

Vision Transformer for Multi-Domain Phase Retrieval in Coherent Diffraction Imaging

Jialun Liu¹, David Yang², Ian Robinson^{1,2}

¹London Centre for Nanotechnology, University College London, London WC1E 6BT, United Kingdom.

²Condensed Matter Physics and Materials Science Department, Brookhaven National Laboratory, Upton, NY 11973, USA.

Abstract

Bragg coherent diffraction imaging (BCDI) phase retrieval becomes rapidly difficult in the strong-phase regime, where a crystal contains distortions beyond half a lattice spacing. An important special case is the phase domain problem, where blocks of a crystal are displaced with sharp jumps at domain walls. The strong-phase, here defined as beyond $\pm\pi/2$, generates split Bragg peaks and dense fringe structure for which classical iterative solvers often stagnate or return different solutions from different initialisations. Here, we introduce an unsupervised Fourier Vision Transformer (Fourier ViT) to solve this block-phase, multi-domain phase-retrieval problem directly from measured 2D Bragg diffraction intensities. Fourier ViT couples reciprocal-space information globally through multiscale Fourier token mixing, while shallow convolutional front and back-ends provide local filtering and reconstruction. We validate the approach on large-scale synthetic datasets of Voronoi multi-domain crystals with strong-phase contrast under realistic noise corruptions, and on experimental diffraction from a $\mathbf{La}_{2-x}\mathbf{Ca}_x\mathbf{MnO}_4$ nanocrystal. Across the regimes considered, Fourier ViT achieves the lowest reciprocal-space mismatch (χ^2) among the compared methods and preserves domain-resolved phase reconstructions for increasing numbers of domains. On experimental data, with the same real-space support, Fourier ViT matches the iterative benchmark χ^2 while improving robustness to random initialisations, yielding a higher success rate of low- χ^2 reconstructions than the complex convolutional neural network baseline.

Keywords: Vision transformer; Bragg coherent diffraction imaging; phase retrieval; multi-domain

1 Introduction

Bragg coherent diffraction imaging (BCDI) has transformed the way we characterise nanoscale materials by revealing the internal structure and lattice distortions of single crystals in three dimensions

(Miao et al., 1999; Robinson and Harder, 2009; Pfeifer et al., 2006). This lensless X-ray technique inverts Bragg peaks into real-space by iteratively applying reciprocal and real-space constraints via fast Fourier transforms (Fienup, 1982). BCDI reconstructions have revealed how edge strain in single Pt nanocatalysts influences methane-oxidation activity (Kim et al., 2018), how nanostrain accumulation drives structural degradation and voltage fade in Li-rich layered oxide cathodes (Liu et al., 2022), and how electric fields can affect domain wall dynamics in BaTiO₃ nanoparticles (Liu et al., 2025). However, its core challenge remains phase retrieval: detectors record only diffraction intensities, so the phase required for real-space reconstruction is missing (Sayre, 1952; Fienup, 1982).

Recovering phase from intensities alone is a non-convex, ill-conditioned inverse problem whose difficulty grows sharply with internal strain and domain complexity. Classical algorithms such as Gerchberg–Saxton error reduction (Gerchberg and Saxton, 1972), hybrid input–output (HIO) (Fienup, 1982), difference map (Elser, 2003) and Relaxed Averaged Alternating Reflections (RAAR) (Luke, 2005) alternate between real and reciprocal-space, enforcing constraints such as support and measured intensities. For “weak-phase” crystals, where phase shifts stay below about $\pi/2$, these methods are often reliable and converge in a few hundred iterations (Robinson and Harder, 2009). Once phase shifts exceed this regime, interference between differently phased regions leads to peak splitting and complex fringe patterns, so called the “strong-phase” problem (Robinson et al., 2020; Wu et al., 2021a; Gao et al., 2021). The central Bragg maximum can split into multiple lobes, producing separated intensity maxima and dense fringes around the Bragg peak, so the measured diffraction becomes extremely sensitive to the internal phase distribution. In this strong-phase, multi-domain regime, iterative solvers stagnate or converge to different solutions under different random seeds, reflecting an apparent non-uniqueness of the reconstruction (Bates, 1984; Ulvestad et al., 2017; Carnis et al., 2019). This slow and fragile convergence limits routine use of BCDI and makes real-time feedback for in situ or operando experiments difficult, especially at X-ray free-electron laser (XFEL) and modern synchrotron sources (Clark et al., 2013; Harder, 2021).

Under idealised conditions, theoretical work on the uniqueness of the phase problem shows that oversampling in two or more dimensions is sufficient to guarantee a unique solution up to trivial symmetries such as real-space translation, conjugate inversion, a global phase shift and Fresnel propagation (Bates, 1984; Bruck and Sodin, 1979; ?; Huang et al., 2011). In contrast, the one-dimensional problem generally admits many non-equivalent solutions that share the same Fourier magnitudes (Beinert and Plonka, 2015; Bendory et al., 2017). Motivated by this, we focus on a two-dimensional formulation in which each synthetic crystal is treated as a projection of a three-dimensional nanocrystal, thereby preserving the essential Fourier-phase structure of the full 3D problem. Following Bates (Bates, 1984), we consider the 2D and 3D phase problems to be equivalent.

We focus on the phase-domain limit relevant to multi-domain crystals. In BCDI, the real-space phase $\phi(\mathbf{r}) = \mathbf{Q} \cdot \mathbf{u}(\mathbf{r})$ is the projection of the lattice displacement field $\mathbf{u}(\mathbf{r})$ onto the reciprocal-lattice vector \mathbf{Q} of the measured Bragg peak (Pfeifer et al., 2006; Robinson and Harder, 2009). We define a phase domain as a connected block-shaped region in which this projected phase is approximately constant, separated from neighbouring regions by sharp phase discontinuities across narrow domain walls. As the number of domains increases, coherent interference between domains produces stronger Bragg-peak splitting and richer fringe structure, and the phase-retrieval landscape becomes more complicated (Robinson et al., 2020). Such phase domain textures occur in real quantum materials, for example, BCDI has been used to image the formation and scaling of low-temperature orthorhombic domains in the cuprate superconductor La_{1.875}Ba_{0.125}CuO₄ (Assefa et al., 2020) and monoclinic domains in Magnetite Fe₃O₄ (?).

In practice, experimental data can be noisy and constraints such as support or non-negativity are only approximate (Fienup, 1987). For noisy data the phase problem no longer has a single exact solution. Instead, one obtains a family of near-solutions whose spread reflects both noise and algorithmic trapping (Bates and Napier, 1972). This has motivated ensemble strategies, in which many independent reconstructions are run and then either selected or averaged according to a figure of merit. Examples include guided HIO procedures (Ulvestad et al., 2017; Chen et al., 2007) and the common-mode (eigen-solution) analysis (Lee and Seung, 1999) implemented in PYNX, which combines multiple BCDI reconstructions using a free log-likelihood metric (Favre-Nicolin et al., 2020b,a). Alternative projection schemes, including Fourier-weighted projections (Guizar-Sicairos and Fienup, 2008), improve robustness for truncated diffraction data but remain iterative projection algorithms and can still be sensitive to initialisation. Overall, phase retrieval for multi-domain crystals remains slow, fragile, and difficult to automate.

In parallel, deep learning is increasingly used to solve coherent imaging problems. Supervised neural networks learn the inverse mapping from diffraction intensities to real-space structure, enabling near real-time reconstructions compared with iterative phase retrieval (Cherukara et al., 2018; Wu et al., 2021b, 2025). CNNs with U-Net architectures have been applied to optical holography (Rivenson et al., 2018), X-ray ptychography (Cherukara et al., 2020), and coherent X-ray diffraction imaging (CXDI) (Vu et al., 2025), including BCDI (Wu et al., 2021a,b; Yu et al., 2024). After training, these networks can return reconstructions orders of magnitude faster than classical iterative solvers (Cherukara et al., 2018, 2020; Yu et al., 2024). The main limitation of supervised approaches is that they learn only what has been given in the training set, and therefore do not guarantee the correct solution for a new object outside the training distribution, particularly when ground-truth labels are unavailable for experimental patterns (Cherukara et al., 2018; Wu et al., 2021b; Yao et al., 2022). Since label generation and verification are most straightforward in simple regimes, many published demonstrations and benchmarks focus on single-phase or few-domain systems (Wu et al., 2021b). Unsupervised and physics-aware variants, such as AutoPhaseNN and deep complex-valued CNNs for ultrafast BCDI, embed the forward imaging model directly into the network and loss function, thereby reducing the dependency on labelled training data (Yao et al., 2022; Yu et al., 2024). This unsupervised setting is particularly natural for BCDI, where ground-truth real-space phases are not available for experimental patterns and where strong-phase, multi-domain structures make the optimisation landscape highly non-convex and sensitive to initialisation (Wu et al., 2021b; Yao et al., 2022). Vision transformers (ViTs) have recently been applied in coherent X-ray imaging, for example in ptychography (Gan et al., 2024), but have not yet been integrated into a fully unsupervised, physics-informed BCDI framework.

Here, we introduce a simple block-phase domain model to access the strong-phase, multi-domain regime, and solve the resulting phase-retrieval problem using an unsupervised Fourier ViT with global reciprocal-space token mixing. The proposed Fourier ViT combines convolutional feature extraction with spectral token mixing, providing global coupling at $\mathcal{O}(N \log N)$ complexity, in contrast to the $\mathcal{O}(N^2)$ cost of full dot-product self-attention in standard Transformers (Vaswani et al., 2017). Fourier token mixing has been explored across modalities, from unparameterised Fourier mixing (FNet) to learnable global frequency filters (GFNet) and adaptive Fourier operators (AFNO/FNO) (Lee-Thorp et al., 2022; Rao et al., 2021; Guibas et al., 2021; Li et al., 2021). Here, we adapt these ideas to oversampled Bragg diffraction through a multiscale spectral design tailored to the structure of BCDI patterns.

We first validated Fourier ViT on large-scale synthetic datasets of Voronoi multi-domain crystals. Over the range of conditions studied, Fourier ViT achieves the lowest reciprocal-space mismatch (χ^2)

among the compared methods, while preserving domain-resolved phase reconstructions under realistic noise degradations. We further apply the method to experimental diffraction from $\text{La}_{2-x}\text{Ca}_x\text{MnO}_4$ (LCMO) (Robinson et al., 2026). Here, Fourier ViT outperforms a C-CNN baseline (Yu et al., 2024) and achieves χ^2 values competitive with the iterative benchmark. The run-to-run χ^2 distribution is broader across random initialisations, highlighting the multi-basin nature of strong-phase retrieval.

