



Phase retrieval via gain-based photonic XY-Hamiltonian optimization



Richard Zhipeng Wang¹, Guangyao Li¹, Silvia Gentilini², Davide Pierangeli²,
Marcello Calvanese Strinati³, Claudio Conti² & Natalia G. Berloff¹ ✉

Phase-retrieval from coded diffraction patterns (CDP) is important to X-ray crystallography, diffraction tomography and astronomical imaging, yet remains a hard, non-convex inverse problem. We show that CDP recovery can be reformulated exactly as the minimization of a continuous-variable XY Hamiltonian and solved by gain-based photonic networks. The coupled-mode equations we exploit are the natural mean-field dynamics of exciton-polariton condensate lattices, coupled-laser arrays and driven photon Bose-Einstein condensates, while other hardware such as the spatial photonic Ising machine can implement the same update rule through high-speed digital feedback, preserving full optical parallelism. Numerical experiments on images, two- and three-dimensional vortices and unstructured complex data demonstrate that the gain-based solver consistently outperforms the state-of-the-art Relaxed-Reflect-Reflect (RRR) algorithm in the medium-noise regime (signal-to-noise ratios 10-40 dB) and retains this advantage as problem size scales. Because the physical platform performs the continuous optimisation, our approach promises fast, energy-efficient phase retrieval on readily available photonic hardware.

Recently, there has been a rising interest in using physics-inspired, physics-based computing systems for solving hard optimization problems, including many that are NP-hard¹⁻⁶. An important example is the minimization of the XY Hamiltonian with sign-varying couplings between spins, where each spin is allowed to take on a continuous value in the vector $\mathbf{s}_i = (\cos \theta_i, \sin \theta_i)$ or complex form $s_i = e^{i\theta_i}$, where $\theta_i \in [0, 2\pi)$. Such continuous-spin systems appear naturally in gain-dissipative photonic lattices, exciton-polariton networks, laser arrays, and related platforms, where each oscillator is represented by a complex variable whose amplitude and phase evolve in time⁷⁻¹².

The XY minimization problem involves the quadratic Hamiltonian

$$H_{XY}(\{s_i\}) = -\frac{1}{2} \sum_{ij} J_{ij} s_i s_j^* + \text{c.c.}, \quad (1)$$

where J_{ij} represents the pairwise couplings. The task is to find the spin configuration $\{s_i\}$ that minimizes H_{XY} . This is a continuous quadratic optimization (QCO) problem, with applications ranging from clustering¹⁰ to portfolio optimization¹³.

One key motivation for studying XY-type physical networks is their ability to perform highly parallel, analog searches for low-energy configurations, thereby offering an alternative to purely digital algorithms.

Recently, there have also been efforts to force an XY-based system into effectively binary (Ising-like) states by introducing a large penalty term. Specifically, one adds

$$H_P(\{s_i\}) = H_{XY} + P \sum_{i=1}^N [s_i^{*2} + \text{c.c.}], \quad (2)$$

where $P > 0$ is chosen so as to penalize any non-zero real part of s_i . If P is sufficiently large, spins tend to align at phases $\pm \pi/2$, effectively reducing the continuous-spin XY problem to the binary Ising problem.

The continuous nature of the XY Hamiltonian also makes it relevant to other QCO tasks. One such problem of practical importance is the phase-retrieval problem. This is typically stated as follows: given a real measurement vector $\mathbf{b} \in \mathbb{R}^M$ and a complex matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$, one seeks to find a complex vector $\mathbf{x} \in \mathbb{C}^N$ satisfying

$$|\mathbf{A}\mathbf{x}| = \mathbf{b}, \quad (3)$$

where $|\cdot|$ denotes element-wise amplitude. Its practical significance arises because this problem is often encountered in applications such as X-ray crystallography^{14,15}, astronomical imaging¹⁶, and diffraction imaging^{17,18}.

In many of these applications, the complex vector \mathbf{x} represents the complete information about the sample and is referred to as the *sample*

¹Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK. ²Department of Physics, Sapienza University of Rome, Rome, Italy. ³Research Center Enrico Fermi, Rome, Italy. ✉e-mail: N.G.Berloff@damtp.cam.ac.uk




Fig. 1 | A schematic diagram for the CDP experiment framework. A light source is diffracted by the sample under investigation, and the diffracted complex-valued signal is split into L identical beams, each of which is directed towards a phase filter that modifies the phase of the incident signal at each spatial location. The phase-modified beams are then directed through a lens system, and their intensities are finally captured.

vector. The matrix \mathbf{A} describes the action of the optical system, often well-approximated by a Fourier transform. While the intensity of the resulting electromagnetic wave can be measured by standard detectors (e.g. charge-coupled devices) to yield the real-valued amplitude \mathbf{b} , the phase component is typically lost in the measurement process. Recovering \mathbf{x} from \mathbf{b} and \mathbf{A} constitutes the *phase-retrieval problem*, which is NP-complete¹⁹, underscoring its computational difficulty.

Moreover, without further constraints, the phase-retrieval problem is frequently *ill-posed* because multiple distinct sample vectors \mathbf{x} can give rise to the same measured amplitude \mathbf{b} . For instance, if \mathbf{A} is a square matrix representing a discrete Fourier transform, any arbitrary phase profile can be applied to \mathbf{b} before performing an inverse Fourier transform, producing infinitely many valid solutions \mathbf{x} . In such scenarios, even a theoretically exact algorithm may yield a recovered vector $\tilde{\mathbf{x}}$ that deviates from the original \mathbf{x} ²⁰. This means that any solution of the form $\mathbf{x} = \mathbf{D}^{(i)}\mathbf{A}^{-1}\mathbf{b}$, where $\mathbf{D}^{(i)}$ is a diagonal matrix with arbitrary diagonal elements of the form $e^{i\theta_i}$, $\theta_i \in [0, 2\pi)$ and \mathbf{A}^{-1} is the inverse discrete Fourier transform matrix, will satisfy the requirement given by Eq. (3). Efforts to mitigate ill-posedness often rely on additional assumptions about the sample vector \mathbf{x} . For instance, some approaches impose a sparsity constraint by specifying the number of non-zero entries in \mathbf{x} ²¹. Others adopt even stricter conditions, stipulating the exact support of \mathbf{x} ; that is, which elements are nonzero^{22–25}. However, none of these methods can guarantee that the resulting phase-retrieval problem is well-posed (i.e. admits a unique solution).

Amongst the works that investigated phase-retrieval problems with support constraint, Tradonsky et al.²⁵ proposed a non-digital solver, namely a cavity laser, to solve small-scale phase-retrieval problems. Their use of an analog physical device to solve phase-retrieval problems is similar to the spirit of this paper. However, the proposed laser solver required an aperture that reproduced the known support of the 2D sample object, which limited the scalability and applicability of the solver to larger-sized and non-trivially shaped sample objects or 3D sample objects.

Candes et al.²⁶ proposed a formulation called the “coded diffraction pattern” (CDP) experiment, which provides some solution uniqueness guarantee. Under this formulation, a sample vector \mathbf{x} is first subjected to the action of L random phase filters, each defined by a diagonal matrix $\mathbf{D}^{(i)}$, $i = 1, \dots, L$. Each filter is defined by diagonal entries that are complex numbers with unit magnitude and randomly distributed phases. Thus, the filter modifies only the phase of the incoming signal while leaving its

amplitude unchanged. This produces L filtered sample vectors, which are Fourier transformed separately. The results are then concatenated, and amplitudes are taken to form the observation vector \mathbf{b} . Hence, in this formulation, \mathbf{A} is an M by N complex matrix, where $M = L \times N$, and represents the combined action of the random filters and the Fourier transforms. Figure 1 provides a schematic diagram showing the experimental setup of this arrangement. This formulation is theoretically appealing because if the distribution of the random diagonal matrix elements satisfies certain conditions, then with sufficiently large L , the solution of the phase-retrieval problem is guaranteed to be unique²⁶. It was also demonstrated numerically that even with filters whose diagonal matrix elements did not follow random distributions that satisfy the conditions laid out in ref. 26, the problem could still become well-posed with increased L . Empirically, the process of applying multiple filters corresponds to the practice of oversampling the target and is reasonably achievable in many imaging applications^{26–28}.

The performance of conventional phase-retrieval algorithms in this setting has been systematically examined in ref. 27. More recent proposals include semidefinite programming^{29,30}, matrix completion³¹, and an alternating-projection method called *Relaxed-Reflect-Reflect* (RRR)³². Although RRR lacks a comprehensive theoretical foundation, subsequent investigations³³ reported robust empirical performance. In particular, ref. 34 demonstrated that RRR outperforms both the “Phase Lift” approach of Candès et al.²⁹ and Wirtinger flow methods³⁵ in a variety of benchmark tests. Another approach, making use of unsupervised learning called “Deep-PhaseCut” was also proposed³⁶, and was shown to also outperform “Phase Lift” and “Phase Cut” in terms of computational speed. This approach demonstrated the possibility of combining machine learning techniques with traditional phase-retrieval algorithms for better computational efficiency.

