Role of Spatial Coherence in Single-Shot Lensless Image Reconstruction: Supplemental Material

Ganesh M. Balasubramaniam[1], Xiao-Liu Chu[2] and Matthew R. Foreman[1,2]*

[1]*School of Electrical and Electronic Engineering,*

Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798

[2]*Institute for Digital Molecular Analytics and Science,*

59 Nanyang Drive, Singapore 636921

(Dated: July 20, 2026)

This supplemental document describes our object generation process and extends the analysis of spatial coherence in lensless image reconstruction to phase objects. The analysis, therefore, considers unit-modulus phase transmittances with phase values ranging from 0 to 2π and evaluates reconstructions from two detector-plane measurements, using the same partially coherent and coherent forward models as in the main article. Representative coherence-dependent measurements, phase reconstructions, and class-resolved quantitative results are presented for high-frequency high-density (HFHD), high-frequency low-density (HFLD), low-frequency high-density (LFHD), and low-frequency low-density (LFLD) objects. Reconstruction performance is assessed using the same metrics as in the main article, namely the Pearson correlation coefficient (PCC), high-frequency retention error (HFRE), low-frequency retention error (LFRE), and edge F_1 score (E_{F1}). The results show that coherence-aware inversion improves phase boundary localisation most consistently, while global and spectral gains depend strongly on object structure and spatial frequency content. These findings provide supporting evidence that spatial coherence is a consequential component of the lensless inverse model for both amplitude and phase objects.

S1. SYNTHETIC OBJECT IMAGE GENERATION

The synthetic object library used in this work was generated from glyph-based binary patterns using MATLAB [1]. Each object was created on a 2000×2000 pixel grid and assigned to one of four structural classes, namely, high-frequency high-density (HFHD), high-frequency low-density (HFLD), low-frequency high-density (LFHD), and low-frequency low-density (LFLD). These classes were designed to independently vary the characteristic spatial scale of the object features and their areal density. High-frequency objects were generated using smaller, sharper glyphs, whereas low-frequency objects were generated using larger glyphs with Gaussian smoothing. High-density objects contained many glyphs distributed across the field of view, whereas low-density objects contained fewer glyphs.

For each class, glyphs were randomly selected from a digit or letter library, resized, rotated within a prescribed angular range, and placed onto the image canvas. An occupancy mask was used during placement to minimise glyph overlap and to enforce class-dependent spacing between neighbouring features. Grid-based placement was used for the high-density classes

* matthew.foreman@ntu.edu.sg

to distribute features broadly across the field, while random placement was used for the low-density classes. After placement, class-dependent Gaussian smoothing was applied to control the effective spatial-frequency content. The final amplitude objects were inverted to produce dark features on a bright background and normalised to the range $[0, 1]$. A total of 25 objects were generated for each class.

Object-class separation was verified using a spatial-frequency score and a feature-density score. For each generated object $I(\mathbf{r})$, the image was first normalised, mean-subtracted, and multiplied by a two-dimensional Hann window $W(\mathbf{r})$. The centred Fourier power spectrum $P(\mathbf{f})$ was then computed as

$$P(\mathbf{f}) = \left|\mathcal{F}\left\{[I(\mathbf{r}) - \langle I(\mathbf{r})\rangle]\, W(\mathbf{r})\right\}\right|^2 . \tag{1}$$

A normalised radial spatial-frequency coordinate was defined as

$$\rho(\mathbf{f}) = \frac{\sqrt{(f_x - f_{x,0})^2 + (f_y - f_{y,0})^2}}{\rho_{\max}}, \qquad 0 \leq \rho \leq 1. \tag{2}$$

The spatial-frequency score was computed as the fraction of spectral power above the radial-frequency cutoff $\rho_c = 0.18$,

$$S_f = \frac{\sum_{\rho(\mathbf{f})\geq 0.18} P(\mathbf{f})}{\sum_{\rho(\mathbf{f})>0.015} P(\mathbf{f})}. \tag{3}$$

The lower cutoff $\rho > 0.015$ removes the near-DC component from the normalisation, so that S_f reflects object structure rather than the average background level.