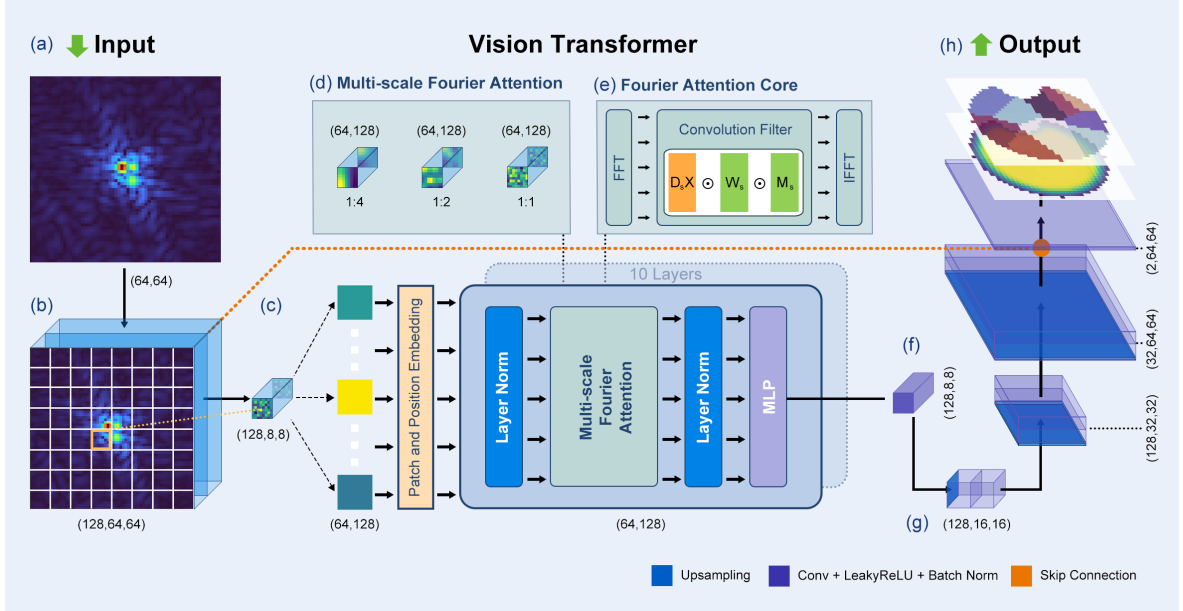


Fig. 1: Architecture of the proposed Fourier ViT model for BCDI phase retrieval. (a) The input diffraction magnitude (64×64 pixels) is passed through (b) a shallow convolutional feature extractor to produce 128 channels ($128 \times 64 \times 64$). (c) Each image is partitioned into patches and embedded as tokens (8×8 patches shown for clarity; all experiments use 16×16 patches), giving a token sequence of shape (64×128). This sequence is processed by (d) a multi-layer Vision Transformer with multi-scale Fourier attention (frequency mixing at 1:4, 1:2 and 1:1 scales) and (e) an FFT-based global convolution block. (f) The transformer output is reshaped into a latent feature map ($128 \times 8 \times 8$) and decoded by (g) a convolutional upsampling path with a skip connection to reconstruct the complex crystal field, yielding (h) 64×64 real-space amplitude and phase maps consistent with the measured diffraction pattern.

2 Results

2.1 Fourier ViT model architecture

2.1.1 Tokenisation and positional embedding

As shown by Fig.1 (a), the input is a 2D diffraction pattern magnitude with 64×64 pixels, normalised to $[0, 1]$. In panel (b), a shallow CNN extracts features at the original 64×64 resolution and outputs

a 128-channel feature map. We retain this as a full-resolution skip map $\mathbf{S} \in \mathbb{R}^{B \times 128 \times 64 \times 64}$, where B denotes the batch size. Patch embedding in Fig. 1(c) uses a single 4×4 convolution with stride 4 applied to \mathbf{S} , partitioning the image into non-overlapping 4×4 patches and projecting each patch to a token of embedding dimension $E=128$. This produces a 16×16 token grid, thus $N=16 \times 16=256$ tokens in total. The typical speckle/fringe spacing in our datasets is about two pixels after binning, so a 4×4 patch spans at least one full fringe period and allows each token to encode a physically meaningful local piece of the diffraction pattern. Following the Vision Transformer formulation (Dosovitskiy et al., 2021; Vaswani et al., 2017), we then flatten the 16×16 grid into a token sequence and refine it with two 1×1 convolutions before adding a learned positional embedding. The resulting sequence $\mathbf{X} \in \mathbb{R}^{B \times N \times E}$ is passed to the stack of Fourier transformer blocks described in Sec. 2.1.2.

2.1.2 Multi-scale Fourier attention

Inspired by Fourier-based transformers and spectral gating mechanisms (Lee-Thorp et al., 2022; Patro and Agneeswaran, 2023), we replace dot product self-attention (Vaswani et al., 2017) with a multi-scale Fourier attention, indicated by the green boxes in Fig.1 (d, e). The tokens are reshaped to a feature map $\mathbf{X} \in \mathbb{R}^{B \times E \times H \times W}$ with $H=W=16$ and processed at three spatial scales $s \in \{1, 2, 4\}$:

$$\mathbf{Z}_s = \mathcal{U}_s \left(\mathcal{F}^{-1} \left(\mathcal{F}(\mathcal{D}_s \mathbf{X}) \odot \mathbf{W}_s \odot \mathbf{M}_s \right) \right), \quad (1)$$

Here \mathcal{D}_s denotes average pooling with stride s , which reduces the spatial size to $h_s \times w_s = (H/s) \times (W/s)$, and \mathcal{U}_s upsamples back to (H, W) . The operator \mathcal{F} is the channel-wise 2D FFT, and \odot denotes element-wise multiplication in frequency space. In practice, we use three mixers at resolutions (H, W) , $(H/2, W/2)$ and $(H/4, W/4)$ and learn their relative importance via the scale weights α_s , giving $\mathbf{Z} = \sum_s \alpha_s \mathbf{Z}_s$ with $\alpha = \text{softmax}(\theta)$. At each scale, we learn per-channel frequency responses $\mathbf{W}_s \in \mathbb{R}^{E \times h_s \times w_s}$ and a channel-shared spectral gate $\mathbf{M}_s \in \mathbb{R}^{1 \times h_s \times w_s}$. This explicit multi-scale design allows the block to respond both to finely-spaced fringes and to broader low frequency components.

We use pre-normalised residual connections in each Fourier transformer block:

$$\mathbf{X}' = \mathbf{X} + \text{Mixer}(\text{LN}(\mathbf{X})), \quad \mathbf{X}_{\text{out}} = \mathbf{X}' + \text{MLP}(\text{LN}(\mathbf{X}')), \quad (2)$$

where LN represents layer normalisation (Ba et al., 2016; Dosovitskiy et al., 2021), Mixer denotes the multi-scale Fourier spectral mixer defined above, and multilayer perceptron (MLP) is a feed-forward projection with hidden dimension d_{mlp} and ReLU nonlinearity. The residual paths improve gradient flow, while pre-normalisation stabilises optimisation for deeper stacks of blocks. Additional details, including the convolution theorem and complexity analysis, are given in the Supplementary Sec. S.1.

2.1.3 Decoder

After the final transformer block, the output sequence $\mathbf{X} \in \mathbb{R}^{B \times N \times E}$ (with $N=16 \times 16$ and $E=128$) is reshaped to an $E \times 16 \times 16$ feature map as shown in Fig.1 (f) and upsampled to 64×64 by a small CNN decoder (Fig.1 (g)). At full resolution the decoder fuses three inputs: (i) the upsampled token features, (ii) the encoder skip map $\mathbf{S} \in \mathbb{R}^{B \times 128 \times 64 \times 64}$, which carries local diffraction pattern context from the CNN feature extractor, and (iii) a frequency-space summary of \mathbf{S} obtained by taking its channel-wise Fourier magnitude, applying $\log(1 + |\mathcal{F}\{\mathbf{S}\}|)$ to compress the dynamic range, and normalising per channel (see Supplementary Sec. S.2 for details). The three streams are concatenated

along the channel dimension and passed through 3×3 convolutions for channel mixing, followed by a short refinement stack. This fusion supplies the decoder with complementary information: (i) non-local diffraction correlations from the transformer, (ii) fine scale speckle features from the encoder, and (iii) an explicit spectral summary of the learned diffraction features.

The decoder heads output a real valued crystal amplitude map $\hat{A}(x, y)$ and two channels $(\hat{c}(x, y), \hat{s}(x, y))$ that parametrise the crystal phase. The amplitude head uses a softplus nonlinearity to enforce non negativity, and both amplitude and phase are multiplied by a fixed real-space support mask $S(x, y)$, in line with standard CXDI support constraints (Fienup, 1982; Marchesini et al., 2003). For experimental data, Fourier ViT uses the same support as the MATLAB iterative reconstruction (see Sec. 4.4) for a fair comparison. Inside the support, the predicted amplitude is blended with a simple prior via a scalar schedule $\alpha(t)$ and then softly rescaled so that its ℓ_2 norm matches that of the prior amplitude; the precise blending and normalisation rules are given in Supplementary Sec. S.2. For the phase, the two decoder outputs are first normalised onto the unit circle, $(\tilde{c}, \tilde{s}) = (\hat{c}, \hat{s})/\sqrt{\hat{c}^2 + \hat{s}^2}$, and the phase is recovered as $\phi(x, y) = \text{atan2}(\tilde{s}(x, y), \tilde{c}(x, y))$.

Given an input diffraction pattern magnitude $M(\mathbf{q})$, the network therefore outputs a complex real-space density $\rho(x, y)$ on a 64×64 grid, with amplitude and phase constrained by the fixed support.

2.1.4 Loss function

For a given input diffraction pattern, the network predicts a complex real-space density $\rho(x, y)$ from the crystal’s amplitude and phase. This estimate is propagated through the BCDI forward model to obtain a predicted diffraction magnitude $M_{\text{pred}}(\mathbf{q})$ (Sec. 2.1.3), which is compared with the measured magnitude $M_{\text{meas}}(\mathbf{q})$ via a hybrid Fourier-space loss. The loss combines Pearson correlation coefficient (PCC), an root mean square (RMS)-normalised χ^2 , a power weighted χ^2 term and a small total variation (TV) regulariser, with epoch-dependent weights that shift the emphasis from global pattern correlation to fine scale intensity agreement as training proceeds.

For the χ^2 -type terms we work with an RMS-normalised magnitude $\hat{M}(\mathbf{q}) = \frac{M(\mathbf{q})}{\sqrt{\langle M(\mathbf{q})^2 \rangle_{\mathbf{q}}}}$, applied separately to M_{meas} and M_{pred} so that any global intensity scale is removed. At training epoch t we minimise,

$$L_{\text{total}}(t) = w_{\text{PCC}}(t) L_{\text{PCC}} + w_{\chi^2}(t) L_{\chi^2} + w_{\text{p}\chi^2}(t) L_{\text{p}\chi^2} + \lambda_{\text{TV}}(t) L_{\text{TV}}, \quad (3)$$

where the data-fidelity weights $w_{\text{PCC}}(t)$, $w_{\chi^2}(t)$, $w_{\text{p}\chi^2}(t)$ and the TV weight $\lambda_{\text{TV}}(t)$ follow simple epoch-dependent schedules designed so that PCC dominates early training, power- χ^2 ramps up over the early epochs to enhance low frequencies, and standard χ^2 is slightly favoured in the late refinement stage, while TV remains a small stabilising term. The full schedule definitions are given Supplementary Sec. S.3.