In this work, we demonstrate that the CDP variant of phase retrieval can be reformulated as an XY Hamiltonian minimization problem, amenable to physical solvers that emulate the same Hamiltonian, for instance, through gain-based oscillator networks^{3,37}. Under moderate noise levels in the measured amplitudes, we show numerically that a gain-based system can outperform the state-of-the-art Relaxed-Reflect-Reflect (RRR) algorithm and reliably reconstruct complex-valued experimental data with high accuracy.

Methods

Mapping phase retrieval into XY problem

The work of Waldspurger et al.³⁰ first recast phase retrieval as a non-convex quadratic program, facilitating its solution via a suitable relaxation method. Following a similar line of reasoning, we show how to reformulate phase retrieval in the CDP experiment setting as an XY Hamiltonian minimization problem.

We start by defining the unknown phase of the observation as \mathbf{p} , where all elements of this vector are complex and have unit amplitude, i.e., $p_i \in \{e^{i\theta_i} | \theta_i \in [0, 2\pi)\}$. By this definition, we have

$$\sum_{j=1}^N A_{ij}x_j = b_i p_i \text{ for all } i = 1, \dots, M. \tag{4}$$

Suppose that the complete observation information including the phase information \mathbf{p} and the amplitude information \mathbf{b} is known, then the process of trying to find \mathbf{x} that satisfy the constraint given by Eq. (4) is equivalent to the minimization of the cost function \mathcal{E} given by

$$\mathcal{E}(\{x_j\}) = \sum_{i=1}^M \left| \sum_{j=1}^N A_{ij}x_j - b_i p_i \right|^2, \tag{5}$$

where the minimization is over all possible sets of $\{x_j\}$ for $j \in \{1, \dots, N\}$. If the observations \mathbf{b} and \mathbf{p} are exact, the minimum value of this expression should be 0, but if uncertainty exists in either of them, then this expression is a least-squares problem over \mathbf{x} . The solution to the least-square problem is

given by the Moore–Penrose inverse, also known as the pseudoinverse of \mathbf{A} , denoted as \mathbf{A}^\dagger :

$$x_i = \sum_j^M A_{ij}^\dagger b_j p_j \tag{6}$$

Hence, for any given set of observation phase $\{p_i\}$, the minimal cost function \mathcal{E} is given by

$$\mathcal{E}(\{p_i\}) = \sum_i^M \left| \sum_j^N \sum_k^M A_{ij} A_{jk}^\dagger b_k p_k - b_i p_i \right|^2 \tag{7}$$

This means that to solve the original phase-retrieval problem, one needs to find a set of $\{p_i\}$ that minimizes this cost function \mathcal{E} , which is essentially a QCO problem over a set of variables $\{p_i\}$ that all have unit amplitude and unconstrained phases, so this problem can be mapped into an XY Hamiltonian minimization problem.

To put it into an explicit XY Hamiltonian form, we rearrange Eq. (7) as follows:

$$\begin{aligned} \mathcal{E}(\{p_i\}) &= \sum_i^M \left| \sum_{j,k}^{N,M} A_{ij} A_{jk}^\dagger b_k p_k - \sum_k^M \delta_{ik} b_k p_k \right|^2 \\ &= \sum_i^M \left| \sum_k^M \left(\sum_j^N A_{ij} A_{jk}^\dagger - \delta_{ik} \right) b_k p_k \right|^2 \\ &= \sum_i^M \left| \sum_k^M G_{ik} b_k p_k \right|^2 \\ &= \sum_i^M \sum_{j,k}^M G_{ij} G_{ki}^* b_j p_j b_k p_k^* \\ &= - \sum_{j,k}^M \tilde{J}_{jk} p_j p_k^*, \end{aligned} \tag{8}$$

where from the second line to the third line for ease of notation we defined $G_{ik} = \sum_j^N A_{ij} A_{jk}^\dagger - \delta_{ik}$, and from the fourth line to the last we identified the coupling matrix elements of the equivalent XY Hamiltonian to be $\tilde{J}_{jk} = - \sum_i^M G_{ij} G_{ki}^* b_j b_k$.

In principle, Eq. (8) is already in the form of an XY Hamiltonian we first introduced in Eq. (1). Noting that $\mathbf{A}\mathbf{A}^\dagger$ is Hermitian and $\mathbf{A}\mathbf{A}^\dagger\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}\mathbf{A}^\dagger$ by property of the pseudoinverse, one can simplify the expression for \tilde{J}_{ij} to

$$\tilde{J}_{ij} = \left(\sum_k^N A_{ik} A_{kj}^\dagger - \delta_{ij} \right) b_i b_j, \tag{9}$$

which leads to an equivalent XY Hamiltonian:

$$\begin{aligned} \mathcal{E}(\{p_i\}) &= - \sum_{ij}^M \left(b_i b_j \sum_k^N A_{ik} A_{kj}^\dagger \right) p_i p_j^* + \sum_{ij}^M \delta_{ij} p_i p_j^* \\ &= H_{XY}(\{p_i\}) + M, \end{aligned} \tag{10}$$

where M is as before the dimension of observation vector \mathbf{b} . Hence, minimizing \mathcal{E} is equivalent to minimizing the simpler XY Hamiltonian H , whose coupling matrix elements are given by

$$J_{ij} = \sum_k^N A_{ik} A_{kj}^\dagger b_i b_j. \tag{11}$$

This coupling matrix was then used as input to the simulated gain-based system, with outcomes presented in the “Results” section.

We note that the Moore–Penrose pseudoinverse \mathbf{A}^\dagger may become ill-conditioned when $\text{rank}(\mathbf{A}) < N$. However, as we will illustrate in the later subsection “Generation of CDP phase-retrieval problems”, the way that matrix \mathbf{A} is built in the CDP phase-retrieval problem via Eq. (15) means that \mathbf{A} is always well-conditioned and its pseudoinverse is exceptionally well-behaved. For completeness, we also tested a Tikhonov-regularized alternative,

$$(c\mathbf{I} + \mathbf{A}\mathbf{A}^*)^{-1} \mathbf{A}^*,$$

with $c > 0$, which can be used to stabilize the inversion. In practice, choosing $c \sim 10^{-3} \|\mathbf{A}\|_2^2$ leaves all numerical results within the error bars reported below.

Simulated gain-based system

To minimize the XY Hamiltonian H_{XY} for a given phase-retrieval problem, we simulate an oscillator network following the gain-based dynamics given by

$$\frac{d\psi_i}{dt} = (\gamma_i - |\psi_i|^2)\psi_i + \sum_j^M J_{ij}\psi_j \tag{12}$$

$$\frac{d\gamma_i}{dt} = \epsilon(1 - |\psi_i|^2), \tag{13}$$

where $\gamma_i \in \mathbb{R}$ is the effective injection rate of oscillator i (gain minus losses), $\psi_i \in \mathbb{C}$ characterize each oscillator, and J_{ij} specifies the coupling strength between oscillators. J_{ij} is calculated from the given phase-retrieval problem according to our previous discussion. ϵ is an externally controlled positive constant, which measures the responsiveness of the gain of each oscillator to the amplitude variations of each oscillator. Eqs. (12) and (13) faithfully reproduce the gain-dissipative evolution observed in networks of exciton–polaritons hosted in semiconductor micro-cavities³⁸, in coupled-laser arrays¹², in driven photonic oscillator lattices³⁹, and spin wave Ising machines⁴⁰. Hardware platforms that do not possess intrinsic gain control, most notably the spatial photonic Ising machine (SPIM)^{41,42}, can nevertheless implement the same update rule by applying digital feedback to the spatial-light modulator after each optical pass, thereby emulating the gain–loss loop in silico while retaining full optical parallelism. This kind of dynamics was formulated and studied in refs. 3,8,43, showing that the dynamics of this system lead to stationary states close to or at the global minimum of the XY Hamiltonian specified by the coupling matrix \mathbf{J} with high probability. Equation (12) encapsulates the main dynamics of ψ_i with the interplay of the effective gain, non-linear loss, and the coupling terms. Eq. (13) provides a feedback mechanism that pushes the amplitudes of all oscillators towards 1. It was previously reported that this feedback mechanism is crucial for the gain-dissipative dynamics to produce good solutions close to the true global minimum of the XY Hamiltonian⁸.

In principle, the gain-based optimizer works as follows. At the start of the dynamical evolution, the oscillator network has a set of highly negative effective gains γ_i , so the system has a stable fixed point at $\psi_i = 0$. Due to the gain dynamics given by Eq. (13), the gains γ_i increase over time and eventually cross a critical value at which supercritical Hopf bifurcation occurs^{3,43}. The $\psi_i = 0$ fixed point becomes unstable, and oscillators spontaneously increase to some non-zero amplitudes and start to have well-defined phases. Over time, all amplitudes $|\psi_i|$ approach 1, while the phases of oscillators also approach their stationary values, which will be read out as the solution to the XY problem.