The feature-density score was computed from the binary object mask $M(\mathbf{r})$ as

$$S_d = \frac{1}{N^2} \sum_{\mathbf{r}} M(\mathbf{r}), \tag{4}$$

where $N = 2000$ is the image size. The raw scores were robustly normalised across the generated library,

$$\widetilde{S} = \mathrm{clip}\left[\frac{S - P_2(S)}{P_{98}(S) - P_2(S)}, 0, 1\right], \tag{5}$$

where P_2 and P_{98} denote the 2nd and 98th percentiles, respectively. The four object classes were then verified in the $(\widetilde{S}_f, \widetilde{S}_d)$ plane using a cutoff of 0.5 along each axis. Thus, HFHD objects occupy the high-$\widetilde{S}_f$, high-$\widetilde{S}_d$ region, HFLD objects occupy the high-$\widetilde{S}_f$, low-$\widetilde{S}_d$ region, LFHD objects occupy the low-$\widetilde{S}_f$, high-$\widetilde{S}_d$ region, and LFLD objects occupy the low-$\widetilde{S}_f$, low-$\widetilde{S}_d$ region.

S2. PHASE OBJECT RECONSTRUCTION UNDER PARTIAL SPATIAL COHERENCE

The phase object is modelled as a unit-modulus transmittance,

$$t(\mathbf{r}_o) = \exp[i\phi(\mathbf{r}_o)]\,, \qquad 0 \leq \phi(\mathbf{r}_o) \leq 2\pi, \tag{6}$$

where $\mathbf{r}_o = (x_o, y_o)$ denotes the transverse object-plane coordinate and $\phi(\mathbf{r}_o)$ is the spatially varying phase delay. The same partially coherent and coherent forward operators used in