(i) **PCC loss.** To enforce global similarity of the diffraction patterns we use

$$L_{\text{PCC}} = 1 - \rho(M_{\text{meas}}, M_{\text{pred}}), \quad (4)$$

with, for a single pattern,

$$\rho(M_{\text{meas}}, M_{\text{pred}}) = \frac{\sum_{\mathbf{q}} (M_{\text{meas}}(\mathbf{q}) - \mu_{\text{meas}})(M_{\text{pred}}(\mathbf{q}) - \mu_{\text{pred}})}{\sqrt{\sum_{\mathbf{q}} (M_{\text{meas}}(\mathbf{q}) - \mu_{\text{meas}})^2} \sqrt{\sum_{\mathbf{q}} (M_{\text{pred}}(\mathbf{q}) - \mu_{\text{pred}})^2}}, \quad (5)$$

where $\mu_{\text{meas}} = \langle M_{\text{meas}}(\mathbf{q}) \rangle_{\mathbf{q}}$, $\mu_{\text{pred}} = \langle M_{\text{pred}}(\mathbf{q}) \rangle_{\mathbf{q}}$.

(ii) RMS-normalised χ^2 loss. To penalise absolute mismatches while remaining insensitive to any residual scale factor, we define

$$L_{\chi^2} = \frac{\sum_{\mathbf{q}} [\hat{M}_{\text{meas}}(\mathbf{q}) - \hat{M}_{\text{pred}}(\mathbf{q})]^2}{\sum_{\mathbf{q}} \hat{M}_{\text{meas}}(\mathbf{q})^2}, \quad (6)$$

where \hat{M}_{meas} and \hat{M}_{pred} are the RMS-normalised magnitudes defined above.

(iii) Power- χ^2 loss. To give additional emphasis to bright and high- q fringes, we include a power-weighted χ^2 term. Starting from the RMS-normalised magnitudes, we form powered magnitude of measured and predicted diffraction pattern, $\hat{M}_{\text{meas}}^{(p)}(\mathbf{q}) = [\hat{M}_{\text{meas}}(\mathbf{q})]^{p(t)}$ and $\hat{M}_{\text{pred}}^{(p)}(\mathbf{q}) = [\hat{M}_{\text{pred}}(\mathbf{q})]^{p(t)}$, then define

$$L_{p\chi^2} = \frac{\sum_{\mathbf{q}} [\hat{M}_{\text{meas}}^{(p)}(\mathbf{q}) - \hat{M}_{\text{pred}}^{(p)}(\mathbf{q})]^2}{\sum_{\mathbf{q}} [\hat{M}_{\text{meas}}^{(p)}(\mathbf{q})]^2}. \quad (7)$$

The exponent $p(t)$ decays smoothly from about 2 at early epochs towards 0.5 later in training, so that optimisation initially focuses on fitting the brightest regions and progressively places more weight on weaker fringes.

(iv) Total-variation amplitude regularisation. Let $S(x, y) \in \{0, 1\}$ be the binary support mask and $N_S = \sum_{x,y} S(x, y)$ the number of support pixels. We define an anisotropic TV penalty

$$L_{\text{TV}} = \frac{1}{N_S} \sum_{x,y} S(x, y) (|\nabla_x A(x, y)| + |\nabla_y A(x, y)|), \quad (8)$$

where $\nabla_x A(x, y) = A(x + 1, y) - A(x, y)$ and $\nabla_y A(x, y) = A(x, y + 1) - A(x, y)$. In implementation, we mask the amplitude by $S(x, y)$ before taking finite differences, so the gradients across the support boundary do not contribute.

2.2 Performance on synthetic data

On synthetic multi-domain crystals, the Fourier ViT recovers the real-space object from reciprocal-space diffraction magnitude, without real-space supervision or support updates. At 64×64 diffraction resolution with known crystal amplitude, the method resolves up to 19 domains with runs that achieve perfect convergence, here defined as $\chi^2 \leq 10^{-5}$. Perfect reconstructions occur 42 times out of 100 runs for 10 domains, 18 times for 15 domains, and 4 times for 19 domains. When crystal amplitude and phase are recovered jointly, convergence is harder but still feasible: for a 10-domain crystal, 12 out of 100 runs reach $\chi^2 \leq 10^{-5}$ after 1000 epochs.

To illustrate object amplitude-phase retrieval, we show a representative 10-domain case in Fig. 2. Figs 2 (a,b) compare the ground truth and Fourier ViT predicted diffraction magnitudes, while Fig. 2(c) shows the radially averaged log-intensity profile of the diffraction pattern as a function of momentum transfer q (radial coordinate in reciprocal-space). The profile matches the ground truth

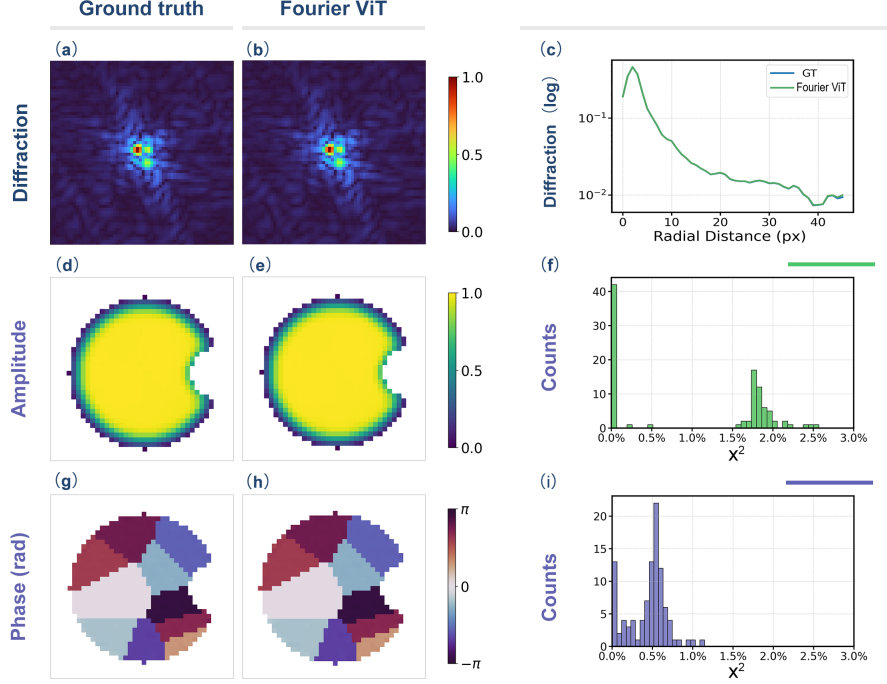


Fig. 2: Fourier ViT reconstruction of a synthetic 10-domain crystal from diffraction only. (a,d,g) Ground truth diffraction magnitude (64×64 pixels), crystal amplitude and phase (40×40 pixels); (b,e,h) corresponding Fourier ViT reconstructions. (c) Radial log intensity against q profiles of ground truth and reconstructed diffraction, showing agreement across the full q -range. (f,i) Histograms of diffraction χ^2 over 100 random initialisations for known crystal amplitude (f) and unknown amplitude (i) phase retrieval.

across the full q range, including the high- q tails produced by strong-phase contrasts and sharp domain walls. Instead of a single central peak, familiar in the weak-phase case, this pattern has multiple centres because the object has highly varying phase values. Figs 2 (d,e) demonstrate that the recovered crystal amplitude is smoothly matching the ground truth, and Figs. 2 (g,h) show that the phase solution recovers the real-space domain configurations with sharp, correctly placed boundaries and no spurious domains. The ground truth panels are used only for quantitative assessment and visualisation, as the network is trained unsupervised from diffraction data without real-space labels. For visual comparison, we fix the global phase offset by setting the phase at the central pixel to zero in all reconstructions.

Fig. 2(f) shows a bimodal distribution over 100 random initialisations in the phase-only setting (known amplitude). Many runs converge to low- χ^2 solutions with $\chi^2 \leq 10^{-5}$, while the remainder settle into higher-error local minima, forming a second cluster at $\chi^2 \approx (1.5-2.5) \times 10^{-2}$. This behaviour reflects the non-convex strong-phase retrieval landscape: small shifts in domain wall placement or topology can produce a structured mismatch in the high- q fringe field, and with fixed amplitude the

optimisation has limited freedom to compensate, leaving stable nonzero residuals. In joint amplitude-phase recovery, the extra degrees of freedom from amplitude prediction reduce the number of perfectly converged runs but concentrate the remaining outcomes into a narrower band of lower χ^2 . Here, partially incorrect phase maps can be offset by amplitude reweighting within the support to better match the diffraction magnitude.

These results show that reciprocal-space modelling with Fourier attention reproduces the diffraction structure and recovers the multi-domain phase topology without real-space labels. Unless otherwise stated, all Fourier ViT models were trained in an unsupervised manner using the optimisation and loss schedule described in Sec. 4.2.

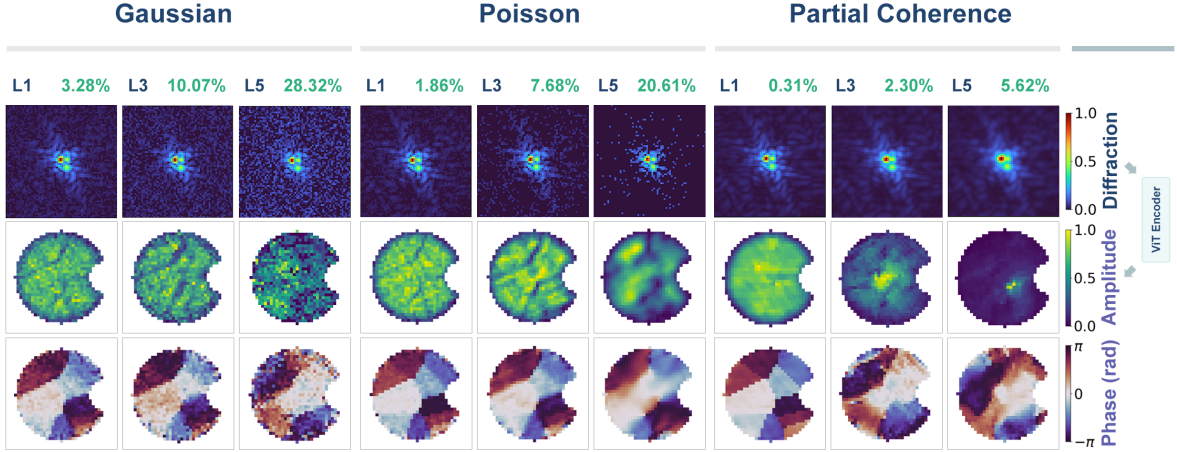


Fig. 3: Noise robustness of the Fourier ViT. Columns show three types of noise applied to the same diffraction pattern: Gaussian noise, Poisson noise, and partial coherence. For each model, L1, L3 and L5 denote increasing noise strengths (levels 1, 3 and 5); green percentages give χ_n^2 (in %), the input diffraction error relative to the clean ground truth. Top row shows the noisy input diffraction amplitudes (64×64 pixels). Middle and bottom rows show the reconstructed crystal amplitude and phase zoomed in 40×40 pixels, respectively; both are masked to the support and phase is in radians.