In our simulations, we first initialized the amplitude of ψ_i to some random small but non-vanishing values uniformly distributed in the range (0, 0.1) and initialized their phase uniformly randomly in the range $[0, 2\pi)$. Initial gains γ_i were initialized uniformly randomly in such a range of values,

so that initial stages of the evolution are below the threshold, which means that oscillator amplitudes will tend to decay to 0 if gains are held constant at this level. In practice, this was achieved by selecting γ_i values uniformly randomly from the range $[-\lambda_{\max}(\mathbf{J})-2, -\lambda_{\max}(\mathbf{J})-1]$ where $\lambda_{\max}(\mathbf{J})$ denotes the maximum eigenvalue of coupling matrix \mathbf{J} . Some preliminary simulations showed that the feedback coefficient ϵ mainly affected the rate of convergence to the final solution, but did not impact the quality of the final solution significantly, as long as the value of ϵ was not so large that the system would diverge. Hence, to show that the proposed gain-based dynamical system can solve a wide range of phase-retrieval problems without considerable hyper-parameter tuning beforehand on a per-problem basis, we simply used the value of $\epsilon = 1$ throughout all simulations carried out in this paper. The system was then evolved with an adaptive timestep Runge–Kutta 4th-order method using the “Scipy” Python package until it reached a stationary state or a maximum simulation time was reached, and the phases of oscillators were used as the spin configurations \mathbf{p} for the XY problem. This can then be substituted into Eq. (6) to find the solution \mathbf{x} to the original phase-retrieval problem.

Generation of CDP phase-retrieval problems

For a given sample vector $\mathbf{x} \in \mathbb{C}^N$, we had to generate an observation vector $\mathbf{b} \in \mathbb{R}^M$ under the CDP experiment framework to serve as the input to our XY minimizer. A set of L random filters was first generated. Each filter could be represented as a diagonal matrix $\mathbf{D}^{(i)}$ where $i = 1, \dots, L$, and each diagonal element $D_{jj}^{(i)}$ was uniformly randomly selected from $\{1, i, -1, -i\}$, corresponding to a phase shift of $0, \pi/2, \pi, 3\pi/2$, respectively.

We represent the action of the Fourier transform on the sample vector by using the discrete Fourier transform (DFT) matrix defined as:

$$\mathbf{F}_N = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & \omega_N & \omega_N^2 & \dots & \omega_N^{N-1} \\ 1 & \omega_N^2 & \omega_N^4 & \dots & \omega_N^{2(N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N^{N-1} & \omega_N^{2(N-1)} & \dots & \omega_N^{(N-1)^2} \end{pmatrix},$$

where $\omega_N = e^{-2\pi i/N}$ is the N th root of unity.

One set of observation from one phase filter can be obtained by multiplying \mathbf{F}_N and $\mathbf{D}^{(i)}$ with \mathbf{x} , so part of the complex-valued matrix \mathbf{A} can be defined as:

$$\mathbf{A}^{(i)} = \mathbf{F}_N \mathbf{D}^{(i)} \quad (14)$$

Matrix $\mathbf{A}^{(i)} \in \mathbb{C}^{N \times N}$ can be stacked columnwise to obtain the complete matrix \mathbf{A} as illustrated below:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}^{(1)} \\ \mathbf{A}^{(2)} \\ \vdots \\ \mathbf{A}^{(L)} \end{pmatrix}. \quad (15)$$

This matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$, together with \mathbf{x} , can be substituted into Eq. (3) to produce the observation vector \mathbf{b} . The observation vector \mathbf{b} and the generated matrix \mathbf{A} define this CDP phase-retrieval problem, and they are supplied to our gain-based optimizer Eqs. (12) and (13), and the phases θ_i are found as the stationary states. The sample vector \mathbf{x} is then reconstructed from Eq. (6).

We note that because each block $\mathbf{A}^{(i)}$ is unitary, the conjugate transpose of \mathbf{A} , denoted as \mathbf{A}^* , has the following property:

$$\mathbf{A}^* \mathbf{A} = \sum_{i=1}^L (\mathbf{A}^{(i)})^* \mathbf{A}^{(i)} = L \mathbb{I}_N \quad (16)$$

where \mathbb{I}_N is the identity matrix of dimension N . Thus, the pseudoinverse of \mathbf{A} simplifies to $\mathbf{A}^\dagger = \frac{1}{L} \mathbf{A}^*$. Hence, \mathbf{A} has full column rank and condition number 1, so its pseudoinverse is always well-behaved.

The block stacking of all Fourier operators \mathbf{F} and diagonal masks $\mathbf{D}^{(i)}$ is mathematically identical to the decomposition of a fully connected interaction matrix into multiple Mattis subproblems recently realized in a fully programmable SPIM via focal-plane division, where the energies of all sub-Hamiltonians are computed in parallel on distinct camera regions⁴². This analogy suggests that an experimental CDP phase-retrieval setup could exploit the same optical parallelism, processing the L masks in a single physical shot and obtaining an L -fold reduction in acquisition time.

Additionally, by substituting $\mathbf{A}^\dagger = \frac{1}{L} \mathbf{A}^*$ into Eq. (11), the coupling matrix \mathbf{J} that needs to be implemented in hardware can be expressed as

$$\mathbf{J} = \frac{1}{L} \text{diag}(\mathbf{b}) \mathbf{A} \mathbf{A}^* \text{diag}(\mathbf{b}) \quad (17)$$

where $\text{diag}(\mathbf{b})$ denotes a diagonal matrix whose elements are elements of vector \mathbf{b} . This implies that although matrix \mathbf{J} has dimension $M \times M$ and is dense, its rank is at most N . This low-rank structure can be exploited in photonic hardware such as SPIM, which natively implements rank-one or low-rank interaction matrices via optical propagation¹³.

In this study, we also considered the case where \mathbf{b} is noisy, which is to be expected in realistic experimental data. In this case, a noisy observation vector $\tilde{\mathbf{b}}$ is produced by adding a normally distributed random noise to each element of the noiseless observation vector \mathbf{b} , i.e. $\tilde{b}_i = b_i + \xi_i$ with $\xi_i \sim \mathcal{N}(0, \sigma^2)$. The variance of noise σ was used to control the magnitude of noise in the observational data so that we could investigate its impact on the performance of the gain-based system in solving the phase-retrieval problem. To quantify the amount of noise in the given noisy observation vector, we define the signal-to-noise ratio (SNR) as follows:

$$\text{SNR} = 10 \log_{10} \frac{\|\mathbf{b}\|_2}{\|\tilde{\mathbf{b}} - \mathbf{b}\|_2}, \quad (18)$$

which is conventionally measured on a logarithmic scale and quoted in units of decibels. In this expression, $\|\cdot\|_2$ denotes the vector 2-norm.

Additionally, we note that beyond Gaussian noise, Poisson noise and systematic errors are also often present in realistic diffraction experiments. To investigate the robustness of our proposed phase-retrieval method against Poisson noise, we also generate noisy observation vector $\tilde{\mathbf{b}}$ whose elements are given by $b_i = B_i/n$, where random variable $B_i \sim P(\lambda = nb_i)$ (i.e. it follows Poisson distribution with λ given by nb_i). The coefficient n controls the strength of the Poisson noise and thus allows us to generate a noisy signal with the required SNR. To understand the impact of systematic error, we generate a noisy observation vector $\tilde{\mathbf{b}}$ with a constant shift from the true value, which means its elements $b_i = b_i + n$, where $n > 0$ is the same for all i and controls the SNR of the resultant observation vector.

Performance evaluation


To measure the quality of the calculated solution, we used metrics that were also used previously in refs. 30,34. The most direct metric to measure the success of the phase-retrieval algorithm is the Euclidean distance between the observation vector calculated from the recovered sample vector $\tilde{\mathbf{x}}$ and the given observation vector \mathbf{b} , normalized over the vector norm of the observation vector \mathbf{b} .

For clearer visual representation, it is more convenient to express this quantity on a logarithmic scale, similar to SNR. Therefore, we define the relative observation error (ROE) in decibels as

$$\text{ROE} = 10 \log_{10} \frac{\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_2}{\|\mathbf{b}\|_2}, \quad (19)$$

where $|\cdot|$ denotes taking the amplitude element-wise, and $\tilde{\mathbf{x}}$ is the sample vector calculated by the phase-retrieval algorithm. In experiments, the true

Fig. 2 | Comparison of phase-retrieval problems with different number of phase masks. **a** Time evolution of errors for the phase-retrieval problem under the CDP experiment framework with $L = 2$ phase masks. The inset gives the final reconstructed image produced by the phase-retrieval algorithm. **b** Time evolution of errors for the phase-retrieval problems with 5 phase masks. In both cases, the problems were solved with the gain-dissipative system with a random initial condition. The original image, which is visually identical to the image shown in the inset of **(b)**, is from the labeled face in the wild (LFW) dataset⁵¹. See table “Fig. 2” in supplementary data.



sample vector \mathbf{x} is typically unavailable, so the ROE would be an appropriate way to evaluate the quality of the recovered solution. However, in generated datasets where the true sample vector \mathbf{x} is known a priori, we can instead measure the Euclidean distance between the recovered sample vector $\tilde{\mathbf{x}}$ and the true sample vector \mathbf{x} . Note that the recovered sample vector $\tilde{\mathbf{x}}$ may have a global phase shift relative to \mathbf{x} while reproducing the exact same observation vector \mathbf{b} . Hence, the error metric should be defined as the minimum Euclidean distance between $e^{i\theta}\tilde{\mathbf{x}}$ and \mathbf{x} for all $\theta \in [0, 2\pi)$. This quantity is also expressed in decibels and is referred to as the relative sample error (RSE), which is defined as

$$\text{RSE} = 10\log_{10} \left(\min_{\theta} \frac{\| e^{i\theta}\tilde{\mathbf{x}} - \mathbf{x} \|_2}{\| \mathbf{x} \|_2} \right). \quad (20)$$

In a phase-retrieval problem, this error is usually the most important quantity because it measures how close the recovered solution is to the original sample.

