

SUB-DIFFRACTION IMAGING USING FOURIER PTYCHOGRAPHY AND STRUCTURED SPARSITY

Gauri Jagatap, Zhengyu Chen, Chinmay Hegde, Namrata Vaswani

Electrical and Computer Engineering, Iowa State University, Ames, IA 50011

ABSTRACT

We consider the problem of super-resolution for sub-diffraction imaging. We adapt conventional Fourier ptychographic approaches, for the case where the images to be acquired have an underlying structured sparsity. We propose some subsampling strategies which can be easily adapted to existing ptychographic setups. We then use a novel technique called CoPRAM with some modifications, to recover sparse (and block sparse) images from subsampled ptychographic measurements. We demonstrate experimentally that this algorithm performs better than existing phase retrieval techniques, in terms of quality of reconstruction, using fewer number of samples.

Index Terms— Phase retrieval, ptychography, sparsity, non-convex algorithms

1. INTRODUCTION

1.1. Motivation

Diffraction blurring is a common optical phenomenon encountered in imaging scenarios due to two reasons: if the aperture of the lens used for imaging is too small; or alternatively, if the lens is placed too far from the object to be imaged. The rays of light that interact with the end points of the camera aperture undergo diffraction, which leads to a diffraction pattern. The central disc of this pattern can have diameter exceeding the spatial resolution of the object or scene to be imaged. In such cases, one observes a diffraction blur. This effect has been well modeled in the context of super-resolution microscopic imaging [1, 2, 3]. Recently, Holloway et al. in [4, 5] have studied this model for long-distance imaging using coherent camera arrays using a technique known as Fourier ptychography [2].

To mitigate the effects of diffraction blur, one approach is to model the image reconstruction in terms of inverting a series of optical operations, given magnitude-only measurements. This formulation fits perfectly into the mold of *phase retrieval* [6, 7, 8]. However, phase retrieval is a highly nonlinear, ill-posed inverse problem with (mostly) heuristic solutions. A recent line of breakthrough results in the phase-retrieval literature [9, 10, 11] have showcased *provably* efficient algorithms for the special case where the measurement vectors arise from certain multi-variate probability distributions. However, the applicability of these methods to Fourier ptychography have not been explored in depth.

A basic challenge for ptychographic reconstruction methods is the requirement of an *overcomplete* set of observations; the “sampling rate” m in order to reconstruct a signal of length n must satisfy $m = \Omega(n)$ or more. A particularly compelling approach to reduce this burden is by leverage the *sparsity* of the target image; if the signal is s -sparse in a given basis (with $s \ll n$), the hope is that far

fewer samples suffice. The assumption of sparsity is natural in several applications in imaging systems, such as sub-diffraction imaging, X-ray crystallography, bio-imaging and astronomical imaging [12, 13, 14]; in these applications, the object to be imaged is often modeled as sparse in the canonical (or wavelet) basis.

Moving beyond sparsity, several algorithms that leverage more refined *structured sparsity* modeling assumptions (such as block sparsity) achieve considerably improved sample-complexity for reconstructing images from compressive measurements [15, 16, 17]. However, analogous algorithms in the context of phase retrieval are not well-studied. In recent previous work [18], we have developed a theoretically-sound algorithmic approach for phase retrieval that integrate sparsity (as well as structured sparsity) modeling assumptions within the reconstruction process. However, that work also assumes certain stringent probabilistic assumptions on the measurement process, similar to the approaches of [9, 10, 11].

1.2. Our Contributions

In this paper, we initiate a new approach for sub-diffraction imaging that combines Fourier ptychography with (structured) sparsity-based phase retrieval methods. In particular, we make the following contributions:

1. We propose suitable “under-sampling” strategies for Fourier ptychography, which can potentially reduce the number of samples required for image reconstruction;
2. We propose and test the efficacy of new sparsity-based algorithms for solving the Fourier ptychography problem;
3. We provide an intuition for a new initialization strategy, and comment on the convergence and sample complexity of the problem;
4. We support these claims via a series of experiments.

In our previous work [18], we have developed an algorithm called *Compressive Phase Retrieval using Alternating Minimization*, or CoPRAM), which significantly lowers the number of samples required for phase retrieval from Gaussian measurements using sparsity modeling assumptions.