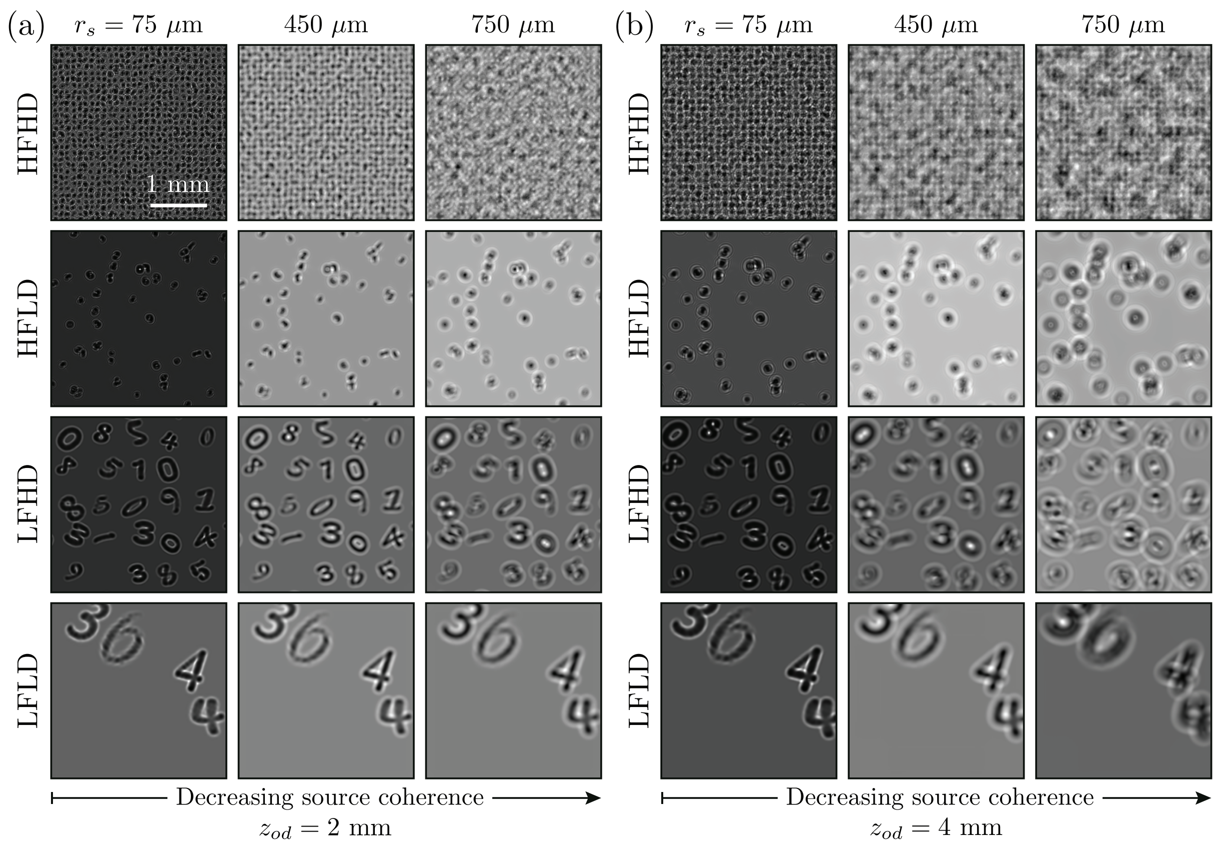

FIG. 1. **Coherence-dependent lensless measurements for phase objects derived from the object classes in Fig. 1 of the main manuscript.** The same object classes used in Fig. 1 of the main manuscript are considered here as phase objects rather than amplitude objects. Rows correspond to high-frequency, high-density (HFHD), high-frequency, low-density (HFLD), low-frequency, high-density (LFHD), and low-frequency, low-density (LFLD) phase objects. The detector intensities correspond to lensless measurements under different source-coherence conditions. (a) and (b) show detector measurements at 2 mm and 4 mm, respectively.

the main text were applied.

Phase recovery was performed using two detector-plane intensities, generated at $z_{od} =$ 2 mm and $z_{od} = 4$ mm. A single-phase estimate was propagated to both detector planes (see Fig. 1), and the plane-resolved loss terms were combined into a common optimisation objective. This two-plane configuration provides additional propagation diversity, reducing ambiguities that are more pronounced with phase-only inversion. All other simulation parameters, including wavelength, source-to-object distance, source radii, and the comparison between partially coherent and coherent forward models, were kept identical to those used in the amplitude analysis in the main text.

S3. REPRESENTATIVE PHASE RECONSTRUCTIONS

Representative phase reconstructions are shown in Fig. 2. In each panel, the left half corresponds to the reconstruction obtained using the partially coherent forward model, while the right half corresponds to the reconstruction obtained using the coherent forward model. The comparison shows that coherence mismatch also affects phase recovery, although the observed improvement differs from that in the amplitude case because phase reconstruction uses two detector planes.

For high-frequency phase objects, the partially coherent model more consistently preserves phase boundaries and local phase structure. In contrast, the coherent model can introduce background texture and oscillatory artefacts, particularly at larger source radii, because it attempts to explain a partially coherent measurement using a fully coherent propagation operator. For low-frequency phase objects, the difference between the two models is less uniform, because the dominant information is carried by broader spatial features and the two-plane configuration partially compensates for forward-model mismatch.

S4. QUANTITATIVE PHASE RECONSTRUCTION ANALYSIS

The quantitative phase analysis used the same four metrics as the main article. Global reconstruction fidelity was quantified using the Pearson correlation coefficient (PCC). Spectral preservation was quantified using high-frequency retention error (HFRE) and low-frequency retention error (LFRE), computed as the logarithmic mismatch between the reconstructed

