) \otimes \mathbf{I}_{N \times N}) \\ &\quad \times ((\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))^T \mathbf{T}^{-1} \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{V}} \\ &= \underline{\mathbf{V}}^H ((\mathbf{T}^{-1} (\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y})) \\ &\quad \times (\mathcal{T}(\mathbf{y}) - \bar{\mathcal{T}}(\mathbf{y}))^T \mathbf{T}^{-1}) \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{V}}. \end{aligned}$$

We now evaluate the expectation which leads to

$$\mathbb{E} \left[\left\| \sum_{m=1}^M t_m \mathbf{V}_m \right\|_F^2 \right] = \text{tr} \left(\underline{\mathbf{V}}^H (\mathbf{T}^{-1} \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{V}} \right)$$

or, equivalently, to

$$\mathbb{E} \left[\left\| \sum_{m=1}^M t_m \mathbf{V}_m \right\|_F^2 \right] = \sum_{m=1}^M \sum_{m'=1}^M [\mathbf{T}^{-1}]_{m,m'} \text{tr}(\mathbf{V}_m^H \mathbf{V}_{m'}).$$

We next will simplify (20). Recall that

$$t_m = \sum_{m'=1}^M [\mathbf{T}^{-1}]_{m,m'} (\mathcal{T}(y_{m'}) - \bar{\mathcal{T}}(y_{m'})),$$

which enables us to write (20) as

$$\begin{aligned} &\mathbb{E} \left[\text{tr} \left((\mathbf{x}\mathbf{x}^H - \mathbf{K}_x)^H \left(\sum_{m=1}^M t_m \mathbf{V}_m \right) \right) \right] \\ &= \sum_{m=1}^M \sum_{m'=1}^M [\mathbf{T}^{-1}]_{m,m'} \\ &\quad \times \text{tr} \left(\mathbb{E} [(\mathbf{x}\mathbf{x}^H - \mathbf{K}_x)^H (\mathcal{T}(y_{m'}) - \bar{\mathcal{T}}(y_{m'}))] \mathbf{V}_m \right) \\ &= \sum_{m=1}^M \sum_{m'=1}^M [\mathbf{T}^{-1}]_{m,m'} \text{tr}(\mathbf{V}_m^H \mathbf{V}_{m'}). \end{aligned} \quad (22)$$

Seeing as (21) is the Hermitian conjugate of (20), we have

$$\begin{aligned} &\mathbb{E} \left[\text{tr} \left(\left(\sum_{m=1}^M t_m \mathbf{V}_m \right)^H (\mathbf{x}\mathbf{x}^H - \mathbf{K}_x) \right) \right] \\ &= \sum_{m=1}^M \sum_{m'=1}^M [\mathbf{T}^{-1}]_{m,m'}^* \text{tr}(\mathbf{V}_m^H \mathbf{V}_{m'}). \end{aligned} \quad (23)$$

Combining all these terms yield the spectral MSE

$$\begin{aligned} S\text{-MSE} &= \mathbb{E} \left[\left\| \sum_{m=1}^M t_m \mathbf{V}_m - (\mathbf{x}\mathbf{x}^H - \mathbf{K}_x) \right\|_F^2 \right] \\ &= \mathbf{C}_{\mathbf{x}\mathbf{x}^H} - \text{tr} \left(\underline{\mathbf{V}}^H (\mathbf{T}^{-1} \otimes \mathbf{I}_{N \times N}) \underline{\mathbf{V}} \right). \end{aligned}$$

with $\mathbf{C}_{\mathbf{x}\mathbf{x}^H} = \mathbb{E} \left[\left\| \mathbf{x}\mathbf{x}^H - \mathbf{K}_x \right\|_F^2 \right]$.

C. Proof of Corollary 1

We bound the estimation error with the spectral MSE of the LSPE as follows. For a given instance, we have

$$\begin{aligned} \|\hat{\mathbf{x}}\hat{\mathbf{x}}^H - \mathbf{x}\mathbf{x}^H\|_F^2 &= \|\hat{\mathbf{x}}\hat{\mathbf{x}}^H - \mathbf{D}_y + \mathbf{D}_y - \mathbf{x}\mathbf{x}^H\|_F^2 \\ &\stackrel{(a)}{\leq} 2\|\hat{\mathbf{x}}\hat{\mathbf{x}}^H - \mathbf{D}_y\|_F^2 + 2\|\mathbf{D}_y - \mathbf{x}\mathbf{x}^H\|_F^2 \\ &\stackrel{(b)}{\leq} 4\|\mathbf{D}_y - \mathbf{x}\mathbf{x}^H\|_F^2, \end{aligned}$$

