

Coherent Inverse Scattering via Transmission Matrices: Efficient Phase Retrieval Algorithms and a Public Dataset

Christopher A. Metzler*
Rice University
chris.metzler@rice.edu

Manoj K. Sharma*
Northwestern University and Rice University
manoj.sharma@rice.edu

Sudarshan Nagesh
Rice University
sn34@rice.edu

Richard G. Baraniuk
Rice University
richb@rice.edu

Oliver Cossairt
Northwestern University
ollie@eecs.northwestern.edu

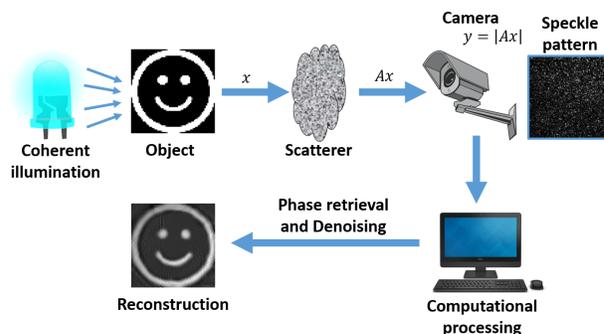
Ashok Veeraraghavan
Rice University
vashok@rice.edu

Abstract

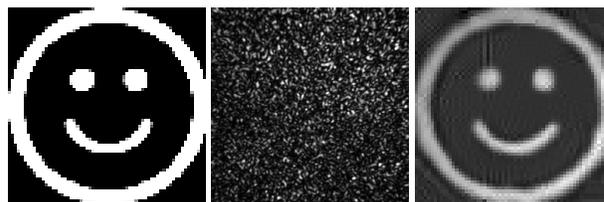
A transmission matrix describes the input-output relationship of a complex wavefront as it passes through/reflects off a multiple-scattering medium, such as frosted glass or a painted wall. Knowing a medium’s transmission matrix enables one to image through the medium, send signals through the medium, or even use the medium as a lens. The double phase retrieval method is a recently proposed technique to learn a medium’s transmission matrix that avoids difficult-to-capture interferometric measurements. Unfortunately, to perform high resolution imaging, existing double phase retrieval methods require (1) a large number of measurements and (2) an unreasonable amount of computation. In this work we focus on the latter of these two problems and reduce computation times with two distinct methods: First, we develop a new phase retrieval algorithm that is significantly faster than existing methods, especially when used with an amplitude-only spatial light modulator (SLM). Second, we calibrate the system using a phase-only SLM, rather than an amplitude-only SLM which was used in previous double phase retrieval experiments. This seemingly trivial change enables us to use a far faster class of phase retrieval algorithms. As a result of these advances, we achieve a $100\times$ reduction in computation times, thereby allowing us to image through scattering media at state-of-the-art resolutions. In addition to these advances, we also release the first publicly available transmission matrix dataset. This contribution will enable phase retrieval researchers to apply their algorithms to real data. Of particular interest to this community, our measurement vectors are naturally *i.i.d.* subgaussian, *i.e.*, no coded diffraction pattern is required.

*Contributed equally.

This work was supported in part by DARPA REVEAL grant HR0011-16-C-0028, ONR grant N00014-15-1-2735, ARO grant W911NF-12-1-0407, and the Big-Data Private-Cloud Research Cyberinfrastructure MRI award funded by NSF under grant CNS-1338099 and by Rice University.



(a) System Diagram



(b) Incident Image x (c) Speckle Pattern y (d) Reconstruction \hat{x}

Figure 1: (Top) A sketch of our experimental setup. (Bottom) A smiley face, the speckle pattern it produces when transmitted through a multiple-scattering media, and our reconstruction of the smiley face from the speckle pattern. Phase retrieval algorithms can recover the original image from the speckle pattern.

1. Imaging through multiple-scattering media

At first glance, seeing through multiple-scattering media seems like an impossible task. Any light incident on the me-

In addition, C. Metzler was supported in part by the NSF GRFP, C. Metzler and R. Baraniuk were supported in part by NSF CCF-1527501 and CCF-1502875, and O. Cossairt was supported in part by NSF CAREER IIS-1453192.

dia will undergo multiple reflections. Thus if we illuminate an object with coherent light, the resulting wavefront will constructively and destructively interfere with itself, and a “speckle pattern” will be produced on the far side of the scatterer. This speckle pattern generally bears no resemblance to the original image. See Fig. 1.

In [11, 31], the authors developed the double phase retrieval method to reconstruct an image on the far side of a scatterer from its speckle pattern. The key enabler of this ability is learning a scattering material’s transmission matrix (TM).

A TM characterizes the input-output relationship of a coherent light source as it passes through a multiple-scattering medium. The scattering process is linear and thus coherent incident light $\mathbf{x} \in \mathbb{C}^N$ will produce an output $\mathbf{z} \in \mathbb{C}^M$, where $\mathbf{z} = \mathbf{A}\mathbf{x} + \epsilon$, \mathbf{A} is the TM, and ϵ represents noise. Typical cameras capture only intensity information, in which case the measurement process becomes

$$\mathbf{y}^2 = |\mathbf{A}\mathbf{x} + \epsilon|^2,$$

with $\mathbf{y} \in \mathbb{R}_+^M$ denoting the measurements and the square taken element-wise. In this work we take the square of the intensity measurements and deal with $\mathbf{y} = |\mathbf{A}\mathbf{x} + \epsilon|$.

If one knows the TM \mathbf{A} then one can use phase retrieval algorithms to reconstruct a signal \mathbf{x} from measurements \mathbf{y} . Thus phase retrieval algorithms can be used to see through multiple-scattering media.

This promise comes at a price; one needs to first learn the TM. Learning a TM can be accomplished by using a spatial light modulator (SLM) to create a series of calibration images and then solving a series of phase retrieval problems [11]. Physically, assuming we want to avoid compressive techniques, we must take over $4\times$ as many measurements as the desired resolution. Computationally, we must then solve a proportional number of phase retrieval problems. Using existing techniques, learning a TM to image at 128×128 resolutions would require capturing over 65,000 images, which takes about 24 hours. Solving the corresponding phase retrieval problems would then require roughly 1 *million* CPU hours.

In this paper we reduce this extreme computational burden in two ways. First, we develop a new phase retrieval algorithm, prVAMP, which runs far faster than existing methods. Second, we propose using a phase-only SLM, rather than an amplitude-only SLM, to learn a TM. This change dramatically eases the calibration procedure, as it allows us to use a far broader class of phase retrieval algorithms.

In addition, we release a dataset to the optics and phase retrieval communities so that TMs can be studied and so that newly developed phase retrieval algorithms can be applied to real-world data.

1.1. What about compressive measurements?

Without prior information, $M \geq 4N - 4$ phaseless measurements are sufficient [3] to reconstruct a signal. However, when prior information, like sparsity, is available potentially far fewer measurements are required [25]. Thus a natural way to reduce the computational cost of learning the TM is to simply reconstruct fewer rows of the TM. This was the technique employed in [31] and [43].

In this work, we chose not to pursue this line of research and instead focused on experiments where $M \geq 4N$. We made this decision because most of the objects we would eventually like to image, for instance a person, are optically rough. That is, the surface depths between adjacent pixels vary by more than a wavelength. As a result, although the amplitude of the signal may be very structured, and thus well captured in a sparse basis, the phase, which depends on surface depth, follows an i.i.d. uniform distribution. Thus the signal is not sparse in any basis [20]. Other regularization techniques, like denoisers, are not easily applied to the optically rough signals either.

2. Related work

2.1. Imaging through scattering media

Starting with [13], a host of techniques have been used to recover signals that were passed through or reflected off of scattering materials. We describe some of the most promising techniques here.

2.1.1 Time-of-flight

Time-of-flight (ToF) imaging systems rely upon the fact that a photon’s propagation time depends on both where the photon came from and how many times it scattered before reaching the detector. This fact allows one to use a pulsed or modulated laser and some sort of temporal gating to accept only ballistic photons (photons that were never scattered) from a region of interest. This allows the system to effectively ignore the scattering material, at the cost of reduced light throughput.

In [39] the authors used a streak-camera based setup to recover the shape of a 3D mannequin placed outside the line-of-sight of a camera. In [26] the authors showed that it is possible to capture the surface reflectance of an object with a similar ToF system. More recently, [17] noted the high-cost of the streak-camera based setup and proposed a cheaper photonic mixer devices based approach which they used to image through water that was contaminated with milk. This method was further improved upon in [37].

ToF systems have a number of advantages. For instance, they allow one to co-locate the detector and the light source and they can handle rapidly changing scattering media, like fog. Unfortunately, ToF systems rely upon ballistic photons and therefore cannot deal with thick scatterers, which are the focus of this work.

2.1.2 Multi-slice light-propagation

As the name implies, the multi-slice light propagation method models a scattering material as a series of 2D scattering slices between which light propagates. Rather than learning a TM, which blindly maps inputs to outputs, the multi-slice propagation method learns a composition of linear transformations.

The multi-slice light propagation method was first described in [38]. In that work the authors use Ptychographic methods to reconstruct a 3D tomogram of algae. The multi-slice light propagation method was subsequently used with holographic measurements in [18]. In that work the model's linear transformations were represented as a neural network and the neural network was optimized to match the measurements. This procedure produced accurate 3D reconstructions of human cells placed between glass slides.

Among the methods presented here, multi-slice light propagation is uniquely suited to perform 3D reconstructions. Moreover, it provides information about the actual structure/composition of the scattering material. However, the multi-slice approach does not model reflections within the scatterer and it is unclear how well this method works with thick scattering materials.

2.1.3 Strong memory effect

A scatterer is said to exhibit the strong memory effect if a translation by the target, e.g. the smiley face in Fig. 1, produces only a translation in the resulting speckle pattern [14]. This property was recently used to enable single shot imaging through a multiple-scattering material [19]. In that work, the authors showed that when the strong memory effect holds the autocorrelation of an object of interest's speckle patterns becomes equal to the autocorrelation of the object. This property allows phase retrieval algorithms to recover the original signal from its speckle pattern.

It is amazing single-shot imaging works at all, and this technique holds great promise for imaging through thin biological tissues. However, the strong memory effect is an unrealistic assumption for imaging at a macroscopic scale. Using a relaxed version of the strong memory effect is an interesting direction for future work.

