




Phasebook: a survey of selected open problems in phase retrieval

Marc Allain¹ · Selin Aslan² · Wim Coene³ · Sjoerd Dirksen⁴ · Jonathan Dong⁵ · Julien Flamant⁶ · Mark Iwen⁷ · Felix Krahmer⁸ · Tristan van Leeuwen^{4,9} · Oleh Melnyk^{10,11} · Andreas Menzel¹² · Allard P. Mosk¹³ · Viktor Nikitin¹⁴ · Palina Salanevich⁴  · Gerlind Plonka¹⁵ · Matthias Wellershoff¹⁶

Received: 23 May 2025 / Accepted: 25 September 2025
© The Author(s) 2025

Abstract

Phase retrieval is an inverse problem that, on one hand, is crucial in many applications across imaging and physics, and, on the other hand, leads to deep research questions in theoretical signal processing and applied harmonic analysis. This survey paper is an outcome of the recent workshop *Phase Retrieval in Mathematics and Applications (PRiMA)* (held on August 5–9 2024 at the Lorentz Center in Leiden, The Netherlands) that brought together experts working on theoretical and practical aspects of the phase retrieval problem with the purpose to formulate and explore essential open problems in the field.

Keywords Phase retrieval · Ptychography · Inverse problem complexity · Generative priors · Random matrices · Quantization · Wigner distribution deconvolution

Mathematics Subject Classification (2020) 42C15 · 65R32 · 78A46 · 94A12 · 65J22 · 15B52


1 Introduction

Phase retrieval is a fundamental problem in signal processing and imaging in which one seeks to recover a signal from magnitude-only measurements. It arises in many areas of the physical sciences—from X-ray crystallography and astronomy to electron microscopy and quantum tomography—whenever detectors record intensity

Communicated by: Akram Aldroubi

Extended author information available on the last page of the article

Published online: 21 October 2025

 Birkhäuser

without phase information. Because the phase information is lost, the inverse reconstruction problem becomes highly ill-posed, raising fundamental questions regarding the uniqueness and stability of solutions and efficient algorithms to find them.

1.1 Purpose of the survey

Despite these inherent challenges, the field of phase retrieval has witnessed significant advances, driven by fundamental theoretical progress and rapid developments in applications across various imaging disciplines. At the same time, a considerable gap remains between the theoretical understanding of the phase retrieval methods and their practical implementation in real-world scenarios. Recognizing this challenge, the recent workshop *Phase Retrieval in Mathematics and Applications (PRiMA)*, held August 5–9 2024 at the Lorentz Center in Leiden, The Netherlands, brought together experts from a diverse set of theoretical and practical backgrounds to explore critical limitations in phase retrieval, and to formulate essential open problems in the field.

As a direct outcome of *PRiMA*, this paper has a twofold purpose: first, to clearly articulate open problems and their practical relevance, thus providing practitioners with insights into theoretical challenges and potential methodological improvements; and, second, to present theorists with concrete, real-world scenarios that need further theoretical investigation and innovation. By bridging this gap, we hope to stimulate productive collaboration and inspire future research directions.

1.2 Organization of the paper

The remainder of this paper is organized as follows. Section 2 introduces the notation and mathematical framework that will be used in the subsequent sections. Section 3 covers the topic of the *Intrinsic Difficulty of Phase Retrieval*, examining complexity-theoretic aspects of the problem. It surveys measures of problem difficulty, theoretical results on uniqueness and algorithmic hardness, and approaches for establishing confidence in reconstructions. Section 4 addresses the problem of *Physics-Informed Modeling of Object and Probe*, discussing how to incorporate prior knowledge into phase retrieval models. It explores object parameterization strategies (including examples like repetitive unit-cell structures in crystallography and patterned components in imaging) and the challenges of jointly estimating object and probe. Section 5, titled *PtyGenography*, presents the use of generative models for phase retrieval. It describes methodologies for leveraging neural-network-based priors, provides examples of reconstructions under varying noise levels, and discusses the observed trade-offs and performance of these methods. Section 6 focuses on *Structured Random Phase Retrieval*, comparing classical random measurement approaches with structured schemes. It highlights recent advances in designing practical measurement patterns (for example, block-coded masks or correlated illuminations) that approximate the advantages of randomness. Section 7 investigates *One-Bit Quantization for Event-Driven 4D STEM*, illustrating how extreme data quantization (down to single-bit per measurement) in high-speed electron microscopy impacts phase retrieval and what strategies might mitigate the information loss in such scenarios. Finally, Section 8 delves into the *Mathematical Foundations of Wigner Distribution Deconvolution*,

an alternative formulation of phase retrieval, and outlines how analyzing the Wigner distribution (or short-time Fourier transform) can provide insight into unique recovery criteria.

2 Notation

We begin by introducing some general notation that is used in later sections.

- We denote the imaginary unit by $i := \sqrt{-1}$.
- Euler's number is denoted $e := 2.718\dots$
- We write $\mathbb{N} = \{0, 1, 2, \dots\}$.
- For a given positive integer n we denote $[n] := \{0, 1, 2, \dots, n-1\}$.
- Vectors are denoted by bold font small letters, e.g. for $\mathbf{x} \in \mathbb{C}^n$. Its j^{th} entry is denoted $x_j \in \mathbb{C}$, that is, $\mathbf{x} = (x_j)_{j \in [n]}$. Note that indexing of vector coordinates starts with 0.
- For $1 \leq p \leq \infty$ we use $\|\mathbf{x}\|_p$ to denote the ℓ^p -norm of $\mathbf{x} \in \mathbb{C}^n$.
- $\mathbb{S}^{n-1} = \{\mathbf{x} \in \mathbb{C}^n : \|\mathbf{x}\|_2 = 1\}$ denotes the unit sphere in \mathbb{C}^n .
- For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$, their inner product is defined as $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{j \in [n]} x_j \overline{y_j}$.
- For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$, their coordinate-wise (or Hadamard) product is defined as $\mathbf{x} \circ \mathbf{y} = (x_j y_j)_{j \in [n]} \in \mathbb{C}^n$.
- Matrices are denoted by capital letters, e.g., $A \in \mathbb{C}^{m \times n}$. Its $(j, k)^{\text{th}}$ entry is denoted by $A_{jk} \in \mathbb{C}$, so that $A = (A_{jk})_{j \in [m], k \in [n]}$.
- For a matrix $A \in \mathbb{C}^{m \times n}$, its adjoint matrix A^* is defined as $A_{jk}^* = \overline{A_{kj}}$.
- $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 , $\text{Unif}[a, b]$ the uniform distribution on the interval $[a, b]$, and $\text{Pois}(\lambda)$ the Poisson distribution with intensity $\lambda > 0$.

While some of the problems consider an infinite-dimensional setup, where the object and the probe are functions, other problems deal with a discretized, finite-dimensional setup that assumes both of them to be vectors. Here, we introduce the notation for both of these setups.

2.1 Finite-dimensional setup

- The input dimension is denoted by $n \in \mathbb{N} \setminus \{0\}$.
- The number of measurements is denoted by $m \in \mathbb{N} \setminus \{0\}$.
- The object to be recovered in a phase retrieval problem is denoted by $\mathbf{f} \in \mathbb{C}^n$.
- A probe (also called mask or window) is denoted by $\mathbf{w} \in \mathbb{C}^n$.
- The $n \times n$ unitary Discrete Fourier Transform matrix is denoted by $F \in \mathbb{C}^{n \times n}$. Its entries are given by $F_{k,j} := e^{-2\pi i k j / n} / \sqrt{n}$, for $k, j \in [n]$.

- For two vectors $x, y \in \mathbb{C}^n$, their convolution is defined as $(x * y)_j = \sum_{k \in [n]} x_k y_{j-k}$, where $j-k$ is understood modulo n .
- For a given $x \in \mathbb{C}^n$, $D_x \in \mathbb{C}^{n \times n}$ denotes the diagonal matrix with x on its diagonal. In particular, $D_x y = x \circ y$ for $y \in \mathbb{C}^n$.
- The shift operator S_r is defined by $(S_r x)_j = x_{j-r}$, for $j \in [n]$.
- The modulation operator M_k is defined by $(M_k x)_j = e^{2\pi i j k/n} x_j$, for $j \in [n]$.
- The measurement matrix is usually denoted by $A \in \mathbb{C}^{m \times n}$ and the corresponding phaseless measurement map is $\mathcal{A}(f) = |Af|^2$, where the absolute value is taken coordinate-wise.

2.2 Infinite-dimensional setup

- The object to be recovered is denoted by $f: \mathbb{R}^d \rightarrow \mathbb{C}$, where d is the domain dimension. In most of the cases considered in this paper, we have $d=2$, but it can also be $d=1$ or $d=3$ in some cases.
- A probe (also called a mask or a window) is denoted by $w: \mathbb{R}^d \rightarrow \mathbb{C}$, where d is the dimension.
- The Fourier transform of a function f is denoted by $\hat{f}(\xi) = \mathcal{F}[f](\xi) = \int_{\mathbb{R}^d} f(x) e^{-i\langle x, \xi \rangle} dx$ for $d \geq 1$.
- For two functions $f, g: \mathbb{R}^d \rightarrow \mathbb{C}$, their convolution is defined as $(f * g)(t) = \int_{\mathbb{R}^d} f(\tau) g(t - \tau) d\tau$.
- The shift operator S_k is defined by $S_k f(x) = f(x - k)$.

We will now begin to present the individual problems proposed and discussed by various working groups of physicists and mathematicians at PRiMA 2024.

3 Problem: intrinsic difficulty of phase retrieval

Authors: Marc Allain, Selin Aslan, Wim Coene, Julien Flamant, Mark Iwen, Oleh Melnyk, Andreas Menzel, Viktor Nikitin, Gerlind Plonka.

In this section, we consider the general continuous phase retrieval problem

$$I = \mathcal{A}(f) = |Af|^2, \quad (1)$$

of recovering f from I and its discretization

$$I = \mathcal{A}(f) = |Af|^2. \quad (2)$$

Note that we use the same notation for measurement operator \mathcal{A} and linear transform A in continuous and discrete cases as they can be distinguished from notation \mathbf{f} or f . We focus mainly on the standard coherent diffraction imaging (CDI) with $A=F$ and

ptychography, where A corresponds to the (subsampled) short-time Fourier transform matrix.

In practice, the expected behavior of phase retrieval algorithms depends largely on the sample properties such as support and (in)homogeneity of the phase of the object. Numerical and experimental evidence shows that all samples are not equally difficult to retrieve. For instance, it is well-known that objects with smooth edges or strong inhomogeneous phase-fields in their domain are generally harder to reconstruct [1–4]. Unfortunately, the specifics of the object are not known prior to the experiment. Thus, establishing complexity measures based solely on the measurements is a crucial issue in phase retrieval. Possible solutions to this open problem are now discussed, which are mainly structured by the following questions.

Questions 3.1

1. What are the current practices and what measures could be used for estimating the object's complexity from the measurements?
2. Which aspects of object complexity are theoretically understood and what open problem should be addressed?
3. How does an object's complexity constitute itself in the reconstruction process? What is its impact on the convergence rate?
4. Can we justify the obtained reconstruction by providing (pixel-wise) confidence intervals?

3.1 Complexity measures

We consider the intensity of the field diffracted by a sample f , i.e., $I = |\mathcal{F}f|^2$, obtained during a standard CDI experiment. From this diffracted intensity, we also introduce $\gamma = \mathcal{F}^{-1}I$, the autocorrelation of the sample. Sampling aliasing is neglected hereafter, and the autocorrelation of the *discretized* problem, denoted γ , is considered consistent with the spatial sampling of γ . In other words, we assume that the following relation holds

$$\gamma = \{\gamma(\mathbf{r}) : \mathbf{r} \in \mathcal{G}\},$$

with \mathcal{G} a rectangular d -dimensional spatial grid. Since the autocorrelation function and its moments are used in both sparse CDI [5, 6] and ptychography [7–11], its use to derive a complexity measure for phase retrieval problems seems appropriate.

A necessary assumption in CDI is that f lives inside a *bounded* spatial domain in \mathbb{R}^d . As a result, the set

$$S_f = \{\mathbf{r} \in \mathcal{G} : f(\mathbf{r}) \neq 0\}$$

is bounded and its cardinal $|S_f|$ corresponds to the (finite) number of complex unknowns to retrieve. Similarly, in the absence of noise, the autocorrelation γ vanishes outside a bounded domain $S_\gamma = S_f \ominus S_f$ with \ominus the Minkowski difference between sets. The (noise-free) intensity measurements \mathbf{I} being in a one-to-one relationship with γ , the number of independent constraints provided by \mathbf{I} is easily derived

from the cardinal of S_γ . More specifically, since γ is centrosymmetric, the number of independent constraints is therefore $\frac{1}{2}|S_\gamma|$.

Recalling that the number of unknowns is $|S_f|$, the ratio of the number of constraints to unknowns called the *overdetermination ratio* [12] is

$$\Omega = \frac{1}{2} \frac{|S_\gamma|}{|S_f|}.$$

It can be shown that $\Omega \geq 2^{d-1}$, where d is the number of spatial dimensions of the problem. In two dimensions, the overdetermination ratio is always greater or equal to 2, the equality being verified when the support is convex and centrosymmetric [12]. Regardless of the structure of the object itself, it can be conjectured that the intrinsic difficulty of phase retrieval depends on this ratio. However, Ω treats all points $\mathbf{r} \in S_f$ equally, independently of $|f(\mathbf{r})|$. An analogous issue arises with the rank of a matrix, in case of which one can use the stable rank as an alternative [13]. We question whether a generalized version of Ω accounting for $|f(\mathbf{r})|$ could be established.

The distribution of autocorrelation values might also provide insight into the difficulty of the phase problem. In particular, the normalized second-order moment (also known as the “radius of gyration”), which we define here in the continuous case [14]

$$B := \frac{\int \|\mathbf{r}\|^2 |\gamma(\mathbf{r})|^2 d\mathbf{r}}{\int |\gamma(\mathbf{r})|^2 d\mathbf{r}},$$

could be considered under the following assumption

“When B decreases, the difficulty of the phase retrieval problem increases.”

