Fast Compressive Phase Retrieval

Aditya Viswanathan and Mark Iwen

Department of Mathematics, Michigan State University
{aditya,markiwen}@math.msu.edu

MICHIGAN STATE UNIVERSITY

The Compressive Phase Retrieval Problem

Let $\mathbf{x} \in \mathbb{C}^n$ be a k-sparse signal, with $k \ll n$. Given squared magnitude measurements

$$\mathbf{y} = |\mathcal{M}\mathbf{x}|^2 + \mathbf{n},$$

where $\mathcal{M} \in \mathbb{C}^{m \times n}$ denotes a measurement matrix and $\mathbf{n} \in \mathbb{R}^m$ denotes measurement noise, the compressive phase retrieval problem seeks to recover the unknown signal $\mathbf{x}$ (upto some global phase offset) using only $m \ll n$ phaseless measurements, $\mathbf{y} \in \mathbb{R}^m$.

We are interested in measurement constructions $\mathcal{M}$ and associated recovery algorithms $\mathcal{A}_\mathcal{M} : \mathbb{R}^m \to \mathbb{C}^n$ which are efficient, use a minimal number of measurements, and are robust to measurement errors.

The phase retrieval problem occurs in several fields of science such as X-ray crystallography, optics, astronomy and quantum mechanics, where, either due to the underlying physics or instrumentation limitations, we are unable to acquire phase information.

Main Result

There exists a deterministic algorithm $\mathcal{A}_\mathcal{M} : \mathbb{R}^m \to \mathbb{C}^n$ for which the following holds: Let $\epsilon \in (0, 1]$, $\mathbf{x} \in \mathbb{C}^n$ with n sufficiently large, and $k \in \{1, 2, \ldots, n\} \subset \mathbb{N}$. Then, one can select a random measurement matrix $\mathcal{M} \in \mathbb{C}^{m \times n}$ such that

$$\min_{\theta \in [0, 2\pi)} \left\| e^{i\theta}\mathbf{x} - \mathcal{A}_\mathcal{M}\left(|\mathcal{M}\mathbf{x}|^2\right)\right\|_2 \leq \left\|\mathbf{x} - \mathbf{x}_k^{\text{opt}}\right\|_2 + \frac{22\epsilon \left\|\mathbf{x} - \mathbf{x}^{\text{opt}}_{(k/\epsilon)}\right\|_1}{\sqrt{k}}$$

is true with probability at least $1 - \frac{1}{C \cdot \log^2(n) \cdot \log^3(\log n)}$. Here, m can be chosen to be $\mathcal{O}\left(\frac{k}{\epsilon} \cdot \log^3\left(\frac{k}{\epsilon}\right) \cdot \log^3\left(\log \frac{k}{\epsilon}\right) \cdot \log n\right)$. Furthermore, the algorithm will run in $\mathcal{O}\left(\frac{k}{\epsilon} \cdot \log^4\left(\frac{k}{\epsilon}\right) \cdot \log^3\left(\log \frac{k}{\epsilon}\right) \cdot \log n\right)$–time.

This is the *first sub-linear time* compressive phase retrieval algorithm.

Both the sampling and runtime complexities are *sub-linear* in the problem size and *(poly)log-linear* in the sparsity.

Proposed Algorithm

Let $\mathcal{P} \in \mathbb{C}^{m \times d}$ denote an admissible phase retrieval matrix associated with the phase retrieval method Δ_P, and let $\mathcal{C} \in \mathbb{C}^{d \times n}$ denote a compressive sensing matrix associated with the sub–linear time compressive sensing algorithm Δ_C.

Construct the measurement matrix $\mathcal{M}$ for the compressive phase retrieval problem as $\mathcal{M} = \mathcal{P}\mathcal{C}$. Now, consider the following simple two stage formulation:

1. Apply a fast phase retrieval method, $\Delta_P : \mathbb{R}^m \to \mathbb{C}^d$, to the phaseless measurements $\mathbf{y}$ and recover an intermediate compressed signal $\mathbf{z} \in \mathbb{C}^d$, where $d = \mathcal{O}(k \log k \cdot \log n)$.
2. Next, use a *sub–linear time* compressive sensing algorithm, $\Delta_C : \mathbb{C}^d \to \mathbb{C}^n$, to recover the unknown signal $\mathbf{x}$.

We can show that $\Delta_C \circ \Delta_P : \mathbb{R}^m \to \mathbb{C}^n$ recovers the unknown signal $\mathbf{x}$ upto a global phase factor accurately and stably.