where (a) follows from the squared triangle inequality and (b) because $\hat{\mathbf{x}}\hat{\mathbf{x}}^H$ is the best rank-1 approximation of \mathbf{D}_y . Averaging over all instances finally yields

$$\mathbb{E} [\|\hat{\mathbf{x}}\hat{\mathbf{x}}^H - \mathbf{x}\mathbf{x}^H\|_F^2] \leq 4 S\text{-MSE}_{\text{LSPE}}.$$

References

- Bahmani, S. and Romberg, J. Phase retrieval meets statistical learning theory: A flexible convex relaxation. In *Proc. Intl. Conf. Artif. Intell. and Stat.*, pp. 252–260, May 2017.
- Bendory, T. and Eldar, Y. C. Non-convex phase retrieval from STFT measurements. *IEEE Trans. Inf. Theory*, Aug. 2017.
- Candès, E. J. and Li, X. Solving quadratic equations via PhaseLift when there are about as many equations as unknowns. *Found. Comput. Math.*, 14(5):1017–1026, Oct. 2014.
- Candès, E. J., Thomas, S., and Voroninski, V. PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. *Commun. Pure Appl. Math.*, 66(8):1241–1274, 2013.
- Candès, E. J., Eldar, Y. C., Strohmer, T., and Voroninski, V. Phase retrieval via matrix completion. *SIAM Rev.*, 57(2): 225–251, Nov. 2015a.
- Candès, E. J., Li, X., and Soltanolkotabi, M. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Trans. Inf. Theory*, 61(4):1985–2007, Feb. 2015b.
- Chandra, R., Zhong, Z., Hontz, J., McCulloch, V., Studer, C., and Goldstein, T. PhasePack: A phase retrieval library. *arXiv preprint: 1711.10175*, Nov. 2017.
- Chen, P., Fannjiang, A., and Liu, G. Phase retrieval with one or two diffraction patterns by alternating projections of the null vector. *arXiv preprint: 1510.07379*, Apr. 2015.
- Chen, Y. and Candès, E. J. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Proc. Adv. Neural Info. Proc. Syst.*, pp. 739–747, Dec. 2015.
- Dhifallah, O. and Lu, Y. Fundamental limits of PhaseMax for phase retrieval: A replica analysis. *arXiv preprint: 1708.03355*, Aug. 2017.
- Dhifallah, O., Thrampoulidis, C., and Lu, Y. Phase retrieval via linear programming: Fundamental limits and algorithmic improvements. *arXiv preprint: 1710.05234*, Oct. 2017.
- Faulkner, H. M. L. and Rodenburg, J. M. Movable aperture lensless transmission microscopy: A novel phase retrieval algorithm. *Phys. Rev. Lett.*, 93(2):023903, July 2004.
- Fienup, J. R. Phase retrieval algorithms: a comparison. *Appl. Opt.*, 21(15):2758–2769, Aug. 1982.
- Fogel, F., Waldspurger, I., and d’Aspremont, A. Phase retrieval for imaging problems. *Math. Prog. Comp.*, 8(3): 311–335, Sept. 2016.
- Gerchberg, R. W. and Saxton, W. O. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, Aug. 1972.
- Ghods, R., Lan, A. S., Goldstein, T., and Studer, C. PhaseLin: Linear phase retrieval. In *Proc. 52nd Ann. Conf. Info. Sci. Systems*, Mar. 2018.
- Goldstein, T. and Studer, C. Convex phase retrieval without lifting via PhaseMax. In *Proc. Intl. Conf. Mach. Learn.*, pp. 1273–1281, Aug. 2017.
- Hand, P. and Voroninski, V. An elementary proof of convex phase retrieval in the natural parameter space via the linear program PhaseMax. *arXiv preprint: 1611.03935*, Nov. 2016.
- Harrison, R. W. Phase problem in crystallography. *J. Opt. Soc. Am. A*, 10(5):1046–1055, May 1993.
- Holloway, J., Asif, M. S., Sharma, M. K., Matsuda, N., Horstmeyer, R., Cossairt, O., and Veeraraghavan, A. Toward long-distance subdiffraction imaging using coherent camera arrays. *IEEE Trans. Comput. Imag.*, 2(3):251–265, Sept. 2016.
- Kan, R. and Robotti, C. On moments of folded and truncated multivariate normal distributions. *J. Comput. Graph Stat.*, 26(4):1–5, Apr. 2017.
- Kou, S. S., Waller, L., Barbastathis, G., and Sheppard, C. J. R. Transport-of-intensity approach to differential interference contrast (TI-DIC) microscopy for quantitative phase imaging. *Opt. Lett.*, 35(3):447–449, Feb. 2010.
- Lu, Y. and Li, G. Phase transitions of spectral initialization for high-dimensional nonconvex estimation. *arXiv preprint: 1702.06435*, Apr. 2017.
- Ma, J., Xu, J., and Maleki, A. Optimization-based AMP for phase retrieval: The impact of initialization and ℓ_2 -regularization. *arXiv preprint: 1801.01170*, Jan. 2018.
- Metzler, C. A., Sharma, M. K., Nagesh, S., Baraniuk, R. G., Cossairt, O., and Veeraraghavan, A. Coherent inverse scattering via transmission matrices: Efficient phase retrieval algorithms and a public dataset. In *Proc. IEEE Intl. Conf. Comput. Photograph.*, pp. 1–16, May. 2017.
- Miao, J., Ishikawa, T., Shen, Q., and Earnest, T. Extending X-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes. *Ann. Rev. Phys. Chem.*, 59:387–410, Nov. 2008.
- Mondelli, M. and Montanari, A. Fundamental limits of weak recovery with applications to phase retrieval. *arXiv preprint: 1708.05932*, Sep. 2017.
- Netrapalli, P., Jain, P., and Sanghavi, S. Phase retrieval using alternating minimization. In *Proc. Adv. Neural Info. Proc. Syst.*, pp. 2796–2804, Dec. 2013.
- Pfeiffer, F., Weitkamp, T., Bunk, O., and David, C. Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources. *Nat. Phys.*, 2(4):258–261, Apr. 2006.