It is worth noting that B can be given an equivalent expression based on the diffracted intensity

$$B = \frac{\int \|\nabla I(\mathbf{q})\|^2 d\mathbf{q}}{\int |I(\mathbf{q})|^2 d\mathbf{q}},$$

with ∇ being the usual gradient operator. Using B as a measure of difficulty is therefore motivated by the following empirical knowledge, disseminated in the CDI community: “the more fringes in the intensity pattern, the easier the problem”. Unlike the overdetermination ratio, this metric depends on the value of the autocorrelation, and not only on its support. As such, it might better capture the intrinsic difficulty of the problem, although it is also more prone to imperfections when the autocorrelation is estimated from data (e.g., the autocorrelation estimate can be severely distorted by the noise and/or by the suppression of the central portion of the intensity pattern I induced by a beam-stop).

3.2 Theoretical background

During the past decade, significant progress was made in understanding the uniqueness and stability of reconstruction from phaseless measurements. Among them a major step forward was the establishment of the local stability bounds [15–17].

Namely, given two objects f and g , it is possible to track their similarity in terms of their diffraction measurements,

$$\min_{|\alpha|=1} \|f - \alpha g\|_{\mathbb{F}} \leq C(g) \|\mathcal{A}(f) - \mathcal{A}(g)\|_{\mathbb{M}}. \quad (3)$$

Here, the left-hand side is the distance between f and g accounting for global phase ambiguity and the right-hand side is the distance between the respective measurements. Typically, $\|\cdot\|_{\mathbb{F}}$ is an \mathbb{L}^{\neq} -norm, possibly restricted to some domain Ω and $\|\cdot\|_{\mathbb{M}}$ is the Sobolev norm, i.e., the sum of \mathbb{L}^{\neq} -norms of the function and its derivative. When we only consider functions on a finite domain Ω , then both norms are the \mathbb{L}^{\neq} -norms in the corresponding spaces [18].

The constant $C(g) > 0$ is called the local stability constant as it depends on g . This dependency is linked to the connectivity of the support of g represented by the function's Cheeger constant, see [19] for more details.

Thus, when the support of g consists of many dis- or weakly connected components, one can construct a new function f by assigning a separate global phase to each “isle”, which will yield similar measurements [15].

In the context of investigating the complexity of the given object, it is interesting whether (3) could be helpful. For instance, by setting $g = \frac{1}{|\Omega|} \int_{\Omega} f(x) dx$ with Ω being the domain of f , the both sides become analogies to variances.

If the above idea is applicable, can we then establish bounds as in (3) in different spaces \mathbb{F} to investigate further properties of f ?

Another interesting direction in understanding stability is the transversality property introduced in [1]. For this, let us look at phase retrieval as the intersection task of finding $f \in \mathbb{A} \cap \mathbb{B}$. In CDI, \mathbb{A} is the set of Fourier magnitudes and \mathbb{B} the support of the object. In ptychography or other diffraction imaging setups with overdetermined linear mapping A , we can instead search for $g \in \mathbb{A} \cap \mathbb{B}$, where

$$\mathbb{A} := \{h : |h| = y\} \quad \text{and} \quad \mathbb{B} := \{h : h = Af \text{ for some } f\},$$

are the set of functions satisfying the magnitude constraints and, respectively, the image of the linear operator A .

When considering \mathbb{B} as the support constraint of the object, one of the claims of [1] is that the difficulty of phase retrieval mainly depends on the geometry near the intersection $\mathbb{A} \cap \mathbb{B}$. In particular, on the dimension of its near-linearization $T_f \mathbb{A} \cap T_f \mathbb{B}$, where $T_f \mathbb{A}$, $T_f \mathbb{B}$ are the tangent spaces at f and since \mathbb{B} is linear, we can write $T_f \mathbb{B} = \mathbb{B}$. The intersection is said to be transversal if $\dim(T_f \mathbb{A} \cap \mathbb{B}) = 0$, and non-transversal otherwise.

One key of the results in [1] Theorem 3.5] is that phase retrieval has a locally Lipschitz inverse mapping at f if and only if the intersection $T_f \mathbb{A} \cap \mathbb{B}$ is transversal. At the present time, it is unclear (at least for us) whether the notion of transversality can be useful in practice. The difficulty in using it seems to be the direct dependency on f that is not available to us.

It would also be beneficial to put transversality into perspective of other uniqueness and stability results known for phase retrieval, such as complementarity property [20, 21] or (robust) rank null space property [22].

3.3 Complexity and reconstruction

Turning to the reconstruction algorithms, we observe that even less is understood about the impact of an object's complexity on the convergence speed. Let us consider a gradient descent method

$$f^{t+1} = f^t - \mu \nabla \mathcal{L}(f^t)$$

applied to the “square-root” objective

$$\mathcal{L}(f) = \sum_{j=1}^m \left| |(Af)_j|^{1/2} - y_j^{1/2} \right|^2,$$

which, from the perspective of maximum likelihood estimation, is an approximation of the Poisson log-likelihood function [23]. With appropriately chosen step size μ [24], its convergence is characterized by

$$\nabla \mathcal{L}(f^t) \rightarrow 0 \text{ as } t \rightarrow \infty, \text{ and } \max_{t=1, \dots, T} \|\nabla \mathcal{L}(f^t)\| \leq \frac{(\mathcal{L}(f^0) - \inf_f \mathcal{L}(f))^{1/2}}{\|A\| \sqrt{T}}, \quad (4)$$

where T is the number of iterations and $\|A\|$ is the spectral norm of the operator A .

A few important remarks are the following:

- By convergence we understand finding a critical point $\nabla \mathcal{L}(f) = 0$, which may not necessarily be the global minimizer. For nonrandom A , this is the best-guaranteed performance up to date as the function \mathcal{L} is nonconvex. On the other hand, for random A the landscape of the (slightly modified) objective is well-understood [25].
- If more than one unknown is being optimized at the time, e.g., probe in blind ptychography [26], the exponent $1/2$ in the denominator is replaced by $p/2$ where p is the number of unknown components, highlighting how reconstruction difficulty scaled with the problem's difficulty.
- It is important to note that (4) is a so-called global bound, which applies for an arbitrary starting point f^0 . Theoretical findings in [27] suggest that for ptychographic reconstruction with the right initialization, it is possible to achieve linear convergence rate, that is to replace $T^{-1/2}$ in (4) with e^{-cT} for $c > 0$. However, how to obtain a sufficiently good initial guess efficiently remains unclear. On the contrary, in [28] for one-dimensional Fourier phase retrieval, the existence of a linearly convergent algorithm is shown to be impossible.¹

¹ unless $p = \text{NP}$.

Prominently, in (4) the impact of noise, measurement operator A , initialization \mathbf{f}^0 and the complexity of the ground truth is hidden in the term $\mathcal{L}(\mathbf{f}^0) - \inf_{\mathbf{f}} \mathcal{L}(\mathbf{f})$.

Due to the empirical evidence on the slower convergence for more complex objects, an important direction of research is to establish a more explicit dependency between the convergence rate (4) and, for instance, the object's smoothness class.

An alternative idea is to perform a numerical study on the conditioning of the Hessian $\nabla^2 \mathcal{L}$ as it determines the (local) convexity of the function and the convergence rate of the gradient method. This leads to a natural follow-up question whether the desired complexity statistics of the object can be based on the Hessian and its conditioning?

3.4 Confidence intervals

Adjacent to the Hessian is the topic of confidence intervals and uncertainty quantification for reconstructed images or their parts. The classical solution to this problem is the construction of confidence intervals, within which the ground truth should lie compared to the obtained estimate. The construction of confidence intervals is particularly crucial for medical applications, where false reconstruction may pose dangers to human health.

To construct the intervals for \mathbf{f} using results for maximum likelihood estimation we require a Fisher information matrix, which is obtained by taking an expectation of the Hessian matrix. This can be done, following the ideas from [29]. However, the performance of the method is limited to dimensions below 100 pixels by the computation time of the inverse of the Fisher information matrix. Whether the computations can be sped up, e.g., using techniques from randomized linear algebra [30], is an open question.

The confidence interval construction is not limited to the maximum likelihood method. Another possibility to compute confidence intervals is using conformal inference techniques [31]. This approach, however, requires a large dataset of possible objects for calibration. Unfortunately, a large realistic dataset of objects is not yet available in diffraction imaging, which significantly limits the applications of conformal inference. The assembly of such a dataset is a high-value future research objective.

4 Problem: physics-informed modeling of object and probe

Authors: Selin Aslan, Oleh Melnyk, Palina Salanevich, Mark Iwen, Felix Krahmer, Gerlind Plonka.

With the advancement of detector capabilities, the dimensionality and complexity of measurement models in diffraction imaging continue to grow, leading to a significant increase in computational demands. Estimating millions of unknown parameters in such high-dimensional models often creates a bottleneck, where iterative methods can stagnate as the updates for each individual pixel are close to zero. To overcome this limitation, reducing the number of unknowns through suitable parameterization of the underlying objects offers a promising solution.

In many applications, the specimen exhibit specific structure that can be represented as a compilation of repeating patterns. For example, in crystallography [32, 33], crystals consist of layers with repeating arrangements of individual atoms. Similarly, in circuit board imaging [34], the focus is on reconstructing patterns formed by numerous individual components, each belonging to a finite set of admissible types, such as chips, resistors, capacitors, and connections. Visualizations of such objects can be found in Fig. 1.

The idea of parametrization is not necessarily restricted to the object itself and can be used for other experimental unknowns. For instance, in electron microscopy, the probe is typically represented by aberrations of circular aperture in terms of Zernike polynomial expansion [35, 36].

Note that the idea of efficient parametrization is not new to imaging and phase retrieval. Ideas from compressed sensing [5, 37–39], dictionary learning [40, 41], and, generative learning [42–44] have been explored in the past. Among these, generative models had the most success in phase retrieval so far, both with pre-trained [42, 43] and zero-shot [44] networks. However, the physical interpretability of such models is limited, which motivates us to consider model-based parametrizations.

4.1 Object parameterization

Motivated by the applications in crystallography and lithography, let us assume that the object f can be modeled as a sum of complex-valued measures

$$f = a_0\lambda + \sum_{j=1}^N a_j S_{t_j} \mu_j(\sigma), \quad (5)$$




Fig. 1 Examples of objects

where λ denotes the Lebesgue measure representing the general object background (i.e., absence of features), and the measures μ_j are (normalized, localized) complex measures representing the features of the object. We assume that $\mu_j \in \mathcal{E}$, where \mathcal{E} with $|\mathcal{E}| \ll N$ is a set of all admissible features. Moreover, the measures μ_j may be influenced by experimental conditions captured by the parameter σ . In this model, the translation operator S_{t_j} centers μ_j at the position t_j . The weights a_j represent the strength of the j th feature.

In the applications above, measures μ_j represent atomic potentials in crystallography and individual chip components in circuit boards. For example, let us consider the crystallographic model from [33], where f represents the Coulomb potential, incorporating temperature effects through the Debye-Waller damping scheme. The model is expressed as

$$f = \exp \left(i\sigma \sum_{j=1}^N \omega_j S_{t_j} v_j \right), \quad (6)$$

where σ is the interaction constant, v_j is the potential function of the j th atom, ω_j is the weight and t_j is the atom's position. If the atoms are sufficiently far apart and the supports of their potentials are approximately disjoint, we can set $\omega_j = 1$ and rewrite (6) in the form of (5),

$$f \approx \lambda + \sum_{j=1}^N S_{t_j} (1 - \exp(i\sigma v_j)) \approx \lambda + \sum_{j=1}^N S_{t_j} (1 - \exp(i\sigma v_j(0))).$$

A crucial detail is that \mathcal{E} is known, as the individual potentials $v \in \mathcal{E}$ of atoms are computed in [45]. In this particular application, the authors initially perform a pixel-based reconstruction, from which they determine N , μ_j and set up initial guesses for σ , a_j , and t_j , which are then optimized.

The above example raises a few important questions.

Question 4.1 The set \mathcal{E} strongly depends on the application and may not be known in advance. What is an efficient strategy to establish it?

Here, we see a potential solution in using patch-based priors [41, 46] to learn \mathcal{E} . However, this approach is prone to reliability and interpretability problems with the learned dictionary.

Question 4.2 How can one select the number of features N and the features μ_j ?

As demonstrated in [33], one approach to this problem is to use the pixel-wise reconstruction as a way to transit to a parametrized model. It can be further developed by considering of a coupled model, where the parametrization model is used as a prior or a regularization term akin to [41]. Studying this idea, as well as alternative meth-

ods for selecting N and μ_j , is an important component in development of an efficient phaseless reconstruction method.

Question 4.3 Can we establish uniqueness and stability guarantees for the model (7)? What are the reasonable sufficient conditions?

In the model (5), the set of parameters is given by (a_j, μ_j, t_j) , $1 \leq j \leq N$ and is generally much smaller than in the pixel-wise representation of the object. While this approach effectively reduces the number of parameters, it comes at the cost of nonlinear interactions between a_j , μ_j and t_j , which makes the analysis of the model more intricate. We discuss Question 4.3 in more detail for a simplified model in Section 4.2.

4.2 Sparse crystallography

In this section, we focus on the following simplification of the crystallographic model (6). We replace the atomic potentials v_j with the Dirac delta measure δ . In this way, we disregard the weight of an atom and rather focus on its position. Then, (6) can be rewritten in terms of distributions as

$$f = \lambda + \sum_{j=1}^N (e^{i\omega_j} - 1) \delta_{t_j} = \lambda + \sum_{j=1}^N a_j \delta_{t_j}, \quad (7)$$

where $a_j = e^{i\omega_j} - 1$ for $j \in \{1, \dots, N\}$.

Problem 4.4 Suppose the signal f is given by the model (7) and let $w \in \mathcal{S}(\mathbb{R}^2)$ be the probe. Given the measurements of the form

$$\mathcal{A}(f)(r, \xi) = |\mathcal{F}(f \cdot S_r w)(\xi)|^2, \quad \xi \in \Omega, \quad r \in R, \quad (8)$$

recover the atom positions t_j and weights a_j for $j \in \{1, \dots, N\}$.

Remark 4.5 In the case of no background, that is, when $f = \sum_{j=1}^N a_j \delta_{t_j}$, the problem can be seen as typical sparse phase retrieval, which can be solved using the Prony's method [47] or lifting with TV-regularization [48].