2.1.4 Holographic interferometry

Holographic interferometry can be used to help learn a material's TM. After one has measured the TM, one can image through multiple-scattering media by solving simple inverse problems.

To learn a TM using holographic interferometry one first measures a large number of the *complex-fields* of responses to calibration signals. These fields are each measured by capturing four interference pattern with a separate reference beam. With these fields in hand, one sets up a matrix inversion problem to learn the TM.

Holographic interferometry was first used to acquire a scattering material's TM in [28]. That work also represents the first paper to compute an optical TM. Using TMs captured with holographic interferometry, [28] and [21] performed imaging through multiple-scattering media.

Unlike the previously described methods, this technique can deal with scattering materials of arbitrary depths and complexities. Unfortunately, this method comes with a major drawback. Because it relies on interferometry, the system requires physical stability and so is very sensitive to perturbations; even minute vibrations, such as those caused by an air conditioning unit turning on and off in a lab, can be enough to change the interference pattern and thereby kill this method. This problem becomes particularly pernicious at higher resolutions where the physical stability requirement becomes more and more demanding.

2.1.5 Temporally modulated phase

By temporally modulating the phase of calibration signals, one can interferometrically measure a TM¹ without a separate, troublesome reference arm.

At a high level, the temporally modulated phase method works as follows. Half the SLM's pixels' phase delays are modulated in time, each at a unique frequency. At the same time, the other half of the SLM's pixels' phase delays are fixed at zero to serve as a reference. The two sets of signals pass through the multiple-scattering material and interfere with one another. The resulting scene on the detector becomes a superposition of many interference patterns which, because the phase delays are modulated in time, are themselves sinusoidal functions of time. A video of the interference patterns is captured and the sequence of frames associated with each detector pixel is Fourier transformed (with respect to time). Because, the interference patterns are multiplexed in frequency the Fourier transform serves to separate them: the phase of the Fourier coefficient associated with each modulation frequency becomes the TM value of the SLM pixel associated with that modulation frequency. The process is repeated with the role of reference and signal groups swapped to learn the other half of the TM.

The aforementioned method was originally developed in [7] to measure a single row of a TM. In [42], the authors extended the method to work with detector arrays and so measured a 2 dimensional TM. Both works validated their matrices by focusing through a scattering media, but did not perform imaging.

Temporal modulation of phase can be used to learn TMs quickly, without a separate reference beam, and without phase retrieval algorithms. However, as mentioned in foot-

¹With phase only modulation (rather than amplitude and phase modulation), the temporally modulated phase method does not measure the TM but rather a very related quantity; the phase modulation patterns that would most excite each pixel on the detector. They would be equivalent if each pixel on the SLM contributed equally to each pixel on the detector.

note 1, this method does not learn the true TM. Furthermore, the SLM and camera refresh rates limit the maximum resolution. In [42] the authors demonstrated an SLM resolution of 33×30 .

2.1.6 Double phase retrieval

Double phase retrieval, the focus of this work, is another method that uses TMs for imaging. The key idea behind this method is that if one measures a sufficient number of the *intensities* of responses to calibration signals, one can use phase retrieval techniques to learn the TM. This method will be described more thoroughly in Section 3.3.

The double phase retrieval method was first proposed in [31] based off of work done in [11]. In [31] the authors performed very low resolution imaging using the prSAMP algorithm [30]. These experiments were repeated successfully at 64×64 resolutions in [43] using the prVBEM algorithm [10]. Following the submission of the current work, [27] learned a 157×100 TM (SLM resolution of 10×10) using a phase-only SLM and the PhaseCut algorithm [41]. They proceeded to use phase-only wavefront shaping to spell “Michigan” on the far side of a scatterer.

Compared to other techniques, the double phase retrieval method is cheap, robust, and easy to setup. Like other TM methods, this technique can deal with scattering materials of arbitrary depths and complexities. Unlike most other methods, this procedure is able to reconstruct signal after capturing only a single image.

This approach does come with two drawbacks however: First, it requires capturing a large number of measurements, and thus is not well suited for rapidly changing scattering materials like fog. Second, at calibration it requires solving a large number of phase retrieval problem, which can induce a large computational cost. This work focuses on solving the latter of these two problems.

2.2. Phase retrieval algorithms

Phase retrieval algorithms reconstruct signals from the magnitudes of linear transformations of the signal. More formally, phase retrieval algorithms solve the following problem.

$$\text{Given observation } \mathbf{y} = |\Phi \mathbf{x}|, \text{ with } \Phi \text{ known, determine } \mathbf{x}. \tag{1}$$

The matrix Φ is known as the measurement matrix. In the context of this work, the measurement matrix is either the TM \mathbf{A} or the calibration matrix \mathbf{X}^H , which will be defined in Section 3.3.

2.2.1 Alternating projection methods

Alternating projection methods, such as the seminal Fienup [12] and Gerchberg-Saxton (GS) [15] algorithms, were the first algorithms to solve the phase retrieval problem. These

Algorithm	$N = 8^2$	$N = 16^2$	$N = 32^2$
prGAMP	0.07	0.39	9.64
prSAMP	3.52	54.80	826.79
prVBEM	0.22	2.66	105.74
prVAMP	0.44	0.48	6.79
GS-50	0.15	1.78	30.56
GS-100	0.26	3.41	60.26
WF-50	0.09	1.50	30.30
WF-500	0.69	14.22	244.54
PhaseLift	5.00	201.00	12321
PhaseMax	0.48	2.85	104.46

Table 1: Running times (in seconds) of various phase retrieval algorithms with $12 \cdot N$ i.i.d. $\mathcal{CN}(0, 1)$ measurements. Unless noted otherwise, all algorithms are run for 50 iterations. Notice at higher resolutions prGAMP and prVAMP are the fastest.

Preprocessing Step	$N = 8^2$	$N = 16^2$	$N = 32^2$
Economical SVD	0.52	0.50	19.92
Pseudoinverse	0.59	0.62	24.80

Table 2: The computation times (in seconds) of one-time preprocessing operations associated with $12 \cdot N \times N$ measurement matrices. prVAMP requires an economical SVD of the measurement matrix and GS requires a pseudoinverse of the measurement matrix. These values can then be reused for all phase retrieval problems using the same measurement matrix.

algorithms alternatively project estimates of the signal from the target domain to the measurement domain such that the estimate eventually lies on the same support as the target and would produce the observed measurements.

Alternating projection methods are often derived and presented with Fourier measurements, as this arises in ptychography, astronomical imaging, and a variety of different imaging modalities that rely upon Fraunhofer diffraction. However, they can be extended to handle arbitrary measurements by recognizing that the pseudoinverse of the measurement matrix serves to both map from the measurement space to the target space and project onto the target’s support.

2.2.2 Convex relaxations and gradient descent

In the last five years, phase retrieval has become the focus of intense renewed interest. A host of recovery algorithms have been developed to deal with general linear transformations, rather than focusing on Fourier measurements. These methods include the well-known PhaseLift [6] and PhaseCut [41] lifted convex relaxations, as well as the stochas-

tic Wirtinger Flow (WF) [5] algorithm. More recently, a convex relaxation of the phase retrieval problem that avoids lifting, termed PhaseMax, was independently developed in [16] and [1].

2.2.3 Approximate message passing

Approximate message passing algorithms form an interesting family of heuristic algorithms that solve inverse problems involving generalized linear measurements,² including the phase retrieval problem (1). They are based off of the approximate message passing (AMP) algorithm [9] and its generalization GAMP [32], which were originally designed to solve the compressed sensing problem. The first AMP algorithm designed for phase retrieval was prGAMP [35]. This was followed by the prVBEM algorithm [10], which is more robust to noise; prSAMP [30], which can handle a broader class of measurement matrices; and most recently D-prGAMP [23], which uses denoisers to impose complex priors on the reconstructed signal. In Section 4 we will introduce a new phase retrieval algorithm, prVAMP.

The AMP algorithms are all heuristic algorithms and so at best offer only asymptotic guarantees. In the case of the phase retrieval problem, most AMP algorithms offer no guarantees at all. Despite this shortcoming, in practice they offer remarkable performance. They are accurate, robust to noise, generally very fast, and can easily incorporate priors into the reconstruction. In Figs. 2 and 3 we see that the AMP algorithms perform phase retrieval as well as or better than competing methods when the measurements follow i.i.d. circularly symmetric complex Gaussian distributions.³ In Table 1 we see that the prGAMP and prVAMP AMP-based algorithms perform these reconstruction far faster than competing methods. Details of the settings used in these simulations are presented in Section 4.3.2.

Unfortunately, until recently, AMP algorithms came with an Achilles’ heel; they required that the measurement matrix has i.i.d. subgaussian, zero-mean elements. When the measurements do not satisfy these assumptions, as is the case in the experimental setup of [11] and in Section 4, the performance of AMP plummets [4]. This can be alleviated somewhat with damping (updating the estimate of the signal gradually with small steps), but this slows convergence considerably [33].

The authors of prSAMP dealt with the challenging measurement matrices from [11] by making the parallel updates with AMP sequential [30], as was first done for the compressed sensing problem in [22]. Thanks to these sequential updates, prSAMP outperforms prGAMP and prVBEM

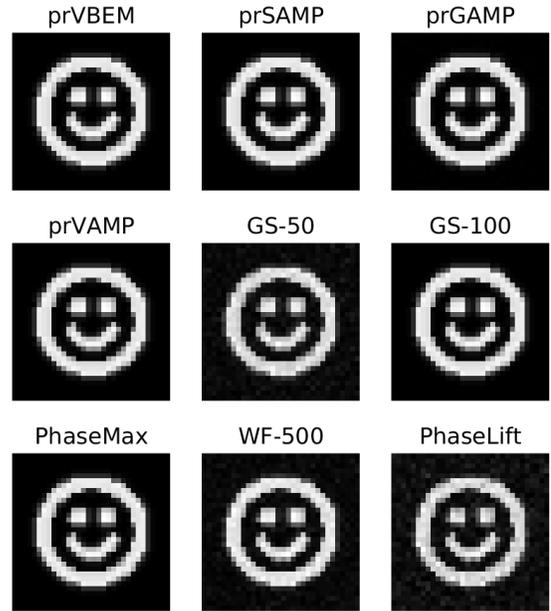


Figure 2: Simulated 32×32 reconstructions with various phase retrieval algorithms from $12 \cdot 32^2$ i.i.d. $\mathcal{CN}(0, 1)$ measurements. GS-50 and GS-100 denote the GS algorithm run for 50 and 100 iterations, respectively. When run for a sufficient number of iterations, all algorithms faithfully reconstruct the signal.

when dealing with poorly conditioned matrices. Unfortunately, the sequential updates mean prSAMP gives up AMP algorithms’ main advantage; fast computational times. In Table 1 we compare the run-times of various phase retrieval algorithms.⁴ At even moderate resolutions prSAMP is nearly one hundred times slower than prGAMP. To deal with this computational complexity, researchers have developed methods that split the overall phase retrieval algorithm into smaller subproblems [29]. While effective, this method degrades the performance of the algorithm and complicates implementation.

