Supplementary Information

(Dated: 19 November 2025)

S1. WEIGHTED COHERENT AVERAGE LOSS

We have implemented an auto-differentiation model to fit a phase 2D phase array to a ground truth one taken from a simulated diffraction pattern. The fit has been performed using the Weighted Coherent Average loss that we have used for the training of the DL model. We could then calculate easily the complex phase-difference vectors at each stage of the fit and extract the gradients from the algorithm to image the direction and the magnitude of the updates. Supplementary Figure. 1 shows the evolution of the predicted phase array over the optimization steps, providing a visual representation of the progressive "alignment" of the complex phase differences vectors from an initially random orientation toward a coherent direction. Recall that the complex phase difference vectors represent the phase offsets for each pixel, and therefore, convergence to a consistent value across the array is essential for a meaningful solution, although the specific value itself is determined freely by the model.

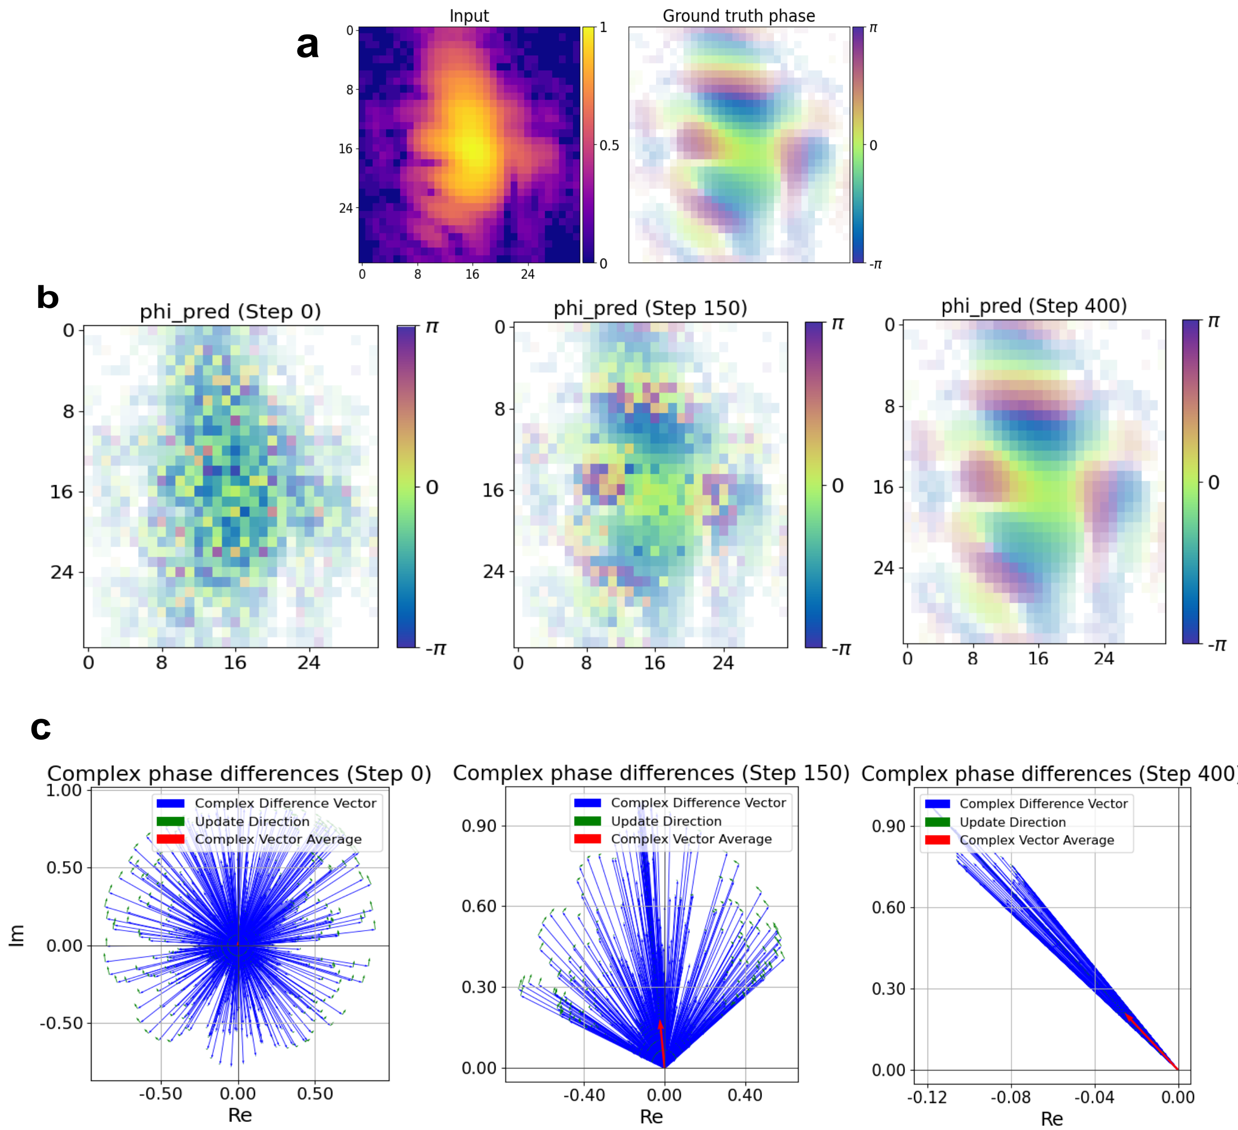

Supplementary Figure 1: Illustration of the functioning of the Weighted Coherent Average (WCA) loss function.