Note that by replacing the atomic potentials with the Dirac delta measure, we transform the object model into a discrete representation. To further simplify the problem, let us now discretize it and consider the finite-dimensional version of (7)

$$f = \mathbb{1} + \sum_{j=1}^N a_j e_{t_j} \in \mathbb{C}^n. \quad (9)$$

Then the measurements (8) become

$$\begin{aligned}
 \mathcal{A}(f)(r, k) &= \left| (FS_r \mathbf{w})_k + \sum_{j=1}^N a_j (S_r \mathbf{w})_{t_j} (F e_{t_j})_k \right|^2 \\
 &= \left| (FS_r \mathbf{w})_k + \sum_{j=1}^N a_j (S_r \mathbf{w})_{t_j} F_{k, t_j} \right|^2 \\
 &= \left| (M_{-r} F \mathbf{w})_k + \sum_{j=1}^N a_j (S_r \mathbf{w})_{t_j} F_{k, t_j} \right|^2 \\
 &= \left| (F \mathbf{w})_k + \sum_{j=1}^N \frac{a_j}{\sqrt{n}} (S_r \mathbf{w})_{t_j} e^{-2\pi i k(t_j - r)/n} \right|^2.
 \end{aligned} \tag{10}$$

Note that the background term $(F \mathbf{w})_k$ here does not depend on the probe shift r . As a result, the above problem could be treated as a special case of affine phase retrieval [49].

This consideration leads to the following open questions, answering which would allow us to understand benefits and limitations of the proposed model-based object parametrization.

Question 4.6 How large the set of the probe shifts R should be to ensure unique reconstruction of $a_j, t_j, j \in \{1, \dots, N\}$ from (10)?

Question 4.7 What methods can be used to solve this problem efficiently and stably, and under what conditions on the probe \mathbf{w} and the set of the probe shifts R ?

To further improve the reconstruction accuracy and/or to reduce the cardinality of R , one can incorporate prior knowledge on a_j, t_j . The examples of such priors are the admissible range of $|a_j|$ and the separation $|t_j - t_\ell| \geq \delta$ for $j \neq \ell$ and known $\delta > 0$. Furthermore, in crystallographic applications, the atom positions t_j often form a superposition of several lattices.

5 PtyGenography: using generative model for regularization with varying noise levels

Authors: Selin Aslan, Wim Coene, Sjoerd Dirksen, Tristan van Leeuwen, Allard Mosk, Palina Salanevich.

In this section, we consider a more general setup that includes the phase retrieval problem as a partial case. Namely, we consider a non-linear inverse problem of retrieving $f \in \mathbb{C}^n$ from noisy measurements $y = \mathcal{A}(f) + \varepsilon$, where $\mathcal{A}: \mathbb{C}^n \rightarrow \mathbb{C}^m$ is the measurement map, and $\varepsilon \in \mathbb{C}^m$ is the noise term. In the case of the phase retrieval

problem, the measurement map is defined as $\mathcal{A}(f) = |Af|^2$, where $A \in \mathbb{C}^{m \times n}$ is a given matrix (see, e.g. [50, 51] for a recent overview). Note that in this section, we only consider the additive noise model, which is typical for detector noise and does not cover Poisson shot noise.

Assuming \mathcal{A} is injective, the conventional formulation of the reconstruction problem is

$$\min_f \|\mathcal{A}(f) - y\|_2^2. \quad (11)$$

To account for instabilities in the presence of noise, one often adds a regularization term. The most common, classical regularization approach is Tikhonov regularization, which is extensively studied and theoretically well-understood [52]. However, it tends to smooth out high-frequency components, which can be problematic when trying to capture detailed features of the signal.

More recently, generative models have emerged as a powerful alternative to capture detailed prior information on the signal, and, thus, enforce reconstruction stability. More specifically, assuming that the signal $f = \mathcal{G}(z)$ for some $z \in \mathbb{C}^k$, where $\mathcal{G}: \mathbb{C}^k \rightarrow \mathbb{C}^n$ is a given generative model and $k < n$, the problem (11) becomes

$$\min_z \|\mathcal{A} \circ \mathcal{G}(z) - y\|_2^2. \quad (12)$$

The rationale here is that the conditioning of $\mathcal{A} \circ \mathcal{G}$ is generally more favorable than that of \mathcal{A} , at the expense of introducing a bias in the reconstruction. It has indeed been observed in numerical experiments that for high signal-to-noise ratio levels formulation (11) performs better, while for in the case of low signal-to-noise ratio formulation (12) outperforms it [43].

This section provides an overview of the recent work by the authors [53] on regularizing an inverse problem using generative priors and discusses the open problems that emerge from it. This work [53] was initiated and partially conducted during the PRiMA workshop at the Lorentz center.

In this section we

1. characterize the reconstruction error of (11) and (12) in terms of the bias and variance;
2. develop a practical method for detecting bias introduced by generative priors;
3. propose a unified variational framework that balances generative and conventional formulations, achieving stable performance across different noise levels.

In Section 5.1, we explore the recovery properties of the two inverse problem formulations (11) and (12). Section 5.1.1 presents the theoretical foundations obtained in [53], including the key results on the reconstruction error and the bias introduced by the generative model that may hinder recovery in the cases when the true signal lies outside the learned distribution. A unified approach that allows to mitigate this problem and balance generalization and stability by combining (11) and (12) is dis-

cussed in Sect. 5.1.4. Section 5.2 summarizes numerical findings that illustrate the trade-offs between generative and classical approaches. Finally, Sect. 5.3 discusses open problems and future research directions, emphasizing challenges in bias detection, theoretical guarantees, and optimization strategies.

5.1 Methods

We assume that both the measurement map \mathcal{A} and the generative model \mathcal{G} are injective and bi-Lipschitz with constants $\alpha, \beta \geq 1$. That is, for any $f, f' \in \mathbb{C}^n$,

$$\alpha^{-1} \|f - f'\|_2 \leq \|\mathcal{A}(f) - \mathcal{A}(f')\|_2 \leq \alpha \|f - f'\|_2,$$

and likewise for \mathcal{G} , with β in place of α .

Remark 5.1 For the phase retrieval problem, where $\mathcal{A}(f) = |Af|^2$ with $A \in \mathbb{C}^{m \times n}$, the measurement map \mathcal{A} satisfies the (bi-)Lipschitz property whenever it is injective [54]. However, for most matrices A , no explicit bound on the Lipschitz constant α is known. An exception is phase retrieval from locally supported measurements [55].

We furthermore assume that the bi-Lipschitz constant γ of $\mathcal{A} \circ \mathcal{G}$ is more favorable than that of \mathcal{A} , that is, $\gamma < \alpha$. In other words, \mathcal{G} effectively regularizes the inverse problem. Moreover, we assume that \mathcal{G} is well-conditioned, in the sense that $0 < \beta - 1 \ll 1$.

5.1.1 Characterizing the reconstruction error

We denote by f_0 the ground-truth signal we aim to recover from the measurements

$$y = \mathcal{A}(f_0) + \varepsilon,$$

where ε is a bounded noise term. The following results characterize the reconstruction error for conventional (11) and generative priors-based (12) reconstruction methods [53].

Lemma 5.2 *Let $\tilde{f} = \arg \min_f \|\mathcal{A}(f) - y\|_2$ with $y = \mathcal{A}(f_0) + \varepsilon$. Then the reconstruction error is given by*

$$\|\tilde{f} - f_0\|_2 \leq 2\alpha \|\varepsilon\|_2.$$

Lemma 5.3 *Let $\tilde{z} = \arg \min_z \|\mathcal{A} \circ \mathcal{G}(z) - y\|_2$ with $y = \mathcal{A}(f_0) + \varepsilon$, and $\tilde{f} = \mathcal{G}(\tilde{z})$. Then the reconstruction error is bounded by*

$$\|\tilde{f} - f_0\|_2 \leq (1 + 2\alpha\beta\gamma) \|\mathcal{G}(z_0) - f_0\|_2 + 2\beta\gamma \|\varepsilon\|_2,$$

where $z_0 = \arg \min_z \|\mathcal{G}(z) - f_0\|_2$.

Note that under the assumptions stated above, it is not unreasonable to have $\beta\gamma < \alpha$, so that the regularized problem indeed leads to less amplification of noise, at the expense of introducing a bias which mainly depends on the expressiveness of the generative model. More specifically, there is a trade-off between the introduced bias and the noise amplification. More restrictive generative model would have smaller β , leading to the smaller variance, but at the same time it would lead to large bias term in out-of-distribution scenarios. This phenomenon is illustrated in Fig. 6.

5.1.2 Detecting bias

A significant challenge in phase retrieval is detecting and correcting bias due to imperfect measurements. In many applications, such as material science [56, 57], lithography [58], and circuits board manufacturing [59], the ground truth signal has the form $f_0 = \mathcal{G}(z_0) + \boldsymbol{\eta}$, where \mathcal{G} models all the “perfect” signals and $\boldsymbol{\eta}$ represents the signal imperfections and manufacturing defects. Detecting these defects is critical for the success of the reconstruction process. Therefore, we would like to answer the following question.

Question 5.4 Given a solution $\tilde{z} = \arg \min_z |\mathcal{A} \circ \mathcal{G}(z) - \mathbf{y}|_2$ to the optimization problem (12), can we detect or perhaps even quantify the bias $|\mathcal{G}(\tilde{z}) - f_0|$?

Answering this question requires more refined exploration of the properties of the measurement map \mathcal{A} . One possible approach to it is using *Vanderlugt correlation* [60], which gives us a way to highlight the defect structure by correlating $\mathcal{A} \circ \mathcal{G}(\tilde{z})$ with \mathbf{y} .

5.1.3 Vanderlugt correlation

In this section, let us assume that the measurements are diffraction patterns in the Fraunhofer regime, so that the measurement map is given by $\mathcal{A}(\mathbf{f}) = |F\mathbf{f}|^2$, where F is a discrete Fourier matrix [61]. In this setup, a procedure inspired by the *VanderLugt filter* provides a simple means to detect and highlight object defects. The analog VanderLugt filter [60] can detect a given object by correlation in the Fourier domain, using an analog or digital hologram of the object to be recognized. Here, we adapt this principle to find and highlight the imperfections $\boldsymbol{\eta}$, assuming these are spatially sparse.

We consider the Fourier transform $F\mathbf{f}$ of the optimized “perfect” object $\mathbf{f} = \mathcal{G}(z_0)$ obtained by fitting the ground truth f_0 to the generative model. For the diffraction pattern intensity, we then have

$$|F\mathbf{f}_0|^2 = |F\mathbf{f}|^2 + |F\boldsymbol{\eta}|^2 + F\boldsymbol{\eta}(F\mathbf{f})^* + (F\boldsymbol{\eta})^*F\mathbf{f}.$$

In the absence of experimental noise, when $\mathbf{y} = |F\mathbf{f}_0|^2$, the solution obtained by the reconstruction method (12) is \mathbf{f} . The term $|F\boldsymbol{\eta}|^2$ is quadratic in the object imperfection $\boldsymbol{\eta}$ that we assume to be small, so that this term can be neglected. Then we have

$$y - |F(f)|^2 \approx F\eta(Ff)^* + (F\eta)^* Ff,$$

which we can rewrite as

$$F\eta \approx \frac{y - |Ff|^2}{(Ff)^*} - \frac{(F\eta)^* Ff}{(Ff)^*}.$$

The first term on the RHS of the equation can be calculated directly, and is a hologram of the misfit between the ground truth measurements and the measurements of the object recovered by (12). The second term is a conjugate hologram. Such terms are common in holography and typically give rise to speckle-like background noise. Using this equation, we can approximate the object imperfection η by taking the inverse Fourier transform.

To illustrate this consideration numerically, let us assume that the generative model approach (12) perfectly reconstructs letters from the standard roman alphabet, but does not allow accents. Let the ground truth image f_0 be as in Fig. 2(a). Since the generative model only reconstructs the letter ‘a’ and not the dot, the procedure outlined above outputs exactly the dot as shown in Fig. 2(b). The conjugate term gives rise to the noisy fringes in the image.

5.1.4 Towards a unified approach

We propose a combined approach to inference using a generative model

$$\min_{z, f} \|\mathcal{A}(f) - y\|_2^2 + \lambda^2 \|\mathcal{G}(z) - f\|_2^2. \quad (13)$$

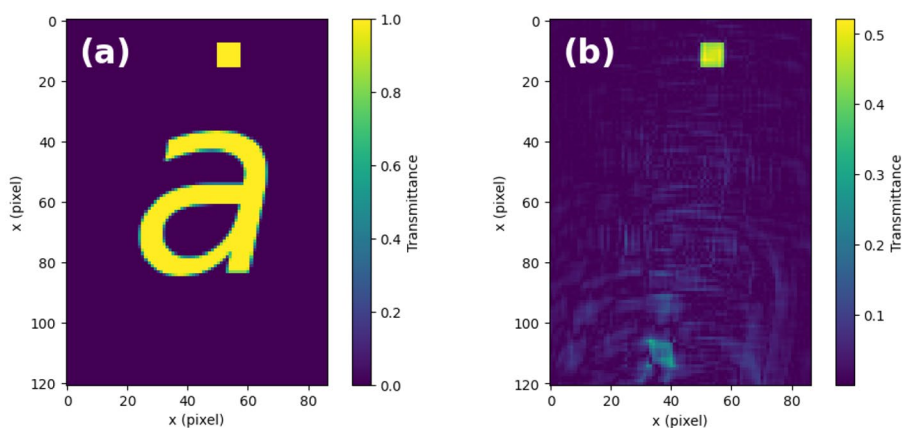


Fig. 2 (a) Example of an object $f_0 = f + \eta$, where f is the letter ‘a’ and the imperfection η is the dot accent. (b) Imperfection η retrieved from correlating the generative reconstruction \mathbf{f} with the given phaseless measurements

Obviously, as $\lambda=0$, we retrieve (11), and as $\lambda \rightarrow \infty$, we retrieve (12). In [53], we show the following reconstruction guarantees for (13).