In this work we avoid the Achilles’ heel of AMP phase retrieval algorithms in two ways. First, we extend the recently developed vector-AMP (VAMP) [34, 36] algorithm, to solve the phase retrieval problem. Our extension, prVAMP, offers the speed of prGAMP with the robustness (in terms of measurement matrices) of prSAMP. Vector-

²Generalized measurements are any measurements of the form $\mathbf{y} = Q(\mathbf{A}\mathbf{x} + \epsilon)$, where $Q(\cdot)$ denotes a simple nonlinearity.

³We use the notation $\mathcal{CN}(0, \sigma^2)$ to denote the circularly symmetric Gaussian distribution. A random variable follows a circularly symmetric complex Gaussian distribution if its real and imaginary parts are each statistically independent and follow identical zero-mean Gaussian distributions.

⁴Table 1 does not include computation times associated with preprocessing operations. prVAMP and GS must perform an economical SVD and a pseudoinverse, respectively, of the measurement matrix. The computations times associated with these operations are presented in Table 2. In the context of double phase retrieval, these operations have a negligible effect on the overall runtime because they can be done once and then reused tens of thousands of times for different problems.

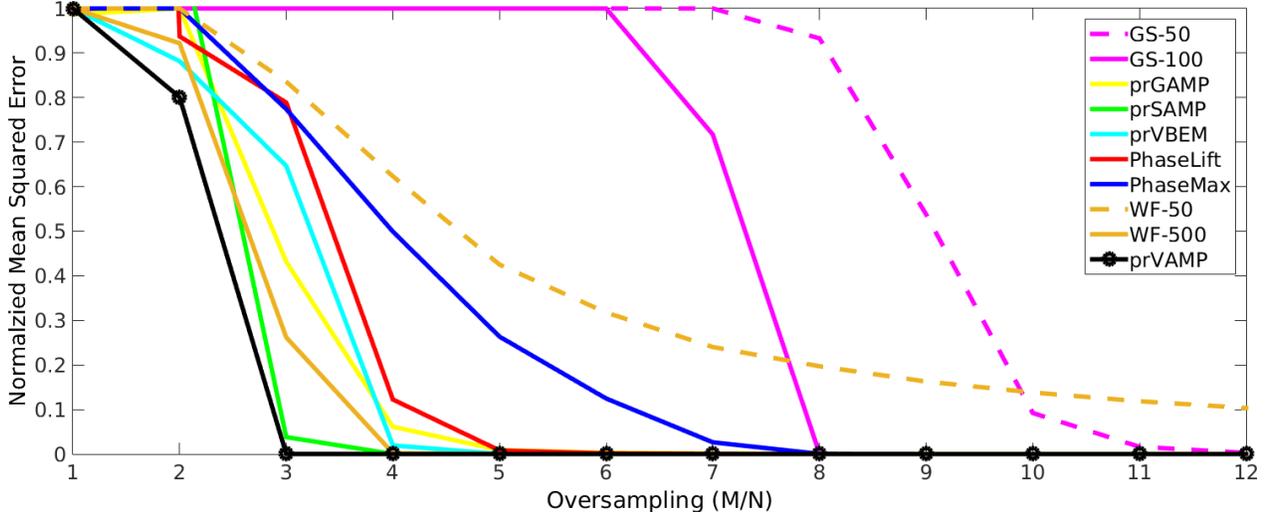


Figure 3: Simulated 16×16 reconstruction errors at various sampling rates with i.i.d. $\mathcal{CN}(0, 1)$ measurements. The AMP-based algorithms reconstruct the signal with fewer measurements than many competing methods.

AMP and our extension will be briefly reviewed and introduced in Section 4. Second, we make a minor change to the experimental setup that lets us avoid nonzero-mean matrices. This allows us to use prGAMP, the simplest AMP-based phase retrieval algorithm.

2.2.4 Reporting Simulations Results

Throughout this work, when we report simulation results we do so in terms of median normalized mean squared error (NMSE). Phase retrieval solutions are ambiguous up to a global phase rotation. That is, if \mathbf{x}_o is a solution, then so is $\mathbf{x}_o e^{j\phi}$ for any ϕ . Therefore, prior to computing the normalized mean squared error, we disambiguate the solution: we determine the global phase rotation that would minimize the error. For the oversampling and the noise-sensitivity experiments, we ran each algorithm for 50 trials and reported the median NMSE. For the smiley face reconstruction experiments, we performed 10 trials and returned the estimate with the median NMSE.

3. Problem Description

3.1. Notation

We will use the following notation throughout the rest of the paper. Matrices will be denoted with bold upper case letters, \mathbf{A} . Vectors will be denoted with bold under case letters, \mathbf{a} . Scalar variables will be denoted with unweighted under case letters, a .

3.2. Physical setup

At a high level, our setup, illustrated in Fig. 1, is as follows. A laser illuminates an SLM and the SLM modulates the light with some pattern \mathbf{x} (the smiley face). The modulated light is incident on a multiple-scattering material. The

scattering material may be transmissive (as shown) or reflective. The scattered light constructively and destructively interferes with itself to form a speckle pattern. This speckle pattern is photographed by a consumer-level camera to produce measurements $\mathbf{y}^2 = |\mathbf{A}\mathbf{x} + \epsilon|^2$, where ϵ represents noise.

3.3. Double phase retrieval approach

This subsection describes the double phase retrieval approach to imaging through complex scattering media. This method gets its name from the fact that phase retrieval is performed twice: once for calibration and once for imaging. [11] was the first work to learn a TM this way. [31] was the first work to then use such a TM for imaging.

3.3.1 Calibration

We first send a series of calibration patterns $\mathbf{x}_p \in \mathbb{R}^N$ with $p = 1, \dots, P$, through the scattering media. For each p , the signal will be transformed by the TM $\mathbf{A} \in \mathbb{C}^{M \times N}$, to produce measurements $\mathbf{y}_p \in \mathbb{R}_+^M$, with

$$\mathbf{y}_p = |\mathbf{A}\mathbf{x}_p + \epsilon_p|,$$

where ϵ_p denotes noise.⁵ Actual calibration patterns and their corresponding measurements are illustrated in Fig. 4.

The sets of calibration, measurement, and noise column vectors are then concatenated with themselves to form $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_P] \in \mathbb{R}^{N \times P}$, $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_P] \in \mathbb{R}_+^{M \times P}$, and $\mathbf{E} = [\epsilon_1, \epsilon_2, \dots, \epsilon_P] \in \mathbb{C}^{M \times P}$.

⁵In this work we have used long exposure times and strong illumination in hopes that the noise vectors ϵ_p follow approximately white circularly-symmetric complex Gaussian distributions.

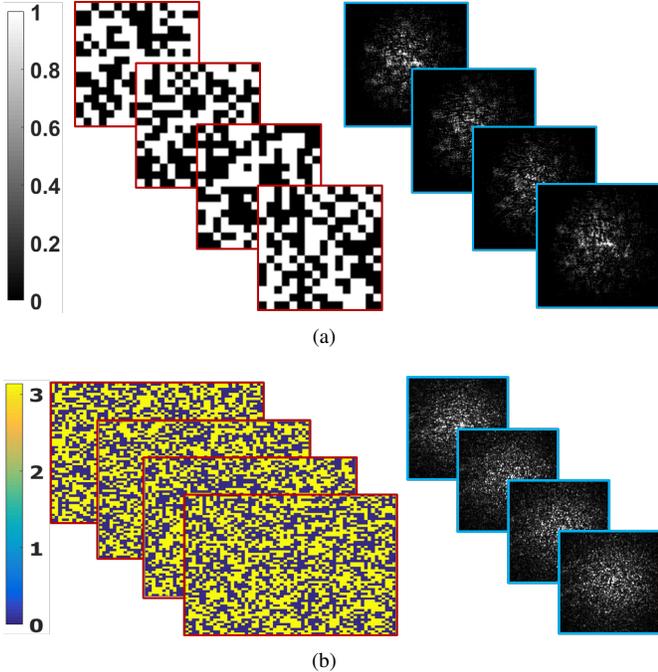


Figure 4: Calibration patterns with amplitude-only modulation (top-left) and phase-only (bottom-left) and the corresponding measured speckle patterns (right).

In this way the entire calibration process can be described by the equation

$$\mathbf{Y} = |\mathbf{A}\mathbf{X} + \mathbf{E}|.$$

Taking the transpose of the above equality, we obtain

$$\mathbf{Y}^H = |\mathbf{X}^H \mathbf{A}^H + \mathbf{E}^H|.$$

Consider the m^{th} column of \mathbf{Y}^H (this consists of the P measurements associated with the m^{th} detector pixel)

$$\mathbf{y}_m^H = |\mathbf{X}^H \mathbf{a}_m^H + \epsilon_m^H|, \quad (2)$$

where \mathbf{a}_m^H and ϵ_m^H denote the m^{th} rows of \mathbf{A} and ϵ .

Assuming P is sufficiently large (generally $P > 4N$), one can apply phase retrieval algorithms to (2) to recover each row \mathbf{a}_m^H ; simply treat \mathbf{y}_m^H as the measurement and \mathbf{X}^H as the measurement matrix. This can be repeated for each of the M rows of \mathbf{A} , potentially in parallel, to learn the entire TM.

3.3.2 Imaging

After one has an estimate $\tilde{\mathbf{A}}$ of the TM, imaging through the scattering medium is straightforward. Given a speckle pattern \mathbf{y} , one needs only apply phase retrieval algorithms to the measurements with $\tilde{\mathbf{A}}$ as the known measurement matrix. Reconstruction results with real-world data are presented in Fig. 5.

