Appl. Comput. Harmon. Anal. • • • (• • • •) • • •–• • •

Contents lists available at ScienceDirect

Applied and Computational Harmonic Analysis

www.elsevier.com/locate/acha

Letter to the Editor

Robust sparse phase retrieval made easy

Mark Iwen [a,*,1], Aditya Viswanathan [b], Yang Wang [c,2]

[a] Department of Mathematics and Department of ECE, Michigan State University, United States
[b] Department of Mathematics, Michigan State University, United States
[c] Department of Mathematics, The Hong Kong University of Science and Technology, Hong Kong

ARTICLE INFO

Article history:
Received 20 October 2014
Received in revised form 25 April 2015
Accepted 19 June 2015
Available online xxxx
Communicated by Thomas Strohmer

ABSTRACT

In this short note we propose a simple two-stage sparse phase retrieval strategy that uses a near-optimal number of measurements, and is both computationally efficient and robust to measurement noise. In addition, the proposed strategy is fairly general, allowing for a large number of new measurement constructions and recovery algorithms to be designed with minimal effort.

1. Introduction

Herein we consider the phase retrieval problem of reconstructing a given vector $\mathbf{x} \in \mathbb{C}^N$ from noisy magnitude measurements of the form

$$b_i := |\langle \mathbf{p}_i, \mathbf{x} \rangle|^2 + n_i, \tag{1}$$

where $\mathbf{p}_i \in \mathbb{C}^N$ is a measurement vector, and $n_i \in \mathbb{R}$ represents arbitrary measurement noise, for $i = 1, \ldots, M$. In particular, we focus on the setting where the dimension N is either very large, or else the number of measurements allowed, M, is otherwise severely restricted. In either case, our inability to gather the $M = \mathcal{O}(N)$ measurements required for the recovery of $\mathbf{x}$ in general [20] forces us to consider the possibility of approximating $\mathbf{x}$ using only $M \ll N$ magnitude measurements, if possible. This is the situation motivating the *compressive phase retrieval problem* (see, e.g., [30,31,26,24,34,15,32,35]), in which one attempts to accurately approximate $\mathbf{x} \in \mathbb{C}^N$ using only $M = o(N)$ magnitude measurements (1) under the assumption that $\mathbf{x}$ is either sparse, or compressible.

* Corresponding author.
E-mail addresses: markiwen@math.msu.edu (M. Iwen), aditya@math.msu.edu (A. Viswanathan), yangwang@ust.hk (Y. Wang).
[1] M.A. Iwen was supported in part by NSF DMS-1416752 and NSA H98230-13-1-0275.
[2] Y. Wang was partially supported by NSF DMS-1043032 and AFOSR FA9550-12-1-0455.

http://dx.doi.org/10.1016/j.acha.2015.06.007
1063-5203/

One question regarding the compressive phase retrieval problem is how many measurements are needed to allow for stable reconstruction of $\mathbf{x}$. Clearly, compressive phase retrieval requires at least as many measurements as the corresponding classical compressive sensing problem since one is given less information. Hence, stable compressive phase retrieval requires at least $\mathcal{O}(s\log(N/s))$ magnitude measurements[3] – but *can it be done with* $M = \mathcal{O}(s\log(N/s))$ *measurements*? It is shown in [15] that stable compressive phase retrieval is indeed achievable with $M = \mathcal{O}(s\log(N/s))$ measurements for *real* $\mathbf{x}$ if the entries of $\mathbf{p}_i$ are real independent and identically distributed (i.i.d.) Gaussians. However, this question was unresolved in the complex case. In this note we extend the result to the complex case. Furthermore, we do so in a constructive way by providing a computational procedure which can stably reconstruct complex $\mathbf{x}$ using only $\mathcal{O}(s\log(N/s))$ magnitude measurements.