The two error metrics are closely related but not equivalent, and both are required to give a complete picture of the well-posedness of the phase-retrieval problem itself and the performance of the phase-retrieval algorithm used. For example, if a solution calculated by a phase-retrieval algorithm yields a small ROE value but a large RSE value, this indicates that the algorithm is performing well, but the problem itself is poorly defined because it has more than one degenerate ground state (i.e., more than one \mathbf{x} can all produce the same observation vector \mathbf{b}). This is because the phase-retrieval algorithm only has access to \mathbf{b} and \mathbf{A} , so any solution $\tilde{\mathbf{x}}$ that can minimize $\| \mathbf{A}\tilde{\mathbf{x}} - \mathbf{b} \|_2$ is equally good to the algorithm, even if it might be far from the true \mathbf{x} from which the problem was first constructed. This situation is illustrated in Fig. 2a, where the simulated gain-based system tries to recover an image from an observation vector produced by 2-phase filters. While ROE keeps decreasing, RSE has remained largely flat, and the system failed to recover the original image. As far as the phase-retrieval method is concerned, it is performing well because it is able to find a vector $\tilde{\mathbf{x}}$ that produces an observation vector $\tilde{\mathbf{b}}$ that is very close to the known observation vector \mathbf{b} . This means that this phase-retrieval problem under the CDP experiment framework with only 2 phase masks is ill-defined, because it has degenerate ground states.

When both RSE and ROE are small, it suggests that the phase-retrieval problem is well defined and the solution found is close to the true solution. This case is shown in Fig. 2b, where the same gain-dissipative system recovers the same image from an observation vector produced by 5 phase filters. RSE and ROE decrease in tandem, indicating that the algorithm is approaching the planted ground state. As far as we know, there is no explicit formula which can give a L value for a given class of sample vector \mathbf{x} such that it is guaranteed that the generated phase-retrieval problem will always have a unique solution³⁶. It is down to empirical trial-and-error to determine what is the minimum value of L such that the generated phase-retrieval problem has a unique solution. Hence, when studying phase-retrieval

algorithms, it is crucial to consider both the ROE and RSE to determine whether the phase-retrieval method itself or the problem at hand is responsible for the failure to recover the original signal. In all of the numerical experiments carried out in this paper, we incremented L until the difference between ROE and RSE remained small throughout the entire simulation duration, and the reconstruction result became visually faithful. We then used the smallest such L in all subsequent tests, but we also always checked both ROE and RSE for each class of problems we investigated to ensure that all random problem instances we generated resulted in a well-posed CDP phase retrieval problem.

Comparison with existing algorithms

To compare the gain-dissipative system with established phase-retrieval methods, we focus on the RRR algorithm for phase retrieval³². Originally designed for the sparse variant of phase retrieval, RRR belongs to the family of alternating projection methods, similar to classic techniques such as the Gerchberg-Saxton (GS) method⁴⁴, Fienup’s hybrid input-output (HIO) scheme²², and the shrinkwrap algorithm⁴⁵. These algorithms operate in an iterative fashion, applying two distinct projection operators in sequence at each iteration. Among them, RRR has shown particularly strong empirical performance, surpassing multiple modern phase-retrieval methods³⁴.

The RRR algorithm starts with a guessed observation vector \mathbf{b}_0 , and employs two projections P_1 and P_2 . When given a (generally complex) estimated observation vector \mathbf{b}_n , P_1 keeps only the S largest elements in the vector and sets other to 0, where S is the given sparsity constraint in the observation data; For P_2 , when given $\mathbf{b}_n \in \mathbb{C}^M$, it keeps the phase of each element but overwrites their amplitudes with the known correct amplitudes in $\mathbf{b} \in \mathbb{R}^M$. Overall, in each iteration, the algorithm applies P_1 and P_2 as follows:

$$\mathbf{b}_{n+1} = \mathbf{b}_n + \beta [P_2(2P_1(\mathbf{b}_n) - \mathbf{b}_n) - P_1(\mathbf{b}_n)], \quad (21)$$

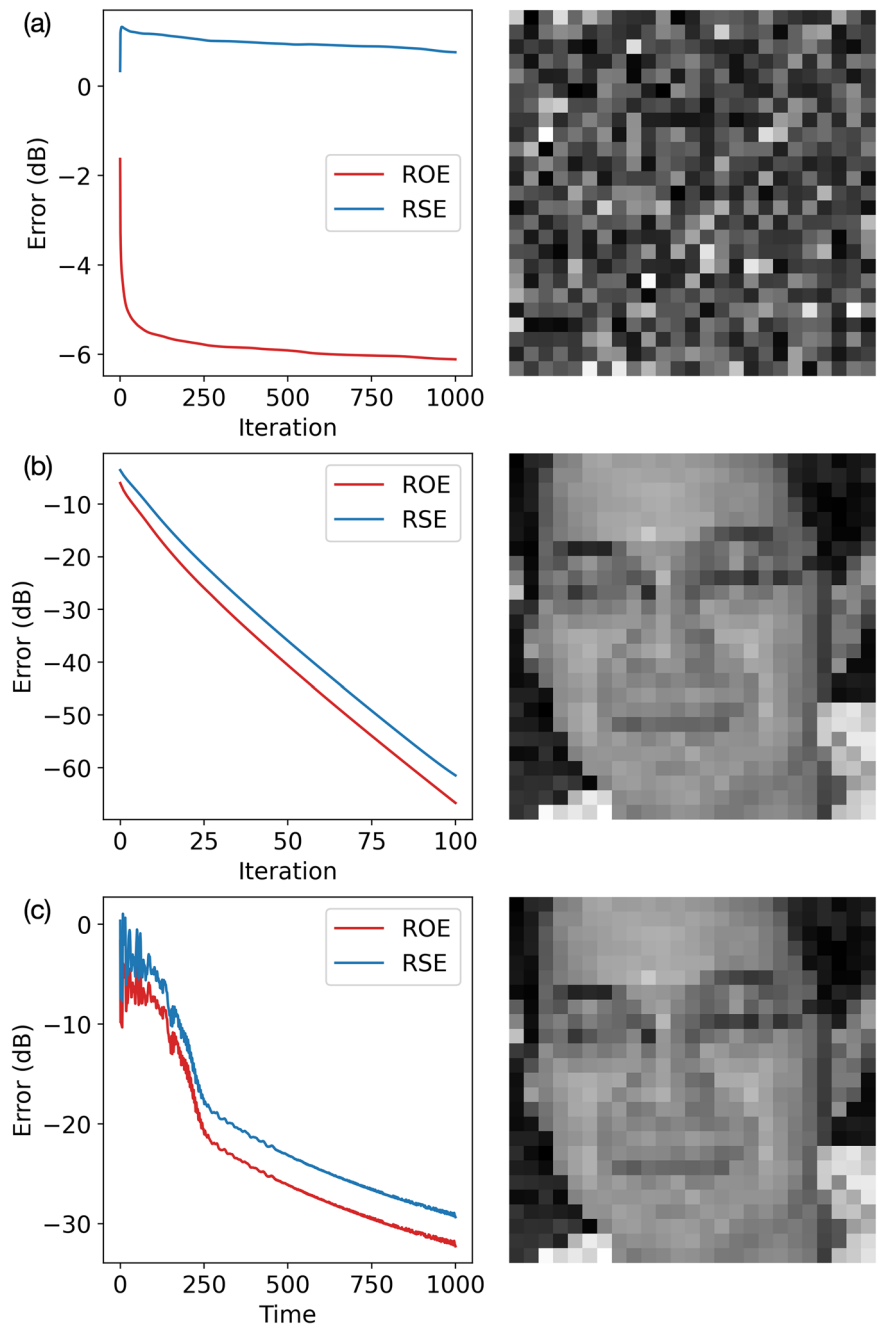
where β is a constant parameter. We found the value $\beta = 0.5$ used by authors of Elser et al.³⁴ generally produced good results.

To adapt RRR for the CDP variant of the phase-retrieval problem, we follow the method proposed in ref. 34 and modify projection P_1 to the following:

$$P_1(\mathbf{b}_n) = \mathbf{A}\mathbf{A}^\dagger \mathbf{b}_n. \quad (22)$$

One can motivate this projection by considering that $\mathbf{A}^\dagger \mathbf{b}_n$ makes use of all L sets of observations, unique to the CDP formulation, to produce an “average” estimated sample vector \mathbf{x}_n , based on all available observations, and then using this best estimation to produce the next \mathbf{b}_{n+1} by calculating $\mathbf{A}\mathbf{x}_n$, which in turn ensures that \mathbf{b}_{n+1} produced this way remains in the range of \mathbf{A} . We then substituted this modified projection P_1 into the iterative scheme of Eq. (21), keeping the original P_2 operator and $\beta = 0.5$, and applied the

Fig. 3 | Comparison of the GS method and the gain-based system performance in recovering a real-valued image. **a** Error evolution and final reconstructed image with the GS method starting from a complex-valued random initial condition whose phase is uniformly randomly distributed in the range $[0, 2\pi)$ and whose amplitude is the known observation vector **b**. **b** Error evolution and final reconstructed image with the GS method starting from an initial condition whose phase is obtained by multiplying **A** with a random real-valued vector \tilde{x} , and whose amplitude is the known observation vector **b**. **c** Error evolution and final reconstructed image with gain-based method starting from the exact same initial condition as used in (a). See “Fig. 3” table in supplementary data.



resulting RRR method to the same phase-retrieval instances used by the gain-based system.