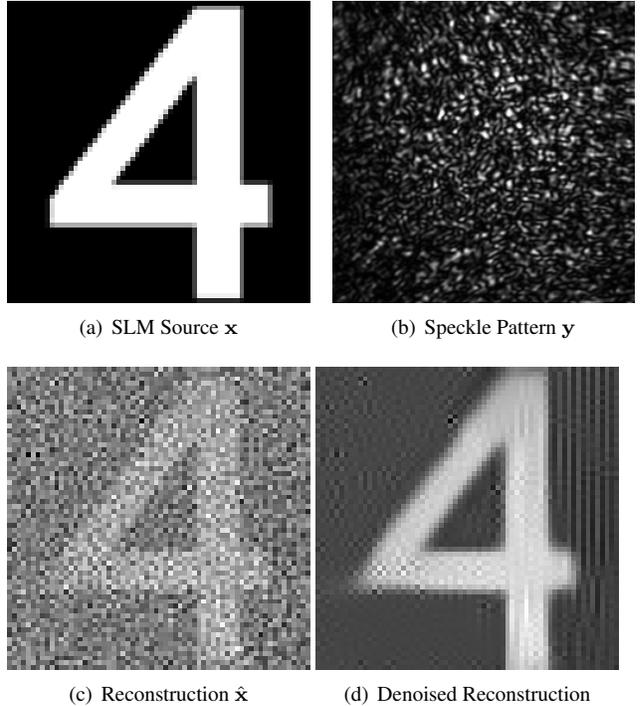


Figure 5: (a) A 64×64 amplitude SLM pattern incident on the scatterer. (b) The corresponding 256×256 speckle pattern captured on the detector. (c) The SLM pattern reconstructed via the TM calibrated using the prVAMP algorithm. (d) The reconstruction after BM3D denoising is applied.

4. A Fast and Robust Phase Retrieval Algorithm

The central algorithmic contribution of our paper is a new AMP-based phase retrieval algorithm that has exceptional run-times and performance, even when the measurement matrix's elements are not zero-mean i.i.d. subgaussian.

4.1. The Algorithm

Rangan, Schniter, and Fletcher recently developed a significant extension to AMP, called vector-AMP (VAMP) [34]. VAMP handles ill-conditioned nonzero-mean matrices without excessive damping nor sequential updates, both of which hurt computation times. In a follow-up work, the authors extended VAMP to solve generalized linear models, creating VAMP-GLM [36]. In this subsection we briefly review VAMP and VAMP-GLM and then introduce our extension, prVAMP, which uses the VAMP-GLM framework to solve the (compressive) phase retrieval problem.⁶

⁶The phase retrieval problem is *compressive* when there are fewer unknowns than required to recover the signal without prior information. It is not the focus of this work; we assume we have high resolution measurements of the signal ($M \gg N$).

4.1.1 VAMP-GLM

VAMP is an evolution of AMP that is able to handle a far broader class of measurement matrices. In particular, whereas AMP has guaranteed asymptotic convergence only when the measurement matrix has i.i.d. zero-mean subgaussian elements [2], VAMP converges asymptotically for all right rotationally invariant matrices⁷ [34]; a far larger set of matrices.

VAMP-GLM is a generalization of VAMP that solves inverse problems involving generalized linear measurements (GLM). GLMs are defined to be any measurement of the form

$$\mathbf{y} = Q(\mathbf{z} + \epsilon) \text{ with } \mathbf{z} = \Phi \mathbf{x}, \quad (3)$$

where \mathbf{x} is our signal of interest, ϵ is noise, and $Q(\cdot)$ denotes a simple non-linearity.

VAMP-GLM works by first splitting the vectors x and z into two sets of identical vectors x_1 and x_2 and z_1 and z_2 . Iterations of the algorithm then broadly consist of four steps: Two denoising steps which impose priors on x_1 and z_1 and two linear minimum mean squared error (LMMSE) estimation steps which ensure x_2 and z_2 are consistent with the measurements y .

VAMP-GLM is presented in Algorithm 1. Within the algorithm $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$ terms are estimates of \mathbf{x} and \mathbf{z} , \mathbf{r} and \mathbf{p} terms act as noisy observations of \mathbf{x} and \mathbf{z} , the γ and τ terms track variances, and the α and β terms are divergence terms used to estimate variances and compute the Onsager correction, which is described later.

On line 4 the algorithm denoises the estimate of the signal \mathbf{x} with the function $g_{x1}(\cdot)$. On line 8 the algorithm denoises the estimate of the signal \mathbf{z} with the function $g_{z1}(\cdot)$. On line 12 it performs LMMSE estimation of \mathbf{x} with the function $g_{x2}(\cdot)$. Finally, on line 17 it perform LMMSE estimation of \mathbf{z} with the function $g_{z2}(\cdot)$.

The LMMSE estimation steps are what differentiate VAMP and VAMP-GLM from previous AMP and GAMP algorithms. Note that the measurement matrix Φ and observations \mathbf{y} do not explicitly show up in the algorithm but are instead modeled within the $g(\cdot)$ functions. See [36] for the full form of $g_{x2}(\cdot)$ and $g_{z2}(\cdot)$. We used a Gaussian prior on \mathbf{x} to set $g_{x1}(\cdot)$; see [32].

Like other AMP algorithms, the exceptional performance of VAMP-GLM is in large part due to the *Onsager correction terms*; the $\alpha \mathbf{r}$ and $\beta \mathbf{p}$ terms in Algorithm 1. These terms debias intermediate results in the algorithm so that at every iteration the denoiser $g_{x1}(\cdot)$ is denoising the true signal plus additive white Gaussian noise. See [9, 24, 32, 34] for more information about the Onsager correction term.

⁷A matrix \mathbf{A} is right-rotationally invariant if it can be written in the form $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^t$, with \mathbf{V} uniformly distributed over the group of orthogonal matrices.

Algorithm 1 VAMP-GLM

```

1: Initialize:  $\mathbf{r}_{10}, \mathbf{p}_{10}, \gamma_{10}, \tau_{10}$ 
2: for  $k=0, 1, \dots, K$  do
3:   Denoise  $x$ :
4:    $\hat{\mathbf{x}}_{1,k} = g_{x1}(\mathbf{r}_{1,k}, \gamma_{1,k}), \alpha_{1,k} = \langle g'_{x1}(\mathbf{r}_{1,k}, \gamma_{1,k}) \rangle$ 
5:    $\mathbf{r}_{2k} = (\hat{\mathbf{x}}_{1,k} - \alpha_{1,k} \mathbf{r}_{1k}) / (1 - \alpha_{1,k})$ 
6:    $\gamma_{2k} = \gamma_{1k} (1 - \alpha_{1,k}) / \alpha_{1k}$ 
7:   Denoise  $z$ :
8:    $\hat{\mathbf{z}}_{1,k} = g_{z1}(\mathbf{p}_{1k}, \tau_{1k}), \beta_{1k} = \langle g'_{z1}(\mathbf{p}_{1k}, \tau_{1k}) \rangle$ 
9:    $\mathbf{p}_{2k} = (\hat{\mathbf{z}}_{1k} - \beta_{1k} \mathbf{p}_{1k}) / (1 - \beta_{1k})$ 
10:   $\tau_{2k} = \tau_{1k} (1 - \beta_{1k}) / \beta_{1k}$ 
11:  LMMSE estimation of  $x$ :
12:   $\hat{\mathbf{x}}_{2k} = g_{x2}(\mathbf{r}_{2k}, \mathbf{p}_{2k}, \gamma_{2k}, \tau_{2k})$ 
13:   $\alpha_{2k} = \langle \gamma'_{x2}(\mathbf{r}_{2k}, \mathbf{p}_{2k}, \gamma_{2k}, \tau_{2k}) \rangle$ 
14:   $\mathbf{r}_{1,k+1} = (\hat{\mathbf{x}}_{2k} - \alpha_{2k} \mathbf{r}_{2k}) / (1 - \alpha_{2k})$ 
15:   $\gamma_{1,k+1} = \gamma_{2k} (1 - \alpha_{2k}) / \alpha_{2k}$ 
16:  LMMSE estimation of  $z$ :
17:   $\hat{\mathbf{z}}_{2k} = g_{z2}(\mathbf{r}_{2k}, \mathbf{p}_{2k}, \gamma_{2k}, \tau_{2k})$ 
18:   $\beta_{2k} = \langle g'_{z2}(\mathbf{r}_{2k}, \mathbf{p}_{2k}, \gamma_{2k}, \tau_{2k}) \rangle$ 
19:   $\mathbf{p}_{1,k+1} = (\hat{\mathbf{z}}_{2k} - \beta_{2k} \mathbf{p}_{2k}) / (1 - \beta_{2k})$ 
20:   $\tau_{1,k+1} = \tau_{2k} (1 - \beta_{2k}) / \beta_{2k}$ 
21: Return  $\hat{\mathbf{x}}_{1K}$ .
```

4.1.2 prVAMP

Our algorithm, prVAMP, is a special case of VAMP-GLM when the output denoiser $g_{z1}(\cdot)$ is designed to handle an absolute value output nonlinearity; $Q(\cdot) = |\cdot|$. In this case, the denoiser $g_{z1}(\cdot)$ is a complicated expression involving the ratio of two Bessel functions. See [35] for its form and a derivation.

To implement prVAMP we modified the original VAMP-GLM code. Code demonstrating prVAMP is available within the GAMP project; <https://sourceforge.net/projects/gampmatlab/files/>.

4.2. prVAMP Simulation Results

In this subsection, we compare the performance of prVAMP to several other phase retrieval algorithms. We restrict ourselves to i.i.d. uniformly distributed $\{0, 1\}^{P \times N}$ measurement matrices, which represent the modulations of an amplitude-only SLM. See Section 5.1 for information about the SLM.

4.2.1 Simulation Settings

For these tests we kept $P = 12N$, which we found offers a reasonable trade-off between runtime and accuracy, and tested at $N = 32^2$ resolutions. See Section 4.3.2 for information about the algorithms' parameters.