For the same simulated diffraction pattern of a 10-domain crystal and fixed support, iterative methods combining Error Reduction (ER) and Hybrid Input-Output (HIO) reach χ^2 values between 10^{-4} and 10^{-3} for approximately half of reconstructions after 500 iterations. No runs reach our “perfect” criterion $\chi^2 \leq 10^{-5}$ with the iterative method, but increasing to 2000 iterations raises the fraction of runs with $10^{-4} \leq \chi^2 \leq 10^{-3}$ to 90%. Under matched initialisation and support constraints, the Complex Convolutional Neural Network (C-CNN) with skip connection achieves much higher χ^2 values of 6×10^{-3} – 1.2×10^{-2} . Among all three methods, the χ^2 outcomes over random initialisations are bimodal; however, only the Fourier ViT reaches the lowest- χ^2 value ($\chi^2 \leq 10^{-5}$), whereas the C-CNN remains confined to a higher-error mode.

Table 1: Noise robustness of the Fourier ViT (magnitude domain). All values are normalised χ^2 errors in %. χ_n^2 compares the noisy input to the clean ground truth; $\chi_{\text{rec},n}^2$ and $\chi_{\text{rec},c}^2$ compare the reconstruction to the noisy input and to the clean ground truth, respectively. Noise strength increases from level 1 to level 5.

Noise	Gaussian			Poisson			Partial coherence		
	χ_n^2	$\chi_{\text{rec},n}^2$	$\chi_{\text{rec},c}^2$	χ_n^2	$\chi_{\text{rec},n}^2$	$\chi_{\text{rec},c}^2$	χ_n^2	$\chi_{\text{rec},n}^2$	$\chi_{\text{rec},c}^2$
Level 1	3.28	1.76	1.54	1.86	1.18	0.92	0.31	0.13	0.25
Level 2	6.55	3.48	2.71	4.15	2.67	1.92	1.50	0.25	1.94
Level 3	10.07	5.16	4.14	7.68	4.69	3.65	2.30	0.23	2.78
Level 4	18.84	9.03	8.21	12.90	8.07	5.82	3.95	0.17	4.66
Level 5	28.32	12.49	14.36	20.61	11.91	10.53	5.62	0.11	6.37

2.3 Noise robustness

To assess robustness under realistic experimental degradations, the same noise-free diffraction pattern used in Fig. 2 was corrupted by three noise models: additive Gaussian noise, accounting for detector background; Poisson counting noise representing counting statistics; and partial coherence of the incident beam as demonstrated in Fig. 3. Strictly, partial coherence is not the same as the statistical noise models because it can be corrected (Clark et al., 2012), but its effects on convergence are worth identifying so they can be used for identification. Each noise model was applied at five strengths, from Level 1 to Level 5, according to its effect on χ^2 . For visual clarity, Fig. 3 shows Levels 1, 3 and 5; Levels 2 and 4 are provided in Supplementary Sec. 5 (Fig. 6).

For each noise type, Table 1 reports three normalised diffraction errors: (i) the noisy input relative to the clean ground truth, χ_n^2 ; (ii) the reconstruction mismatch to the noisy input, $\chi_{\text{rec},n}^2$; and (iii) the reconstruction error relative to the clean ground truth, $\chi_{\text{rec},c}^2$. In our unsupervised setting, we select the best run across random initialisations by minimising $\chi_{\text{rec},n}^2$, and report the corresponding $\chi_{\text{rec},c}^2$ as evaluation against the clean ground truth.

For Gaussian noise, χ_n^2 increases from 3.28% to 28.32% with noise strength. The Fourier ViT reduces the error relative to the clean diffraction to $\chi_{\text{rec},c}^2 = 1.54$ –14.36%, corresponding to a 49–59% reduction relative to χ_n^2 . For Poisson noise, χ_n^2 increases from 1.86% to 20.61%, and the model reduces this to $\chi_{\text{rec},c}^2 = 0.92$ –10.53%, giving a comparable 49–55% reduction. Across both stochastic noise processes, the reconstruction remains closer to the clean diffraction than the corrupted input itself, demonstrating genuine denoising rather than simply reproducing the noise statistics.

Fig. 3 illustrates the effects of Gaussian and Poisson noise on the real-space images. At low noise levels, the fringes of the diffraction remain sharp and the reconstructions show clean, domain-resolved crystal phase maps, whereas stronger noise introduces pixel-to-pixel fluctuations and visibly corrupts the crystal phase, eventually losing the clear boundaries of the domains. Under Gaussian noise, pixel-scale intensity fluctuations developed in both the reconstructed crystal’s amplitude and phase, while under Poisson statistics the amplitude develops grain boundary contrast that mirrors the underlying domain structure. In both cases, the Fourier ViT suppresses unstructured speckle and background fluctuations while preserving the physically meaningful fringe pattern around the centre of the Bragg

peak, thereby enabling robust phase retrieval even under strongly degraded conditions. Consistent with Table 1, the Fourier ViT therefore acts like a filter, reducing the diffraction domain error by roughly a factor of two in χ^2 relative to the noisy input.

For the partial coherence condition, χ_n^2 increases from 0.31% to 5.62% with increasing blur. The reconstruction matches the blurred measurement extremely closely as $\chi_{\text{rec},n}^2 = 0.11\text{--}0.25\%$, but the error relative to the clean diffraction lies in the range $\chi_{\text{rec},c}^2 = 0.25\text{--}6.37\%$. The ratio $\chi_{\text{rec},c}^2/\chi_n^2$ is ≈ 0.81 at the weakest level, but increases to $\approx 1.13\text{--}1.29$ at stronger blur. The diffraction fringes are progressively blurred and power is concentrated near the central Bragg peak. The reconstructed amplitude shows an over-localised central intensity (“hot spot” region) and a smoothed phase, reflecting the loss of high- q information. Partial coherence is therefore a natural physical origin for the central hot spot in the reconstructed crystal amplitude, a behaviour documented by (Vartanyants and Robinson, 2001) and also observed in the experimental reconstructions discussed in Sec. 2.4.

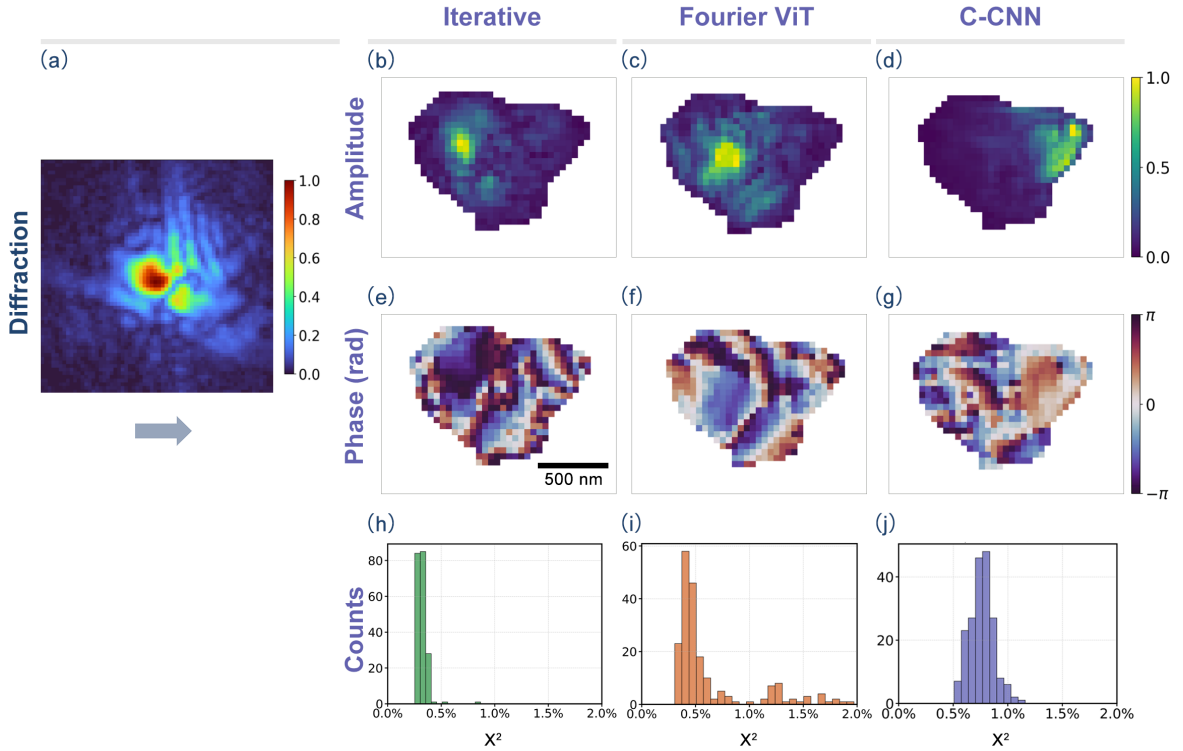


Fig. 4: Comparison study of phase retrieval on experimental data from a multi-domain LCMO-500 nanocrystal. (a) Measured BCDI diffraction magnitude (64×64 pixels). (b–d) Reconstructed crystal’s amplitude and (e–g) phase from the iterative method, Fourier ViT, and C-CNN. Scale bar is 500 nm. (h–j) Corresponding χ^2 histograms over 200 runs per method, each run with a random initial phase. χ^2 range is between 0.0% and 2.0%.

2.4 Experimental verification

As a preliminary experimental check, we tested the Fourier ViT on a 2D diffraction pattern from a weak-phase, single-domain 450 nm SrTiO₃ (STO) nanocrystal; results are shown in Supplementary Sec. S.5. In Fig. 7(a,b), the reconstructions recover a single-domain morphology across runs. Across 100 random initialisations, the final reciprocal-space mismatch χ^2 falls within a narrow band ($\sim 0.35\%$ – $\sim 0.68\%$; Fig. 7(c)), indicating reliable convergence in the weak-phase regime. This contrasts with the strong-phase, multi-domain case discussed next, where multiple local minima lead to a substantially broader spread of χ^2 values across runs.

We then consider the strong-phase regime using an experimental BCDI pattern from a strongly distorted, multi-domain LCMO-500 nanocrystal (Sec. 4.3). Fig. 4 compares reconstructions from a traditional iterative method (ER/HIO), the Fourier ViT, and a C-CNN baseline, using an identical real-space support. For each method, we report the best run (lowest χ^2) over 200 random initialisations.

For the raw measured diffraction magnitude in Fig. 4(a), the iterative reconstruction (Fig. 4(b,e)) recovers the multi-domain phase structure, but shows isolated hot spots within the amplitude map. The best result reaches $\chi^2 = 0.25\%$ with PCC = 99.77% against the measured magnitude. Trained unsupervised on the same single diffraction pattern and constrained by the same support, the Fourier ViT returns the amplitude and phase in Fig. 4(c,f). The recovered domain pattern is consistent with the iterative solution, with clearer domain boundaries and more spatially coherent phase regions, and fewer isolated high-contrast features. The best Fourier ViT run attains $\chi^2 = 0.30\%$ and PCC = 99.79%. By contrast, the best C-CNN result in Fig. 4(d,g) satisfies the support constraint but converges to a different solution: the amplitude concentrates near the support boundary and the phase topology deviates from both the iterative and Fourier ViT reconstructions. This behaviour is accompanied by a higher $\chi^2 = 0.50\%$, suggesting a poorer match to the measured diffraction magnitude under the same forward model.