At a high level, CoPRAM uses a thresholding-based spectral procedure to provide a coarse initial estimate of the underlying signal, and successfully refines this estimate by a variant of (classical) alternating minimization [19]. Our approach follows the same procedure in the context of Fourier ptychography, albeit with a somewhat simpler initialization procedure which we describe below. This paper focuses on static images that obey (structured) sparsity assumptions. In a companion paper [20], we develop algorithms for Fourier ptychography for dynamic scenes that obey low-rank assumptions.

1.3. Prior Work

Phase retrieval is a long-standing challenge in optical signal processing, dating back to the early work of [19]. However, the classical approach in phase retrieval is to alternately estimate the phase and the

This work is supported in parts by the National Science Foundation under the grants CCF-1566281, CCF-1526870 and IIP-1632116




Fig. 1: Sampling procedure, using operator \mathcal{A}_i . The green box indicates extra subsampling step. Camera index is denoted by $i = [N]$.

signal in an iterative fashion, and convergence of such an approach is not always guaranteed. In the seminal paper [11], Netrapalli et al. proposed the first rigorous theoretical claims for analyzing alternating minimization for the special case of Gaussian measurements. Their results were improved upon by Candes et al. in [10] and subsequent works [21, 22, 23].

Related work on integrating sparsity assumptions within phase retrieval includes a variant of alternating minimization [11], methods based on convex relaxation [24, 25, 26] and iterative thresholding-based techniques [27, 28, 18]. For recovering s -sparse signals of length n , all of the above techniques incur a sample complexity of $\mathcal{O}(s^2 \log n)$ for stable recovery, which is an improvement upon the standard limit of $\mathcal{O}(n)$ for $s \ll n$. Our recent work [18] suggests that it is possible to reduce the sample complexity even further (in fact matching the optimal limit $\mathcal{O}(s \log n)$) if the target signal is known a priori to exhibit block sparsity with sufficiently large block size. Moreover, some algorithms which use carefully designed measurements such as constrained sensing vectors, or Fourier-like measurements show a complexity of $\mathcal{O}(s \log n)$ [29, 30, 31].

The Fourier and ptychographic literature, on the other hand, has primarily focused on experimental advantages [1, 2, 3, 32]; however, few rigorous theoretical guarantees exist. Recently the works [33, 34, 35] have demonstrated theoretical guarantees for the convergence of phase retrieval algorithms for Short Time Fourier Transform (STFT) measurements; while these methods could be plausibly used in ptychography applications but these works only focus on simple synthetic test cases.

In [4], the authors have demonstrated a method for super-resolution of diffraction-blurred images using alternating minimization. Moreover, they demonstrate a number of experiments on how the quality of reconstruction is affected by parameters such as amount of overlap between consecutive cameras, aperture size, noise, etc. However, they do not explore the effect of integrating sparsity (or structured sparsity) models in order to improve reconstruction quality and/or reduce measurement rates. Our work in this paper demonstrates clear numerical evidence towards this objective.

2. FORMULATION

2.1. Optical Setup

A Fourier ptychographic setup, such as in [4], involves imaging a long-distance object using a series of sensing operations. The object is said to be illuminated using coherent light. A thin lens placed in front of the object results in a phase shift, which transforms the image from spatial to Fourier domain. This Fourier domain image is acquired via an camera array with aperture pupils placed in a square grid, with significant overlap between consecutive lenses (such an arrangement is called a coherent camera array). This has the effect of simulating a large effective “synthetic” aperture. The signal then

undergoes an inverse Fourier transform (due to a second phase shift from the acquisition camera lens). This complex, spatial domain image is captured by the optical sensor, which is only capable of recording magnitudes of the image pixels (and loses phase information).

The entire sensing procedure is described in Fig. 1. To recover the original image from such magnitude-only measurements, [4] uses a variant of the alternating minimization approach of [19] with an extra regularization mechanism that limits the norm of the signal estimate. They call this the Iterative Error Reduction Algorithm (IERA), which serves as our primary baseline for comparisons.