4.2.2 Simulation Results

Fig. 7 demonstrates that with an i.i.d. uniformly distributed $\{0, 1\}^{P \times N}$ (amplitude-only) measurement matrix

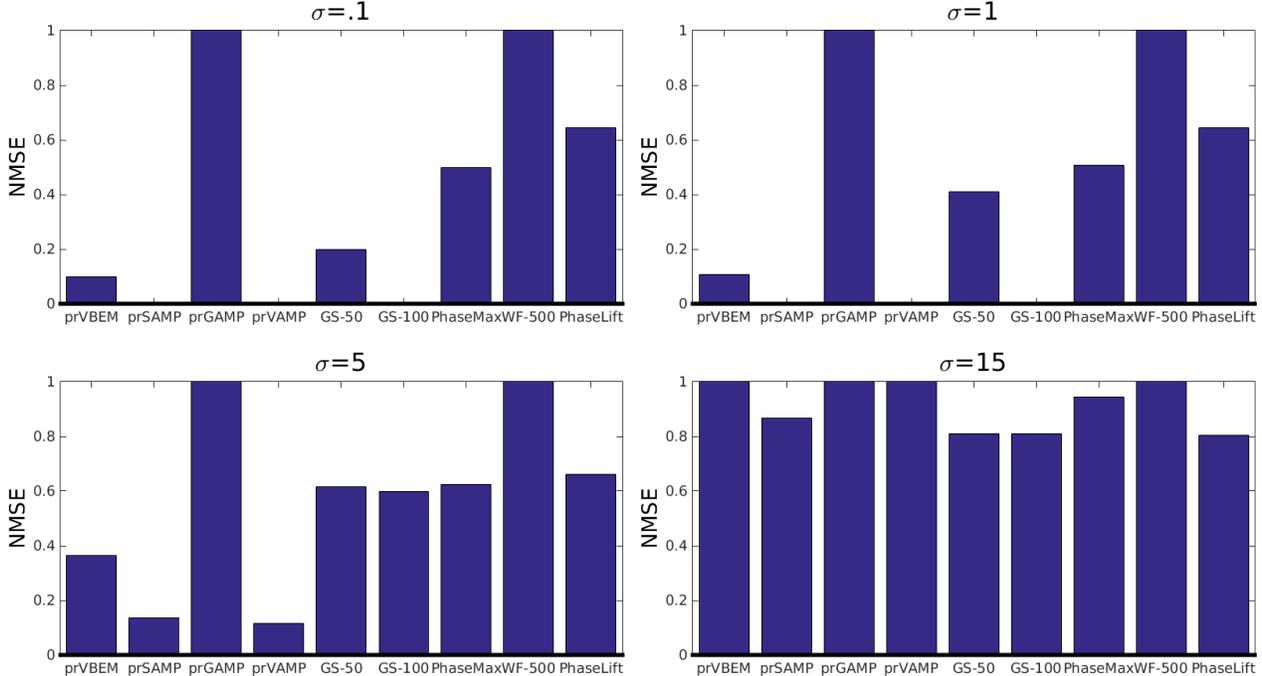


Figure 6: Simulated reconstruction errors with varying amounts of measurement noise from $12 \cdot 16^2$ i.i.d. uniformly distributed $\{0, 1\}$ measurements. prVAMP and prSAMP are the most robust to noise.

prVBEM, prSAMP, prVAMP, and GS are far more accurate than the other methods. Fig. 6 demonstrates that prVAMP and prSAMP are the most robust to noise.

Recall from Table 1 in Section 2.2.3 that prGAMP, prVAMP, GS, and Wirtinger Flow are significantly faster than other algorithms. Thus prVAMP is the only fast, accurate, and robust algorithm we tested.

4.3. prVAMP Experimental Results

In this subsection we provide a detailed description of our experimental setup and demonstrate imaging through a multiple-scattering material.

4.3.1 Physical Setup

Our physical setup is illustrated in Fig. 8. As shown in the figure, a spatially filtered and collimated laser beam ($\lambda = 632.8$ nm) illuminates an SLM from Holoeye, which is a transmissive type display (LC 2012) with 1024×768 resolution and 36 micrometer size square pixels, which modulates the amplitude of the beam before the lens L ($f = 150$ mm), which focuses it onto the scattering medium (a holographic 5 degree diffuser from Thorlabs). Then a microscope objective (Newport, X10, $NA: 0.25$) is used to image the SLM calibration pattern onto the sensor (Point Grey Grasshopper2, pixel size 6.45 micrometer). Using this setup, we learned a $256^2 \times 64^2$ TM using $12 \cdot 64^2$ calibration patterns.

4.3.2 Algorithmic Parameters

Throughout this paper we compare the performance of eight phase retrieval algorithms; Gerchberg Saxon (GS), Wirtinger Flow (WF), PhaseMax, Phase Lift, prSAMP, prVBEM, prGAMP, and prVAMP. All of these algorithms require parameter tuning of some sort and we have done our best to maximize each of their performance with respect to both accuracy and computation time. We describe the parameter tuning in this section. The parameter settings are summarized in Table 3.

We initially ran all algorithms for 50 iterations. We found that GS and WF benefited from more iterations and so also ran GS for 100 and WF for 500.

We initialized each of the algorithms as follows. prGAMP, prVAMP, GS, PhaseLift, and PhaseMax were initialized with random vectors. WF was fed its spectral initializer, described in [5]. prVBEM came with its own initializer which we left intact. prSAMP was sensitive to initialization and was much slower than prVBEM. Therefore it was initialized with the solution produced by prVBEM, as this greatly improved its performance without significantly changing its runtime.

Unlike other phase retrieval algorithms, AMP-based algorithms can easily incorporate, and in fact require, priors about the reconstructed signals. For prSAMP, prGAMP, and prVAMP we tuned the following three parameters for each problem: the signal mean, the signal variance, and the noise variance. We experimented with tuning parameters

Algorithm	Iterations	Initialization	Signal Priors	Other Parameters
prGAMP	50	Random vector	Signal mean and variance and noise variance	Damping = .8
prSAMP	50	prVBEM solution	Signal mean and variance and noise variance	Damping = .9
prVBEM	50	Initializes itself	Noise variance, otherwise defaults	Defaults
prVAMP	50	Random vector	Signal mean and variance and noise variance	Damping = .8
GS	50-100	Random vector	None	Defaults
WF	50-500	Spectral initializer	None	Defaults
PhaseLift	50	Random vector	None	Defaults
PhaseMax	50	Random vector	None	Defaults

Table 3: The parameters used for testing various phase retrieval algorithms.

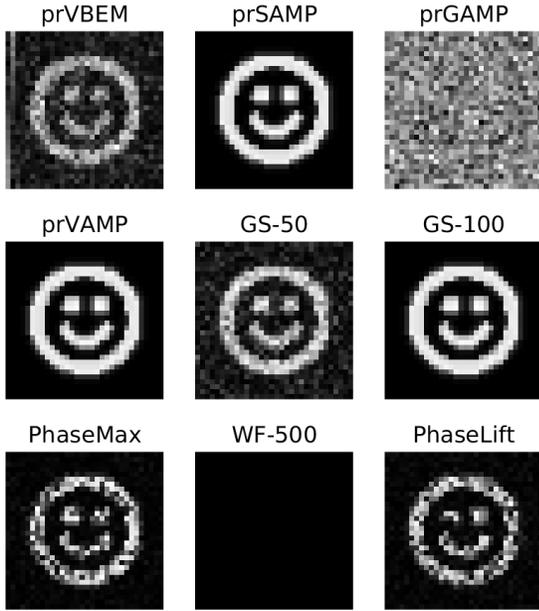
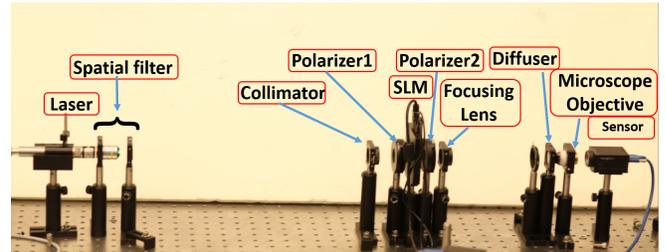


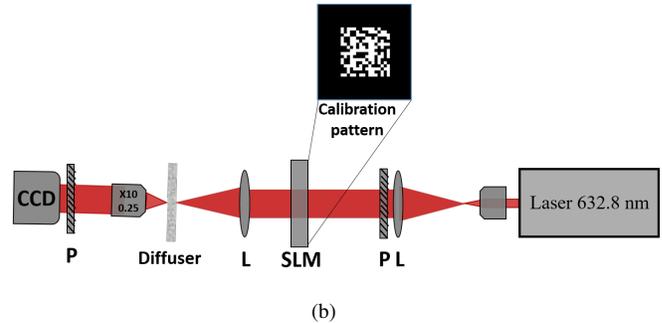
Figure 7: Simulated 32×32 reconstructions with various phase retrieval algorithms from $12 \cdot 32^2$ i.i.d. uniformly distributed $\{0, 1\}$ measurements. Many algorithms suffer, and prGAMP and WF fail, with this type of measurement matrix. prVAMP successfully reconstructs the signal.

for prVBEM as well but found that doing so had negligible effect on the reconstruction performance, so we used its defaults. In all tests the signal mean was set to zero. In simulations the signal and noise variances were set to the oracle truths. When dealing with real data, we used the prVBEM algorithm to calculate an estimate of the signal and noise variance for one row of the TM. These values were then used to tune the other three AMP algorithms for all other rows of the TM.

While using oracle means and variance for one set of al-



(a)



(b)

Figure 8: Experimental setup with an amplitude-only SLM.

gorithms and not for another is in some sense cheating, we feel it is justified in the context of this work. To reconstruct a $256^2 \times 64^2$ transmission matrix with the double phase retrieval approach, one solves over 65000 statistically similar phase retrieval problems. Therefore one has ample data to use for parameter tuning.

In addition to signal priors, prGAMP, prSAMP, and prVAMP have a damping parameter. Limited experimentation found a parameter of .8 worked well for prGAMP and prVAMP and a parameter of .9 worked well for prSAMP. These damping parameters are excessive when dealing with Gaussian measurements, however we found they worked well when reconstructing real-world data.

All other algorithm parameters were set to their default values.