To remove the hot spot in the crystal amplitude, we additionally tested a mild partial-coherence correction by Gaussian deconvolution of the measured diffraction magnitude (Vartanyants and Robinson, 2001). While the deconvolution visibly sharpens the fringes, it also amplifies other detector noises, and the reconstruction χ^2 increases relative to the raw-data case: iterative gives $\chi^2 = 0.61\%$, Fourier ViT has $\chi^2 = 0.57\%$, without significantly improving the hot spot. We therefore use the raw measured pattern as the primary basis for comparison in Fig. 4.

The χ^2 histograms in Fig. 4(h–j), plotted on a common axis, summarise run-to-run variability over 200 independent runs per method. Iterative reconstructions (h) cluster tightly within $\chi^2 \approx 0.25$ – 0.87% , indicating stable convergence. The Fourier ViT (i) achieves comparable best-case values $\chi^2 = 0.3\%$ but shows a broader spread (~ 0.3 – 2.0%), while the C-CNN distribution (j) is shifted to higher errors (~ 0.5 – 1.1%). Supplementary Movies M1–M3 visualise the per-epoch evolution for the three methods and highlight distinct optimisation trajectories.

Overall, on this experimental dataset, the Fourier ViT matches the measured diffraction at least as well as the best iterative reconstruction and outperforms the C-CNN baseline in χ^2 , while producing multi-domain crystal’s amplitude and phase that remain consistent with the iterative solution under identical support constraints.

3 Discussion

We developed a Vision Transformer with Fourier attention (Fourier ViT) for multi-domain phase retrieval in BCDI. On synthetic 64×64 diffraction patterns, the model recovers strong-phase domain

structure directly from diffraction in an unsupervised setting. For phase-only prediction with known amplitude, runs can reach essentially perfect agreement with the target diffraction ($\chi^2 \leq 10^{-5}$) while still resolving complex domain structure up to 19 domains (Fig. 2f). Perfect reconstructions occur 42/100 times for 10 domains, 18/100 for 15 domains, and 4/100 for 19 domains, indicating a low- χ^2 solution class and competing local minima in the strong-phase regime. For joint amplitude-phase prediction, a representative 10-domain case reproduces the ground-truth diffraction magnitude over the full q range, including the high- q tails (Fig. 2c), and the reconstructed crystal phase recovers the same domain layout with sharp domain walls (Fig. 2h); 12/100 runs reach $\chi^2 \leq 10^{-5}$ after 1000 epochs. Under the same fixed support on the same pattern, ER/HIO reaches $\chi^2 = 10^{-4}$ – 10^{-3} for approximately half of reconstructions after 500 iterations, and with sufficient iterations a larger fraction of runs fall into this range. The complex convolutional neural network (C-CNN) remains at higher χ^2 error (6×10^{-3} – 1.2×10^{-2}) under the same initial setup and does not reach the $\chi^2 \leq 10^{-5}$ regime achieved by Fourier ViT with sufficient iterations. Across all three methods, random initialisation in the strong-phase regime consistently separates runs into two outcome families under the same constraints: a low- χ^2 set and a second, higher- χ^2 set.

Robustness tests in Fig. 3 and Table 1 show that, under Gaussian and Poisson noise, the reconstructions remain closer to the clean diffraction than the noisy input across all levels ($\chi_{\text{rec,c}}^2 < \chi_{\text{n}}^2$), with $\chi_{\text{rec,c}}^2$ reduced by roughly a factor of two relative to χ_{n}^2 . This indicates genuine denoising in the diffraction magnitude. By contrast, partial coherence primarily blurs the diffraction and reduces fringe visibility. Although the model can fit the blurred measurement extremely closely, the reconstructed diffraction departs from the clean target as the blur increases. Partial coherence can also produce an over-localised central amplitude “hot spot”, a behaviour documented previously (Vartanyants and Robinson, 2001) and observed in our experimental reconstructions (Sec. 2.4). Therefore, an explicit coherence model is needed when the blur is significant. All three distortions appear differently in the predicted images: Gaussian noise introduces noisy local features, Poisson noise produces grainy fluctuations, and partial coherence blurs the fringes and produces less well-defined domain boundaries.

We then evaluated the method on measured diffraction from the LCMO-500 nanocrystal against ER/HIO and the C-CNN, using the same support constraint and the same FFT-based forward model in the loss. The three approaches show distinct optimisation behaviour as demonstrated in Supplementary Movies M1–M3. In ER/HIO, the negative feedback step helps the iterations escape local minima and reach low- χ^2 solutions. The C-CNN tends to build the phase map from smooth initial fields toward sharper domain boundaries, adding spatial detail gradually during training; in our tests, it more often converges to solutions with edge-localised amplitude and a phase topology that differs from the other methods. By contrast, the Fourier ViT couples the diffraction-pattern features globally through token mixing and produces multi-domain configurations early in training. This reduces time spent in over-smoothed intermediate states and moves the optimisation quickly into the relevant strong-phase solution class. In Fig. 4, the Fourier ViT reaches χ^2 values comparable to the best iterative reconstruction (iterative: $\chi^2 \approx 0.25\%$, Fourier ViT: $\chi^2 \approx 0.30\%$) and improves over the C-CNN baseline ($\chi^2 \approx 0.50\%$). Under the same support constraint, the recovered domain morphology is consistent with the ER/HIO result.

The reconstruction error χ^2 histograms in Fig. 4 (h-j) also highlight a practical aspect of the strong-phase regime. With realistic noise and unavoidable model mismatch under a magnitude-only data constraint, many near-degenerate minima can fit the measured diffraction almost equally well. As a result, reconstructions may be nearly indistinguishable in reciprocal-space while remaining meaningfully different in real-space. Therefore, reducing χ^2 is necessary but not sufficient to identify a

unique domain topology. The broader spread of Fourier ViT outcomes is thus more naturally interpreted as rapid access to multiple acceptable strong-phase solutions enabled by global token mixing, rather than as optimisation instability. Practically, we report the best- χ^2 reconstructions and verify that key domain-level features are reproducible across independent runs when the measured intensity and imposed constraints do not enforce a unique solution.

Computationally, ER/HIO is dominated by repeated FFT projections, typically one forward and one inverse FFT per iteration (Fienup, 1982; Duhamel and Vetterli, 1990). In the unsupervised learning condition, the FFT forward model is embedded in the loss alongside network forward/backward propagation for both the C-CNN and Fourier ViT. Within the Fourier ViT, Fourier token mixing enables global coupling while avoiding the quadratic scaling of dot-product self-attention (Vaswani et al., 2017; Lee-Thorp et al., 2022).

Overall, Fourier ViT offers a practical route to strong-phase, multi-domain BCDI reconstructions. In near-ideal conditions, all methods improve with additional iterations or epochs, but for a fixed optimisation budget the Fourier ViT reaches structured multi-domain solutions earlier and outperforms the C-CNN baseline, while remaining comparable to the best ER/HIO reconstructions on experimental data. Future work will incorporate explicit noise and partial-coherence models in the forward operator, develop uncertainty-aware selection from solution ensembles, learn how to derive the support from the diffraction data and extend the approach to fully 3D experimental datasets.

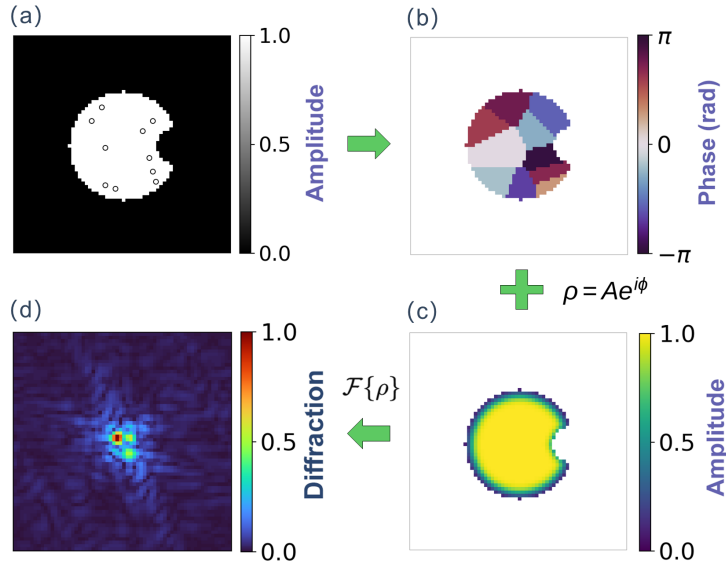


Fig. 5: Simulation pipeline for synthetic crystals and diffraction patterns. (a) Random seed placement within the support. (b) Voronoi domain generation with assigned phases, wrapped to $[-\pi, \pi]$. (c) Construction of the crystal amplitude mask, normalised to $[0, 1]$. (d) Corresponding diffraction magnitude computed by fast Fourier transform (FFT), normalised to $[0, 1]$. All images are 64×64 pixels.

4 Methods

4.1 Synthetic Data and Noise

4.1.1 Voronoi multi-domain crystals

To simulate the heterogeneous structure of functional nanocrystals, we constructed two-dimensional crystals made up of multiple phase domains with sharp, straight boundaries. In Fig. 5(a), seed points were randomly placed inside a soft-edged circular support with a crescent cutout, representing the particle shape, and each seed was assigned a random phase in $[-\pi, \pi]$. The soft support suppresses artificial high-frequency components in Fourier space that would arise from box-shaped boundaries, while the crescent cut breaks symmetry and helps to avoid twin solutions. A minimum spacing of 2 pixels was enforced between seeds to prevent clustering. Each seed then generated a Voronoi cell whose pixels take the phase of the closest seed, producing a tessellation of non-overlapping domains with straight boundaries. This Voronoi-based construction mimics the abrupt domain walls typically observed in ferroelectric and ferroelastic materials (Setter et al., 2006).

In Fig. 5(b), we set the central domain to $\phi = 0$ to remove the global phase offset. We required $\Delta\phi > 0.5$ rad between neighbouring domains to enforce strong-phase contrast and avoid inversion-symmetric cases. Fig. 5(c) shows the crystal amplitude $A(x, y)$, chosen to represent the 2D projected amplitude of a finite 3D particle. We use a flat-core radial profile with Gaussian edge smoothing, modulated by the crescent mask, to produce a confined but soft-edged support. The final real-space object is defined as the complex density

$$\rho(x, y) = A(x, y) e^{i\phi(x, y)}, \quad (9)$$

with $A(x, y)$ fixed and $\phi(x, y)$ set by the Voronoi partition.

In Fig. 5(d), the diffraction intensity is obtained as the modulus squared of the Fourier transform of the complex density,

$$I_{\text{pred}}(\mathbf{q}) = |\mathcal{F}\{\rho(x, y)\}(\mathbf{q})|^2, \quad (10)$$

on a 64×64 grid containing a 32-pixel-diameter support, corresponding to an oversampling ratio of 2 in reciprocal-space. Outside the support, $\rho(x, y)$ is set to zero so that only physically allowed regions contribute to the scattering. The crystal amplitude, phase and diffraction pattern are stored as 16-bit TIFF images. Overall, this pipeline produces synthetic data with a controlled number of domains, strong-phase boundaries and realistic reciprocal-space structure, which we use to benchmark Fourier ViT phase retrieval networks.

4.1.2 Noise tests

We assess robustness in the regime where the network predicts both crystal amplitude and phase with noise added to the diffraction pattern. Starting from a normalised diffraction pattern intensity $I(q_x, q_y) \in [0, 1]$, we construct three families of corrupted data: Poisson shot noise, Gaussian read noise and partial coherence.