Unlike previous sparse phase retrieval approaches, we propose a generic two-stage solution technique consisting of (i) using the phase retrieval technique of one's choice to recover compressive sensing measurements of $\mathbf{x}$, $\mathcal{C}\mathbf{x} \in \mathbb{C}^m$, followed by (ii) utilizing the compressive sensing method of one's choice in order to approximate $\mathbf{x}$ from the recovered measurements $\mathcal{C}\mathbf{x}$. As we shall see, the generic nature of the proposed sparse phase retrieval procedure not only allows for a relatively large number of measurement matrices and recovery algorithms to be used, but also allows robust recovery guarantees for the sparse phase retrieval problem to be proven in the complex setting essentially "for free" by combining existing robust recovery results from the compressive sensing literature with robust recovery results for the standard phase retrieval setting. As a result, we are able to show that $\mathcal{O}(s\log(N/s))$ magnitude measurements suffice in order to recover a large class of compressible vectors with the same quality of error guarantee as commonly achieved in the compressive sensing literature. Finally, numerical experiments demonstrate that the proposed approach is also both efficient and robust in practice.

2. Background

In this section we briefly recall selected results from the existing literature on compressive sensing [14,17] and phase retrieval [3,2,12,11,1,16]. Let $\|\mathbf{x}\|_0$ denote the number of nonzero entries in a given $\mathbf{x} \in \mathbb{C}^N$, and $\|\mathbf{x}\|_p$ denote the standard ℓ_p-norm of $\mathbf{x}$ for all $p \geq 1$, i.e., $\|\mathbf{x}\|_p := \left(\sum_{n=1}^N |x_n|^p\right)^{1/p}$ for all $\mathbf{x} \in \mathbb{C}^N$.

2.1. Compressive sensing

Compressive sensing methods deal with the construction of an $m \times N$ measurement matrix, $\mathcal{C}$, with m minimized as much as possible subject to the constraint that an associated approximation algorithm, $\Delta_{\mathcal{C}} : \mathbb{C}^m \to \mathbb{C}^N$, can still accurately approximate any given vector $\mathbf{x} \in \mathbb{C}^N$. More precisely, compressive sensing methods allow one to minimize m, the number of rows in $\mathcal{C}$, as a function of s and N such that

$$\left\|\Delta_{\mathcal{C}}\left(\mathcal{C}\mathbf{x}\right) - \mathbf{x}\right\|_p \leq C_{p,q} \cdot s^{\frac{1}{p}-\frac{1}{q}} \left(\inf_{\mathbf{z}\in\mathbb{C}^N, \|\mathbf{z}\|_0 \leq s} \|\mathbf{x}-\mathbf{z}\|_q\right) \tag{2}$$

holds for all $\mathbf{x} \in \mathbb{C}^N$ in various fixed ℓ_p,ℓ_q norms, $1 \leq q \leq p \leq 2$, for an absolute constant $C_{p,q} \in \mathbb{R}$ (e.g., see [13,17]). Note that this implies that $\mathbf{x}$ will be recovered exactly if it contains only s nonzero entries. Similarly, $\mathbf{x}$ will be accurately approximated by $\Delta_{\mathcal{C}}\left(\mathcal{C}\mathbf{x}\right)$ any time its ℓ_q-norm is dominated by its largest s entries.

There are a wide variety of measurement matrices $\mathcal{C} \in \mathbb{C}^{m\times N}$ with $m = \mathcal{O}(s\log(N/s))$ that have associated approximation algorithms, $\Delta_{\mathcal{C}}$, which are computationally efficient, numerically robust, and able to achieve error guarantees of the form (2) for all $\mathbf{x} \in \mathbb{C}^N$. For example, this is true of "most" random matrices $\mathcal{C} \in \mathbb{C}^{m\times N}$ with i.i.d. subgaussian random entries [4,17]. Similarly, one may construct such a $\mathcal{C} \in \mathbb{C}^{m\times N}$ with

[3] See, e.g., Chapter 10 of [17] concerning the minimal number of measurements required for stable compressive sensing.