The classical GS method, which was one of the first proposed heuristic methods for solving phase-retrieval problems, used the same P_1 and P_2 projections, but a simpler iterative formula:

$$\mathbf{b}_{n+1} = P_2(P_1(\mathbf{b}_n)). \quad (23)$$

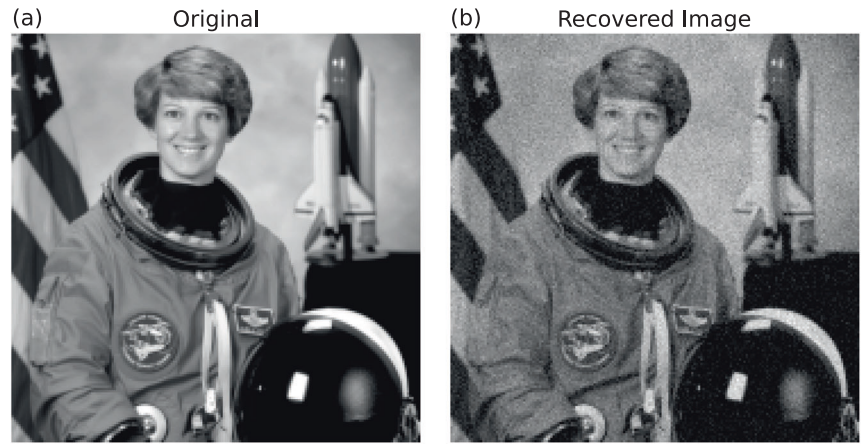
This iterative heuristic can also be applied to CDP phase-retrieval problems with modified projection P_1 given by Eq. (22). The quality of solutions found by these methods was measured by the performance metrics presented in the “Performance evaluation” section.

In summary, for each class of sample vectors **x** considered in the “Results” section, we first generated matrix **A** with Eqs. (14) and (15), and then generated observation vector **b** with Eq. (3). Matrix **A** and vector **b** completely define the CDP phase-retrieval problem. To solve this problem

with the simulated gain-based method, matrix **A** and observation vector **b** were then substituted into Eq. (11) to obtain the coupling matrix **J**, and it was used to construct the dynamical system defined by Eqs. (12) and (13). After the dynamical system evolves until it reaches a stationary state or a maximum simulation duration, the final phases of the complex-valued state parameters $\{\psi_i\}$ can be read out. The phases can then finally be substituted into Eq. (6) to produce the reconstructed sample vector **x**. This solution was compared against the solution obtained by presenting matrix **A** and observation vector **b** to the RRR and GS methods.

To ensure fairness in each set of comparisons, initial conditions for the RRR method and the gain-based method were always generated from the same initial random guess of phases θ_i , which were uniformly randomly chosen from the range $[0, 2\pi)$. Angles $\{\theta_i\}$ became the initial phases of the complex-valued state $\{\psi_i\}$ for the gain-based method, and $\{b_i e^{i\theta_i}\}$ became the initial guess for the RRR and GS methods. Simulation results are shown in the “Results” section.

Fig. 4 | Phase retrieval with a large-scale sample vector using the gain-based system. **a** A 180×180 pixel grayscale image downsampled from the “astronaut” image of the “scikit-image” Python image dataset package⁵². This image is used as the sample vector in a CDP-based phase-retrieval setup with 8 phase filters, yielding an observation vector **b** of length 259,200. **b** The final reconstruction after the gain-based system evolves for $t = 5$ from a random initial condition. The resulting RSE is -9.4 and the ROE is -12 .



Results

Performance with real-valued sample vectors

Many studies^{18,24,34} have evaluated phase-retrieval algorithms using images that are real and strictly positive, represented by two-dimensional matrices. Although such images offer a convenient test bed, they restrict **x** to real values only, whereas real-world diffraction data are generally complex. Nevertheless, these simple test cases already reveal key limitations of traditional phase-retrieval algorithms.

For instance, we used the positive real-valued image depicted in Fig. 2 to construct a CDP-based phase-retrieval problem with five phase filters. The resulting observation vector **b** was then presented to different solvers. As shown in Fig. 3a, when the GS method began from a random complex initialization for **b**, it quickly became stuck in a local minimum after about 1000 iterations, failing to reproduce any recognizable features of the original image. By contrast, starting from the same initial condition, the gain-dissipative system (Fig. 3c) followed a markedly different trajectory and produced significantly lower errors, ultimately reconstructing the image with high fidelity.

We further observed that a priori knowledge of **x** being real and positive can substantially simplify phase retrieval. For example, by multiplying **A** with a random positive real-valued vector $\tilde{\mathbf{x}}$ and extracting its phase as the initial guess for GS, we obtained the error evolution and final reconstruction depicted in Fig. 3b. This carefully chosen initialization, which already exhibits a lower error than a random complex guess, allowed the GS method to recover the underlying image. Clearly, such information (i.e., that **x** is a non-negative real) makes the phase-retrieval problem more tractable. However, real-world experimental data generally produce complex-valued **x** without providing a straightforward initialization. Consequently, the GS method often encounters difficulty in practical scenarios. In the following section, we therefore focus on more general, complex-valued samples and benchmark the gain-based solver in that setting.

As a final demonstration using real-valued data, Fig. 4a shows a high-resolution grayscale image of size 180×180 , totaling 32,400 real-valued pixels. Despite this large problem dimension, the gain-based method successfully reconstructs the main features of the image after a short simulation (duration $t = 5$), as shown in Fig. 4b. Notably, this was achieved without leveraging the fact that the target sample vector is purely real. While the final RSE remains moderately high at -9.4 , the essential image structure is clearly recognizable, albeit with visible background noise.

Although the gain-based method easily solves real-valued phase-retrieval problems, most practical applications involve recovering complex-valued data. Consequently, the remaining sections focus on benchmark cases with complex-valued samples to more accurately reflect real-world experimental conditions.

Performance with complex-valued sample vectors

We begin by investigating the reconstruction of a complex-valued image that can potentially be produced in an experiment: a two-dimensional vortex, where the vortex flow is given by the gradient of phase at each point in a plane. The field value at each point can be approximated by

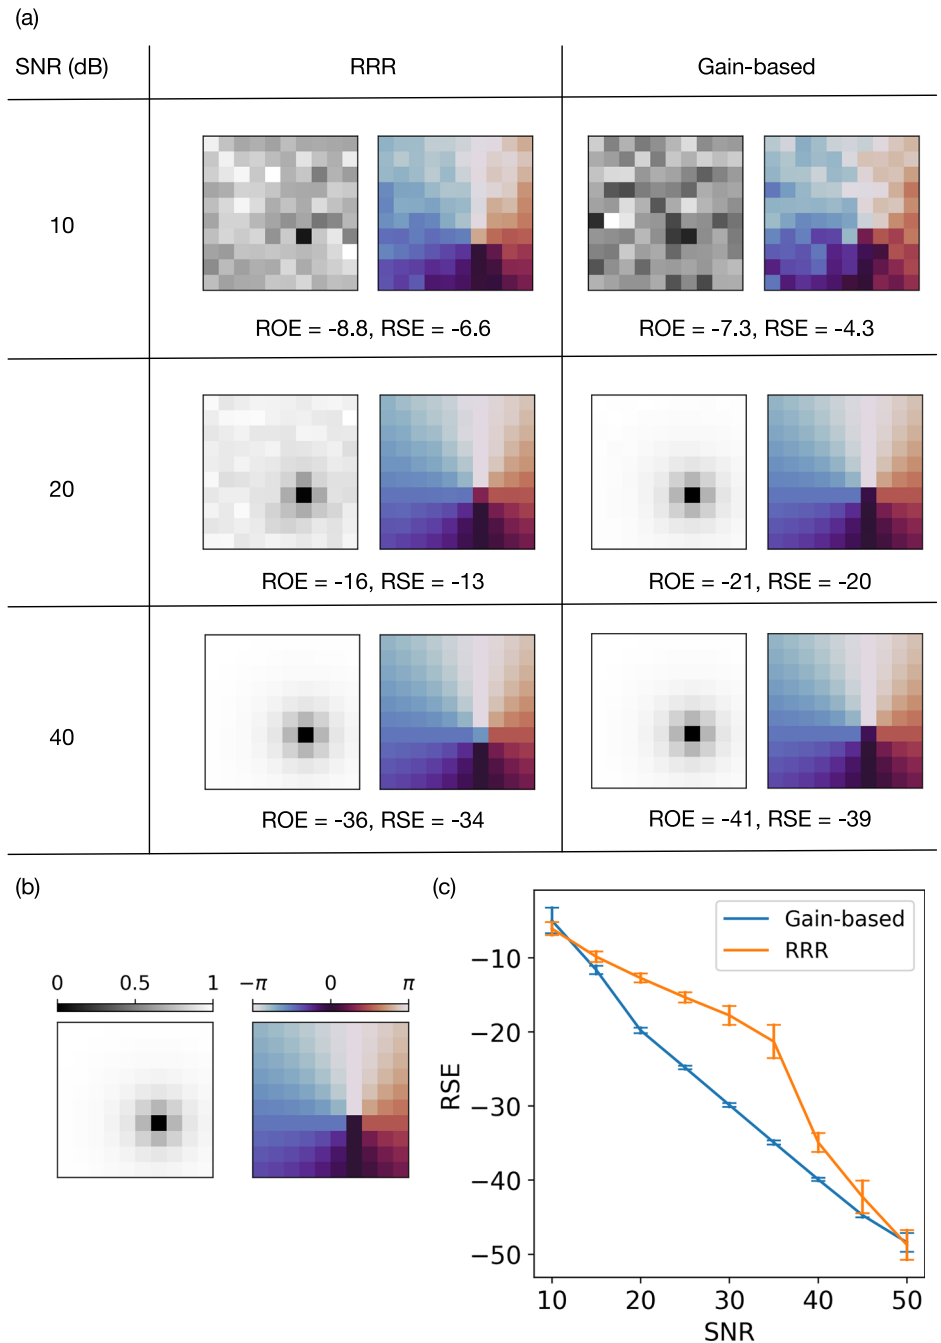
$$v(x, y) = \frac{(x - x_0) + i(y - y_0)}{\sqrt{(x - x_0)^2 + (y - y_0)^2 + \xi^2}}, \quad (24)$$