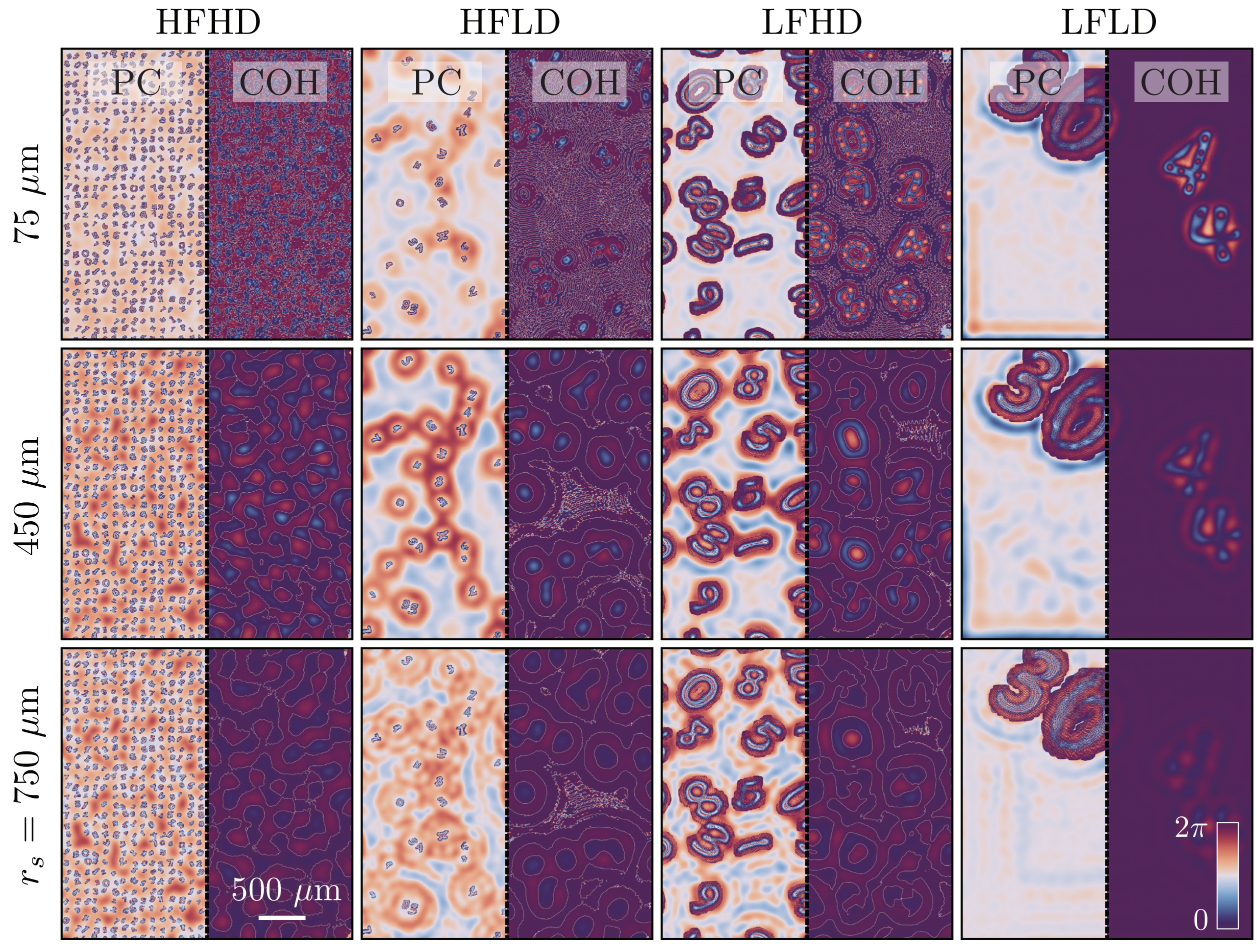

FIG. 2. **Representative phase reconstruction comparison across object classes and source coherence.** Phase-object reconstructions for HFHD, HFLD, LFHD, and LFLD object classes. Rows correspond to source radii $r_s = 75$ μm, 450 μm, and 750 μm. In each image, the left half shows the reconstruction obtained using the partially coherent (PC) forward model, while the right half shows the reconstruction obtained using the coherent (COH) forward model. The dashed line marks the boundary between the two reconstructions.

and reference spectral energy fractions above and below a fixed radial-frequency cutoff. Boundary recovery was quantified using the edge F_1 score, E_{F1}, computed from boundary precision and recall with a small spatial tolerance.