2.2. Mathematical Model

We model the Fourier ptychography phase retrieval problem as follows. A signal (or vectorized image) $\mathbf{x} \in \mathbb{C}^n$ is acquired in the form of measurements $\mathbf{y}_i \in \mathbb{C}^n$. The linear operators $\mathcal{A}_i : \mathbb{C}^n \rightarrow \mathbb{C}^n$ represents the traversal of the optical signal through the measurement system prior to the sensor measurement step. The measured images are given by:

$$\mathbf{y}_i = |\mathcal{A}_i(\mathbf{x})|, \quad (1)$$

where i is an index that spans the number of cameras in a camera array grid ($i = 1, 2, \dots, N$). We can equivalently denote this as:

$$\begin{aligned} \mathbf{y} &= |\mathcal{A}(\mathbf{x})|, \quad \text{where} \\ \mathcal{A} &= [\mathcal{A}_1^\top \dots \mathcal{A}_i^\top \dots \mathcal{A}_N^\top], \\ \mathbf{y} &= [\mathbf{y}_1^\top \dots \mathbf{y}_i^\top \dots \mathbf{y}_N^\top], \end{aligned}$$

and $\mathbf{y} \in \mathbb{C}^{nN}$ and $\mathcal{A} : \mathbb{C}^n \rightarrow \mathbb{C}^{nN}$. The optical setup in [4] can be described mathematically as:

$$\mathcal{A}_i = \mathcal{M}_i \mathcal{F}^{-1} \mathcal{P}_i \circ \mathcal{F} \quad \text{and} \quad \mathcal{A}_i^\top = \mathcal{F}^{-1} \mathcal{P}_i \circ \mathcal{F} \mathcal{M}_i, \quad (2)$$

where \mathcal{P}_i is a pupil mask corresponding to the i^{th} camera. The operators \mathcal{P}_i constitute a series of bandpass filters which cover different parts of the Fourier domain image, and the symbol \circ represents the Hadamard product. The subsampling masks \mathcal{M}_i resembles the operation of an *identity*, in the conventional setup (i.e. all measurements are retained). The number of cameras and overlap between subsequent cameras is chosen such that the entire Fourier spectrum is captured by the camera array.

We note that this setup is not designed to benefit from a sparsity constraint. Below, we discuss two subsampling procedures which, when coupled with appropriate sparsity-driven reconstruction procedures, provide better images at comparable sampling rates than the standard IERA. Specifically, we propose to design subsampling masks \mathcal{M}_i , as illustrated in the form of green boxes in Fig. 1. The subsampling mask \mathcal{M}_i can be designed in different ways, and we discuss these techniques in Section 4.

3. RECONSTRUCTION ALGORITHM

We seek the signal \mathbf{x} given magnitude-only measurements \mathbf{y} . Suppose we enforce a constraint on the sparsity of the signal in a given basis (assumed to be the canonical basis for simplicity). Then, the signal estimate can be posed as the solution to the non-convex optimization problem:

$$\min_{\mathbf{x}} \sum_{i=1}^N \|\mathcal{A}_i(\mathbf{x}) - \mathbf{y}_i\|_2^2, \quad \text{s.t. } \mathbf{x} \in \mathfrak{M}_s^b, \quad (3)$$

where \mathbf{x} is the signal in the sparse domain, where \mathfrak{M}_s^b denotes the model of the signal, consisting of a set of s -sparse signals with uniform block length $b \in \mathbb{Z}$. For the standard sparsity model $b = 1$, whereas for the block sparse model $b > 1$. Here, \mathcal{A} is the modified

Algorithm 1 Model-based CoPRAM for Ptychography

input $\mathcal{A}, \mathbf{y}, s, b, t_0$

1: Initialize \mathbf{x}^0 according to: $\sqrt{\frac{1}{N} \sum_{i=1}^N \mathbf{y}_i^2}$.
2: **for** $t = 0, \dots, t_0 - 1$ **do**
3: $\mathbf{P}^{t+1} \leftarrow \text{diag}(\text{sign}(\mathcal{A}(\mathbf{x}^t)))$,
4: $\mathbf{x}^{t+1} = \text{MODELCoSAMP}\left(\frac{\mathcal{A}}{\sqrt{nN}}, \frac{\mathbf{P}^{t+1}\mathbf{y}}{\sqrt{nN}}, s, b, \mathbf{x}^t\right)$.
5: **end for**
output $\mathbf{z} \leftarrow \mathbf{x}^{t_0}$.

measurement operator, which accounts for the domain transformation.

To solve the problem (3), we employ the technique called Compressive Phase Retrieval with Alternating Minimization (CoPRAM) [18] (and its structured sparsity variant, Block CoPRAM). This procedure is described in Alg. 1.