where (x_0, y_0) is the center of the vortex and ξ is the size of the vortex core⁴⁶. The amplitude, $|v|$ (greyscale), and phase of the vortex $\arg(v)$ (color) can be visualized as shown in Fig. 5b.

In most practical experiments, the observed amplitudes $\tilde{\mathbf{b}}$ deviate from the ideal measurements **b** due to noise. As illustrated in Fig. 5a, both the RRR and gain-based solvers were presented with noisy observations injected with Gaussian noises of varying magnitudes. In the high-noise regime (small SNR values), neither method recovered a solution close to the original sample **x**, shown in Fig. 5b. Under moderate noise ($\text{SNR} \approx 20$), the RRR reconstruction exhibited pronounced amplitude distortions, whereas the gain-based solver produced a visually accurate approximation of **x**. This improvement is reflected quantitatively by the lower RSE values achieved by the gain-dissipative method, a performance gap that persists until around $\text{SNR} \approx 40$. Above this threshold, the visual differences between the two solvers diminish, although residual discrepancies in ROE and RSE remain. Figure 5c summarizes the impact of noise on reconstruction accuracy, revealing that the gain-based approach and RRR perform similarly in the high-noise ($\text{SNR} < 10$) and low-noise ($\text{SNR} > 40$) regimes. Notably, however, the gain-dissipative solver holds a distinct advantage in the intermediate range $10 < \text{SNR} < 40$, where the difference in RSE is both statistically and visually significant. In these conditions, the gain-dissipative system often recovers important structural details that RRR loses in noise.

While the two-dimensional vortex example provides instructive insight into solver behavior at various noise levels, it remains a highly structured sample. To assess robustness on less structured data, we next consider *unstructured* samples **x**, whose amplitudes are drawn uniformly from $[0, 1)$ and phases from $[0, 2\pi)$. Keeping the dimensionality fixed and varying the noise level, we compared the RRR and gain-based methods, as summarized in Fig. 6a. The results largely mirror the vortex case: in the medium-noise regime, the gain-based solver attains markedly lower RSE than RRR, whereas both methods perform comparably under very high ($\text{SNR} < 10$) or very low ($\text{SNR} > 40$) noise. These findings indicate that the gain-based approach handles both structured and unstructured samples effectively, consistently outperforming RRR in the medium-noise band. Moreover, its RSE increases in tandem with SNR, suggesting stable performance across different noise levels.

Fig. 5 | Phase retrieval of a two-dimensional vortex in the presence of noise, comparing RRR and the gain-based system. The phase-retrieval problems were constructed with $L = 5$ phase filters. The size of the vortex core $\xi = 1$, and the vortex field were discretised into a 10×10 grid. **a** Each panel displays the reconstructed sample vector $\tilde{\mathbf{x}}$, where the grayscale image encodes amplitude, and the color image encodes phase. For each noise level, both RRR and the gain-based method start from the same initial condition. Here, RRR runs for 10,000 iterations, while the gain-based system is evolved to $t = 1000$. **b** The ground-truth sample vector \mathbf{x} that describes a 2D vortex, showing amplitude (left) and phase (right). **c** Phase-retrieval error (RSE) versus the signal-to-noise ratio (SNR). In this test, the noisy observation vectors $\tilde{\mathbf{b}}$ presented to the solvers contained Gaussian noise. Each data point represents the average of 20 random instances, where the vortex core is placed at different positions, and each algorithm is initialized randomly. Error bars denote the standard deviation of the final RSE values. See “Fig. 5” table in supplementary data.



All of the above tests involved samples of size 100. However, realistic applications typically require recovering much larger vectors. To explore how solution accuracy scales with dimensionality, Fig. 6b plots RSE against problem size for medium-noise (SNR = 30, in blue) and low-noise (SNR = 50, in orange) conditions, comparing the gain-based method (solid lines) and RRR (dashed lines). In both noise regimes, the RSE remains nearly constant even as the number of sample elements grows by a factor of four. This suggests that noise level influences the gain-based solver’s performance more than the underlying problem dimension. At very low RSE (~50 dB), the gain-based solver exhibits a slightly larger spread of RSE than RRR in Fig. 6. The spread can be explained by the following two effects. Firstly, the gain-based dynamics are integrated using an adaptive time-step ODE solver, so tiny variations in step acceptance or in the stopping time lead to small changes in the final state; when plotted in decibels, these tiny differences are amplified at very low decibel levels. Secondly, different initial conditions can relax

into slightly different low-lying minima of the XY Hamiltonian. RRR iterations are discrete and extremely numerically stable, so the spread for RRR comes almost entirely from the choice of initial phases. Hence, the increased spread in RSE produced by the gain-based solver only appears deep in the regime of extremely small error, and does not indicate instability at moderate noise levels.

In addition, noisy observation vectors $\tilde{\mathbf{b}}$ with Poisson noise or systematic errors, instead of Gaussian noise, were also presented to the solvers to reveal their robustness to different types of noise. Figure 7a shows the trend of solution quality measured by RSE as a function of Poisson noise magnitude measured by SNR, while Fig. 7b shows the RSE trend for observation vectors with constant systematic shifts to each of their elements. In both cases, phase-retrieval problems were generated from complex-valued sample vectors with random amplitudes in the range $[0, 1)$ and phases in the range $[0, 2\pi)$. Just like the Gaussian noise case shown in Fig. 6a, when the noisy observation vector $\tilde{\mathbf{b}}$ contained Poisson noise or systematic

Fig. 6 | Phase-retrieval problems generated from random complex-valued samples with $L = 5$ phase filters. To solve each problem, the gain-based solver evolved to $t = 1000$, and RRR ran for 10,000 iterations. All noisy observation vectors $\tilde{\mathbf{b}}$ contained Gaussian noise. **a** Phase-retrieval error (RSE) produced by the RRR method (dashed lines) and the gain-based system (solid lines) as a function of Gaussian noise in the measured amplitudes. At each noise level, 20 random complex samples were generated, each with 100 elements whose amplitudes are uniformly distributed in $[0, 1)$ and phases in $[0, 2\pi)$. The resulting observation vectors were then used for both methods. Vertical error bars indicate the standard deviation in RSE across the 20 trials. **b** Phase-retrieval error (RSE) versus the dimensionality of the sample vectors, comparing medium noise (SNR = 30, shown in blue) and low noise (SNR = 50, shown in orange). Solid lines again correspond to the gain-based system, and dashed lines correspond to RRR. Each data point represents the average RSE over 20 distinct random instances of the specified dimension, with error bars indicating the standard deviation. See “Fig. 6a” and “Fig. 6b” tables in supplementary data.