- Salehi, F., Abbasi, E., and Hassibi, B. A precise analysis of PhaseMax in phase retrieval. *arXiv preprint: 1801.06609*, Jan. 2018.
- Schniter, P. and Rangan, S. Compressive phase retrieval via generalized approximate message passing. *IEEE Trans. Sig. Process.*, 63(4):1043–1055, Feb. 2015.
- Sun, J., Qu, Q., and W., J. A geometric analysis of phase retrieval. *IEEE Intl. Symp. Info. Th.*, July 2016.
- Waldspurger, I., d’Aspremont, A., and Mallat, S. Phase recovery, maxcut and complex semidefinite programming. *Math. Prog.*, 149(1-2):47–81, Feb. 2015.
- Wang, G., Giannakis, G. B., and Eldar, Y. C. Solving systems of random quadratic equations via truncated amplitude flow. *arXiv preprint: 1605.08285*, Aug. 2017a.
- Wang, G., Giannakis, G. B., Saad, Y., and Chen, J. Solving almost all systems of random quadratic equations. *arXiv preprint: 1705.10407*, May 2017b.
- Wei, K. Solving systems of phaseless equations via Kaczmarz methods: A proof of concept study. *Inverse Probl.*, 31(12):125008, Nov. 2015.
- Yeh, L., Dong, J., Zhong, J., Tian, L., Chen, M., Tang, G., Soltanolkotabi, M., and Waller, L. Experimental robustness of Fourier ptychography phase retrieval algorithms. *Opt. Express*, 23(26):33214–33240, Dec. 2015.
- Yuan, Z. and Wang, H. Phase retrieval via reweighted Wirtinger flow. *Appl. Op.*, 56(9):2418, Mar. 2017.
- Zeng, W. and So, H. C. Coordinate descent algorithms for phase retrieval. *arXiv preprint: 1706.03474*, June 2017.
- Zhang, H. and Liang, Y. Reshaped Wirtinger flow for solving quadratic system of equations. In *Proc. Adv. Neural Info. Proc. Syst.*, pp. 2622–2630, Dec. 2016.
- Zhang, H., Chi, Y., and Liang, Y. Provable non-convex phase retrieval with outliers: Median truncated Wirtinger flow. In *Proc. Intl. Conf. Mach. Learn.*, pp. 1022–1031, June 2016.