Ingredients: (I) Fast (Non-Sparse) Phase Retrieval

Use measurement constructions $\mathcal{P}$ arising from local correlation-based measurements. For example, with noiseless measurements $\mathbf{y} \in \mathbb{R}^{12}$, $\mathbf{z} \in \mathbb{C}^4$, and $\mathcal{P} \in \mathbb{C}^{12 \times 4}$, we have

$$\mathcal{P} = \begin{pmatrix} \mathcal{P}_1 \\ \mathcal{P}_2 \\ \mathcal{P}_3 \end{pmatrix}, \quad \mathcal{P}_i \in \mathbb{C}^{4 \times 4},\ i \in \{1,2,3\}, \quad \mathcal{P}_i = \begin{pmatrix} (\mathbf{p}_i)_1^* & (\mathbf{p}_i)_2^* & 0 & 0 \\ 0 & (\mathbf{p}_i)_1^* & (\mathbf{p}_i)_2^* & 0 \\ 0 & 0 & (\mathbf{p}_i)_3^* & (\mathbf{p}_i)_4^* \\ (\mathbf{p}_i)_2^* & 0 & 0 & (\mathbf{p}_i)_1^* \end{pmatrix}.$$

This corresponds to (squared magnitude) *correlation* measurements of the intermediate compressed signal $\mathbf{z} \in \mathbb{C}^4$ with three *local* masks $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{C}^4$, where $(p_i)_\ell = 0$ for $\ell > 2, i \in \{1,2,3\}$. Writing out the correlation sum explicitly and setting $\delta = 2$, we obtain

$$(y_i)_\ell = \left|\sum_{k=1}^{\delta} (\mathbf{p}_i)_k^* \cdot z_{\ell+k-1}\right|^2 = \sum_{j,k=1}^{\delta} (\mathbf{p}_i)_j\, (\mathbf{p}_i)_k^*\, z_{\ell+j-1}\, z_{\ell+k-1}^* := \sum_{j,k=1}^{\delta} (\mathbf{p}_i)_{j,k}\, z_{\ell+j-1}\, z_{\ell+k-1}^*,$$

where we have used the notation $(\mathbf{p}_i)_{j,k} := (\mathbf{p}_i)_j (\mathbf{p}_i)_k^*$. The resulting *linear* system of equations for the (scaled) phase differences $\{z_i z_j^*\}$ may be written as

$$\begin{pmatrix} (y_1)_1 \\ (y_2)_1 \\ (y_3)_1 \\ (y_1)_2 \\ (y_2)_2 \\ (y_3)_2 \\ (y_1)_3 \\ (y_2)_3 \\ (y_3)_3 \\ (y_1)_4 \\ (y_2)_4 \\ (y_3)_4 \end{pmatrix} = \begin{pmatrix}
(\mathbf{p}_1)_{1,1} & (\mathbf{p}_1)_{1,2} & (\mathbf{p}_1)_{2,1} & (\mathbf{p}_1)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
(\mathbf{p}_2)_{1,1} & (\mathbf{p}_2)_{1,2} & (\mathbf{p}_2)_{2,1} & (\mathbf{p}_2)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
(\mathbf{p}_3)_{1,1} & (\mathbf{p}_3)_{1,2} & (\mathbf{p}_3)_{2,1} & (\mathbf{p}_3)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & (\mathbf{p}_1)_{1,1} & (\mathbf{p}_1)_{1,2} & (\mathbf{p}_1)_{2,1} & (\mathbf{p}_1)_{2,2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & (\mathbf{p}_2)_{1,1} & (\mathbf{p}_2)_{1,2} & (\mathbf{p}_2)_{2,1} & (\mathbf{p}_2)_{2,2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & (\mathbf{p}_3)_{1,1} & (\mathbf{p}_3)_{1,2} & (\mathbf{p}_3)_{2,1} & (\mathbf{p}_3)_{2,2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_1)_{1,1} & (\mathbf{p}_1)_{1,2} & (\mathbf{p}_1)_{2,1} & (\mathbf{p}_1)_{2,2} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_2)_{1,1} & (\mathbf{p}_2)_{1,2} & (\mathbf{p}_2)_{2,1} & (\mathbf{p}_2)_{2,2} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_3)_{1,1} & (\mathbf{p}_3)_{1,2} & (\mathbf{p}_3)_{2,1} & (\mathbf{p}_3)_{2,2} & 0 & 0 \\
(\mathbf{p}_1)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_1)_{1,1} & (\mathbf{p}_1)_{1,2} & (\mathbf{p}_1)_{2,1} \\
(\mathbf{p}_2)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_2)_{1,1} & (\mathbf{p}_2)_{1,2} & (\mathbf{p}_2)_{2,1} \\
(\mathbf{p}_3)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_3)_{1,1} & (\mathbf{p}_3)_{1,2} & (\mathbf{p}_3)_{2,1}
\end{pmatrix} \begin{pmatrix} |z_1|^2 \\ z_1 z_2^* \\ z_2 z_1^* \\ |z_2|^2 \\ z_2 z_3^* \\ z_3 z_2^* \\ |z_3|^2 \\ z_3 z_4^* \\ z_4 z_3^* \\ |z_4|^2 \\ z_4 z_1^* \\ z_1 z_4^* \end{pmatrix}.$$