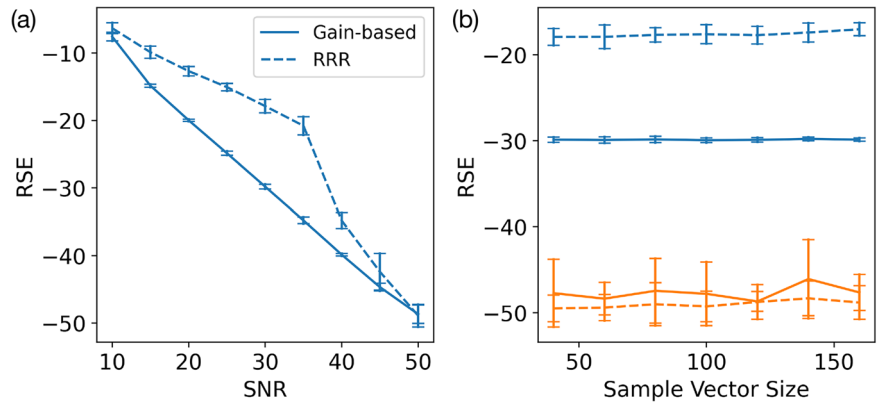



Fig. 7 | Phase-retrieval problems with Poisson or systematic errors. These problems were generated from random complex-valued samples with $L = 5$ phase filters and with Poisson noise or systematic errors added to observation vectors. For each problem, the gain-based solver evolved to $t = 1000$, and the RRR ran for 10,000 iterations. **a** Phase-retrieval error (RSE) produced by the RRR method (dashed line) and the gain-based system (solid line) as a function of Poisson noise in the measured amplitudes. **b** RSE produced by RRR and the gain-based methods as a function of systematic errors introduced into the measured amplitudes. In both panels, vertical error bars are the standard deviations in RSE across 20 different CDP phase-retrieval problems generated from random complex-valued sample vectors. These sample vectors had amplitudes uniformly randomly distributed in $[0, 1)$ and phases uniformly randomly distributed in $[0, 2\pi)$. See “Fig. 7a” and “Fig. 7b” tables in supplementary data.



errors, the gain-based method still clearly outperformed the RRR method under moderate noise, and performed similarly to the RRR method under low and high noise. This shows that the advantage of the gain-based method over the RRR method is robust to different kinds of noises that can be found in real diffraction experiments.

To illustrate how the gain-based method can recover phase information in cold-atomic Bose-Einstein condensate (BEC) experiments, where topological defects like solitons⁴⁷, vortex lines⁴⁸, and vortex rings^{49,50} naturally appear, we consider reconstructing a three-dimensional complex vortex ring. The sample vector is specified by

$$v(r, \theta, z) = \frac{(r - r_0) + iz}{\sqrt{(r - r_0)^2 + z^2 + \xi^2}}, \quad (25)$$

where (r, θ, z) are cylindrical coordinates, r_0 is the ring’s radius, and ξ is the vortex-core size. Figure 8a shows an isosurface at 30% of the maximum amplitude, with phase isolines superimposed. The vortex flow, which is orthogonal to these lines, winds along the ring.

We construct a CDP phase-retrieval problem from this 3D complex-valued vector (3087 elements) and $L = 8$ phase filters, giving an observation vector \mathbf{b} of size 24,696. We solved it using the gain-based system and Fig. 8b charts the time evolution of the retrieval error, while Fig. 8c–f depict snapshots of the reconstructed three-dimensional field at various stages. By the final state, the key vortex-ring features closely match the original, indicating that, with non-destructive CDP measurements, the gain-based approach can accurately recover the phase of 3D wavefunctions in BEC systems.

Conclusion

We have shown that the coded diffraction pattern variant of phase-retrieval can be rigorously reformulated as an XY Hamiltonian minimization problem, paving the way for direct solution by gain-based oscillator networks and related physics-inspired systems. Through numerical tests, we demonstrated that such gain-based dynamics significantly outperforms the state-of-the-art Relaxed-Reflect-Reflect algorithm, particularly under medium-level noise (SNR values between 10 and 40 dB). Our findings hold for both structured data (e.g., two-dimensional

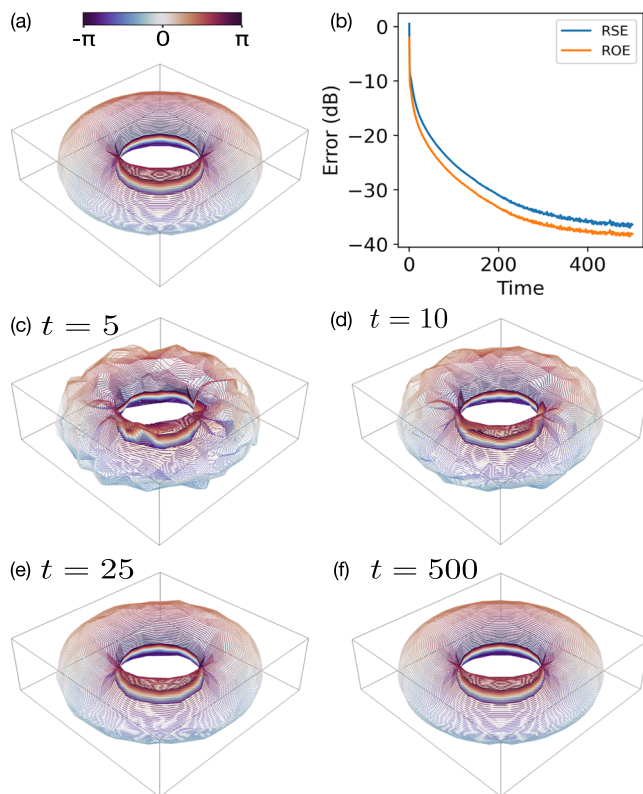


Fig. 8 | Three-dimensional vortex ring reconstructed using the gain-based phase-retrieval method. The phase-retrieval problem was constructed with $L = 8$ phase filters, and the field was discretised into a $21 \times 21 \times 7$ grid, resulting in a sample vector \mathbf{x} of size 3087, and an observation vector \mathbf{b} of size 24,696. **a** Original vortex ring, visualized as an isosurface at 30% of its maximum amplitude. Phase isolines on this surface are colored according to their phase values, ranging from $-\pi$ to π . **b** Time evolution of the phase-retrieval error (RSE and ROE) under gain-based dynamics. **c–f** Snapshots of the reconstructed vortex ring at $t = 5$, $t = 10$, $t = 25$, and $t = 500$, respectively, illustrating the progressive refinement of amplitude and phase in three dimensions. See “Fig. 8” table in supplementary data.

vortices and three-dimensional vortex rings) and unstructured complex-valued data with random amplitudes and phases, and they are robust to Gaussian and Poisson noises as well as systematic errors in the observation vector.

Critically, we observed that the superior accuracy of the gain-based solver remains robust even as problem sizes grow. This scalability, combined with its noise resilience, indicates strong potential for large-scale real-world imaging tasks. Moreover, the gain-based approach can be physically realized in optical, polaritonic, or other nonlinear oscillator networks, thereby offering a hardware platform for rapid, energy-efficient phase retrieval. Such physical devices could perform continuous, parallel searches for global minima in the XY energy landscape, transforming this theoretical advantage into practical gains for real-time imaging and beyond.

Our gain-based algorithm requires only one continuous evolution of the oscillator network to a steady state for each CDP phase-retrieval problem, so the wall-clock time is essentially the physical relaxation time of the hardware. A summary of reported performance metrics of experimental implementations of various unconventional physical platforms could be found in ref. 13, and from there, we could provide some order-of-magnitude estimate of the expected computational speed of our gain-based algorithm when implemented in that hardware. In SPIMs, measured performance for $\sim 10^2$ spins with 60 Hz liquid-crystal SLMs corresponds to a time-to-solution of order 2 s, dominated by the SLM refresh rate. Next-generation electro-optic SLMs with frame rates >1 GHz are projected to reduce this to ~ 10 μ s for problems of similar logical size, since the optical evaluation of the XY

energy is fully parallel in the number of spins. The SPIM platform has demonstrated all-to-all connectivity for spin counts up to $\sim 10^5$ and projected capabilities near $\sim 10^7$, making the CDP problem sizes studied here ($M = LN \sim 2.5 \times 10^4$ – 2.6×10^5) well within the expected scaling regime. Coherent Ising machines (CIMs) report time-to-solution of a few milliseconds for problems of ~ 100 spins, suggesting that an XY implementation of Eqs. (12) and (13) would converge on sub-ms to ms timescales for the dimensionalities considered in this work.

Our results open a promising direction in the development of continuous-variable “ XY machines”, enabling them to tackle large and noisy phase-retrieval instances that arise in a variety of scientific and industrial settings.

Data availability

The authors declare that all data supporting the findings of this study are available within the supplementary information files of this paper.

Code availability

The code that supports the findings of this study is available from the first author (zw321@cam.ac.uk) upon request.