For metrics where larger values indicate better reconstruction, namely PCC and E_{F1}, the reconstruction advantage was defined as

$$\Delta Q = Q_{\mathrm{pc}} - Q_{\mathrm{coh}}, \tag{7}$$

where Q_{pc} and Q_{coh} are the metric values obtained from the partially coherent and coherent

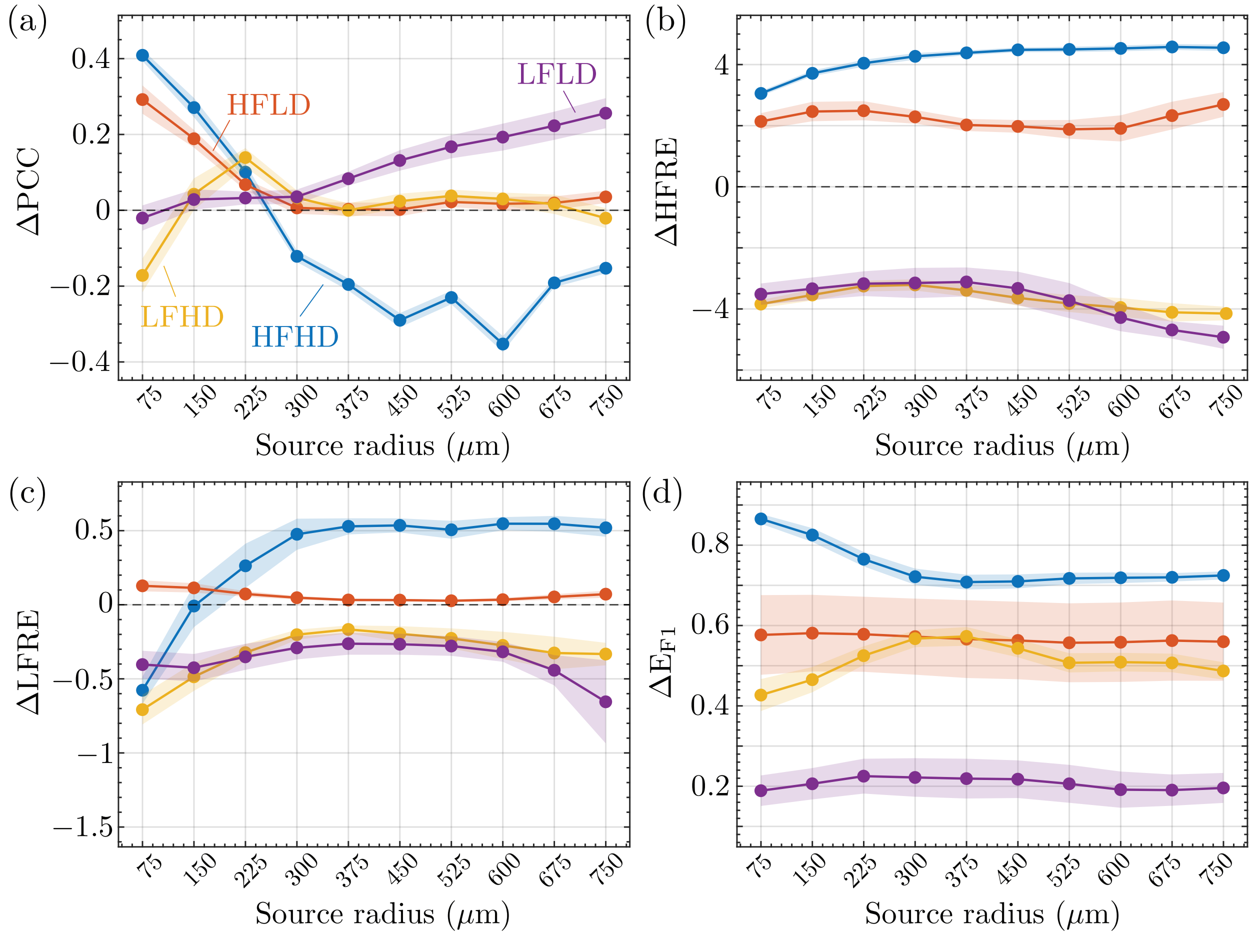

FIG. 3. **Class-resolved reconstruction advantage under varying source coherence.** Paired reconstruction differences are plotted as a function of source radius for the four object classes. Panels show differences in (a) PCC, (b) HFRE, (c) LFRE, and (d) E_{F1}. Solid markers denote the class-averaged value at each source radius, and shaded bands denote standard deviation across objects within that class.

forward models, respectively. For the error metrics HFRE and LFRE, the sign was reversed,

$$\Delta Q = Q_{\mathrm{coh}} - Q_{\mathrm{pc}}. \tag{8}$$

Thus, positive values always indicate an advantage for the partially coherent forward model.

The phase reconstruction results in Fig. 3 show that coherence-aware inversion produces the most consistent improvement in boundary recovery. All four object classes show positive ΔE_{F1} across the full range of source radii, indicating that the partially coherent forward model localises phase boundaries more accurately than the coherent forward model. The magnitude of this improvement depends strongly on object class. HFHD phase objects show the largest boundary advantage, with ΔE_{F1} remaining high across all source radii. HFLD and LFHD objects show intermediate positive boundary gains, whereas LFLD objects show the weakest but still positive E_{F1} imrovement.

The PCC is more class-dependent. HFHD objects show a positive ΔPCC at small source radii, but this advantage decreases with increasing source radius and becomes negative at larger source radii. This indicates that the partially coherent model improves boundary localisation for dense, high-frequency phase objects, even when full-image correlation is no longer favourable. HFLD objects show a positive ΔPCC at small source radii and a near-zero improvement at larger source radii. LFHD objects show only a weak PCC advantage over most of the source-radius range. In contrast, LFLD objects show an increasing positive ΔPCC with source radius, indicating that the partially coherent model most clearly improves global phase recovery for sparse low-frequency phase objects under lower spatial coherence.