This is a *block-circulant* system which can be inverted efficiently using FFTs. Moreover, both random and *deterministic* prescriptions for the measurement masks $\mathbf{p}_i$ are available, and we can show that the resulting system is well-conditioned (see [1] for details). Note that by solving this linear system, we automatically obtain $|\mathbf{z}|$. Moreover, we can solve an *angular synchronization* problem using an eigenvector method to recover $\arg \mathbf{z}$.

$$\underbrace{\begin{pmatrix} |z_1|^2 & z_1 z_2^* & 0 & z_1 z_4^* \\ z_2 z_1^* & |z_2|^2 & z_2 z_3^* & 0 \\ 0 & z_3 z_2^* & |z_3|^2 & z_3 z_4^* \\ z_4 z_1^* & 0 & z_4 z_3^* & |z_4|^2 \end{pmatrix}}_{2\delta - 1 \text{ entries in band}} \xrightarrow{\text{normalize}} \begin{pmatrix} 1 & e^{i(\phi_1-\phi_2)} & 0 & e^{i(\phi_1-\phi_4)} \\ e^{i(\phi_2-\phi_1)} & 1 & e^{i(\phi_2-\phi_3)} & 0 \\ 0 & e^{i(\phi_3-\phi_2)} & 1 & e^{i(\phi_3-\phi_4)} \\ e^{i(\phi_4-\phi_1)} & 0 & e^{i(\phi_4-\phi_3)} & 1 \end{pmatrix} \xrightarrow[\text{eigenvector}]{\text{leading}} \begin{pmatrix} e^{i\phi_1} \\ e^{i\phi_2} \\ e^{i\phi_3} \\ e^{i\phi_4} \end{pmatrix}.$$

Note that the leading eigenvector may be computed using the power method in essentially linear-time (see [3] for details). The above framework recovers "flat" (i.e., non-sparse) $\mathbf{z}$. To recover arbitrary vectors, we multiply $\mathcal{P}$ with a random unitary matrix (a fast Johnson-Lindenstrauss transform) to "flatten" $\mathbf{z}$.

The sampling and runtime complexities for this method are $\mathcal{O}(d \cdot \log^2 d \cdot \log^3(\log d))$ and $\mathcal{O}(d \cdot \log^3 d \cdot \log^3(\log d))$ respectively.

Ingredients: (II) Sub-Linear Time Compressive Sensing

Choose the measurement matrix $\mathcal{C}$ to be a random sparse binary matrix obtained by randomly sub-sampling rows of a well-chosen incoherent matrix (for example, the adjacency matrix of certain unbalanced expander graphs). In [2], it is shown that these matrices satisfy certain combinatorial properties which permit the use of fast compressed sensing recovery algorithms.

The recovery algorithm then proceeds in two phases:

1. Identify the k largest magnitude entries of $\mathbf{x}$ using standard bit-testing techniques.
2. Estimate these k largest entries using median estimates and techniques from computer science streaming literature.

The sampling and runtime complexities of this method are both $\mathcal{O}(k \cdot \log k \cdot \log n)$.

Numerical Results

Left panel figure shows execution time as a function of sparsity for various problem dimensions. We observe that the overall execution time is sub-linear in the problem size n and (poly) log-linear in the sparsity k.

The right panel figure illustrates robustness of the method to (i.i.d. Gaussian) measurement noise. It plots the reconstruction error in dB as a function of the added noise level (in dB) for a length $n = 2^{20}$ signal using less than 10% of measurements.

In both cases, complex sparse test signals with i.i.d complex Gaussian non-zero entries were used, with non-zero index locations chosen by k-permutations.