Under the measurement setup explained in the previous section, we introduce our strategy to reconstruct the diffraction-blurred image from subsampled measurements. We invoke a slight modification of the algorithm CoPRAM as discussed in [18]. The algorithm proceeds in two stages — an *Initialization* stage and a *Descent* stage — which we describe as follows.

3.1. Initialization

For the initialization stage, we improve upon the one given in [4], by using root-mean-squared measurements as the estimator (line 1 of Alg. 1). We establish experimentally that this initialization is superior to that in [4] (experiments show an average of 2% improvement in SSIM, for the same number of outer iterations of CoPRAM). We also deviate from the conventional spectral initialization for Gaussian measurements as in [11, 10, 18]. While a spectral initial estimate succeeds for Gaussian measurements, both in theory and practice, it fails for the Fourier ptychographic setup, because there is no clear separation between the first and second largest singular values of the estimator matrix \mathbf{M} :

$$\mathbf{M} = \frac{1}{nN} \mathbf{A}^\top \text{diag}(\mathbf{y}^2) \mathbf{A} = \frac{1}{nN} \sum_{i=1}^N \mathbf{A}_i^\top \text{diag}(\mathbf{y}_i^2) \mathbf{A}_i.$$

Here $\mathbf{A}_i = \mathcal{A}_i(\mathbb{I})$. In the case of Gaussian measurements, $\mathbb{E}[M] = 2\mathbf{x}\mathbf{x}^\top + \mathbb{I} \|\mathbf{x}\|_2^2$ [11], and subsequently, if all elements of \mathbf{x} are positive (which is true in the case of real images), then the diagonal of this matrix $M_{j,j} = 2x_j^2 + \|\mathbf{x}\|_2^2$, with $j = [n]$ is sufficient to estimate \mathbf{x} . Now, if one assumes high overlap between cameras, and big enough aperture diameter, then \mathbf{A}_i 's are nearly diagonal (hence approximately commuting with $\text{diag}(\mathbf{y}_i^2)$) and $\mathbf{A}_i^\top \mathbf{A}_i \approx \mathbb{I}$, which suggests that $\mathbf{M} \approx \frac{1}{nN} \sum_{i=1}^N \text{diag}(\mathbf{y}_i^2) \mathbf{A}_i^\top \mathbf{A}_i \approx \frac{1}{nN} \sum_{i=1}^N \text{diag}(\mathbf{y}_i^2)$. Evaluating the diagonal terms of this matrix gives us the intuition behind using the root-mean-squared measurements (line 1 of Alg. 1), as a coarse initial estimate. We are in the process of developing a more concrete proof for the initialization based on this intuition.

3.2. Descent

Once we have a coarse estimate for the initialization of the CoPRAM algorithm, we then refine this estimate using a variant of alternating minimization. Specifically, at any given iteration, we first estimate the phase (line 3 of Alg. 1) by applying the forward operator \mathcal{A} to the signal estimate \mathbf{x}^t . Next, we impute this estimated phase into our observed intensity measurements, and subsequently obtain the next signal estimate \mathbf{x}^{t+1} using a sparse recovery algorithm (line 4 of Alg. 1) such as CoSaMP [36]. Moreover, in order to incorporate structural assumptions beyond sparsity, the only modification is to




Fig. 2: (a) Resolution chart, used as ground truth (b) simulated block sparse image, used as ground truth for experimental analysis.

replace the sparse recovery method by any other model-based recovery method, such as model CoSaMP [16] (line 4 of Alg. 1).

In [18] we have demonstrated (both theoretically and numerically) that the estimates (\mathbf{x}^{t+1}) of the above alternating minimization technique converges to the solution (\mathbf{x}) at a linear rate, using an appropriate termination condition. The basic idea is that the “phase noise” induced due to the estimation error can be suitably bounded provided the initial estimate is good enough. Below, we empirically demonstrate that for the case of Fourier ptychographic measurements, similar gains can be achieved using our algorithm.

4. EXPERIMENTAL RESULTS

We now demonstrate the efficacy of Algorithm 1 via a number of simulated image reconstruction experiments. (Note: for all of the subsequent experiments, we set the outer number of iterations of all algorithms to 10).

4.1. Uniform random subsampling


We first establish a subsampling strategy, which will be kept fixed for both types of signal models - one with sparsity constraint and one without.