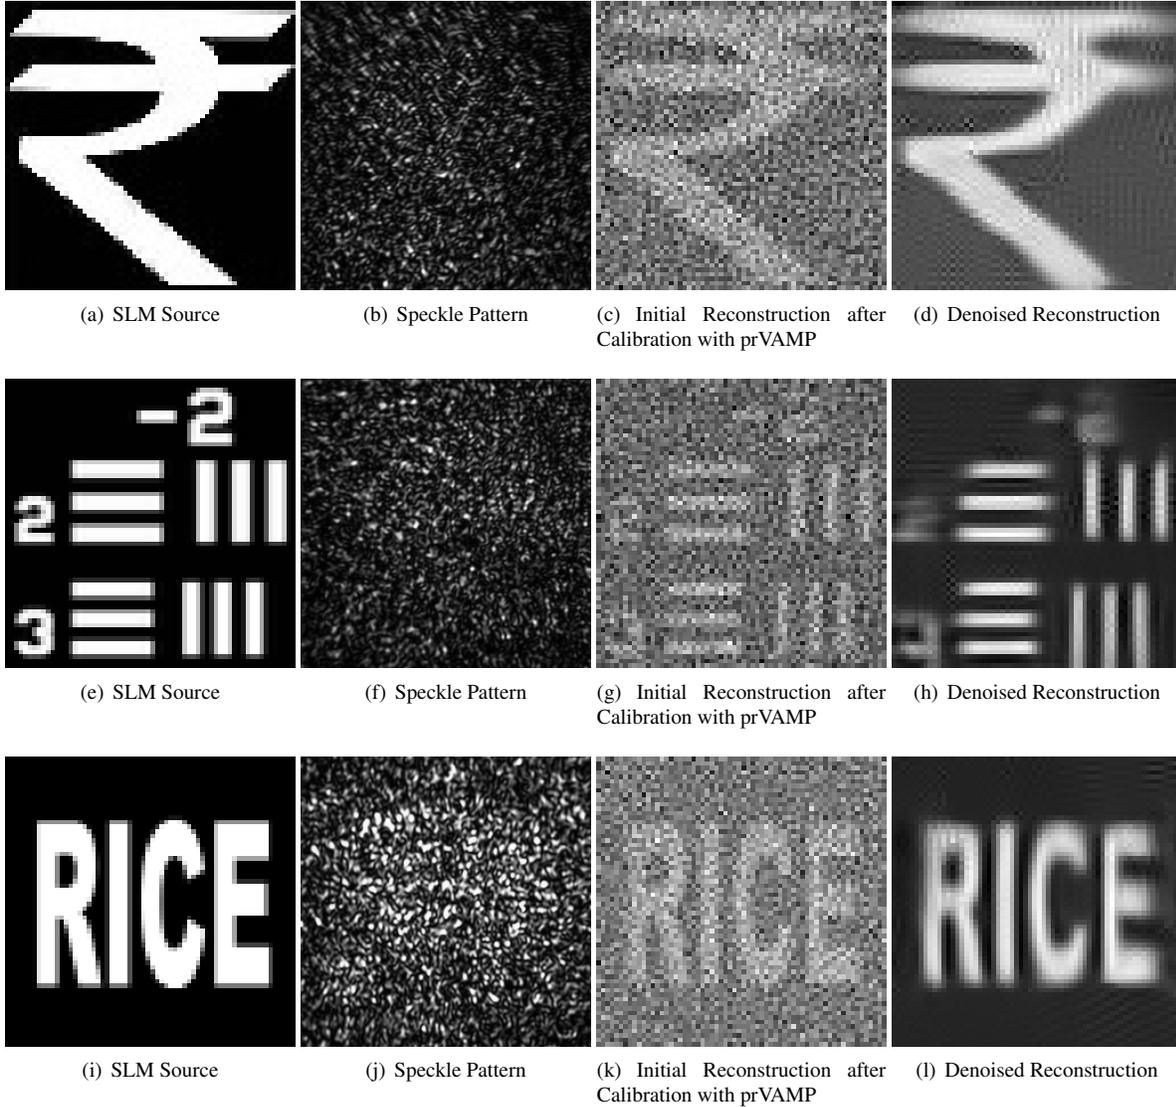


Figure 9: (a,e,i) The first column consists of 64×64 amplitude SLM patterns incident on the scatterer. (b,f,j) The second column consists of the corresponding 256×256 speckle patterns captured on the detector. (c,g,k) The third column consists of the SLM patterns reconstructed via the TM calibrated using the prVAMP algorithm. (d,h,l) The fourth column corresponds to the reconstructions after BM3D denoising is applied. The reconstructions demonstrate prVAMP was able to accurately reconstruct the TM with an amplitude-only SLM.

When estimating the rows of the TM, we observed that the solutions to the phase retrieval problem (2) had residuals varying between 0.1 and 1. A residual greater than 0.4 generally implied that the estimate for the row was inaccurate. Therefore, during imaging, we discarded the rows of the TM whose corresponding phase retrieval problems had a residual higher than 0.4. In general, we found that the algorithm had a residual smaller than 0.4 for most (about 98%) of the rows of the TM. After learning the TM we used it to reconstruct images from their speckle patterns. Lastly,

we cleaned up the initial reconstructions with the BM3D denoiser [8].

4.3.3 Computational Resources

Since the computations of each row of the TM are fully parallelizable, we used a high-throughput computing (HTC) cluster to compute the TMs of the scattering media. The cluster consists of 80 dual processor compute nodes, each of which have two 8-core Intel E5 Ivy Bridge processors running at 2.6 GHz. The compute nodes are each equipped

with between 32 GB and 128 GB of memory.

4.3.4 Results

Our reconstruction results can be found in Fig. 9. Using an amplitude-only SLM and prVAMP for calibration we were able to image through a multiple-scattering material at 64×64 resolutions.

5. Algorithm-friendly TM Learning

Our physical contribution is a simple one; we switch from an amplitude-only SLM to a phase-only SLM. In this section, we explain why this minor change has profound impacts on our ability to learn TMs. In particular, it enables us to reconstruct TMs with the fast and simple prGAMP algorithm [35], as well as a host of other phase retrieval algorithms, like Wirtinger-Flow [5], that rely upon i.i.d. subgaussian measurements.

5.1. Current: Amplitude-only SLM

As the name implies, an amplitude-only SLM modulates the amplitude of the light passing through/off it. To maximize light throughput, we set each source pixel as either completely off, 0, or completely on, 1. Thus with an SLM with N source pixels, each of the calibration signal \mathbf{x}_p is an element of $\{0, 1\}^N$. Likewise, our “measurement matrix” for calibration \mathbf{X}^H is an element of $\{0, 1\}^{P \times N}$.

AMP methods, introduced in Section 2.2.3, were derived under the assumption that the measurement matrix consisted of i.i.d. subgaussian *zero-mean* entries. Our measurement matrix is clearly not zero-mean. Furthermore, because we observe the absolute value of $\mathbf{A}\mathbf{x}$, we cannot apply mean removal [40]. Consequently we have no convergence guarantees and by default most AMP algorithms will diverge.

As described in Section 2.2.3, a number of methods have already been developed to deal with these poorly behaved matrices. These methods work by trading off convergence speed for increased stability. Switching to a phase-only SLM lets us avoid this trade-off.

5.2. Proposed: Phase-only SLM

The physical contribution of our work is changing from an amplitude-only SLM to a phase-only SLM. At each source pixel, the 8-bit phase-only SLM modulates the incoming uniform complex wavefront by an element of $\{\exp 0, \exp 2\pi i \frac{1}{256}, \exp 2\pi i \frac{2}{256}, \dots, \exp 2\pi i \frac{255}{256}\}$. We restricted ourselves to modulations of $\exp 0$ and $\exp \pi i$, which means $\mathbf{X}^H \in \{-1, 1\}^{P \times N}$.

The elements of our matrix $\mathbf{X}^H \in \{-1, 1\}^{P \times N}$ follow a zero-mean i.i.d. subgaussian distribution. Thus the matrix satisfies the AMP assumptions and minimal damping is required. This means that we can apply the low cost prGAMP algorithm with a large step size to perform fast and effective phase retrieval. We can also apply other phase retrieval algorithms that rely upon i.i.d. subgaussian measurements.

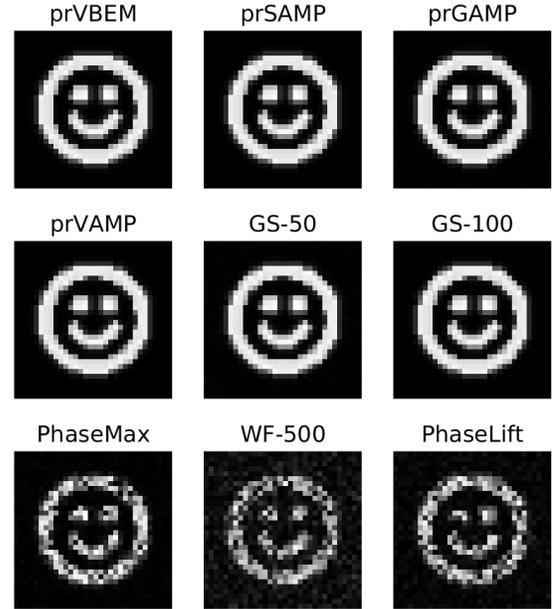


Figure 10: Simulated 32×32 reconstructions with various phase retrieval algorithms from $12 \cdot 32^2$ i.i.d. uniformly distributed $\{-1, 1\}$ measurements. These measurements are zero-mean i.i.d. subgaussian and so all algorithms reconstruct the signal to some degree.

5.3. Phase-only SLM Simulation Results

In this section, we validate our physical contributions in simulation. In particular, we compare the performance of various phase retrieval algorithm when used with uniformly distributed $\{-1, 1\}^{P \times N}$ measurement matrices. For these tests we again kept $P = 12N$ and kept resolutions at $N = 32^2$.

Fig. 10 demonstrates that with uniformly distributed $\{-1, 1\}^{P \times N}$ measurement matrices, all tested algorithms can perform reconstruction successfully, albeit with slightly more error than when the elements of the measurement matrix are i.i.d. $\mathcal{CN}(0, 1)$. Fig. 11 compares the performance of the algorithms at different noise levels when the measurement matrix is uniformly distributed $\{-1, 1\}^{P \times N}$. In this case, the performance of prGAMP is similar or better than the other algorithms. At the same time, prGAMP, GS, WF, and prVAMP enjoy the low computations times observed in Table 1.

5.4. Phase-only SLM Experimental Results

In this section we provide a detailed description of our experimental setup and demonstrate imaging through a multiple-scattering medium.