Lemma 5.5 *The estimate \hat{f} resulting from (13) satisfies*

$$\|\hat{f} - f_0\|_2 \leq \lambda \alpha \|G(z_0) - f_0\|_2 + 2\alpha \|\varepsilon\|_2$$

We note that can view (11), (12), and (13) as partial cases of the generic non-linear optimization problem of the form

$$\min_{x \in \mathbb{C}^d} \|\mathcal{A} \circ \mathcal{B}(x) - y\|_2^2 + \lambda^2 \|w \odot x\|_2^2. \quad (14)$$

We retrieve the specific instances as follows:

- (11) by letting $d=n$, $x := f$, $\mathcal{B} = \mathcal{I}$, and $\lambda=0$;
- (12) by letting $d=k$, $x := z$, $\mathcal{B} = \mathcal{G}$, and $\lambda=0$;
- (13) by letting $d = k + n$ with $x = (x_1, x_2)$, where $x_1 \in \mathbb{C}^k$ and $x_2 \in \mathbb{C}^n$. We define $\mathcal{B} : \mathbb{C}^{n+k} \rightarrow \mathbb{C}^n$ as $\mathcal{B}(x_1, x_2) = \mathcal{G}(x_1) + x_2$, and $w \in \mathbb{C}^{k+n}$ is defined by $w(t) = 0$ for $t \in [k]$ and $w(t) = 1$ for $t \in [k+n] \setminus [k]$. The optimization problem (13) then follows by taking $x_1 := z$ and $x_2 := f - \mathcal{G}(z)$.

To obtain the object reconstruction $\tilde{f} = \mathcal{B}(\tilde{z})$, we solve (14) using a Quasi-Newton method, such as limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm [62]. The reconstruction algorithm we employ thus has parameters \mathcal{B} , λ , w , and a stopping tolerance for L-BFGS.

5.2 Numerical results

For our numerical experiment, we define the measurement map as

$$\mathcal{A}(f) = |Af|^2,$$

where matrix A represents a masked Fourier transform, corresponding to the measurements with ℓ different probes

$$A = \begin{pmatrix} F \text{diag}(a_1) \\ F \text{diag}(a_2) \\ \vdots \\ F \text{diag}(a_\ell) \end{pmatrix}.$$

Here, $a_i \in \mathbb{R}^n$ are random binary probes, that is, their entries are independent identically distributed Bernoulli random variables. The number of measurements is thus $m = n \cdot \ell$.

We define the generative model as

Fig. 3 Samples of the data set on which the generative model was trained ($n=64$). The top row displays the real part while the bottom row displays the imaginary part




Fig. 4 Samples generated by the generative model ($k=30$) the top row displays the real part while the bottom row displays the imaginary part





Fig. 5 Examples of the binary masks used to generate the measurements



$$\mathcal{G}(z) = Gz + b,$$

where $b \in \mathbb{C}^n$ and $G \in \mathbb{R}^{n \times k}$ are obtained by principle component analysis of a data set of handwritten digits [63]. The elements of the data set are 8×8 images, so that $n=64$. For our experiment, we choose $k=30$.

Examples of the elements of the data set on which the generative model \mathcal{G} is trained are shown in Fig. 3. Figure 4 shows examples of the signals $\mathcal{G}(z)$, $z \in \mathbb{C}^{30}$, obtained from the trained generative model. To define the measurement map \mathcal{A} , we use $\ell = 100$ randomly generated binary probes a_j , $j \in [\ell]$. Examples of the probes are shown in Fig. 5. For the numerical experiment, we generate measurements with additive Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_m)$.

To make the reconstruction methods (11) and (12) more stable to the measurement noise, we introduced additional Tikhonov regularization term by setting $\lambda = \sigma^2$ and $w = 1$ in (14). For the reconstruction method (13), we set $\lambda = 10 \cdot \sigma^2$ and $w = 1$, which also introduces additional Tikhonov regularization term for \mathbf{z} . The reconstruction errors for the three methods, tested on both in-distribution (that is, the ground truth is generated by the generative model) and out-of-distribution (where we use samples from the original data set as ground truth) data are shown in Fig. 6.

We see that for high signal-to-noise ratio levels, all methods perform well on in-distribution data, and that for low signal-to-noise ratio levels the generative model shows a slight advantage. On out-of-distribution data, the generative approach (12)




Fig. 6 Relative reconstruction error for the three methods for varying signal-to-noise ratio levels, tested on both in-distribution and out-of-distribution scenarios

clearly shows the bias in the error for high signal-to-noise ratio regime. The combined method (13) achieves the best result for both low and high signal-to-noise ratio levels.

5.3 Conclusion and discussion

In this note, we explored the use of generative models to regularize certain inverse problems, such as phase retrieval. These preliminary results we showed indicate that generative priors can indeed improve the robustness of the inverse problem solution to measurement noise, at the expense of introducing a bias in the reconstruction. To mitigate this issue, we propose a method that aims to combine the best characteristics of both conventional and regularized methods by interpolating between them. Numerical results on phase retrieval from masked Fourier measurements show that the combined method can indeed achieve the best results. However, the presented error bounds are rather crude and can probably be improved with more careful analysis methods. Further research is needed to solidify our understanding of the combined method, to refine it, and to make it feasible for high-dimensional problems. In particular, the following problems should be addressed.

Problem 5.6 Improve the theoretical guarantees on the reconstruction error for a unified approach (13) under appropriate assumptions on the measurement map \mathcal{A} , generative model \mathcal{G} , and the ground truth f_0 . What if \mathcal{A} does not have a (global) bi-Lipschitz property? How the bounds in Lemmas 5.3 and 5.5 can be improved for specific classes of learned and model-based generative models \mathcal{G} ?

Problem 5.7 Devise an efficient algorithm for solving (14) with provable convergence guarantees. Study its robustness under different assumptions on the noise ε and bias $\mathcal{G}(z_0) - f_0$. In particular, as the signal-to-noise ratio level is not known in advance, the regularization parameter λ needs to be adjusted throughout iterations of the reconstruction algorithm.

6 Problem: structured random phase retrieval

Authors: Sjoerd Dirksen, Jonathan Dong, Felix Krahmer, Palina Salanevich.

In this section, we consider the probabilistic formulation of the noiseless phase retrieval problem. More specifically, we make the following assumptions on the object and the measurement map.

- The object $f \in \mathbb{C}^n$ is assumed to be a random vector uniformly distributed over the complex unit sphere, that is, $f \sim \text{Unif}(\mathbb{S}^{n-1})$.
- The measurement map is given by $\mathcal{A}(f) = |Af|^2$, where $A \in \mathbb{C}^{m \times n}$ is a random matrix. The reconstruction guarantees will be drawn depending on its statistical properties.
- Finally, we will consider an asymptotic regime, when $m, n \rightarrow \infty$, while the oversampling ratio $\alpha = \frac{m}{n}$ stays constant.

For simplicity of the discussion, we will focus on the noiseless case, when we are aiming to recover (an estimate) of the ground truth object f_0 from its phaseless measurements

$$y = \mathcal{A}(f_0). \quad (15)$$

Remark 6.1 This phase retrieval problem formulation (15) belongs to a wide class of generalized linear estimation problems. In these problems, the measurements are assumed to have the form

$$y = \varphi_{\text{out}}(Af_0),$$

where $A \in \mathbb{C}^{m \times n}$ is a random measurement matrix and φ_{out} is a stochastic function. All the distributions are assumed to be known. Let $P_{\text{out}}(\cdot|z)$ be the probability density function associated to $\varphi_{\text{out}}(z)$. We obtain the noiseless phase retrieval problem (15), by setting $P_{\text{out}}(y|z) = \delta(y - |z|^2)$, where δ denotes the Dirac delta measure. For more details, we refer an interested reader to [64].

While larger values of the oversampling ratio α imply higher numbers of measurements and, thus, make the recovery of the object f_0 a simpler task, the number of measurements available in practice is often limited. For this reason, we aim to understand the limitations of the phaseless reconstruction in the case when α is small.

More specifically, the goal is to answer the following question for different classes of random sensing matrices A .

Question 6.2 Depending on statistical properties of the random measurement matrix A , what is the smallest value of α for which the reconstructed approximation of f_0 is better than a random guess? What is the smallest value of α for which the average reconstruction error is as small as can be achieved?

In order to formalize this question, we consider the following two reconstruction regimes and define the corresponding threshold values for α .

Let \tilde{f} be the Bayes optimal estimator of f_0 minimizing the mean square error for a given α . Following [64], we define

$$\text{MMSE}(\alpha) = \lim_{n \rightarrow \infty} \mathbb{E} \|\tilde{f} - f_0\|_2^2 \quad (16)$$

to be the asymptotic mean squared approximation error. Here, the expectation is taken with respect to the distribution of the random sensing matrix A . We furthermore define the best possible asymptotic mean squared approximation error as

$$\text{MMSE}_0 = \inf_{\alpha > 0} \text{MMSE}(\alpha). \quad (17)$$

- *Weak recovery.* In this regime, we are not looking to reconstruct f_0 precisely, but rather to retrieve some information about it. We define the weak recovery threshold as

$$\alpha_{\text{WR}} = \inf \{ \alpha > 0 : \text{MMSE}(\alpha) < \mathbb{E} \|f_0\|_2^2 \}. \quad (18)$$

Below α_{WR} , any estimate will inevitably be as good as a random guess, whereas reconstruction can start to return an informed guess for $\alpha > \alpha_{\text{WR}}$.

- *Full recovery.* In this regime, we seek to achieve the best possible reconstruction in expectation with respect to the ground truth object distribution. We define the full recovery threshold as

$$\alpha_{\text{FR}} = \inf \{ \alpha > 0 : \text{MMSE}(\alpha) = \text{MMSE}_0 \}. \quad (19)$$

Note that, depending on the rank of A , full recovery can be perfect, when $\text{MMSE}_0 = 0$, or imperfect, when $\text{MMSE}_0 > 0$. While full reconstruction is theoretically possible for any values of $\alpha > \alpha_{\text{FR}}$, it is not clear it can be obtained algorithmically.

Remark 6.3 We note that the notion of full recovery defined above is different from the injectivity (up to a global phase factor) of the measurement map \mathcal{A} . While injectivity requires that any (deterministic) object can be uniquely reconstructed from its phaseless measurements, perfect full recovery only implies that $\lim_{n \rightarrow \infty} \mathbb{E} \|\tilde{f} - f_0\|_2^2 = 0$. In other words, for high enough input dimension n , the reconstruction error for almost all signals $f_0 \in \mathbb{S}^{n-1}$ is close to zero. Thus, the notion of full recovery is much weaker than the notion of injectivity. Since it is known that for injectivity of \mathcal{A} it is required to have $\alpha \geq 4$ [65, 66], it is expected that $\alpha_{\text{FR}} \leq 4$.

Remark 6.4 Alternatively to the probabilistic viewpoint discussed above, also an algebraic viewpoint has been proposed in the literature. Rather than recoverability with high probability, this viewpoint considers the space of measurement maps and asks if the subset where injectivity fails is of measure zero (typically a low-dimensional algebraic variety). There are also two versions of injectivity in this framework [67]. Strong identifiability requires that almost all measurement maps have the property that no two inputs yield the same measurements, while weak identifiability is satisfied if for a fixed input x , almost all measurement maps do not admit any alternative input x' that yields the same measurements as x . While the focus of our agenda is on the probabilistic notions discussed above, it would be interesting to investigate also the corresponding questions for the algebraic formulation.

6.1 Results for a random measurement matrix

The case when the entries of the measurement matrix A are i.i.d. complex Gaussian random variables, both weak and full recovery regimes are well understood. The following results were shown in [64, 68]:

- The weak recovery threshold satisfies $\alpha_{\text{WR}} = 1$ [68]. Furthermore, the authors also showed that for any $\alpha > 1$, weak recovery of f_0 can be achieved using a spectral method.
- The full recovery threshold satisfies $\alpha_{\text{FR}} = 2$ [64]. To obtain this result, the authors utilized an information theory approach. Furthermore, since in this case matrix A is full-rank with probability 1, for any $\alpha > 2$ perfect full recovery is theoretically guaranteed. At the same time, it is still unknown how to achieve it algorithmically when $\alpha \approx 2$. The closest (polynomial time) algorithmic full recovery threshold is achieved by the Approximate Message Passing (AMP) algorithm, which provably achieves perfect recovery for $\alpha > 2.03$ (see Fig. 7).

These results have been extended to a larger class of random measurement matrices in [64]. Similar thresholds have been derived for Haar-distributed (random unitary) matrices, products of i.i.d. Gaussian matrices, and right-unitarily invariant distributions. For instance, for a random unitary ensemble, both weak and perfect recovery thresholds satisfy $\alpha_{\text{WR}} = \alpha_{\text{FR}} = 2$.

These findings establish a foundation for a robust, practical phase retrieval framework with theoretical reconstruction guarantees at the (close to) minimal number of measurements. This would be in sharp contrast with current approaches in Fourier-based or ptychographic phase retrieval settings where such theoretical guarantees are lacking. However, these large-scale random models are hard to physically implement, making them impractical for applications. Thus, an important problem is to construct a new class of (random) matrices that would be easily implementable while enjoying similar statistical properties.

In practice, randomized models have been introduced in both coded diffraction imaging (CDI) [69] and random-probe ptychography [70]. These randomized physical models improve the reconstruction quality over conventional implementations [70] and reconstruction guarantees for a particular convex relaxation algorithm were

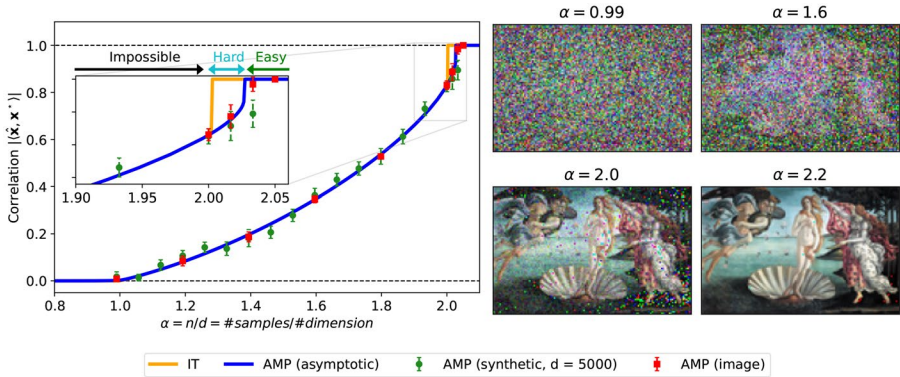


Fig. 7 (Left) correlation $|\langle \hat{f}, f_0 \rangle|$ (higher is better) achieved by the AMP algorithm in noiseless random phase retrieval as a function of the oversampling ratio α . We compare it to the asymptotic theoretical prediction, known as “state evolution” (blue curve) and to the information-theoretic (IT) performance, which is the optimal error that any algorithm can reach, no matter its running time. (right) examples of reconstruction from the AMP algorithm for different oversampling ratios α . Figure from [50]

first derived for $m \sim n \log^4 n$ [69], later improved to $m \sim n \log^2 n$ [71]. This shows the gap in our understanding of these models, as for the random model, the required number of measurements m scales linearly with the problem dimension n .