We construct a subsampling mask, in which the elements of the mask are picked up from a continuous *standard uniform* distribution $u_i \in \mathbb{R}^n$, with elements u_j^i being independent standard uniform random variables. The mask resembles the operation of a matrix with 1s and 0s, where pixels corresponding to 1s are retained and those corresponding to 0s are discarded. A total of $m = f \times (nN)$ measurements are retained, where f denotes the fraction of samples (or pixels). The schematic diagram in Fig. 1 describes this sampling procedure. In this case, for an input vector $\mathbf{v} \in \mathbb{C}^n$, the subsampling mask operates as

$$\mathcal{M}_i(\mathbf{v})_j = \begin{cases} 0 & u_j^i > f, \\ v_j & u_j^i < f. \end{cases} \quad (4)$$

We describe the effect of enforcing the sparsity constraint in various domains as follows. We use two different datasets: (i) the USAF resolution chart as shown in Figure 2 (a), and (ii) a simulated image which is specifically block sparse as shown in Figure 2 (b). The parameters fed to the main algorithm are as follows: we used a $n = 256^2$ (256×256) image of the Resolution Chart (resChart) as the ground truth. The camera array consists of $N = 81$ (9×9) cameras, each with aperture diameter 72.75 pixels and overlap of 0.72 between consecutive cameras. A subsampling factor of $f = 0.3$ picks up 30% of the original number of measurements. To implement this, we generated masks \mathcal{M}_i as in (4). For the sparse phase retrieval algorithm CoPRAM, we enforce a sparsity of $s = 0.25n$.

The reconstruction procedure relies heavily on the extent of overlap, hence the norm of the reconstructed images is not pre-



(a) Initial center, (b) IERA, (c) Fourier, (d) Spatial, SSIM=0.3517 SSIM=0.3369 SSIM=0.5544 SSIM=0.8740

Fig. 3: 30% samples, (a) center image, reconstruction using (b) IERA (c) CoPRAM (Fourier sparse) (d) CoPRAM (spatially sparse).

served. We use Structural Similarity Index (SSIM) [37] as a metric to appropriately capture the quality of reconstruction, as it compares the two images in terms of luminance, contrast and structure, instead of utilizing a straightforward distance measure.

We employ CoPRAM by enforcing sparsity in both spatial and Fourier bases and compare the reconstruction from subsampled magnitude-only measurements to those from IERA. These results are displayed in Figure 3 for the input image in Figure 2. It can be noted that we can also impose sparsity in a wavelet basis (such as Haar) and we expect to achieve similar improvements in the SSIM.

We have also analyzed the variation of the SSIM with different undersampling rates. For this, we used CoPRAM while assuming sparsity in the Fourier and spatial bases for the input image in Fig. 2. We also invoked Block CoPRAM, (refer Sec. 4.4 for details) which assumes block sparsity in the spatial domain. For comparison, we used IERA and also a modified version of another sparse phase retrieval algorithm called SPARTA [28], where we have used the same initialization as in line 1 of Alg. 1. These results can be found in Figure 4.





Fig. 4: Variation of SSIM with undersampling ratio, with sparsity $s = 0.25n$, (block size $b = 4 \times 4$ for Block CoPRAM).

4.2. Subsampling using uniform random camera pattern

Another subsampling strategy is to turn some cameras “on” or “off”. We use sampling masks \mathcal{M}_i , which are picked up from a continuous *standard uniform* distribution $u \in \mathbb{R}^N$, with elements u_i being independent standard uniform random variables. In terms of the sampling mask, for a vector input $\mathbf{v} \in \mathbb{C}^n$, the subsampling mask,

$$\mathcal{M}_i(\mathbf{v}) = \begin{cases} 0 & u_i > f, \\ \mathbf{v} & u_i < f. \end{cases} \quad (5)$$

We utilize this feature to test the robustness of CoPRAM against IERA, under the sparsity assumption. We switch off $\approx 50\%$ of the cameras (for this experiment, 38 cameras are active, from 81 total), where the camera locations are picked according to (5) (the central



(a) Initial center, (b) IERA, (c) Fourier, (d) Spatial, SSIM=0.3927 SSIM=0.4225 SSIM=0.5613 SSIM=0.9053

Fig. 5: 50% cameras, (a) center image, reconstruction using (b) IERA (c) CoPRAM (Fourier sparse) (d) CoPRAM (spatially sparse).


camera is kept “on” by default). The results are displayed in Figure 5 for the input image in Figure 2. We observed that enforcing sparsity in the spatial domain gives a better reconstruction (Fig. 5 (d)).