Received: 19 August 2025; Accepted: 22 January 2026;

Published online: 03 February 2026

References

1. Stroeve, N. & Berloff, N. G. Analog photonics computing for information processing, inference, and optimization. *Adv. Quantum Technol.* **6**, 2300055 (2023).
2. Lucas, A. Ising formulations of many NP problems. *Front. Phys.* **2**, 5 (2014).
3. Syed, M. & Berloff, N. G. Physics-enhanced bifurcation optimisers: all you need is a canonical complex network. *IEEE J. Sel. Top. Quantum Electron.* **29**, 7400406 (2023).
4. Honjo, T. et al. 100,000-spin coherent Ising machine. *Sci. Adv.* **7**, eabh0952 (2021).
5. Honari-Latifpour, M., Mills, M. S. & Miri, M.-A. Combinatorial optimization with photonics-inspired clock models. *Commun. Phys.* **5**, 104 (2022).
6. Inagaki, T. et al. A coherent Ising machine for 2000-node optimization problems. *Science* **354**, 603–606 (2016).
7. Kalinin, K. P. & Berloff, N. G. Polaritonic network as a paradigm for dynamics of coupled oscillators. *Phys. Rev. B* **100**, 245306 (2019).
8. Kalinin, K. P. & Berloff, N. G. Networks of non-equilibrium condensates for global optimization. *N. J. Phys.* **20**, 113023 (2018).
9. Parto, M., Hayenga, W., Marandi, A., Christodoulides, D. N. & Khajavikhan, M. Realizing spin Hamiltonians in nanoscale active photonic lattices. *Nat. Mater.* **19**, 725–731 (2020).
10. Kim, K., Kumagai, M. & Yamamoto, Y. Combinatorial clustering with a coherent XY machine. *Opt. Express* **32**, 33737 (2024).
11. Pal, V., Mahler, S., Tradonsky, C., Friesem, A. A. & Davidson, N. Rapid fair sampling of the xy spin Hamiltonian with a laser simulator. *Phys. Rev. Res.* **2**, 033008 (2020).
12. Nixon, M., Ronen, E., Friesem, A. A. & Davidson, N. Observing geometric frustration with thousands of coupled lasers. *Phys. Rev. Lett.* **110**, 184102 (2013).
13. Wang, R. Z. et al. Efficient computation using spatial-photonic Ising machines with low-rank and circulant matrix constraints. *Commun. Phys.* **8**, 86 (2025).
14. Harrison, R. W. Phase problem in crystallography. *J. Opt. Soc. Am. A* **10**, 1046 (1993).
15. Barnett, M. J., Millane, R. P. & Kingston, R. L. Analysis of crystallographic phase retrieval using iterative projection algorithms. *Acta Crystallogr. Sect. D Struct. Biol.* **80**, 800–818 (2024).
16. Krist, J. E. & Burrows, C. J. Phase-retrieval analysis of pre- and post-repair Hubble Space Telescope images. *Appl. Opt.* **34**, 4951 (1995).

17. Miao, J., Charalambous, P., Kirz, J. & Sayre, D. Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature* **400**, 342–344 (1999).
18. Shechtman, Y. et al. Phase retrieval with application to optical imaging: a contemporary overview. *IEEE Signal Process. Mag.* **32**, 87–109 (2015).
19. Huang, M. & Xu, Z. No existence of a linear algorithm for the one-dimensional Fourier phase retrieval. *J. Complex.* **86**, 101886 (2025).
20. Eldar, Y. C. & Mendelson, S. Phase retrieval: stability and recovery guarantees. *Appl. Comput. Harmon. Anal.* **36**, 473–494 (2014).
21. Bendory, T., Khoo, Y., Kileel, J., Mickelin, O. & Singer, A. Autocorrelation analysis for cryo-EM with sparsity constraints: improved sample complexity and projection-based algorithms. *Proc. Natl. Acad. Sci. USA* **120**, e2216507120 (2023).
22. Fienup, J. R. Phase retrieval algorithms: a comparison. *Appl. Opt.* **21**, 2758 (1982).
23. Chen, C.-C., Miao, J., Wang, C. W. & Lee, T. K. Application of optimization technique to noncrystalline x-ray diffraction microscopy: guided hybrid input–output method. *Phys. Rev. B* **76**, 064113 (2007).
24. Lатычевская, T. Iterative phase retrieval in coherent diffractive imaging: practical issues. *Appl. Opt.* **57**, 7187 (2018).
25. Tradonsky, C. et al. Rapid laser solver for the phase retrieval problem. *Sci. Adv.* **5**, eaax4530 (2019).
26. Candès, E. J., Li, X. & Soltanolkotabi, M. Phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.* **39**, 277–299 (2015).
27. Fannjiang, A. & Liao, W. Phase retrieval with random phase illumination. *J. Opt. Soc. Am. A* **29**, 1847 (2012).
28. Fannjiang, A. & Strohmer, T. The numerics of phase retrieval. *Acta Numer.* **29**, 125–228 (2020).
29. Candès, E. J., Strohmer, T. & Voroninski, V. PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming. *Commun. Pure Appl. Math.* **66**, 1241–1274 (2013).
30. Waldspurger, I., d’Aspremont, A. & Mallat, S. Phase recovery, MaxCut and complex semidefinite programming. *Math. Program.* **149**, 47–81 (2015).
31. Candès, E. J., Eldar, Y. C., Strohmer, T. & Voroninski, V. Phase retrieval via matrix completion. *SIAM J. Imaging Sci.* **6**, 199–225 (2013).
32. Elser, V. The complexity of bit retrieval. *IEEE Trans. Inf. Theory* **64**, 412–428 (2018).
33. Levin, E. & Bendory, T. A note on Douglas–Rachford, gradients, and phase retrieval. arXiv: 1911.13179 (2020).
34. Elser, V., Lan, T.-Y. & Bendory, T. Benchmark problems for phase retrieval. *SIAM J. Imaging Sci.* **11**, 2429–2455 (2018).
35. Chen, Y. & Candès, E. J. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Commun. Pure Appl. Math.* **70**, 822–883 (2017).
36. Cha, E., Lee, C., Jang, M. & Ye, J. C. Deepphasecut: deep relaxation in phase for unsupervised Fourier phase retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 9931–9943 (2022).
37. Cummins, J. S., Salman, H. & Berloff, N. G. Ising Hamiltonian minimization: gain-based computing with manifold reduction of soft spins vs. quantum annealing. *Phys. Rev. Res.* **7**, 013150 (2025).
38. Berloff, N. G. et al. Realizing the classical xy Hamiltonian in polariton simulators. *Nat. Mater.* **16**, 1120–1126 (2017).
39. Toebes, C., Vretenar, M. & Klaers, J. Dispersive and dissipative coupling of photon Bose–Einstein condensates. *Commun. Phys.* **5**, 59 (2022).
40. Litvinenko, A. et al. A spinwave Ising machine. *Commun. Phys.* **6**, 227 (2023).
41. Pierangeli, D., Marcucci, G. & Conti, C. Large-scale photonic Ising machine by spatial light modulation. *Phys. Rev. Lett.* **122**, 213902 (2019).
42. Veraldi, D. et al. Fully programmable spatial photonic Ising machine by focal plane division. *Phys. Rev. Lett.* **134**, 063802 (2025).
43. Kalinin, K. P. & Berloff, N. G. Global optimization of spin Hamiltonians with gain-dissipative systems. *Sci. Rep.* **8**, 17791 (2018).
44. Gerchberg, R. W. Holography without fringes in the electron microscope. *Nature* **240**, 404–406 (1972).
45. Marchesini, S. et al. X-ray image reconstruction from a diffraction pattern alone. *Phys. Rev. B* **68**, 140101 (2003).
46. Berloff, N. G. Padé approximations of solitary wave solutions of the Gross–Pitaevskii equation. *J. Phys. A: Math. Gen.* **37**, 1617 (2004).
47. Yefsah, T. et al. Heavy solitons in a fermionic superfluid. *Nature* **499**, 426–430 (2013).
48. Serafini, S. et al. Vortex reconnections and rebounds in trapped atomic Bose–Einstein condensates. *Phys. Rev. X* **7**, 021031 (2017).
49. Ku, M. J. H. et al. Motion of a solitonic vortex in the BEC–BCS crossover. *Phys. Rev. Lett.* **113**, 065301 (2014).
50. Bulgac, A., Forbes, M. M., Kelley, M. M., Roche, K. J. & Wlazłowski, G. Quantized superfluid vortex rings in the unitary fermi gas. *Phys. Rev. Lett.* **112**, 025301 (2014).
51. Huang, G. B., Mattar, M., Berg, T. & Learned-Miller, E. Labeled faces in the wild: a database for studying face recognition in unconstrained environments. In *Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition, Marseille, France* (eds Learned-Miller, E., Ferencz, A. & Jurie, F.) (2008).
52. van der Walt, S. et al. scikit-image: image processing in Python. *PeerJ* **2**, e453 (2014).

Acknowledgements

The authors acknowledge the support from HORIZON EIC-2022-PATHFINDERCHALLENGES-01 HEISINGBERG Project 101114978. R.Z.W. and N.G.B. acknowledge the support from the Julian Schwinger Foundation Grant No. JSF-19-02-0005. N.G.B. also acknowledges support from Weizmann-UK Make Connection Grant 142568 and the EPSRC UK Multidisciplinary Centre for Neuromorphic Computing (grant UKRI982).

Author contributions

R.Z.W. developed the simulation code and performed numerical simulations. G.L. performed large-scale numerical simulations, produced Fig. 4, and provided comments during manuscript preparation. S.G., D.P., M.C.S., and C.C. contributed implementation considerations for physical hardware, including expected operating timescales, and provided input on experimental feasibility. R.Z.W. wrote the initial manuscript draft. N.G.B. substantially revised and edited the manuscript. N.G.B. supervised the study.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-026-02525-7>.

Correspondence and requests for materials should be addressed to Natalia G. Berloff.

Peer review information *Communications Physics* thanks Mostafa Honari, Roman Khymyn and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

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/>.

© The Author(s) 2026