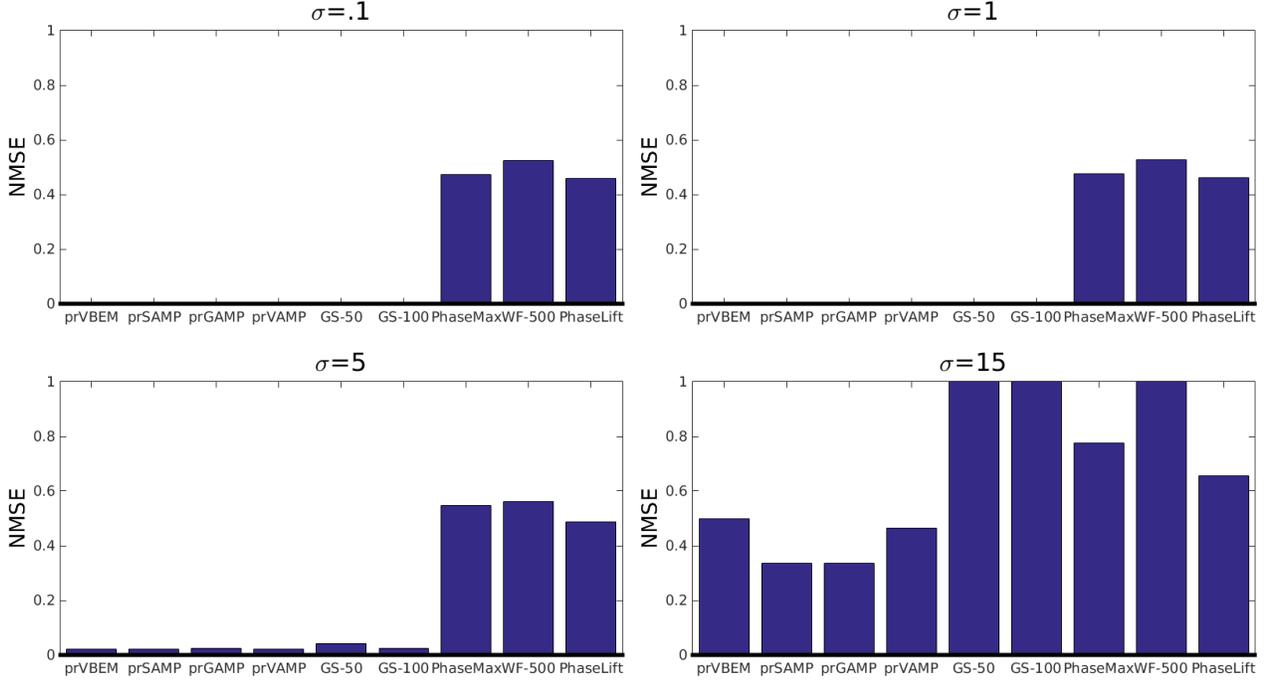
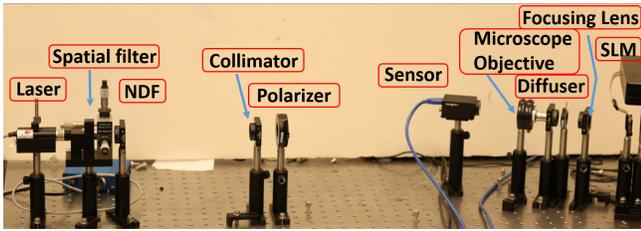
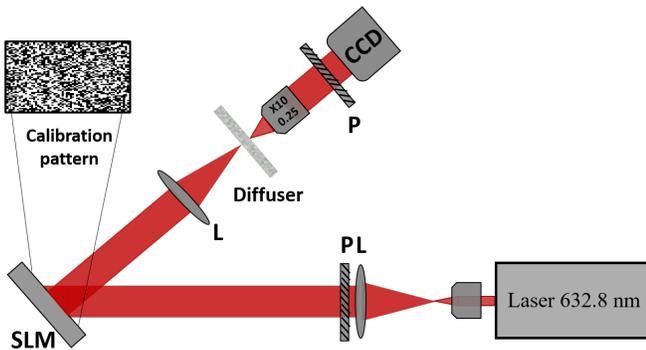


Figure 11: Simulated reconstruction errors with varying amounts of measurement noise from $12 \cdot 16^2$ i.i.d. uniformly distributed $\{-1, 1\}$ measurements. Notice prGAMP now performs comparably to the other AMP algorithms.



(a)



(b)

Figure 12: Experimental setup with a phase-only SLM.

5.4.1 Setup

The physical setup for the phase-only SLM is illustrated in Fig. 12. It is very similar to the phase-only SLM setup describe in Section 4.3.1. The biggest difference between the two is that the phase-only SLM is reflective, whereas the amplitude-only SLM is transmissive.

The algorithms were setup the same as they were in the amplitude-only SLM tests. See Section 4.3.2 for more information. We again applied BM3D denoising [8] on the reconstructed signals.

5.4.2 Results

Our reconstruction results can be found in Fig. 13. Using a phase-only SLM and prGAMP for calibration we were able to image through a multiple-scattering material at 40×40 resolutions.

6. A Public Transmission Matrix Dataset

Our third and most important contribution is the public dissemination of our dataset. Our dataset, which includes the transmission matrices presented here, as well as data captured for $N = 128^2$ transmission matrices that we did not have time to process, has been posted to the following website: <http://dsp.rice.edu/research/transmissionmatrices/>.

This dataset has at least two potential uses. First, the high resolution TM, visualized in Fig. 14, may allow optics

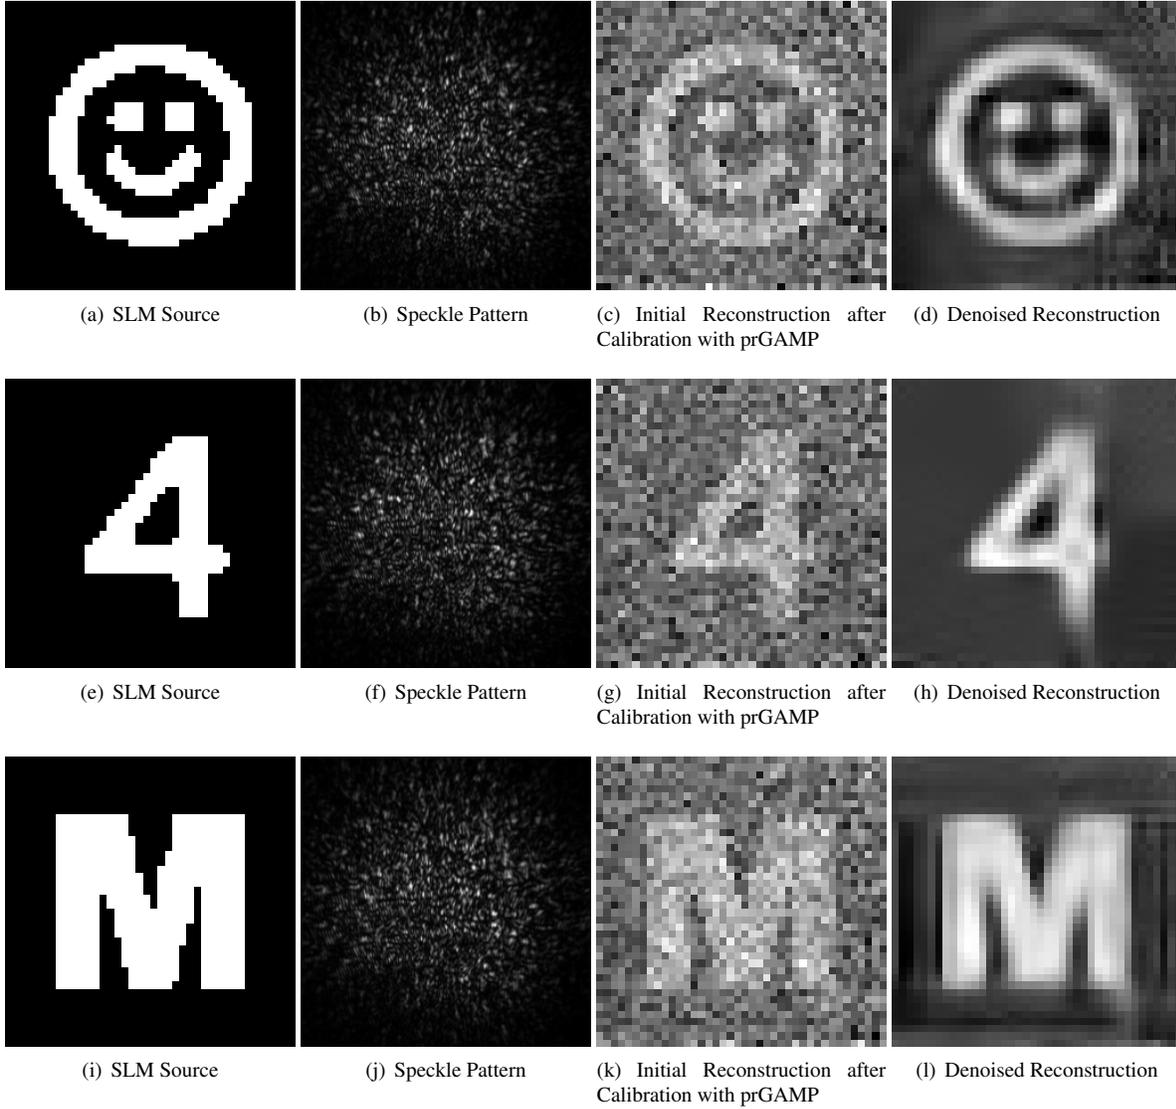


Figure 13: (a,e,i) The first column consists of 40×40 phase SLM patterns incident on the scatterer. (b,f,j) The second column consists of the corresponding 256×256 speckle patterns captured on the detector. (c,g,k) The third column consists of the SLM patterns reconstructed via the TM calibrated using the prGAMP algorithm. (d,h,l) The fourth column corresponds to the reconstructions after BM3D denoising is applied to the reconstructed SLM patterns. The reconstructions demonstrate prGAMP was able to accurately reconstruct the TM with a phase-only SLM.