4.3. Effect of decreased overlap

One of the issues of the implementation in [4] is that they require consecutive camera arrays to have overlap with each other. This is physically impractical if one wants to implement a camera array in the same plane. However, with no camera overlap, their experiments perform poorly (oversampling is imperative for standard phase retrieval strategies). On the other hand CoPRAM uses a sparsity constraint to improve quality of reconstruction (Note: for this setup $f = 1$). For this experiment, we changed the amount of overlap between two cameras from 0.72 to 0.12. The results of this experiment suggest a superior reconstruction when CoPRAM is invoked, with sparsity in spatial basis (SSIM=0.6124) as compared to IERA (SSIM=0.3088).

4.4. Extension to block sparsity

Since we were able to demonstrate the advantage of sparse modeling to reduce number of samples required for good reconstruction, we also applied CoPRAM to images with block sparsity (in the spatial domain). Instead of using CoSaMP (line 4 of Alg. 1), we use a block variant of model-based CoSaMP [16] (we call this Block CoPRAM). For this experiment, we synthetically generated a block sparse image (Fig. 2 (b)), and measured it using the uniform random subsampling pattern described in (4), with a low overlap of 0.12 between adjacent cameras. The reconstructions are displayed in Fig. 6, showing pronounced improvement when Block CoPRAM is used.



(a) Initial, (b) IERA, (c) Spatial, (d) Block, SSIM=0.99687 SSIM=0.99665 SSIM=0.99995 SSIM=0.99998

Fig. 6: Using 0.12 overlap and 30% samples (a) center image, reconstructed image using (b) IERA (c) CoPRAM (spatially sparse) (d) Block CoPRAM (spatially block sparse).

5. CONCLUSION AND DISCUSSION

In conclusion, we demonstrate that by enforcing sparsity constraints in the reconstruction procedure, we are able to reproduce good quality high-resolution images from diffraction-blurred images, using far fewer samples than pre-existing methods. This also translates to lower operational costs for a faster imaging procedure. In future

work, we aim to (i) establish a formal analysis on the sample complexity of Fourier ptychography; and also (ii) explore better initialization mechanisms, and their effect on reconstruction performance.