In coded diffraction imaging and ptychography, the measurement matrix A has a prescribed structure

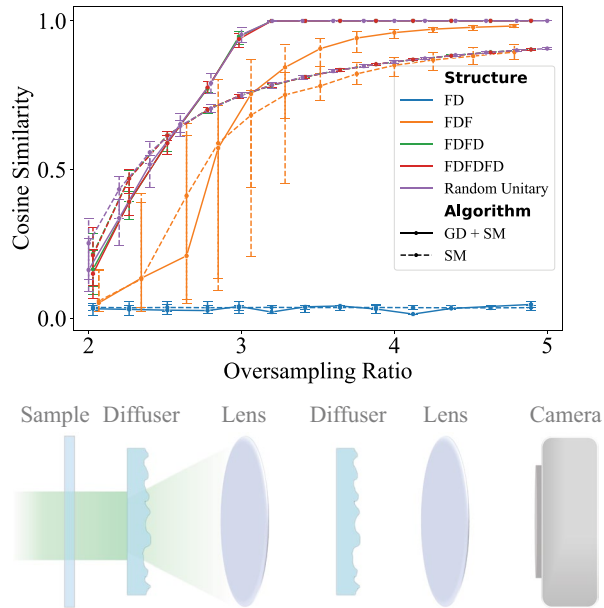
$$A = \begin{bmatrix} F D_{w_1} \\ \vdots \\ F D_{w_\ell} \end{bmatrix}, \quad (20)$$

where the diagonal matrices D_{w_k} correspond to the element-wise multiplication by probes $w_k \in \mathbb{C}^n$, i.i.d. random independent vectors for random CDI or, shifted versions of an initial probe w in the ptychography case. Oversampling corresponds to the number of images ℓ .

6.2 Structured random measurement matrices

To bridge the gap between realistic physical measurement models and the i.i.d. random case with strong recovery guarantees, structured random measurement matrices have recently been proposed [72]. They are built using two main optical components: optical lenses and diffusers (see Fig. 8 (Bottom)). Lenses perform a Fourier transform F on the electric field (neglecting physical scaling factors), which is similar to the operation performed by free-space propagation in ptychography and coded diffraction imaging. Optical diffusers are devices that scatter light by using a rough glass or plastic surface. Neglecting spatial correlations, they can be modeled by a random elementwise multiplication by w with entries $w(j) = e^{iu_j}$ where $u_j \sim \text{i.i.d. Unif}[0, 2\pi]$ [73].

Fig. 8 (Top) performance of structured random models with several architectures and random unitary models. Two structured layers are sufficient to achieve the best results. (bottom) optical system implementing the structured random model. Image adapted from [72]



To emulate the properties of i.i.d. random matrices, several lenses and diffusers can be stacked sequentially:

- In the measurement setup when $A = FDF$, we observe a phase transition from weak to full recovery regime after $\alpha = 3$.
- Consider $A = FD_1FD_2 \cdots FD_\ell$ with $\ell \geq 2$ and $D_k = \text{diag}(w_k)$, where $w_k(j) = e^{iu_{k,j}}$, $u_{k,j} \sim \text{i.i.d. Unif.}[0, 2\pi]$. For such measurement matrices A , we observe a sharp phase transition from weak to full recovery at $\alpha = 2$, suggesting the behavior similar to the Gaussian case discussed in Sect. 6.1.

The numerical findings, including those presented on Fig. 8 (Top), motivate the following open problems, addressing which will lead to rigorous theoretical reconstruction guarantees for this class of physically implementable structured random measurement matrices with optimal oversampling ratio.

Question 6.5 Can we derive (information theory and algorithmic) bounds on the threshold values α_{WR} and α_{FR} for the measurement matrices of the form $A = FDF$ and $A = FD_1FD_2 \cdots FD_\ell$?

In all our considerations above, we assumed that the object f_0 is uniformly distributed on the complex unit sphere. However, in practice, it is often the case that there is prior information available about the ground truth object. This prior information can come in the form of positivity or support constraints, or in form of (deep) generative priors $f_0 = \mathcal{G}(z_0)$ for some $z \in \mathbb{C}^d$ with $d < n$. Such prior information can make the problem of phaseless reconstruction easier, influencing the values of α_{WR} and α_{FR} .

Question 6.6 What is the dependence of α_{WR} and α_{FR} on the object priors?

7 Problem: one-bit quantization for event-driven 4D STEM acquisition in low-dose applications

Authors: Wim Coene, Felix Krahmer, Allard Mosk, Palina Salanevich.

In many imaging experiments, the measured objects, such as proteins or various cell materials, may be highly sensitive to the received dose (see, e.g., [74]). In these cases, the measurements are performed in low-dose regime, and each detector pixel essentially measures the received illumination particle count. Due to the discrete nature of this process, the measurements at the detector pixels can be modeled by independent Poisson random variables.

More formally, let $f \in \mathbb{C}^n$ be the measured object, and let the phaseless measurement map be given by $\mathcal{A}(f) = |Af|^2$, where $A \in \mathbb{C}^{m \times n}$. For instance, in ptychography measurement setup with probe $w \in \mathbb{C}^n$ and ℓ probe positions, the measurement matrix A is given by

$$A = \begin{bmatrix} FD_1 \\ \vdots \\ FD_\ell \end{bmatrix}, \quad (21)$$

where $D_k = \text{diag}(S_k w)$ and F is the discrete Fourier transform matrix. In the Poisson noise model, we assume that the entries of the measurement vector y are given by

$$y_j \sim \text{i.i.d. Poisson}(\mathcal{A}(f)(j)), \quad j \in \{1, \dots, m\}.$$

A commonly used approach to the phase retrieval problem with Poisson noise is negative log-likelihood minimization

$$\tilde{f} = \arg \min_{f \in \mathbb{C}^n} \sum_{j=1}^m (y_j \log \mathcal{A}(f)(j) - \mathcal{A}(f)(j)).$$

The resulting non-convex optimization problem can be solved, for instance, using ADMM [75] or Wirtinger flow algorithm [76]. While these and similar methods show high values of the signal-to-noise ratio in the case when the values $\mathcal{A}(f)(j)$ are large, their performance is limited in the extremely low dose setting [70, 77, 78].

The recent work [79] proposes a different approach to Poisson phase retrieval under extremely low dose assumption by considering this problem from the perspective of *one-bit quantization*. Indeed, already in the case when $\mathcal{A}(f)(j) = 0.2$, $\mathbb{P}(y_j > 1) < 0.018$, so in low-dose regime the probability of the measurements y_j being outside of the set $\{0, 1\}$ is very small. Furthermore, in the case of acquisition with event-driven detectors, the detectors have capacity to record only the first particle arrival, and the measurements are naturally quantized. We discuss this measurement approach in more detail in Section 7.1.

To connect the considered phase retrieval problem to one-bit quantization schemes, let us consider the following. Let, $\hat{y}_j = \min\{y_j, 1\}$, $j \in \{1, \dots, m\}$, be the random quantized measurements obtained by truncating the Poisson measurements y_j with threshold 1. We have that are independent random variables with

$$\hat{y}_i \sim \text{Berboulli}\left(e^{-\mathcal{A}(\mathbf{f})(j)}\right).$$

The idea explored in [79] is to compare them to component-wise deterministic quantization of the phaseless measurements $\mathcal{A}(\mathbf{f})$ with threshold $\delta \in (0, 1)$

$$q_j = \begin{cases} 0, & \mathcal{A}(\mathbf{f})(j) < \delta \\ 1, & \mathcal{A}(\mathbf{f})(j) \geq \delta \end{cases} = \frac{1}{2} (\text{sign}(\mathcal{A}(\mathbf{f})(j) - \delta) + 1).$$

We can now interpret the quantized measurements $\hat{\mathbf{y}}$

as one-bit quantization \mathbf{q} of $\mathcal{A}(\mathbf{f})$ corrupted by independent random bit flip noise.

For this problem reformulation, the authors in [79] proposed a hybrid reconstruction method that uses regularized nuclear norm minimization with appropriate constraints as an initialization for the gradient descent minimization of a loss function incorporating the one-bit quantization assumption. In numerical experiments, where the measurement matrix A is chosen to be a random complex Gaussian matrix with independent entries and the number of measurements is $m = 4n$, the proposed hybrid method shows performance that is superior to the previously studied methods [77, 78]. In these experiments, the low dose regime is simulated via scaling the variance of the entries of A . We note that in this measurement scenario, the number of measurements $m = 4n$ is not sufficient for full reconstruction, thus the algorithms performance is measured in terms of *weak recovery*. In this regime, we aim for the relative reconstruction error (RRE) to be smaller than one, implying that the reconstructed approximation of the object is “better than a random guess”. This is empirically achieved by the method proposed in [79]. At the same time, it is remarkable that for zero initialization (with $\text{RRE} = 1$), the methods from [77, 78] converge to an approximation with $\text{RRE} > 1$, suggesting that their iterations converge to suboptimal local minima.

To make a transition from weak to full recovery (with $\text{RRE} \ll 1$), one needs to increase the number of measurements. To avoid the damage and preserve the overall dose received by an object, this would lead to reducing the dose per pixel value. In this setting, most of the pixels will receive 0 particles, and only a few measurements will be 1. This measurement scenario is realized in event-driven 4D STEM methods.

7.1 Event-driven 4D STEM acquisition

The imaging method we discuss in this section is known as *4D-scanning transmission electron microscopy* (4D-STEM) with event driven detectors, which are originally developed at CERN for detection of elementary particles. In this method, a probe beam of high-energy electrons is scanned over the object on a 2D grid, and the transmitted electrons are detected on a 2D detector working in *event mode* [80].

That is, it detects every incident electron and outputs the time-code, as well as its x and y position. Such an event-driven detector can register a diffraction pattern for each of 10^6 probe positions in 1 second, and the detector consist of 256×256 pixels, with about 50 particle counts per probe position. This imaging method, among other benefits, offers high reconstruction precision and high dose efficiency, with spatial resolution below 0.05 nm. This makes this technology very appealing for applications in material science, including newly emerging 2D materials.

In this measurement setup, for each probe position, we have on average $50/256^2 \approx 7.6 \times 10^{-4}$ detections per pixel of the detector, that is we operate in extremely low dose regime. However, close to the center of the measured diffraction patterns (where the so-called 0^{th} diffraction order resides), the intensity is typically a factor 10^4 higher than in the other (possibly diffuse) diffraction orders. Hence in this area the probability of having $y_j > 1$ is not negligible, but can still be assumed small (estimated to be of order 0.2).

Problem 7.1 Design an efficient reconstruction method for the setting of 4D STEM event-driven acquisition with $m \gg n$ and low overall dose received by the object. Prove theoretical reconstruction and stability guarantees in this measurement setup.

As in this event-driven detection scenario the measurements are naturally quantized, it is particularly interesting to see if the method proposed in [79] can be adapted to the setting with $m \gg n$, with even lower dose per measurement, and the measurement matrix A given by (21).

8 Problem: mathematical foundations of Wigner distribution deconvolution

Authors: Mark Iwen, Felix Krahmer, Oleh Melnyk, Matthias Wellershoff.

One well-known algorithm for solving ptychographic reconstruction is Wigner Distribution Deconvolution (WDD) introduced by Rodenburg and Bates in [11]. It is based on the following identity originally derived by Leon Cohen in a slightly more general form [81]. The intensity measurements²

$$I(x, \xi) = |\mathcal{F}[f \cdot S_x w](\xi)|^2, \quad x, \xi \in \mathbb{R}, \quad (22)$$

with object f and window w , can be rewritten in the form of a convolution

$$I = \mathcal{W}_\zeta * \mathcal{K}_\sqsupset, \quad (23)$$

where \mathcal{W}_ζ is the Wigner distribution of f given by

$$\mathcal{W}_\zeta(x, \xi) := q \int_{\mathbb{R}} f(x + \tau/2) \overline{f(x - \tau/2)} e^{-i\langle \xi, \tau \rangle} d\tau \quad \text{and} \quad \mathcal{K}_\sqsupset(x, \xi) := q \mathcal{W}_\sqsupset(-x, \xi).$$

²Note that these are identical to form to, e.g., (8) up to $\Omega \times R \neq \mathbb{R}^d \times \mathbb{R}^d$.

Applying both spatial and temporal Fourier transforms separates the object term from the window term,

$$\mathcal{F}^{-1}I(\xi, x) = \mathcal{A}_f(x, -\xi) \cdot \mathcal{A}_w(x, \xi), \quad (24)$$

where the *ambiguity function* is given by

$$\mathcal{A}_f(x, \xi) := \mathcal{F}^{-1}\mathcal{W}_f(-\xi, x) = e^{i/2\langle x, \xi \rangle} \cdot \mathcal{F}[f \cdot S_x \bar{f}](\xi). \quad (25)$$

If w is known and $\mathcal{A}_w \neq 0$ almost everywhere, we can recover the ambiguity function of f and then reconstruct f from \mathcal{A}_f .

Since its introduction in 1992, WDD has been used both for the reconstruction from measurements [7, 82, 83] and the analysis of uniqueness and stability of the ptychographic phase retrieval problem [84–86]. A particular interest was dedicated to the discretization of the procedure, which was independently established in the acoustics [87] and imaging communities [88, 89].

For the discrete measurements

$$I_{\xi, x} = |F[f \circ S_x w]_{\xi}|^2, \quad x \in \mathcal{X} \subset [n], \quad \xi \in [n],$$

an analogy of (24) applies

$$(F^{-1}IF)_{\xi, x} = F[f \circ S_{\xi} \bar{f}]_x \cdot \overline{F[\bar{w} \circ S_{\xi} w]_x}, \quad (26)$$

if all scan positions of the window are provided, i.e., $\mathcal{X} = [n]$ [10, 90].