For the two detector-noise models we work on the exposure (photon-count) scale. In this setting, read noise from the sensor chain is well approximated by additive, zero-mean Gaussian noise in counts, whereas photon shot noise is described by Poisson statistics (Goodman, 2015; Janesick, 2001, 2007). Given a conversion factor $S > 0$ from normalised intensity to expected counts, the ideal noiseless

counts are $C_0(q_x, q_y) = S I(q_x, q_y)$. Gaussian read noise is then modelled as

$$C_g(q_x, q_y) = C_0(q_x, q_y) + \eta(q_x, q_y), \quad (11)$$

where the random variable $\eta(q_x, q_y)$ is drawn independently at each pixel from a normal distribution with zero mean and variance σ_c^2 . Poisson shot noise is modelled as

$$C_p(q_x, q_y) \sim \text{Poisson}(C_0(q_x, q_y)), \quad (12)$$

reflecting the discrete, independent arrival of photons at each pixel. The function of Poisson is generated using the Poisson random number generator in the NumPy array library (Harris et al., 2020). In both cases we enforce the physical non-negativity of the recorded counts by truncating $C(q_x, q_y) \geq 0$. For analysis we return to an intensity scale via $\tilde{I}(q_x, q_y) = C(q_x, q_y)/S$, and use the corresponding diffraction magnitude $\tilde{M}(q_x, q_y) = \sqrt{\tilde{I}(q_x, q_y)}$ for all reported image space metrics and for the reconstruction loss χ_{rec}^2 in Table 1.

Partial coherence is treated as a degradation of the diffraction intensity rather than a stochastic detector noise. We model it as a contrast-reducing blur

$$I_{\text{pc}}(q_x, q_y) = \mathcal{G}_\sigma * I(q_x, q_y), \quad (13)$$

where \mathcal{G}_σ approximates the far-field mutual-coherence function as a Gaussian in reciprocal space, so that partial coherence is modelled as a convolution of the coherent intensity (Williams et al., 2007; Clark et al., 2012). For simulated data, we convolve the coherent diffraction with \mathcal{G}_σ and then renormalise I_{pc} to match the clean pattern’s overall scaling, so that changes in reconstruction quality primarily reflect the loss of high- q fringe visibility rather than intensity rescaling. For experimental data (Sec. 2.4), we apply the same Gaussian model in reverse using a regularised deconvolution to partially undo the blurring and obtain a more nearly coherent effective diffraction pattern.

4.2 Fourier ViT-network training

Fourier ViT operates on 64×64 max-normalised diffraction magnitudes. With a patch size of 4, this gives 16×16 tokens and an embedding dimension of 128 processed through 10 Fourier transformer blocks. The networks are trained for 1000 epochs with batch size 1 using AdamW (Loshchilov and Hutter, 2019) (initial learning rate 5×10^{-4} , weight decay 10^{-5} , $\beta_1 = 0.9$, $\beta_2 = 0.999$) with a cosine-annealed learning rate schedule down to 10^{-6} . Training uses automatic mixed precision on an NVIDIA Quadro RTX 5000 (Max-Q).

The hybrid loss function of Sec. 2.1.4 is used throughout. A fixed binary support is imposed; the amplitude is initialised from a Gaussian prior within the support and kept fixed for the first 10 epochs, before being learned jointly with the phase via the blending schedule $\alpha(t)$ (Supplementary Sec. S.2). Dropout is set to 0, and we apply gradient clipping with a maximum gradient norm of 1.0 to stabilise optimisation.

For each domain count, we perform 100 independent runs with different random seeds and report statistics over these runs. We select the best run across random initialisations by minimising the diffraction mismatch χ^2 between the predicted and measured diffraction magnitudes, and report the corresponding reconstruction metrics, including χ^2 against the clean target when available.

4.3 BCDI experiment

We measured BCDI data from a crystal of $\text{La}_{0.5}\text{Ca}_{0.5}\text{MnO}_3$ (LCMO). LCMO is an important complex transition-metal oxide showing colossal magnetoresistance (CMR) close to room temperature, in which large changes in electrical resistance are induced by small changes in magnetic field. We are interested in whether the phase domains play a role in the CMR behaviour. LCMO was prepared by chemical synthesis in the form of nanocrystals with sizes in the range of 300 nm, aggregated into clusters (Robinson et al., 2026). These were attached to silicon wafer substrates for measurement by BCDI using the tetraethyl orthosilicate (TEOS) bonding method (Monteforte et al., 2016): a dilute solution of TEOS was used to coat the dry powder, and it was then calcined at 500 °C in air using a furnace for several hours.

X-ray measurements were made at beamline 34-ID-C of the Advanced Photon Source (APS) (Robinson et al., 2026). The monochromatic X-ray beam, with wavelength $\lambda = 0.138$ nm, was made coherent with a $50\ \mu\text{m}$ (vertical) \times $20\ \mu\text{m}$ (horizontal) slit, 48 m from the source, then focused with Kirkpatrick–Baez (KB) mirrors to a beam size of $600\ \text{nm} \times 800\ \text{nm}$ at the centre of a precise air-bearing rotation axis, θ . A single nanocrystal was selected by rotating θ until an isolated peak appeared on the Medipix area detector, with $55\ \mu\text{m}$ pixels, positioned at the Bragg angle of the (110) Bragg peak, at a distance $D = 2.2$ m away from the sample. A single frame from that 3D data set, LCMO-500, was selected for the study. The corresponding image is a phased projection of the 3D nanocrystal structure, whose Fourier transform magnitude squared (plus noise) should match exactly the measured frame, according to the Fourier-slice theorem.

4.4 Iterative method

We implement a conventional 2D iterative phase retrieval algorithm in MATLAB following standard BCDI practice (Robinson and Harder, 2009; Clark et al., 2012). The initial support is defined relative to the reconstruction grid, with $s_x = 0.5 N_x$ and $s_y = 0.5 N_y$. We then use a predefined switching schedule that alternates between 20 iterations of error-reduction (ER) and 80 iterations of hybrid input-output (HIO) algorithms, with HIO relaxation parameter $\beta = 0.9$ (Fienup, 1982). The support was refined using shrink-wrap updates (Marchesini et al., 2003) with a normalised amplitude threshold = 0.15. During ER, we enforce the crystal’s amplitude to zero outside of the support and keep phase within $[-\pi, \pi]$. Reconstructions of LCMO-500 used 300 iterations with shrinkwrap starting at iteration 200 to establish a support, which was then used to compare Fourier ViT, CNN and iterative results. Fixed-support reconstructions of simulated data were performed for up to 2000 iterations to establish convergence statistics.

4.5 Complex Convolutional Neural Network

For experimental comparison with the proposed Fourier ViT, we implement a complex-valued CNN baseline following the complex-network recipe used for BCDI phase retrieval (Yu et al., 2024). The model has a U-Net encoder-decoder architecture. In our implementation, the encoder uses complex channel widths of 32, 64, 128, 256 and 512 (with a decoder of 256, 128, 64 and 32), followed by a final 1×1 complex convolution that outputs a single complex channel. Feature maps are represented as real and imaginary channels, and all convolutions are performed via complex operations by mixing the two channels according to complex multiplication. Skip connections bridge the encoder and decoder stages to preserve fine spatial detail. For the LCMO-500 measurement, we initialised the network input as two channels, with $\Re\{x\}$ set to the normalised measured diffraction magnitude and $\Im\{x\} = 0$,

corresponding to a zero initial phase. The network outputs a complex object $\rho(x) = \rho_{\Re}(x) + i\rho_{\Im}(x)$, from which amplitude and phase are obtained by $|\rho|$ and $\arg(\rho)$. Training and evaluation follow the same procedure as Sec. 4.2.

Code availability

The code used to train and evaluate the Fourier ViT and to reproduce all figures is available at https://github.com/JialunSimonLiu/Fourier-Attention_Vision-Transformer.

Acknowledgements

We would like to thank Xi Yu, Wenxuan Fang, Jack Griffiths, Yue Dong, Ruiyang Duan, Seungjun Lee, Yuewei Lin, Shinjae Yoo and Longlong Wu for helpful discussions, and Weibin Peng for advice on figure design and presentation. We would like to thank Ross Harder, Tadesse Assefa, Yue Cao, Xiaojing Huang and Jesse Clark for help with the measurements of LCMO and Josh Turner and Arup Kumar Raychaudhuri for providing the sample. The measurements were carried out at the Advanced Photon Source (APS) beamline 34-ID-C, which was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357. The beamline 34-ID-C was built with U.S. National Science Foundation grant DMR-9724294. Work performed at UCL was supported by EPSRC and ERC. Work at Brookhaven National Laboratory was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-SC0012704. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising.

5 Supplementary

S.1 Details of multi-scale Fourier attention

The spectral mixer introduced in Sec. 2.1.2 is defined here. Let $\mathbf{X} \in \mathbb{R}^{B \times E \times H \times W}$ with $H=W=16$ denote the token feature map inside a transformer block. At scale s we downsample by average pooling, apply a per-channel Fourier gate, and upsample back:

$$\mathbf{Z}_s = \mathcal{U}_s \left(\mathcal{F}^{-1} \left(\mathcal{F}(\mathcal{D}_s \mathbf{X}) \odot \mathbf{W}_s \odot \mathbf{M}_s \right) \right),$$

where \mathcal{D}_s is average pooling with stride s to $h_s \times w_s = (H/s) \times (W/s)$, \mathcal{U}_s upsamples back to (H, W) , \mathcal{F} is the channel-wise 2D FFT, $\mathbf{W}_s \in \mathbb{R}^{E \times h_s \times w_s}$ are learned per-channel frequency responses, and $\mathbf{M}_s \in \mathbb{R}^{1 \times h_s \times w_s}$ is a shared gate. The multi-scale output is

$$\mathbf{Z} = \sum_s \alpha_s \mathbf{Z}_s, \quad \alpha = \text{softmax}(\boldsymbol{\theta}),$$

with scale weights α_s parameterised by $\boldsymbol{\theta} \in \mathbb{R}^{|\{s\}|}$.

By the convolution theorem,

$$\mathbf{Z}_s = \mathcal{U}_s \left(\mathcal{D}_s \mathbf{X} * \mathbf{g}_s \right), \quad \mathbf{g}_s = \mathcal{F}^{-1}(\mathbf{W}_s \odot \mathbf{M}_s),$$

with $*$ denoting circular convolution on the $h_s \times w_s$ grid. Each scale applies the learned kernel \mathbf{g}_s , upsamples the result to (H, W) , and combines the scales using the weights α_s . This gives global token mixing via FFT-based convolution. In frequency space, $\mathbf{W}_s \odot \mathbf{M}_s$ is a multiplicative gate on the Fourier modes.

Let $N = HW$ denote the number of tokens. For each transformer block, the spectral mixer requires $\mathcal{O}(|\{s\}| B E N \log N)$ time and $\mathcal{O}(|\{s\}| B E N)$ memory, compared with $\mathcal{O}(B E N^2)$ time and $\mathcal{O}(B N^2)$ memory for dot-product attention. This follows Fourier-based token mixing (FNet) (Lee-Thorp et al., 2022) and spectral parameterisations used in Fourier neural operators (Li et al., 2021), but in our model it appears only as a compact multi-scale gate inside each transformer block.