6. REFERENCES

- [1] S. Dong, R. Horstmeyer, R. Shiradkar, K. Guo, X. Ou, Z. Bian, H. Xin, and G. Zheng. Aperture-scanning fourier ptychography for 3d refocusing and super-resolution macroscopic imaging. *Optics express*, 22(11):13586–13599, 2014.
- [2] G. Zheng, R. Horstmeyer, and C. Yang. Wide-field, high-resolution fourier ptychographic microscopy. *Nature photonics*, 7(9):739–745, 2013.
- [3] A. Maiden, M. Humphry, F. Zhang, and J. Rodenburg. Superresolution imaging via ptychography. *J. Opt. Soc. Am. A*, 28(4):604–612, Apr 2011.
- [4] J. Holloway, M. S. Asif, M. K. Sharma, N. Matsuda, R. Horstmeyer, O. Cossairt, and A. Veeraraghavan. Toward long-distance subdiffraction imaging using coherent camera arrays. *IEEE Trans. Comp. Imag.*, 2(3):251–265, 2016.
- [5] J. Holloway, Y. Wu, M. K. Sharma, O. Cossairt, and A. Veeraraghavan. Savi: Synthetic apertures for long-range, subdiffraction-limited visible imaging using fourier ptychography. *Science Advances*, 3(4):e1602564, 2017.
- [6] J. Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.
- [7] S. Marchesini. Phase retrieval and saddle-point optimization. *JOSA A*, 24(10):3289–3296, 2007.
- [8] K. Nugent, A. Peele, H. Chapman, and A. Mancuso. Unique phase recovery for nonperiodic objects. *Physical review letters*, 91(20):203902, 2003.
- [9] E. Candes, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.*, 66(8):1241–1274, 2013.
- [10] E. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Trans. Inform. Theory*, 61(4):1985–2007, 2015.
- [11] P. Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. In *Adv. Neural Inf. Proc. Sys. (NIPS)*, pages 2796–2804, 2013.
- [12] Y. Shechtman, Y. Eldar, O. Cohen, H. Chapman, J. Miao, and M. Segev. Phase retrieval with application to optical imaging: a contemporary overview. *IEEE Sig. Proc. Mag.*, 32(3):87–109, 2015.
- [13] R. Millane. Phase retrieval in crystallography and optics. *JOSA A*, 7(3):394–411, 1990.
- [14] A. Maiden and J. Rodenburg. An improved ptychographical phase retrieval algorithm for diffractive imaging. *Ultra-microscopy*, 109(10):1256–1262, 2009.
- [15] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. Royal Stat. Soc. Stat. Meth.*, 68(1):49–67, 2006.
- [16] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Info. Theory*, 56(4):1982–2001, Apr. 2010.
- [17] C. Hegde, P. Indyk, and L. Schmidt. Fast algorithms for structured sparsity. *Bul. of the EATCS*, 1(117):197–228, Oct. 2015.
- [18] G. Jagatap and C. Hegde. Fast, sample-efficient algorithms for structured phase retrieval. (To appear in) *Adv. Neural Inf. Proc. Sys. (NIPS)*, 2017.
- [19] R. Gerchberg and W. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35(237), 1972.
- [20] Z. Chen, G. Jagatap, S. Nayer, C. Hegde, and N. Vaswani. Low rank fourier ptychography, 2017. Project website: <https://saranayer.github.io/Low-Rank-Phase-Retrieval/>.
- [21] E. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Ap. Comp. Har. An.*, 39(2):277–299, 2015.
- [22] Y. Chen and E. Candes. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Adv. Neural Inf. Proc. Sys. (NIPS)*, pages 739–747, 2015.
- [23] G. Wang and G. Giannakis. Solving random systems of quadratic equations via truncated generalized gradient flow. In *Adv. Neural Inf. Proc. Sys. (NIPS)*, pages 568–576, 2016.
- [24] H. Ohlsson, A. Yang, R. Dong, and S. Sastry. Cprl—an extension of compressive sensing to the phase retrieval problem. In *Adv. Neural Inf. Proc. Sys. (NIPS)*, pages 1367–1375, 2012.
- [25] Y. Chen, Y. Chi, and A. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Trans. Inform. Theory*, 61(7):4034–4059, 2015.
- [26] K. Jaganathan, S. Oymak, and B. Hassibi. Sparse phase retrieval: Convex algorithms and limitations. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, pages 1022–1026. IEEE, 2013.
- [27] T. Cai, X. Li, and Z. Ma. Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *Ann. Stat.*, 44(5):2221–2251, 2016.
- [28] G. Wang, G. Giannakis, J. Chen, and M. Akçakaya. Sparta: Sparse phase retrieval via truncated amplitude flow. In *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc. (ICASSP)*, pages 3974–3978. IEEE, 2017.
- [29] M. Iwen, A. Viswanathan, and Y. Wang. Robust sparse phase retrieval made easy. *Ap. Comp. Har. An.*, 42(1):135–142, 2017.
- [30] S. Bahmani and J. Romberg. Efficient compressive phase retrieval with constrained sensing vectors. In *Adv. Neural Inf. Proc. Sys. (NIPS)*, pages 523–531, 2015.
- [31] H. Qiao and P. Pal. Sparse phase retrieval using partial nested Fourier samplers. In *Proc. IEEE Global Conf. Signal and Image Processing (GlobalSIP)*, pages 522–526. IEEE, 2015.
- [32] Y. Shechtman, A. Beck, and Y. C. Eldar. Gespar: Efficient phase retrieval of sparse signals. *IEEE Trans. Sig. Proc.*, 62(4):928–938, 2014.
- [33] K. Jaganathan, Y. Eldar, and B. Hassibi. Stft phase retrieval: Uniqueness guarantees and recovery algorithms. *IEEE J. Select. Top. Sig. Proc.*, 10(4):770–781, 2016.
- [34] L. Li, C. Cheng, D. Han, Q. Sun, and G. Shi. Phase retrieval from multiple-window short-time fourier measurements. *IEEE Sig. Proc. Lett.*, 24(4):372–376, 2017.
- [35] T. Bendory, Y. Eldar, and N. Boumal. Non-convex phase retrieval from stft measurements. *IEEE Trans. Info. Theory*, 2017.

- [36] D. Needell and J. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Ap. Comp. Har. An.*, 26(3):301–321, 2009.
- [37] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Proc.*, 13(4):600–612, 2004.