If the window is localized, i.e., $\text{supp}(w) = [\delta]$ for some $\delta < n/2$ and we assume that

$$F[\bar{w} \circ S_{\xi} w]_x \neq 0 \quad \text{for } x \in [d], \quad \xi \in [\delta], \quad (27)$$

then all Fourier coefficients $F[f \circ S_{\xi} \bar{f}]_x$, $x \in [d]$, $\xi \in [\delta]$, and respective vectors $f \circ S_{\xi} \bar{f}$, $\xi \in [\delta]$ are recovered in $\mathcal{O}(n\delta \log n)$ operations.

8.1 Vanishing ambiguity function and subspace completion

In stark contrast to the continuous case, where vanishing of the ambiguity function $\mathcal{F}^{-1}\mathcal{K}_{\square}$ is allowed on a set of measure zero, (27) does not allow any zeros. This is particularly problematic as it excludes symmetric windows persistent in applications [91]. Therefore, a natural question arises:

Question 8.1 How many zero coefficients $F[\bar{w} \circ S_{\xi} w]_x$ are allowed so that \mathbf{f} may be recovered uniquely?

Recent results provide an example with a single zero that allows for unique reconstruction [92], while losing all zero-frequency components $F[\bar{w} \circ S_{\xi} w]_0$, $\xi \in [\delta]$,

leads to ambiguities [93]. As for symmetric windows, they fall in between the two cases above by having at least one zero coefficient in every second vector $\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}$, $\xi \in [\delta]$.

Since $F[\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}]$ are all constructed from the same vector \mathbf{f} , there is a redundancy in the number of the coefficients and recovery of n unknowns should be still possible even if a few out of $n\delta$ coefficients are lost.

By viewing $\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}$ as diagonals of rank-one matrix $\mathbf{f}\mathbf{f}^*$, we can link the Fourier coefficients $F[\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}]_x$ to one another. One possibility is to use the low-rank structure of $\mathbf{f}\mathbf{f}^*$ to formulate recovery as a low-rank recovery problem [94]. In [91], this relation is used in the case when one diagonal $\mathbf{f} \circ S_{\xi_0} \bar{\mathbf{f}}$ for some $\xi_0 \in [\delta]$ is completely recovered and the rest of the diagonals are missing at most one Fourier coefficient, similarly to the symmetric window case. Then, the missing Fourier coefficients are estimated by solving a specially constructed linear system. Although it is yet unclear, whether this procedure, named Subspace Completion (SC), theoretically allows for a unique recovery, its empirical performance suggests so.

The main drawback of SC is that its theoretical underpinning is limited to date to the specific case where at most one coefficient is lost, without a clear way to extend it for scenarios with more missing coefficients. Such settings arise, when the contribution of the Fourier coefficients $F[\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}]$ falls below the noise level. Then, (26) can be regularized by truncation for better noise robustness, i.e., all coefficients $F[\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}]_x$ are set to zero whenever $|F[\bar{\mathbf{w}} \circ S_{\xi} \bar{\mathbf{w}}]_x|$ is small. As typically $|F[\bar{\mathbf{w}} \circ S_{\xi} \bar{\mathbf{w}}]_x|$ are small for x and ξ close to $n/2$, this procedure is similar to applying a low-pass filter.

In turn, the equation

$$F[\mathbf{f} \circ S_{\xi} \bar{\mathbf{f}}]_x = \frac{1}{n} e^{-\frac{2\pi i x \xi}{n}} F[F\mathbf{f} \circ S_{-x} \bar{F\mathbf{f}}]_{\xi},$$

derived in [10] Lemma 3], links bandlimitedness of the object and diagonals of its rank-one representation. This suggests that the uniqueness of reconstruction could be achieved for band-limited objects even if some Fourier coefficients of diagonals are not recovered via (26).

The following three open questions summarize these considerations.

Question 8.2 Can a version of SC be designed to handle an arbitrary number of missing coefficients with a reasonable operational cost?

Question 8.3 Can the performance of SC be theoretically justified?

Question 8.4 Which band limit has to be assumed to guarantee unique recovery for symmetric windows and, more generally, what is the quantitative relation between vanishing $|F[\bar{\mathbf{w}} \circ S_{\xi} \bar{\mathbf{w}}]_x|$ and band-limited \mathbf{f} in terms of WDD?

8.2 Scanning patterns in the discrete setting

As mentioned before, $\mathcal{X} = [n]$ is a crucial assumption for (26) to hold. However, all possible scans are typically not available in practice and (26) has to be adapted to other scanning patterns. The only scenario, in which a version of (26) is established so far, is when the scan positions from a grid, $\mathcal{X} = \{sk : k \in [n/s]\}$. Then, bandlimited [10] or piecewise constant [9] objects can be recovered via (26).

At the same time, working with grid scans is undesirable due to the ambiguities arising in blind ptychography, when the window is unknown [95] and, instead, spiral scanning patterns [96] or randomly perturbed grids [97] are used. This leads to the following question.

Question 8.5 Can WDD be derived for arbitrary scanning patterns with sufficient overlap and which object assumptions are necessary?

We conjecture that the answer is positive for two reasons. First of all, Eq. (26) can be viewed as a discretization of (24) and, therefore, with a proper quadrature rule, one should be able to derive (26) for arbitrary \mathcal{X} . Secondly, it seems possible to apply deconvolution techniques for partially known measurements.

Scanning patterns can also be investigated for the continuous problem (22) with subsampled measurements $(x, \xi) \in \Omega$. Due to the higher complexity of the problem, much more fundamental questions are still open, even the question of the uniqueness of the solutions is still of mathematical interest in many cases. We will now briefly review relevant work in this setting.

8.3 Sampled STFT phase retrieval with compactly supported functions

Consider the recovery of an object $f \in L^2(\mathbb{R})$ from the magnitude of its *short-time Fourier transform (STFT)*,

$$\mathcal{V}_w f(x, \xi) := \int_{\mathbb{R}} f(t) w(t - x) e^{-it\xi} dt = \mathcal{F}[f \cdot S_x w](\xi), \quad x, \xi \in \mathbb{R},$$

where $w \in L^2(\mathbb{R})$ is the window. The STFT is also called the *Gabor transform* when w is chosen to be the Gaussian $t \mapsto 2^{1/4} \exp(-\pi t^2)$. In this case, we denote $\mathcal{G} := \mathcal{V}_w$.

In *sampled STFT phase retrieval*, one is interested in the recovery of f from $|\mathcal{V}_w f|_{\Omega}^2 := \{I(x, \xi) \mid (x, \xi) \in \Omega\}$ with I as per (22) for $d=1$, and $\Omega \subset \mathbb{R}^2$ a, e.g., countable *lattice*,

$$\Omega := \{k v_1 + \ell v_2 \mid k, \ell \in \mathbb{Z}\} \subset \mathbb{R}^2,$$

for a basis $\{v_1, v_2\}$ of \mathbb{R}^2 . Note that the forward operator $f \mapsto |\mathcal{V}_w f|_{\Omega}^2$ is invariant under the action of $\mathbb{S}^0 = \{e^{i\alpha} \mid \alpha \in \mathbb{R}\} \subset \mathbb{C}$ by scalar multiplication on $L^2(\mathbb{R})$. Hence, as above, one instead attempts to recover the orbits of this action, which is referred to as recovery *up to a global phase* in the literature. To describe this the equivalence relation

$$f \sim g \iff \exists \alpha \in \mathbb{R} : f = e^{i\alpha} g$$

is commonly used.

Let $B > 0$ and suppose that f is supported in $[-B, B]$. It is known that $\exists a, b > 0$ such that, for $\Omega = a\mathbb{Z} \times b\mathbb{Z}$, all sufficiently smooth and compactly supported functions $f \in L^2([-B, B])$ whose Fourier series coefficients (after periodically extending f to all of \mathbb{R}) decay monotonically in an averaged sense can be recovered up to a global phase using WDD with windows w that consist of the product of a trigonometric polynomial with a characteristic function whose support includes the support of f as a sufficiently small subset [98, 99]. In the Gabor setting it is also known that compactly supported functions $f \in L^p([-B, B])$, where $p \in [1, \infty]$, can be recovered up to a global phase from their sampled Gabor transform magnitudes on $\Omega = \mathbb{N} \times 2\pi b\mathbb{Z}$ if $b < (4B)^{-1}$ is just slightly smaller than half the Nyquist–Shannon sampling rate [100, 101]. The same applies to bandlimited functions f with $\Omega = 2\pi b\mathbb{Z} \times \mathbb{N}$. We also note that the result extends to sampling on $\{0, 1\} \times \frac{\pi}{2B}\mathbb{Z}$ when $p \geq 2$ [102]. So, only two time bands are actually necessary for uniqueness up to a global phase in the Gabor setting.

Question 8.6 What happens when w is not, e.g., the Gaussian or a simple window of the type utilized in [98, 99]? What other window functions $w \in L^2(\mathbb{R})$ allow recovery up to a global phase from STFT magnitude measurements sampled on a lattice Ω (cf. also [103])?

Question 8.7 Fix any sufficiently “nice” window function w . Do uniqueness results hold for spiral-like sampling sets $\Omega \subset \mathbb{R}^2$? More generally, what exactly determines when a given sampling set Ω is sufficiently informative to allow for recovery of f up to a global phase from $|\mathcal{V}_w f|_\Omega^2$?

A natural question related to the paragraph above is whether it is possible to remove the assumption that our signals $f \in L^2(\mathbb{R})$ are compactly supported (or bandlimited). More precisely, can we recover any given $f \in L^2(\mathbb{R})$ up to a global phase from $|\mathcal{V}_w f|_\Omega^2$? And, if so, with what window function(s) w and sampling set(s) $\Omega \subset \mathbb{R}^2$? A related result [104] shows that $f \in L^2(\mathbb{R})$ cannot be recovered up to a global phase from $|\mathcal{G}f|_\Omega^2$ measured on any set Ω of infinitely many equidistant parallel lines in \mathbb{R}^2 , and this remains true for all other windows $w \in L^2(\mathbb{R})$ as well [105]. Therefore, the recovery of all $f \in L^2(\mathbb{R})$ up to a global phase is impossible when Ω is any lattice (since lattices are always contained in such a set of equidistant parallel lines).

Question 8.8 A more tractable problem than considering the recovery of all $f \in L^2(\mathbb{R})$ up to a global phase using *some* window function w might be to instead focus on related problems in the Gabor setting. For instance, there is a dense set of functions in $L^2(\mathbb{R})$ that *cannot* be recovered up to a global phase from $|\mathcal{G}f|_{a\mathbb{Z} \times b\mathbb{Z}}^2$ while the Gaussian w *can* be recovered up to a global phase from $|\mathcal{G}w|_{a\mathbb{Z} \times b\mathbb{Z}}^2$ [106]. Is there a dense set of other functions that can also be recovered up to a global phase from Gabor magnitude measurements on $\Omega = a\mathbb{Z} \times b\mathbb{Z}$? If so, is this set of the first

or second category? This final question arises from the desire to understand whether we are more “likely” to encounter a function that can be recovered from Gabor magnitude measurements or one that cannot.

In summary, the known results concerning the recovery of f from a *countably infinite* set of measurements $|\mathcal{V}_w f|_\Omega^2$ are currently of the following three types:

1. $f \in L^2(\mathbb{R})$ can be recovered up to a global phase from measurements $\bigcup_{i=1}^4 |\mathcal{V}_{w_i} f|_\Omega^2$ for sufficiently fine lattices Ω provided that the four windows w_i are suitable linear combinations of the Gaussian and the first Hermite function $t \mapsto 2^{5/4} \pi t \exp(-\pi t^2)$ [107],
2. $f \in L^2(\mathbb{R})$ can be recovered up to a global phase from Gabor magnitude measurements $|\mathcal{G}f|_\Omega^2$ if Ω is a square-root lattice [108] or Ω is a union of three separated sets (where Ω is not separated itself) [109], and
3. as previously mentioned, compactly supported functions f (as well as bandlimited functions) can be recovered from measurements $|\mathcal{V}_w f|_\Omega^2$ with special windows w on sufficiently fine lattices Ω [98–102].

Of all of these results, only [98, 99] also considers the *approximate* recovery of f up to a global phase from a *finite* set Ω of (lattice) measurements $|\mathcal{V}_w f|_\Omega^2$. Indeed, such results provide a bridge linking the problems discussed in this section to the discrete WDD algorithms and theory discussed above. As a result, the development of more results along these lines is encouraged.

Question 8.9 Develop practical phase retrieval algorithms $\mathcal{A} : [0, \infty)^k \rightarrow L^2(\mathbb{R}^d)$ together with window functions w and finite sampling schemes $\Omega_k = \{(x_j, \xi_j) \mid 1 \leq j \leq k\} \subset \mathbb{R}^d \times \mathbb{R}^d$ such that, e.g.,

$$\inf_{\alpha \in S^1} \|\alpha f - \mathcal{A}(|\mathcal{V}_w f|_{\Omega_k}^2)\|_{L^2} \rightarrow 0$$

as $k \rightarrow \infty$ for all f in a dense set of functions of $L^2(\mathbb{R}^d)$. We note here that the case $d > 1$ is of particular interest given that [98, 99] focus on $d = 1$.

Acknowledgements The Phase Retrieval in Mathematics and Applications (PRiMA) workshop, Aug. 5–9 2024, which resulted in this manuscript was supported by a Lorentz Center conference grant (funded by Leiden University & Dutch Research Council). In addition, the authors would also like to acknowledge the following additional support: Julien Flamant acknowledges the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-23-CE48-0007 (project ATEMPORAL). Mark Iwen was supported in part by the US National Science Foundation grants NSF DMS 2108479 and NSF EDU DGE 2152014. Palina Salanevich is supported by NWO Talent program Veni ENW grant, file number VI.Veni.212.176. The authors would also like to acknowledge Mahdi Soltanolkotabi, Rima Alaifari, Francesca Bartolucci, and Mitchell Taylor for their valuable insights and discussions that contributed to the development of this manuscript.