The spectral error metrics distinguish between high- and low-frequency phase classes. HFHD and HFLD objects show positive ΔHFRE across all source radii, demonstrating that the partially coherent model substantially reduces high-frequency spectral error for phase objects containing fine spatial structure. This effect is strongest for HFHD objects and remains large as the source radius increases. Conversely, LFHD and LFLD objects show negative ΔHFRE, indicating that the coherent model gives lower high-frequency retention error for these low-frequency phase classes under this metric. This behaviour is expected because low-frequency objects contain limited reference high-frequency energy, so small differences in reconstructed edge sharpness or residual texture can dominate the high-frequency error measure.

The LFRE improvement shows a different pattern. HFHD objects show negative or near-zero ΔLFRE at small source radii, followed by positive values at larger source radii, indicating improved low-frequency spectral preservation by the partially coherent model once coherence mismatch becomes stronger. HFLD objects show a small but consistently positive ΔLFRE. In contrast, LFHD and LFLD objects show negative ΔLFRE across the source-radius range, indicating that the coherent model gives lower low-frequency spectral error for these broader phase structures. The negative LFRE improvement is especially pronounced for LFLD objects at large source radii.

In summary, the phase analysis shows that partial-coherence modelling improves phase reconstruction in an object-dependent manner. The most consistent benefit is observed in boundary localisation, with positive ΔE_{F1} across all phase-object classes and source radii. Spectral advantages are strongest for high-frequency phase objects, whereas low-frequency phase objects show weaker or negative spectral-error enhancement despite improvements in boundary recovery. These results support the conclusion of the main article that spatial coherence must be treated as an explicit component of the lensless inverse model, while showing that its effect on phase reconstruction depends on object structure and on the reconstruction property being evaluated.

---

[1] Y. LeCun, C. Cortes, and C. J. Burges, "MNIST handwritten digit database," ATT Labs [Online] **2**, 5 (2010).