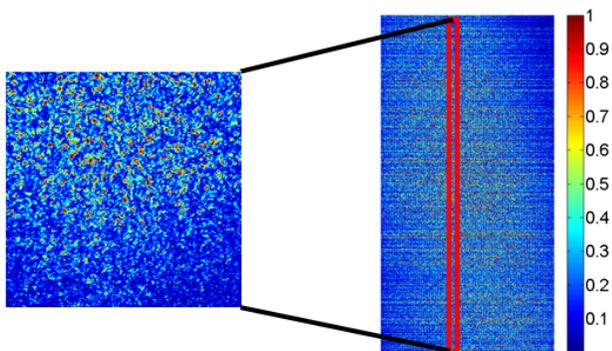
researchers to test new ideas and search for new properties, without themselves having to measure a TM. The transmission matrix clearly has structure, and it is our hope future researchers will figure out how to take advantage of it.

Second, this release provides algorithms researchers with real-world data with which they can apply and test new phase retrieval algorithms. The phase-SLM data may be of particular interest. Many recently developed phase retrieval algorithms rely upon i.i.d. subgaussian measurements. As a result, they do not apply to most real-world datasets. Al-

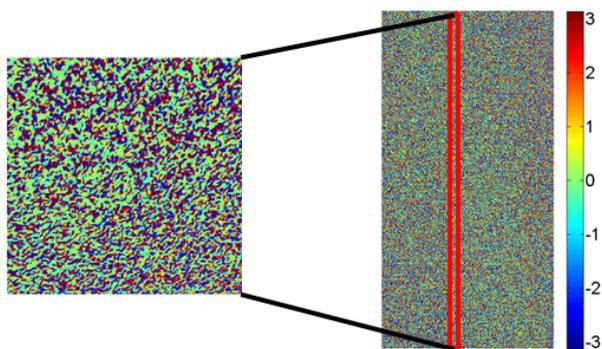
though many of these methods could work with a coded diffraction microscope [5], to our knowledge no such microscope yet exists. Our phase-SLM dataset represents a real-world dataset for which the subgaussian measurements requirement holds.

7. Discussion

In this work we have made three contributions; a new algorithm, a new experimental design, and a new public dataset. Our algorithm prVAMP, a special case of VAMP-



(a) Amplitude of TM



(b) Phase of TM

Figure 14: A visualization of the amplitude (a) and phase (b) of a reconstructed TM. (Left) A column of the TM displayed as $M \times M$ matrices. (Right) The entire transmission matrix. The amplitude visualization of the TM demonstrates some structure.

GLM [36], works with $\{0, 1\}$ measurements and runs hundreds of times faster than competing methods. Our new experimental design learns TMs using $\{-1, 1\}$ i.i.d. subgaussian measurements, thereby enabling a host of different phase retrieval algorithm to be applied to the TM recovery problem. Finally, the public release of our dataset will enable optics researchers to study TMs and will enable phase-retrieval researchers to apply their algorithms to real-world data.

Looking to the future, our high resolution TMs could have a host of interesting uses. The unprecedented resolution of our dataset might help researchers uncover new scattering media properties, e.g., the strong memory effect, which could in turn enable new imaging modalities. Alternatively, this public dataset could serve as a benchmark for newly developed phase retrieval algorithms.

Acknowledgments

Thanks to Phil Schniter for suggesting we extend VAMP-GLM to prVAMP and for providing useful tips along the way. Thanks also to Laurent Daudet and Boshra Rajaei for sharing their prSAMP code.

References

- [1] S. Bahmani and J. Romberg. Phase retrieval meets statistical learning theory: A flexible convex relaxation. *arXiv preprint arXiv:1610.04210*, 2016.
- [2] M. Bayati and A. Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.
- [3] B. G. Bodmann and N. Hammen. Stable phase retrieval with low-redundancy frames. *Advances in Computational Mathematics*, 41(2):317–331, 2015.
- [4] F. Caltagirone, L. Zdeborová, and F. Krzakala. On convergence of approximate message passing. In *2014 IEEE International Symposium on Information Theory*, pages 1812–1816. IEEE, 2014.
- [5] E. J. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- [6] E. J. Candes, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [7] M. Cui. Parallel wavefront optimization method for focusing light through random scattering media. *Optics letters*, 36(6):870–872, 2011.
- [8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007.
- [9] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [10] A. Drémeau and F. Krzakala. Phase recovery from a bayesian point of view: the variational approach. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3661–3665. IEEE, 2015.
- [11] A. Drémeau, A. Liutkus, D. Martina, O. Katz, C. Schülke, F. Krzakala, S. Gigan, and L. Daudet. Reference-less measurement of the transmission matrix of a highly scattering material using a dmd and phase retrieval techniques. *Opt. Express*, 23(9):11898–11911, May 2015.
- [12] J. R. Fienup. Reconstruction of an object from the modulus of its fourier transform. *Optics letters*, 3(1):27–29, 1978.
- [13] I. Freund. Looking through walls and around corners. *Physica A: Statistical Mechanics and its Applications*, 168(1):49–65, 1990.
- [14] I. Freund, M. Rosenbluh, and S. Feng. Memory effects in propagation of optical waves through disordered media. *Physical Review Letters*, 61(20):2328, 1988.

- [15] R. W. Gerchberg. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237, 1972.
- [16] T. Goldstein and C. Studer. Phasemax: Convex phase retrieval via basis pursuit. *arXiv preprint arXiv:1610.07531*, 2016.
- [17] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich. Low-budget transient imaging using photonic mixer devices. *ACM Transactions on Graphics (TOG)*, 32(4):45, 2013.
- [18] U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, and D. Psaltis. Learning approach to optical tomography. *Optica*, 2(6):517–522, 2015.
- [19] O. Katz, P. Heidmann, M. Fink, and S. Gigan. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature Photonics*, 8(10):784–790, 2014.
- [20] S. I. Kelly, C. Du, G. Rilling, and M. E. Davies. Advanced image formation and processing of partial synthetic aperture radar data. *IET Signal Processing*, 6(5):511–520, 2012.
- [21] A. Liutkus, D. Martina, S. Popoff, G. Chardon, O. Katz, G. Lerosey, S. Gigan, L. Daudet, and I. Carron. Imaging with nature: Compressive imaging using a multiply scattering medium. *Scientific Reports*, 4, 2014.
- [22] A. Manoel, F. Krzakala, E. Tramel, and L. Zdeborova. Swept approximate message passing for sparse estimation. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1123–1132, 2015.
- [23] C. A. Metzler, A. Maleki, and R. G. Baraniuk. Bm3d-prgamp: Compressive phase retrieval based on bm3d denoising. In *Image Processing (ICIP), 2016 IEEE International Conference on*, pages 2504–2508. IEEE, 2016.
- [24] C. A. Metzler, A. Maleki, and R. G. Baraniuk. From denoising to compressed sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016.
- [25] M. L. Moravec, J. K. Romberg, and R. G. Baraniuk. Compressive phase retrieval. In *Optical Engineering+ Applications*, pages 670120–670120. International Society for Optics and Photonics, 2007.
- [26] N. Naik, S. Zhao, A. Velten, R. Raskar, and K. Bala. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. In *ACM Transactions on Graphics (TOG)*, volume 30, page 171. ACM, 2011.
- [27] M. N’Gom, M.-B. Lien, T. B. Norris, E. Michielssen, and R. R. Nadakuditi. Controlling light transmission through highly scattering media using semi-definite programming as a phase retrieval computation method. *arXiv preprint arXiv:1612.08214*, 2016.
- [28] S. Popoff, G. Lerosey, R. Carminati, M. Fink, A. Boccarda, and S. Gigan. Measuring the transmission matrix in optics: an approach to the study and control of light propagation in disordered media. *Physical Review Letters*, 104(10):100601, 2010.
- [29] B. Rajaei, S. Gigan, F. Krzakala, and L. Daudet. Fast phase retrieval for high dimensions: A block-based approach. *IEEE Signal Processing Letters*, 23(9):1179–1182, 2016.
- [30] B. Rajaei, S. Gigan, F. Krzakala, and L. Daudet. Robust phase retrieval with the swept approximate message passing (prsamp) algorithm. *arXiv preprint arXiv:1605.07516*, 2016.
- [31] B. Rajaei, E. W. Tramel, S. Gigan, F. Krzakala, and L. Daudet. Intensity-only optical compressive imaging using a multiply scattering material and a double phase retrieval approach. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4054–4058, March 2016.
- [32] S. Rangan. Generalized approximate message passing for estimation with random linear mixing. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pages 2168–2172. IEEE, 2011.
- [33] S. Rangan, P. Schniter, and A. Fletcher. On the convergence of approximate message passing with arbitrary matrices. In *2014 IEEE International Symposium on Information Theory*, pages 236–240. IEEE, 2014.
- [34] S. Rangan, P. Schniter, and A. Fletcher. Vector approximate message passing. *arXiv preprint arXiv:1610.03082*, Oct. 2016.
- [35] P. Schniter and S. Rangan. Compressive phase retrieval via generalized approximate message passing. *IEEE Transactions on Signal Processing*, 63(4):1043–1055, 2015.
- [36] P. Schniter, S. Rangan, and A. K. Fletcher. Vector Approximate Message Passing for the Generalized Linear Model. *arXiv preprint arXiv:1612.01186*, Dec. 2016.
- [37] R. Tadano, A. Kumar Pediredla, and A. Veeraraghavan. Depth selective camera: A direct, on-chip, programmable technique for depth selectivity in photography. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3595–3603, 2015.
- [38] L. Tian and L. Waller. 3d intensity and phase imaging from light field measurements in an led array microscope. *Optica*, 2(2):104–111, 2015.
- [39] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3:745, 2012.
- [40] J. Vila, P. Schniter, S. Rangan, F. Krzakala, and L. Zdeborová. Adaptive damping and mean removal for the generalized approximate message passing algorithm. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2021–2025. IEEE, 2015.
- [41] I. Waldspurger, A. d’Apremont, and S. Mallat. Phase recovery, maxcut and complex semidefinite programming. *Mathematical Programming*, 149(1-2):47–81, 2015.
- [42] J. Yoon, K. Lee, J. Park, and Y. Park. Measuring optical transmission matrices by wavefront shaping. *Optics Express*, 23(8):10158–10167, 2015.
- [43] Y. Zhao, Q. Chen, X. Sui, S. Zhou, and H. Gao. Parallel compressed sensing super-resolution imaging via using multiply scattering medium. In *SPIE Optical Engineering+ Applications*, pages 994813–994813, 2016.