Author contributions Marc Allain - Sect. 3. Selin Aslan - Sect. 3, 4, 5. Wim Coene - Sect. 3, 5, 7. Sjoerd Dirksen - Sect. 2, 5, 6. Jonathan Dong - Sect. 6. Julien Flamant - Sect. 3. Mark Iwen - Sect. 1, 2, 3, 4, 8. Felix Krahmer - Sect. 4, 6, 7, 8. Tristan van Leeuwen - Sect. 5. Oleh Melnyk - Sect. 3, 4, 8. Andreas Menzel - Sect. 1, 2, 3. Allard P. Mosk - Sect. 5, 7. Viktor Nikitin - Sect. 3. Gerlind Plonka - Sect. 3, 4. Palina Salanevich - Sect. 1, 2, 4, 5, 6, 7. Matthias Wellershoff - Sect. 8.

Data availability The article presents a survey on the selected open problems in phase retrieval and related research. The contributions are included in the article; further inquiries can be directed to the corresponding author.

Declarations

Competing interest The authors have no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Barnett, A.H., Epstein, C.L., Greengard, L., Magland, J.: Geometry of the phase retrieval problem: graveyard of algorithms. In: Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press (2022).
2. Fienup, J.R.: Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint. *J. Opt. Soc. Am. A* **4**(1), 118–123 (1987).
3. Kim, C., Chamard, V., Hallmann, J., Roth, T., Wei, L., Boesenberg, U., Zozulya, A., Leake, S., Madsen, A.: Three-dimensional imaging of phase ordering in an Fe-Al alloy by Bragg ptychography. *Phys. Rev. Lett.* **121**(25), 6101 (2018).
4. Wang, Z., Gorobtsov, O., Singer, A.: An algorithm for Bragg coherent X-ray diffractive imaging of highly strained nanocrystals. *New J. Phys.* **22**(1), 013021 (2020).
5. Beinert, R., Plonka, G.: Sparse phase retrieval of one-dimensional signals by Prony's method. *Front. Appl. Math. Stat.* **3** (2017).
6. Novikov, A., White, S.: Support recovery for sparse multidimensional phase retrieval. *IEEE Trans. On Signal Process.* **69**, 4403–4415 (2021).
7. Chapman, H.N.: Phase-retrieval X-ray microscopy by Wigner-distribution deconvolution. *Ultramicroscopy* **66**(3–4), 153–172 (1996).
8. Lazić, I., Bosch, E.G.T., Lazar, S.: Phase contrast STEM for thin samples: integrated differential phase contrast. *Ultramicroscopy* **160**, 265–280 (2016).
9. Melnyk, O.: Phase Retrieval from Short-Time Fourier Measurements and Applications to Ptychography. PhD thesis, 2023, Technical University of Munich, Munich, Germany.
10. Perlmutter, M., Merhi, S., Viswanathan, A., Iwen, M.: Inverting spectrogram measurements via aliased Wigner distribution deconvolution and angular synchronization. *Inf. And Inference: A J. IMA* (2020).
11. Rodenburg, J.M., Bates, R.H.T.: The theory of super-resolution electron microscopy via Wigner-distribution deconvolution. *Phil. Trans. of the R. Soc. of Lond. Ser. A: Phys. And Eng. Sci.* **339**(1655), 521–553 (1992).
12. Elser, V., Millane, R.P.: Reconstruction of an object from its symmetry-averaged diffraction pattern. *Acta Crystallogr. A* **64**(2), 273–279 (2008).
13. Rudelson, M., Vershynin, R.: Sampling from large matrices: an approach through geometric functional analysis. *J. ACM (JACM)* **54**(4), 21–es (2007).
14. Cohen, L. Time-frequency analysis, volume 778. Prentice Hall PTR, New Jersey (1995).
15. Alaifari, R., Daubechies, I., Grohs, P., Yin, R.: Stable phase retrieval in infinite dimensions. *Found. Comput. Math.* **19**(4), 869–900 (2019).


16. Cheng, C., Daubechies, I., Dym, N., Jianfeng, L.: Stable phase retrieval from locally stable and conditionally connected measurements. *Appl. Comput. Harmon. Anal.* **55**, 440–465 (2021).
17. Grohs, P., Rathmair, M.: Stable Gabor phase retrieval for multivariate functions. *J. Retailing The Eur. Math. Soc.* **24**(5), 1593–1615 (2022).
18. Grohs, P., Rathmair, M.: L^2 -stability analysis for Gabor phase retrieval. *arXiv preprint arXiv:2108.06154*, 2021.
19. Grohs, P., Rathmair, M.: Stable Gabor phase retrieval and spectral clustering. *Commun. Pure Appl. Math.* **72**(5), 981–1043 (2019).
20. Alaifari, R., Grohs, P.: Phase retrieval in the general setting of continuous frames for Banach spaces. *SIAM J. On Math. Anal.* **49**(3), 1895–1911 (2017).
21. Bandeira, A.S., Cahill, J., Mixon, D.G., Nelson, A.A.: Saving phase: injectivity and stability for phase retrieval. *Appl. Comput. Harmon. Anal.* **37**(1), 106–125 (2014).
22. Kabanava, M., Kueng, R., Rauhut, H., Terstiege, U.: Stable low-rank matrix recovery via null space properties. *Inf. And Inference: A J. IMA* **5**(4), 405–441 (2016).
23. Thibault, P., Guizar-Sicairos, M.: Maximum-likelihood refinement for coherent diffractive imaging. *New J. Phys.* **14**(6), 063004 (2012).
24. Xu, R., Soltanolkotabi, M., Haldar, J.P., Unglaub, W., Zusman, J., Levi, A.F.J., Leahy, R.M.: Accelerated Wirtinger flow: a fast algorithm for ptychography. *arXiv preprint arXiv:1806.05546*, (2018).
25. Cai, J.-F., Huang, M., Dong, L., Wang, Y.: The global landscape of phase retrieval I: perturbed amplitude models. *arXiv preprint arXiv:2112.07993*, (2021).
26. Melnyk, O.: Convergence properties of epic and stochastic gradient methods for blind ptychography. *Anal. And Appl.*, accepted, (2025). <https://doi.org/10.1142/S0219530525500216>.
27. Bendory, T., Eldar, Y.C., Boumal, N.: Non-convex phase retrieval from STFT measurements. *IEEE Trans. Intell Transp. Syst. On Inf. Theory* **64**(1), 467–484 (2018).
28. Huang, M., Xu, Z.: No existence of a linear algorithm for the one-dimensional Fourier phase retrieval. *J. Complex.* **86**(101886) (2025). <https://doi.org/10.1016/j.jco.2024.101886>
29. Wei, X., Urbach, H.P., Coene, W.M.J.: Cramér-Rao lower bound and maximum-likelihood estimation in ptychography with Poisson noise. *Phys. Rev., A* **102**(43516) (2020). <https://doi.org/10.1103/PhysRevA.102.043516>
30. Martinsson, P.-G., Tropp, J.A.: Randomized numerical linear algebra: foundations and algorithms. *Acta Numerica* **29**, 403–572 (2020).
31. Angelopoulos, A.N., Bates, S.: A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*, (2021).
32. Chen, Z., Jiang, Y., Shao, Y.-T., Holtz, M.E., Odstrčil, M., Guizar-Sicairos, M., Hanke, I., Ganschow, S., Schlom, D.G., Muller, D.A.: Electron ptychography achieves atomic-resolution limits set by lattice vibrations. *Science* **372**(6544), 826–831 (2021). <https://doi.org/10.1126/science.abg2533>.
33. Diederichs, B., Herdegen, Z., Strauch, A., Filbir, F., Müller-Caspary, K.: Exact inversion of partially coherent dynamical electron scattering for picometric structure retrieval. *Nat. Commun.* **15**(1), 101 (2024).
34. Shao, Y., Weerdenburg, S., Jacob Seifert, H.P.U., Mosk, A.P., Coene, W.: Wavelength-multiplexed multi-mode EUV reflection ptychography based on automatic differentiation. *Light Sci. Appl.* **13**(1), 196 (2024).
35. Mahajan, V.N.: Zernike circle polynomials and optical aberrations of systems with circular pupils. *Appl. Opt.* **33**(34), 8121–8124 (1994). <https://doi.org/10.1364/AO.33.008121>
36. Ou, X., Zheng, G., Yang, C.: Embedded pupil function recovery for Fourier ptychographic microscopy. *Opt. Express.* **22**(5), 4960–4972 (2014). <https://doi.org/10.1364/OE.22.004960>
37. Iwen, M., Viswanathan, A., Wang, Y.: Robust sparse phase retrieval made easy. *Appl. Comput. Harmon. Anal.* **42**(1), 135–142 (2017). <https://doi.org/10.1016/j.acha.2015.06.007>.
38. Jaganathan, K., Oymak, S., Hassibi, B.: Sparse phase retrieval: convex algorithms and limitations. In: 2013 IEEE International Symposium on Information Theory, pp. 1022–1026. (2013). <https://doi.org/10.1109/ISIT.2013.6620381>.
39. Wang, G., Zhang, L., Giannakis, G.B., Akçakaya, M., Chen, J.: Sparse phase retrieval via truncated amplitude flow. *IEEE Trans. On Signal Process.* **66**(2), 479–491 (2018). <https://doi.org/10.1109/TSP.2017.2771733>.
40. Krishnan, J.P., Bioucas-Dias, J.M., Katkovnik, V.: Dictionary learning phase retrieval from noisy diffraction patterns. *Sensors* **18**(11) (2018). <https://doi.org/10.3390/s18114006>.
41. Tillmann, A.M., Eldar, Y.C., Mairal, J.: DOLPHIn—dictionary learning for phase retrieval. *IEEE Trans. On Signal Process.* **64**(24), 6485–6500 (2016). <https://doi.org/10.1109/TSP.2016.2607180>

42. Aslan, S., Liu, Z., Nikitin, V., Bicer, T., Leyffer, S., Gürsoy, D.: Joint ptycho-tomography with deep generative priors. *Mach. Learn.: Sci. Technol.* **2**(4), 045017 (2021).
43. Seifert, J., Shao, Y., Mosk, A.P.: Noise-robust latent vector reconstruction in ptychography using deep generative models. *Opt. Express.* **32**(1), 1020–1033 (2024). <https://doi.org/10.1364/OE.513556>.
44. Shamshad, F., Ahmed, A.: Compressed sensing-based robust phase retrieval via deep generative priors. *IEEE Sensors J.* **21**(2), 2286–2298 (2021). <https://doi.org/10.1109/JSEN.2020.3018751>.
45. Lobato, I., Van Dyck, D.: An accurate parameterization for scattering factors, electron densities and electrostatic potentials for neutral atoms that obey all physical constraints. *Acta Crystallogr. A* **70**(6), 636–649 (2014). <https://doi.org/10.1107/S205327331401643X>.
46. Altekürger, F., Denker, A., Hagemann, P., Hertrich, J., Maass, P., Steidl, G.: PatchNR: learning from very few images by patch normalizing flow regularization. *Inverse Probl.* **39**(6), 064006 (2023).
47. Beinert, R., Rezaei, S.: Prony-based super-resolution phase retrieval of sparse, multivariate signals. *arXiv preprint arXiv:2301.07696*, (2023).
48. Abouserie, A., et al.: Phase Retrieval for Finitely-Supported Complex Measures via the Fourier Transform. PhD thesis, 2022, University of Houston.
49. Gao, B., Sun, Q., Wang, Y., Zhiqiang, X.: Phase retrieval from the magnitudes of affine linear measurements. *Adv Appl Math.* **93**, 121–141 (2018). <https://doi.org/10.1016/j.aam.2017.09.004>.
50. Dong, J., Valzania, L., Maillard, A., Pham, T.A., Gigan, S., Unser, M.: Phase retrieval: from computational imaging to machine learning: a tutorial. *IEEE Signal Process. Mag.* **40**, 45–57 (2023). <https://doi.org/10.1109/MSP.2022.3219240>.
51. Grohs, P., Koppensteiner, S., Rathmair, M.: Phase retrieval: Uniqueness and stability. *SIAM Rev.* **62**, 301–350 (2020). <https://doi.org/10.1137/19M1256865>.
52. Vasin, V.V.: Some tendencies in the Tikhonov regularization of ill-posed problems. *J. Inverse. Ill. Posed. Probl.* **14**(8) (2006).
53. Aslan, S., van Leeuwen, T., Mosk, A., Salanevich, P.: PtyGenography: using generative models for regularization of the phase retrieval problem. *arXiv:2502.01338*, (2025).
54. Balan, R., Zou, D.: On Lipschitz analysis and Lipschitz synthesis for the phase retrieval problem. *Linear Algebra Appl.* **496**, 152–181 (2016).
55. Iwen, M.A., Merhi, S., Perlmutter, M.: Lower Lipschitz bounds for phase retrieval from locally supported measurements. *Appl. Comput. Harmon. Anal.* **47**(2), 526–538 (2019).
56. Agour, M., Falldorf, C., Bergmann, R.B.: Investigation of composite materials using SLM -based phase retrieval. *Optics Lett.* **38**(13), 2203–2205 (2013).
57. Lehmann, L., Bouillet, S., Leymarie, C., Ameil-Schuh, C., Benoit, M., Rouyer, C.: Characterization of surface defects using a phase retrieval technique in a high-power laser system. *Appl. Opt.* **61**(6), 1545–1551 (2022).
58. Mochi, I., Goldberg, K.A., Huh, S.: Actinic imaging and evaluation of phase structures on extreme ultraviolet lithography masks. *J. Vac. Sci.amp; Technol. B* **28**(6), C6E11–C6E16, (2010).
59. Kaya, G.U.: Development of hybrid optical sensor based on deep learning to detect and classify the micro-size defects in printed circuit board. In: *Measurement* **206**, 112247 (2023).
60. Lugt, A.V.: Signal detection by complex spatial filtering. *IEEE Trans. Intell Transp. Syst. On Inf. Theory* **10**, 139–145 (1964). <https://doi.org/10.1109/TIT.1964.1053650>.
61. Goodman, J.W.: *Introduction to Fourier Optics*. McGraw-Hill, New York (1988).
62. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. *Math. Program.* **45**(1), 503–528 (1989).
63. Deng, L.: The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Process. Mag.* **29**(6), 141–142 (2012).
64. Maillard, A., Loureiro, B., Krzakala, F., Zdeborová, L.: Phase retrieval in high dimensions: statistical and computational phase transitions. *Adv. Neural Inf. Process. Syst.* **33** (2020).
65. Fickus, M., Mixon, D.G., Nelson, A.A., Wang, Y.: Phase retrieval from very few measurements. *Linear Algebra Appl.* **449**, 475–499 (2014).
66. Heinosaari, T., Mazzarella, L., Wolf, M.M.: Quantum tomography under prior information. *Commun. Math. Phys.* **318**(2), 355–374 (2013).
67. Li, Y., Lee, K., Bresler, Y.: Identifiability and stability in blind deconvolution under minimal assumptions. *IEEE Trans. Intell Transp. Syst. On Inf. Theory* **63**(7), 4619–4633 (2017). <https://doi.org/10.1109/TIT.2017.2689779>
68. Mondelli, M., Montanari, A.: Fundamental limits of weak recovery with applications to phase retrieval. *Found. Comput. Math.* **19**(3) (2019).