S.2 Amplitude prior and normalisation

The decoder predicts a real-space amplitude $\hat{A}(x, y)$ on the 64×64 grid, together with a phase field as described in Sec. 2.1.3. Before forming the complex density we combine $\hat{A}(x, y)$ with a simple prior amplitude $A_{\text{prior}}(x, y)$ defined on the real-space support $S(x, y)$, where $S(x, y) = 1$ inside the support region and $S(x, y) = 0$ outside. From this support we construct one of three priors: (i) a flat prior with $A_{\text{prior}}(x, y) = S(x, y)$; (ii) a blurred prior obtained by repeated local averaging of $S(x, y)$; or (iii) a Gaussian prior centred on the support and truncated by $S(x, y)$.

During training we blend the prior and predicted amplitudes via a scalar weight $\alpha(t)$:

$$A_{\text{blend}}(x, y; t) = \alpha(t) A_{\text{prior}}(x, y) + (1 - \alpha(t)) \hat{A}(x, y).$$

This keeps early epochs guided by the support-based prior and later epochs dominated by the network prediction. We use a cosine schedule with a lock epoch t_{lock} and a decay-end epoch t_{end} :

$$\alpha(t) = \begin{cases} 1, & t < t_{\text{lock}}, \\ \frac{1}{2} \left(1 + \cos \frac{\pi(t - t_{\text{lock}})}{t_{\text{end}} - t_{\text{lock}}} \right), & t_{\text{lock}} \leq t < t_{\text{end}}, \\ 0, & t \geq t_{\text{end}}. \end{cases}$$

For the experimental single-pattern runs, we use $t_{\text{lock}} = 10$ and $t_{\text{end}} = 50$. For the synthetic dataset runs, we use a slower schedule ($t_{\text{lock}} = 50$, $t_{\text{end}} = 200$) to reduce early amplitude drift in joint amplitude-phase recovery; in all cases the prior is removed once $\alpha(t) = 0$.

To stabilise the overall scale of the reconstructed amplitude, we apply a soft energy-matching step inside the support. Let $A_{\text{blend}}(x, y; t)$ denote the blended amplitude at epoch t . We compute the total amplitude ‘‘energy’’ inside the support for the blended and prior amplitudes as

$$E_{\text{cur}}(t) = \sqrt{\sum_{x,y} (A_{\text{blend}}(x, y; t) S(x, y))^2}, \quad E_{\text{tgt}} = \sqrt{\sum_{x,y} (A_{\text{prior}}(x, y) S(x, y))^2}.$$

where E_{cur} is the current energy, E_{tgt} is the targeted energy. We then rescale

$$\tilde{A}(x, y; t) = \gamma(t) A_{\text{blend}}(x, y; t), \quad \gamma(t) = \text{clip}(E_{\text{tgt}}/E_{\text{cur}}(t), \gamma_{\text{min}}, \gamma_{\text{max}}),$$

with $\gamma_{\text{min}}=0.8$ and $\gamma_{\text{max}}=1.2$ in all experiments. This constrains the support energy to remain close to the prior value. In practice, it prevents global amplitude drift while allowing local deviations from A_{prior} .

S.3 Loss weights and exponents

The hybrid loss combines Pearson correlation coefficient (PCC), standard χ^2 , and power- χ^2 ,

$$L_{\text{total}}(t) = w_{\text{PCC}}(t) L_{\text{PCC}} + w_{\chi^2}(t) L_{\chi^2} + w_{\text{p}\chi^2}(t) L_{\text{p}\chi^2},$$

with epoch index $t \geq 0$. We set the raw weight schedules as

$$\begin{aligned} \text{raw}_{\text{PCC}}(t) &= 200 \exp(-t/15) + 1, \\ \text{raw}_{\chi^2}(t) &= 1, \\ \text{raw}_{\text{p}\chi^2}(t) &= 0.5 r(t), \end{aligned}$$

where $r(t)$ is a sine ramp from 0 to 1 over the first 20 epochs. For $t \geq 200$, we rescale the raw weights as

$$\text{raw}_{\text{PCC}}(t) = 0.6 \text{raw}_{\text{PCC}}(t), \quad \text{raw}_{\text{p}\chi^2}(t) = 0.6 \text{raw}_{\text{p}\chi^2}(t), \quad \text{raw}_{\chi^2}(t) = 1.1 \text{raw}_{\chi^2}(t).$$

We then normalise by the sum of raw weights at epoch t to obtain $w_{\text{PCC}}(t)$, $w_{\chi^2}(t)$, and $w_{\text{p}\chi^2}(t)$.

The exponent in the power- χ^2 term decays from 2.0 to 0.5,

$$p(t) = p_{\text{end}} + (p_{\text{start}} - p_{\text{end}}) \exp(-\lambda t),$$

with $p_{\text{start}} = 2.0$, $p_{\text{end}} = 0.5$, and $\lambda = 0.005$. We include a total-variation term on the amplitude with weight $\lambda_{\text{TV}}(t) = 0.01$ for $t < 200$ and $\lambda_{\text{TV}}(t) = 0$ thereafter.

S.4 Additional noise test

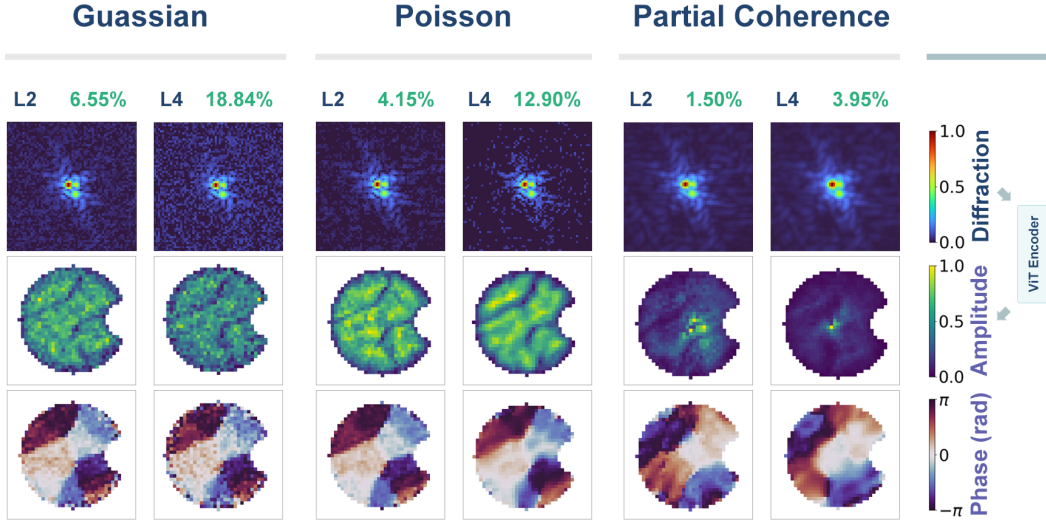


Fig. 6: Additional noise robustness visualisations for the Fourier ViT at Levels 2 and 4 (Levels 1, 3 and 5 are shown in Fig. 3). Columns show Gaussian noise, Poisson noise, and partial coherence applied to the same reference diffraction pattern. For each noise model, the green percentages report χ_n^2 (in %), the input diffraction error relative to the clean ground truth. The top row shows the noisy input diffraction magnitude (64×64 pixels). The middle and bottom rows show the reconstructed crystal amplitude and phase (40×40 pixels), respectively; both are masked to the support and the phase is shown in radians.

S.5 Single-domain strontium titanate (STO)




Fig. 7: Single-domain SrTiO_3 (STO) nanocrystal with a cross-section size of $\sim 460 \times 450 \text{ nm}^2$. (a) Measured STO diffraction magnitude, shown as a 64×64 pixel cropped region centred on the Bragg peak and normalised to $[0, 1]$. (b) Distribution of final reciprocal-space mismatch χ^2 over 100 random initialisations, ranging from $\sim 0.35\%$ to $\sim 0.68\%$. (c) Reconstructed real-space amplitude, normalised to $[0, 1]$. (d) Reconstructed real-space phase, wrapped to $[-\pi, \pi]$; scale bar represents 200 nm.

6 Supplementary Movies

M.1

ER/HIO

M.2

Fourier ViT

M.3

C-CNN

References

- Assefa TA, Cao Y, Diao J, et al (2020) Scaling behavior of low-temperature orthorhombic domains in the prototypical high-temperature superconductor $\text{La}_{1.875}\text{Ba}_{0.125}\text{CuO}_4$. *Phys Rev B* 101:054104. <https://doi.org/10.1103/PhysRevB.101.054104>, URL <https://link.aps.org/doi/10.1103/PhysRevB.101.054104>
- Ba JL, Kiros JR, Hinton GE (2016) Layer normalization. arXiv preprint <https://doi.org/10.48550/arXiv.1607.06450>, URL <https://arxiv.org/abs/1607.06450>
- Bates RHT (1984) Uniqueness of solutions to two-dimensional fourier phase problems for localized and positive images. *Computer Vision, Graphics, and Image Processing* 25(2):205–217. [https://doi.org/10.1016/0734-189X\(84\)90103-8](https://doi.org/10.1016/0734-189X(84)90103-8)
- Bates RHT, Napier PJ (1972) Identification and removal of phase errors in interferometry. *Monthly Notices of the Royal Astronomical Society* 158(4):405–424. <https://doi.org/10.1093/mnras/158.4.405>, URL <https://doi.org/10.1093/mnras/158.4.405>
- Beinert R, Plonka G (2015) Ambiguities in one-dimensional discrete phase retrieval from fourier magnitudes. *Journal of Fourier Analysis and Applications* 21(6):1169–1198. <https://doi.org/10.1007/s00041-015-9405-2>, URL <https://doi.org/10.1007/s00041-015-9405-2>
- Bendory T, Beinert R, Eldar YC (2017) Fourier phase retrieval: Uniqueness and algorithms. In: Boche H, Caire G, Calderbank R, et al (eds) *Compressed Sensing and its Applications. Applied and Numerical Harmonic Analysis*, Springer, p 55–91, https://doi.org/10.1007/978-3-319-69802-1_2, URL https://doi.org/10.1007/978-3-319-69802-1_2
- Bruck YM, Sodin LG (1979) On the ambiguity of the image reconstruction problem. *Optics Communications* 30(3):304–308. [https://doi.org/10.1016/0030-4018\(79\)90358-4](https://doi.org/10.1016/0030-4018(79)90358-4), URL [https://doi.org/10.1016/0030-4018\(79\)90358-4](https://doi.org/10.1016/0030-4018(79)90358-4)
- Carnis J, Gao L, Labat S, et al (2019) Towards a quantitative determination of strain in bragg coherent x-ray diffraction imaging: artefacts and sign convention in reconstructions. *Scientific Reports* 9:17357. <https://doi.org/10.1038/s41598-019-53774-2>
- Chen C, Miao J, Wang CW, et al (2007) Application of optimization technique to noncrystalline x-ray diffraction microscopy: Guided hybrid input-output method. *Phys Rev B* 76:064113. <https://doi.org/10.1103/PhysRevB.76.064113>
- Cherukara MJ, Nashed YSG, Harder RJ, et al (2018) Real-time coherent diffraction inversion using deep generative networks. *Scientific Reports* 8:16520. <https://doi.org/10.1038/s41598-018-34525-1>
- Cherukara MJ, Zhou T, Nashed YSG, et al (2020) Ai-enabled high-resolution scanning coherent diffraction imaging. *Applied Physics Letters* 117(4):044103. <https://doi.org/10.1063/5.0013065>
- Clark JN, Huang X, Harder R, et al (2012) High-resolution three-dimensional partially coherent diffraction imaging. *Nature Communications* 3:993. <https://doi.org/10.1038/ncomms1994>
- Clark JN, Beitra L, Xiong G, et al (2013) Ultrafast three-dimensional imaging of lattice dynamics in individual gold nanocrystals. *Science* 341(6141):56–59. <https://doi.org/10.1126/science.1236034>, URL <https://www.science.org/doi/10.1126/science.1236034>
- Dosovitskiy A, Beyer L, Kolesnikov A, et al (2021) An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (ICLR)*, <https://doi.org/10.48550/arXiv.2010.11929>
- Duhamel P, Vetterli M (1990) Fast fourier transforms: A tutorial review and a state of the art. *Signal Processing* 19(4):259–299. [https://doi.org/10.1016/0165-1684\(90\)90158-U](https://doi.org/10.1016/0165-1684(90)90158-U)
- Elser V (2003) Phase retrieval by iterated projections. *Journal of the Optical Society of America A* 20(1):40–55. <https://doi.org/10.1364/JOSAA.20.000040>