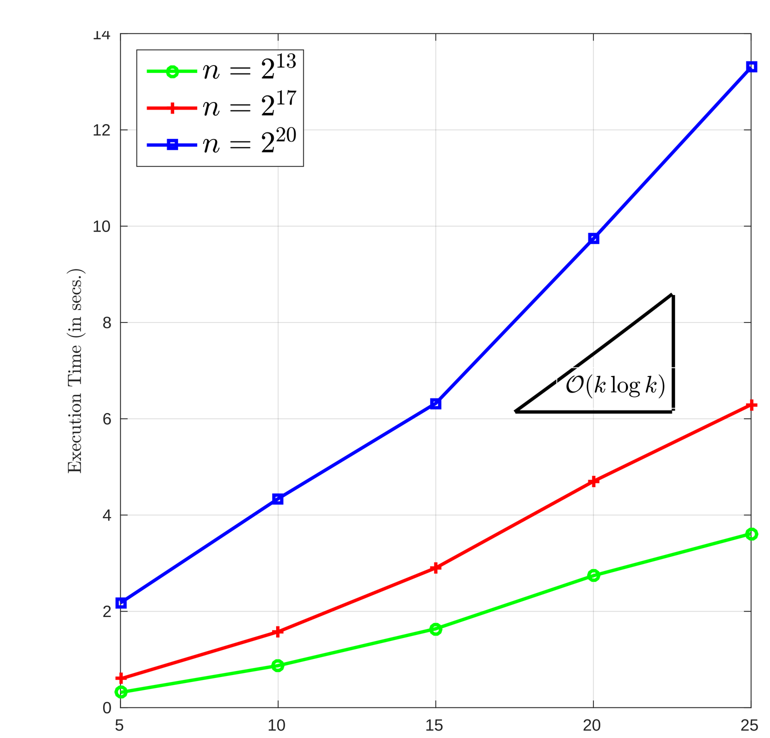

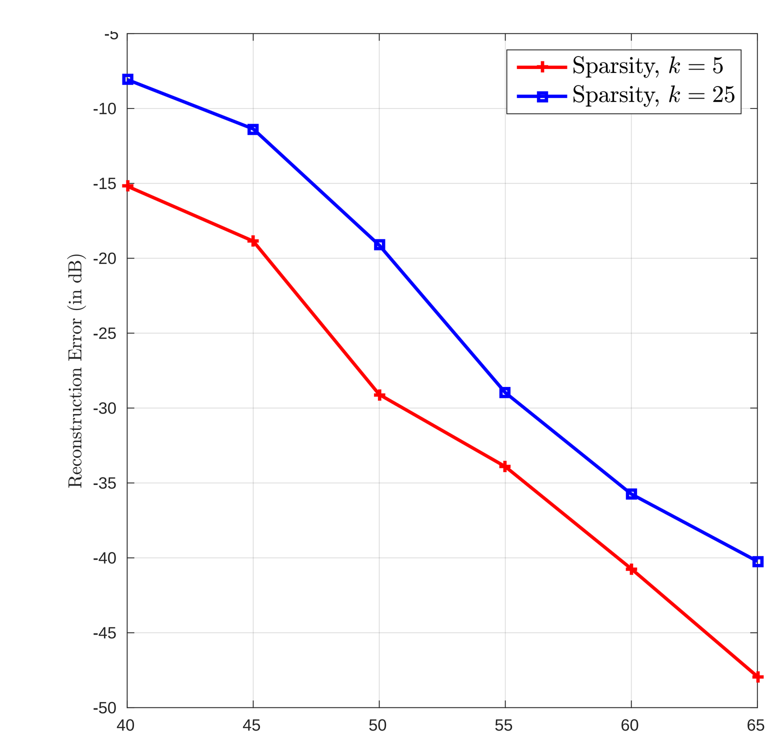

References and Acknowledgement

1. M. Iwen, A. Viswanathan and Yang Wang. *Fast Phase Retrieval for High-Dimensions*. arXiv:1501.02377, 2015.
2. M. Iwen. *Compressed Sensing with Sparse Binary Matrices: Instance Optimal Error Guarantees in Near-Optimal Time*. J. Complexity, Vol. 30, Issue 1, pp. 1 – 15, 2014.
3. A. Viswanathan and Mark Iwen. *Fast Angular Synchronization for Phase Retrieval via Incomplete Information*. Proc. SPIE 9597, Wavelets and Sparsity XVI, 959718, Aug. 2015.

This work was supported in part by NSF grant DMS-1416752 and NSA grant H98230-13-1-0275.