69. Candes, E.J., Xiaodong, L., Soltanolkotabi, M.: Phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.* **39**(2), 277–299 (2015).
70. Pelz, P.M., Qiu, W.X., Bücker, R., Kassier, G., Miller, R.D.: Low-dose cryo electron ptychography via non-convex bayesian optimization. *Sci. Rep.* **7**(1), 9883 (2017).
71. Gross, D., Krahmer, F., Kueng, R.: Improved recovery guarantees for phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.* **42**(1), 37–64 (2017).
72. Hu, Z., Tachella, J., Unser, M., Dong, J.: Structured random model for fast and robust phase retrieval. *arXiv preprint arXiv:2409.05734*, (2024).
73. Goodman, J.W.: *Statistical Optics*. John Wiley & Sons (2015).
74. Glaeser, R.M.: Limitations to significant information in biological electron microscopy as a result of radiation damage. *J. Ultrastruct. Res.* **36**(3–4), 466–482 (1971).
75. Chang, H., Lou, Y., Duan, Y., Marchesini, S.: Total variation–based phase retrieval for Poisson noise removal. *SIAM J Imag. Sci.* **11**(1), 24–55 (2018).
76. Chen, Y., Candes, E.: Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Adv. Neural Inf. Process. Syst.* **28** (2015).
77. Diederichs, B., Filbir, F., Römer, P.: Wirtinger gradient descent methods for low-dose Poisson phase retrieval. *Inverse Probl.* **40**(12), 125030 (2024).
78. Li, Z., Lange, K., Fessler, J.A.: Poisson phase retrieval in very low-count regimes. *IEEE. Trans. Comput. Imag.* **8**, 838–850 (2022).
79. Krahmer, F., Römer, P.: A one-bit quantization approach for low-dose Poisson phase retrieval. In: 6th International Workshop on the Theory of Computational Sensing and its applications to Radar, Multimodal Sensing, and Imaging (CoSeRa), IEEE (2024).
80. Jannis, D., Hofer, C., Gao, C., Xie, X., Béché, A., Pennycook, T.J., Verbeeck, J.: Event driven 4D STEM acquisition with a Timepix3 detector: microsecond dwell time and faster scans for high precision and low dose applications. *Ultramicroscopy* **233**(113423) (2022).
81. Cohen, L.: Generalized phase-space distribution functions. *J Math Phys.* **7**(5), 781–786 (1966, 05). <https://doi.org/10.1063/1.1931206>.
82. Mitsuishi, K., Nakazawa, K., Sagawa, R., Shimizu, M., Matsumoto, H., Shima, H., Takewaki, T.: Direct observation of Cu in high-silica chabazite zeolite by electron ptychography using Wigner distribution deconvolution. *Sci. Rep.* **13**(1), 316 (2023).
83. Yang, H., MacLaren, I., Jones, L., Martinez, G.T., Simson, M., Huth, M., Ryll, H., Soltau, H., Sagawa, R., Kondo, Y., Ophus, C., Ercius, P., Jin, L., Kovács, A., Nellist, P.D.: Electron ptychographic phase imaging of light elements in crystalline materials using Wigner distribution deconvolution. *Ultramicroscopy* **180**, 173–179 (2017). Ondrej Krivanek: A research life in EELS and aberration corrected STEM. <https://doi.org/10.1016/j.ultramicroscopy.2017.02.006>
84. Alaifari, R., Wellershoff, M.: Stability estimates for phase retrieval from discrete Gabor measurements. *J Fourier Anal Appl.* **27**(2), 2021. <https://doi.org/10.1007/s00041-020-09802-1>
85. Bojarovska, I., Flinth, A.: Phase retrieval from Gabor Measurements. *J Fourier Anal Appl.* **22**(3), 542–567 (2016). <https://doi.org/10.1007/s00041-015-9431-0>.
86. Grohs, P., Liehr, L.: Stable Gabor phase retrieval in Gaussian shift-invariant spaces via biorthogonality. *Constr. Approx.* **59**(1), 61–111 (2024).
87. Richman, M.S., Parks, T.W., Shenoy, R.G.: Discrete-time, discrete-frequency, time-frequency analysis. *IEEE Trans. On Signal Process.* **46**(6), 1517–1527 (1998). <https://doi.org/10.1109/78.678465>.
88. Iwen, M.A., Viswanathan, A., Wang, Y.: Fast phase retrieval from local correlation measurements. *SIAM J Imag. Sci.* **9**(4), 1655–1688 (2016). <https://doi.org/10.1137/15M1053761>.
89. Li, L., Cheng, C., Han, D., Sun, Q., Shi, G.: Phase retrieval from multiple-window short-time Fourier measurements. *IEEE Signal Process. Lett.* **24**(4), 372–376 (2017). <https://doi.org/10.1109/LSP.2017.2663668>
90. Cordor, C., Williams, B., Hristova, Y., Viswanathan, A.: Fast 2D phase retrieval using bandlimited masks. In: Marques, A., Hunyadi, B. (eds.): 28th European Signal Processing Conference (EUSIPCO 2020), pp. 980–984. IEEE, [Piscataway, NJ] (2020). <https://doi.org/10.23919/Eusipco47968.2020.9287439>.
91. Forstner, A., Krahmer, F., Melnyk, O., Sissouno, N.: Well-conditioned ptychographic imaging via lost subspace completion. *Inverse Probl.* **36**(10), 105009 (2020). <https://doi.org/10.1088/1361-6420/abaf3a>.
92. Bartusel, D.: Injectivity conditions for STFT phase retrieval on Z , Z^d and R^d . *J Fourier Anal Appl.* **29**(4), 53 (2023).

93. Melnyk, O., Römer, P.: Background removal for ptychography via Wigner distribution deconvolution. *SIAM J Imag. Sci.* **17**(3), 1978–2014 (2024). <https://doi.org/10.1137/24M1642433>
94. Lee, J., Barbastathis, G.: Denoised Wigner distribution deconvolution via low-rank matrix completion. *Opt. Express.* **24**(18), 20069–20079 (2016). <https://doi.org/10.1364/OE.24.020069>
95. Bendory, T., Edidin, D., Eldar, Y.C.: Blind phaseless short-time Fourier transform recovery. *IEEE Trans. Intell Transp. Syst. On Inf. Theory* **66**(5), 3232–3241 (2020). <https://doi.org/10.1109/TIT.2019.2947056>
96. Huang, X., Yan, H., Harder, R., Hwu, Y., Robinson, I.K., Chu, Y.S.: Optimization of overlap uniformness for ptychography. *Optics Express* **22**(10), 12634–12644 (2014). <https://doi.org/10.1364/OE.22.012634>
97. Fannjiang, A., Chen, P.: Blind ptychography: uniqueness and ambiguities. *Inverse Probl.* **36**(4), 045005 (2020). <https://doi.org/10.1088/1361-6420/ab6504>
98. Iwen, M., Perlmutter, M., Sissouno, N., Viswanathan, A.: Phase retrieval for $L^2([-\pi, \pi])$ via the provably accurate and noise robust numerical inversion of spectrogram measurements. *J Fourier Anal Appl.* **29**(8) (2023). <https://doi.org/10.1007/s00041-022-09988-6>
99. Perlmutter, M., Sissouno, N., Viswanathan, A., Iwen, M.: A provably accurate algorithm for recovering compactly supported smooth functions from spectrogram measurements. In: 2020 28th European Signal Processing Conference (EUSIPCO), pp. 970–974. (2021). <https://doi.org/10.23919/Eusipco47968.2020.9287698>
100. Grohs, P., Liehr, L.: Injectivity of Gabor phase retrieval from lattice measurements. *Appl. Comput. Harmon. Anal.* **62**, 173–193 (2023). <https://doi.org/10.1016/j.acha.2022.09.001>
101. Wellershoff, M.: Injectivity of sampled Gabor phase retrieval in spaces with general integrability conditions. *J. Math. Anal. Appl.* **530**(2), 127692 (2024). <https://doi.org/10.1016/j.jmaa.2023.127692>
102. Wellershoff, M.: Sampling at twice the Nyquist rate in two frequency bins guarantees uniqueness in Gabor phase retrieval. *J Fourier Anal Appl.* **29**(7) (2023). <https://doi.org/10.1007/s00041-022-09990-y>
103. Grohs, P., Liehr, L., Shafkulovska, I.: From completeness of discrete translates to phaseless sampling of the short-time Fourier transform. (2022). <https://doi.org/10.48550/arXiv.2211.05687>
104. Alaifari, R., Wellershoff, M.: Phase retrieval from sampled Gabor transform magnitudes: counterexamples. *J Fourier Anal Appl.* **28**(9) (2022). <https://doi.org/10.1007/s00041-021-09901-7>
105. Grohs, P., Liehr, L.: On foundational discretization barriers in STFT phase retrieval. *J Fourier Anal Appl.* **28**(39) (2022). <https://doi.org/10.1007/s00041-022-09935-5>
106. Alaifari, R., Bartolucci, F., Steinerberger, S., Wellershoff, M.: On the connection between uniqueness from samples and stability in Gabor phase retrieval. *Sampling Theory, Signal Process., And Data Anal.* **22**(6) (2024). <https://doi.org/10.1007/s43670-023-00079-1>
107. Grohs, P., Liehr, L., Rathmair, M.: Multi-window STFT phase retrieval: lattice uniqueness. *J. Educ. Chang. Funct. Anal.* **288**(3), 110733 (2025).
108. Grohs, P., Liehr, L.: Phaseless sampling on square-root lattices. In: *Found. Comput. Math.* 1–24 (2024).
109. Grohs, P., Liehr, L., Rathmair, M.: Phase retrieval in Fock space and perturbation of Liouville sets. *Revista Matemática Iberoam.* **41**(3), 969–1008 (2025). <https://doi.org/10.4171/RMI/1517>

Authors and Affiliations

Marc Allain¹ · Selin Aslan² · Wim Coene³ · Sjoerd Dirksen⁴ · Jonathan Dong⁵ · Julien Flamant⁶ · Mark Iwen⁷ · Felix Krahmer⁸ · Tristan van Leeuwen^{4,9} · Oleh Melnyk^{10,11} · Andreas Menzel¹² · Allard P. Mosk¹³ · Viktor Nikitin¹⁴ · Palina Salanevich⁴  · Gerlind Plonka¹⁵ · Matthias Wellershoff¹⁶

✉ Palina Salanevich
p.salanevich@uu.nl

Marc Allain
marc.allain@fresnel.fr

Selin Aslan
saslan@ku.edu.tr

Wim Coene
w.m.j.m.coene-1@tudelft.nl

Sjoerd Dirksen
s.dirksen@uu.nl

Jonathan Dong
jonathan.dong@epfl.ch

Julien Flamant
julien.flamant@cnrs.fr

Mark Iwen
iwenmark@msu.edu

Felix Krahmer
felix.krahmer@tum.de

Tristan van Leeuwen
t.van.leeuwen@cw.nl

Oleh Melnyk
oleh.melnik@tu-berlin.de

Andreas Menzel
andreas.menzel@psi.ch

Allard P. Mosk
a.p.mosk@uu.nl

Viktor Nikitin
vnikitin@anl.gov

Gerlind Plonka
plonka@math.uni-goettingen.de

Matthias Wellershoff
wellersm@umd.edu

¹ Aix Marseille Université, CNRS, Centrale Marseille, Institut Fresnel, Marseille, France

² Department of Mathematics, College of Science, Koc University, Sariyer Rumeli Feneri Yolu, Istanbul, Sariyer 34450, Turkey

³ Applied Sciences, Imaging Physics, Delft University of Technology, Mekelweg 5, Delft 2628 CD, The Netherlands

- ⁴ Mathematical Institute, Utrecht University, Budapestlaan 6, Utrecht 3584 CD, The Netherlands
- ⁵ Ecole Polytechnique Fédérale de Lausanne, Biomedical Imaging Group, Lausanne 1015, Switzerland
- ⁶ Université de Lorraine, CNRS, CRAN, Nancy F-54000, France
- ⁷ Department of Mathematics, and Department of Computational Mathematics, Science and Engineering (CMSE), Michigan State University, East Lansing, Michigan, USA
- ⁸ Department of Mathematics, Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany
- ⁹ Centrum Wiskunde & Informatica, Computational Imaging, Science Park 123, Amsterdam, Netherlands
- ¹⁰ Institute of Mathematics, Technische Universität Berlin, Straße des 17. Juni 136, D-10623 Berlin, Germany
- ¹¹ Ludwig-Maximilians-Universität München, Mathematical Foundation of Artificial Intelligence, Akademiestr. 7, 80799 Munich, Germany
- ¹² Paul Scherrer Institute, Center for Photon Science, Villigen, PSI 5232, Switzerland
- ¹³ Debye Institute for Nanomaterials Science, Utrecht University, Princetonplein 1, Utrecht, Netherlands
- ¹⁴ Advanced Photon Source, Argonne National Laboratory, 9700 S Cass Ave, Lemont, IL 60439, USA
- ¹⁵ Institute for Numerical and Applied Mathematics, University of Goettingen, Lotzestrasse 16-18, 37083 Goettingen, Germany
- ¹⁶ Department of Mathematics, University of Maryland, 4176 Campus Drive, College Park, MD 20742, USA