- Favre-Nicolin V, Girard G, Leake S, et al (2020a) Pynx: high-performance computing toolkit for coherent x-ray imaging based on operators. *Journal of Applied Crystallography* 53(5):1404–1413. <https://doi.org/10.1107/S1600576720010985>, URL <https://doi.org/10.1107/S1600576720010985>
- Favre-Nicolin V, Leake S, Chushkin Y (2020b) Free log-likelihood as an unbiased metric for coherent diffraction imaging. *Scientific Reports* 10:2664. <https://doi.org/10.1038/s41598-020-57561-2>
- Fienup JR (1982) Phase retrieval algorithms: a comparison. *Applied Optics* 21(15):2758–2769. <https://doi.org/10.1364/AO.21.002758>
- Fienup JR (1987) Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint. *Journal of the Optical Society of America A* 4(1):118–123. <https://doi.org/10.1364/JOSAA.4.000118>, URL <https://doi.org/10.1364/JOSAA.4.000118>
- Gan W, Zhai Q, McCann MT, et al (2024) PtychoDV: Vision transformer-based deep unrolling network for ptychographic image reconstruction. *IEEE Open Journal of Signal Processing* 5:539–547. <https://doi.org/10.1109/OJSP.2024.3375276>
- Gao Y, Huang X, Yan H, et al (2021) Bragg coherent diffraction imaging by simultaneous reconstruction of multiple diffraction peaks. *Phys Rev B* 103:014102. <https://doi.org/10.1103/PhysRevB.103.014102>, URL <https://link.aps.org/doi/10.1103/PhysRevB.103.014102>
- Gerchberg RW, Saxton WO (1972) A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* 35(2):237–246
- Goodman JW (2015) *Statistical Optics*, 2nd edn. John Wiley & Sons, Hoboken, NJ
- Guibas J, Mardani M, Li Z, et al (2021) Adaptive fourier neural operators: Efficient token mixers for transformers. arXiv preprint arXiv:2111.13587 <https://doi.org/10.48550/arXiv.2111.13587>
- Guizar-Sicairos M, Fienup JR (2008) Phase retrieval with fourier-weighted projections. *Journal of the Optical Society of America A* 25(3):701–709. <https://doi.org/10.1364/JOSAA.25.000701>
- Harder R (2021) Deep neural networks in real-time coherent diffraction imaging. *IUCrJ* 8(Pt 1):1–3. <https://doi.org/10.1107/S2052252520016590>
- Harris CR, Millman KJ, van der Walt SJ, et al (2020) Array programming with NumPy. *Nature* 585:357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- Huang X, Harder R, Xiong G, et al (2011) Propagation uniqueness in three-dimensional coherent diffractive imaging. *Phys Rev B* 83:224109. <https://doi.org/10.1103/PhysRevB.83.224109>, URL <https://link.aps.org/doi/10.1103/PhysRevB.83.224109>
- Janesick JR (2001) *Scientific Charge-Coupled Devices*. SPIE Press, Bellingham, WA
- Janesick JR (2007) *Photon Transfer*. SPIE Press, Bellingham, WA
- Kim D, Chung M, Carnis J, et al (2018) Active site localization of methane oxidation on pt nanocrystals. *Nature Communications* 9(1):3422. <https://doi.org/10.1038/s41467-018-05464-2>
- Lee DD, Seung HS (1999) Learning the parts of objects by non-negative matrix factorization. *Nature* 401(6755):788–791. <https://doi.org/10.1038/44565>
- Lee-Thorp J, Ainslie J, Eckstein I, et al (2022) Fnet: Mixing tokens with fourier transforms. In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, pp 4296–4313, <https://doi.org/10.18653/v1/2022.naacl-main.319>
- Li Z, Kovachki NB, Azizzadenesheli K, et al (2021) Fourier neural operator for parametric partial differential equations. In: *International Conference on Learning Representations*, URL <https://openreview.net/forum?id=c8P9NQVtmnO>
- Liu J, Yang D, Suzana AF, et al (2025) Electric field driven domain wall dynamics in batio_3 nanoparticles. *Physical Review B* 111:054101. <https://doi.org/10.1103/PhysRevB.111.054101>

- Liu T, Liu J, Li L, et al (2022) Origin of structural degradation in li-rich layered oxide cathode. *Nature* 606(7913):305–312. <https://doi.org/10.1038/s41586-022-04689-y>
- Loshchilov I, Hutter F (2019) Decoupled weight decay regularization. In: International Conference on Learning Representations (ICLR), <https://doi.org/10.48550/arXiv.1711.05101>, URL <https://arxiv.org/abs/1711.05101>
- Luke DR (2005) Relaxed averaged alternating reflections for diffraction imaging. *Inverse Problems* 21(1):37–50. <https://doi.org/10.1088/0266-5611/21/1/004>
- Marchesini S, He H, Chapman HN, et al (2003) X-ray image reconstruction from a diffraction pattern alone. *Physical Review B* 68(14):140101. <https://doi.org/10.1103/PhysRevB.68.140101>
- Miao J, Charalambous P, Kirz J, et al (1999) Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature* 400(6742):342–344. <https://doi.org/10.1038/22498>
- Monteforte M, Estandarte AK, Chen B, et al (2016) Novel silica stabilization method for the analysis of fine nanocrystals using coherent X-ray diffraction imaging. *Journal of Synchrotron Radiation* 23(4):953–958. <https://doi.org/10.1107/S1600577516006408>
- Patro BN, Agneeswaran VS (2023) Scattering vision transformer: Spectral mixing matters. arXiv preprint <https://doi.org/10.48550/arXiv.2311.01310>, URL <https://arxiv.org/abs/2311.01310>
- Pfeifer MA, Williams GJ, Vartanyants IA, et al (2006) Three-dimensional mapping of a deformation field inside a nanocrystal. *Nature* 442(7098):63–66. <https://doi.org/10.1038/nature04867>
- Rao Y, Zhao W, Zhu Z, et al (2021) Global filter networks for image classification. arXiv preprint <https://doi.org/10.48550/arXiv.2107.00645>, URL <https://arxiv.org/abs/2107.00645>
- Rivenson Y, Zhang Y, Günaydin H, et al (2018) Phase recovery and holographic image reconstruction using deep learning in neural networks. *Light: Science & Applications* 7:17141. <https://doi.org/10.1038/lsa.2017.141>
- Robinson I, Assefa TA, Cao Y, et al (2020) Domain texture of the orthorhombic phase of $\text{La}_{2-x}\text{Ba}_x\text{CuO}_4$. *Journal of Superconductivity and Novel Magnetism* 33:99–106. <https://doi.org/10.1007/s10948-019-05252-z>
- Robinson IK, Harder R (2009) Coherent x-ray diffraction imaging of strain at the nanoscale. *Nature Materials* 8(4):291–298. <https://doi.org/10.1038/nmat2400>
- Robinson IK, Assefa T, Cao Y, et al (2026) Phase domains in crystals
- Sayre D (1952) Some implications of a theorem due to shannon. *Acta Crystallographica* 5:843–843. <https://doi.org/10.1107/S0365110X52002276>
- Setter N, Damjanovic D, Eng L, et al (2006) Ferroelectric thin films: Review of materials, properties, and applications. *Journal of Applied Physics* 100(5):051606. <https://doi.org/10.1063/1.2336999>
- Ulvestad A, Nashed Y, Beutier G, et al (2017) Identifying defects with guided algorithms in bragg coherent diffractive imaging. *Scientific Reports* 7:9920. <https://doi.org/10.1038/s41598-017-09582-7>
- Vartanyants IA, Robinson IK (2001) Partial coherence effects on the imaging of small crystals using coherent x-ray diffraction. *Journal of Physics: Condensed Matter* 13(47):10593–10611. <https://doi.org/10.1088/0953-8984/13/47/305>
- Vaswani A, Shazeer N, Parmar N, et al (2017) Attention is all you need. In: *Advances in Neural Information Processing Systems*, <https://doi.org/10.48550/arXiv.1706.03762>, URL <https://arxiv.org/abs/1706.03762>
- Vu TS, Ha MQ, Dam HC (2025) Pid3net: a deep learning approach for single-shot coherent x-ray diffraction imaging of dynamic phenomena. *npj Computational Materials* 11:66. <https://doi.org/10.1038/s41524-025-01549-x>

- Williams GJ, Quiney HM, Peele AG, et al (2007) Coherent diffractive imaging and partial coherence. *Physical Review B* 75(10):104102. <https://doi.org/10.1103/PhysRevB.75.104102>
- Wu L, Juhas P, Yoo S, et al (2021a) Complex imaging of phase domains by deep neural networks. *IUCrJ* 8(1):12–21. <https://doi.org/10.1107/S2052252520013780>
- Wu L, Yoo S, Suzana AF, et al (2021b) Three-dimensional coherent x-ray diffraction imaging via deep convolutional neural networks. *npj Computational Materials* 7:175. <https://doi.org/10.1038/s41524-021-00644-z>
- Wu L, Yang D, Wang W, et al (2025) Unveiling nano-scale crystal deformation using coherent x-ray dynamical diffraction. *npj Computational Materials* 11:379. <https://doi.org/10.1038/s41524-025-01858-1>, URL <https://doi.org/10.1038/s41524-025-01858-1>
- Yao Y, Chan H, Sankaranarayanan S, et al (2022) Autophasenn: unsupervised physics-aware deep learning of 3d nanoscale bragg coherent diffraction imaging. *npj Computational Materials* 8:124. <https://doi.org/10.1038/s41524-022-00803-w>
- Yu X, Wu L, Lin Y, et al (2024) Ultrafast bragg coherent diffraction imaging of epitaxial thin films using deep complex-valued neural networks. *npj Computational Materials* 10:24. <https://doi.org/10.1038/s41524-024-01208-7>