**a** Simulated 2D BCDI diffraction pattern with corresponding reciprocal space phase. **b** Evolution over three different optimization steps of the result of the fit computed with an automatic differentiation model. The WCA was used as cost function to minimize. **c** Corresponding representation of the complex phase-differences vectors extracted during the optimization. The images show the updates (in green, rescaled) for each complex phase-difference vector and the complex average vector.

S2. COMPARISON WITH OTHER MODELS

Here we compare the results of our model with the ones obtained from existing architectures proposed in other works on Deep Learning for BCDI phase retrieval. We have adapted the model proposed in[1] for 3D images and we have trained it with the same dataset we have used for ours. We have trained the network as the authors describe in the paper, using a Mean Absolute Value on both the real and imaginary part of the predicted objects. Completed the training, the model was tested on the same data we showed in Fig.3 of the main text and obtained the following predicted objects. Despite the cleaner objects' amplitudes, obtained thanks to the loss function calculated in real space, the predictions do not correspond to the ground truth, both in shape and phase.

We have also downloaded AutoPhaseNN from the GitHub repository and computed the phase retrieval with it on the same diffraction patterns. The results in Supplementary Figure 3 appear to be very far from the expected outcome.

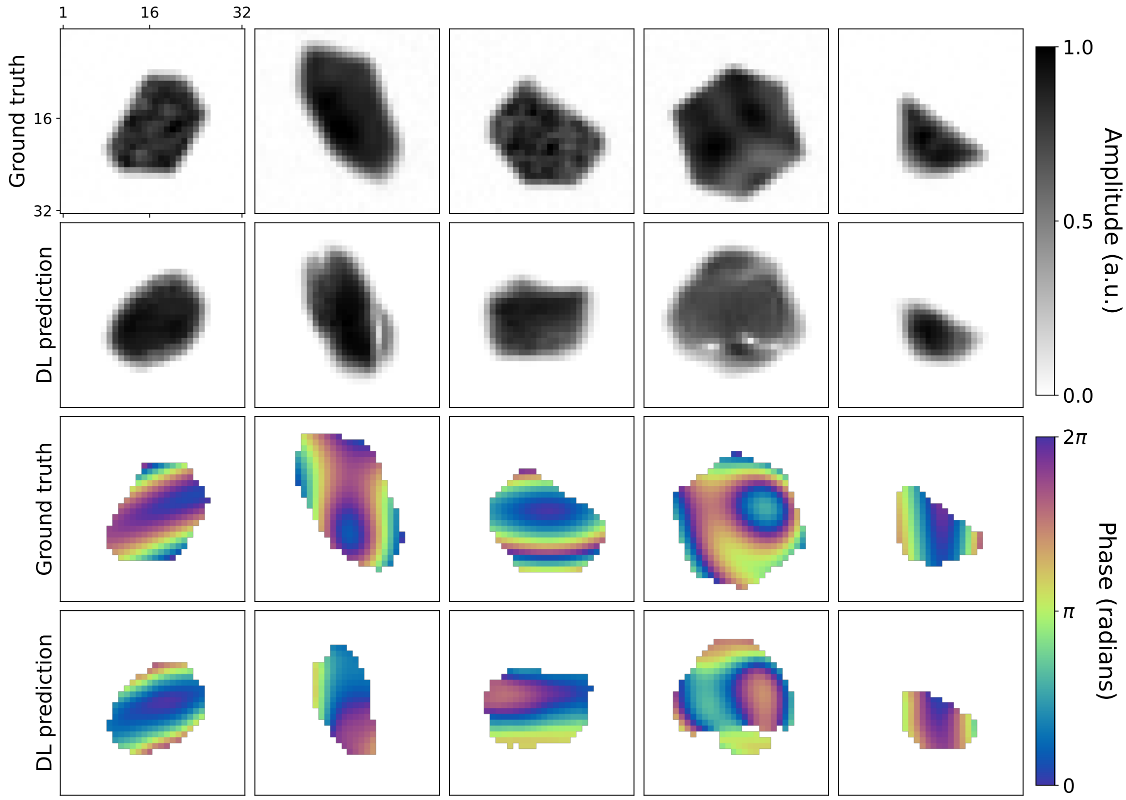

Supplementary Figure 2: Results of the 3D Complex-CNN

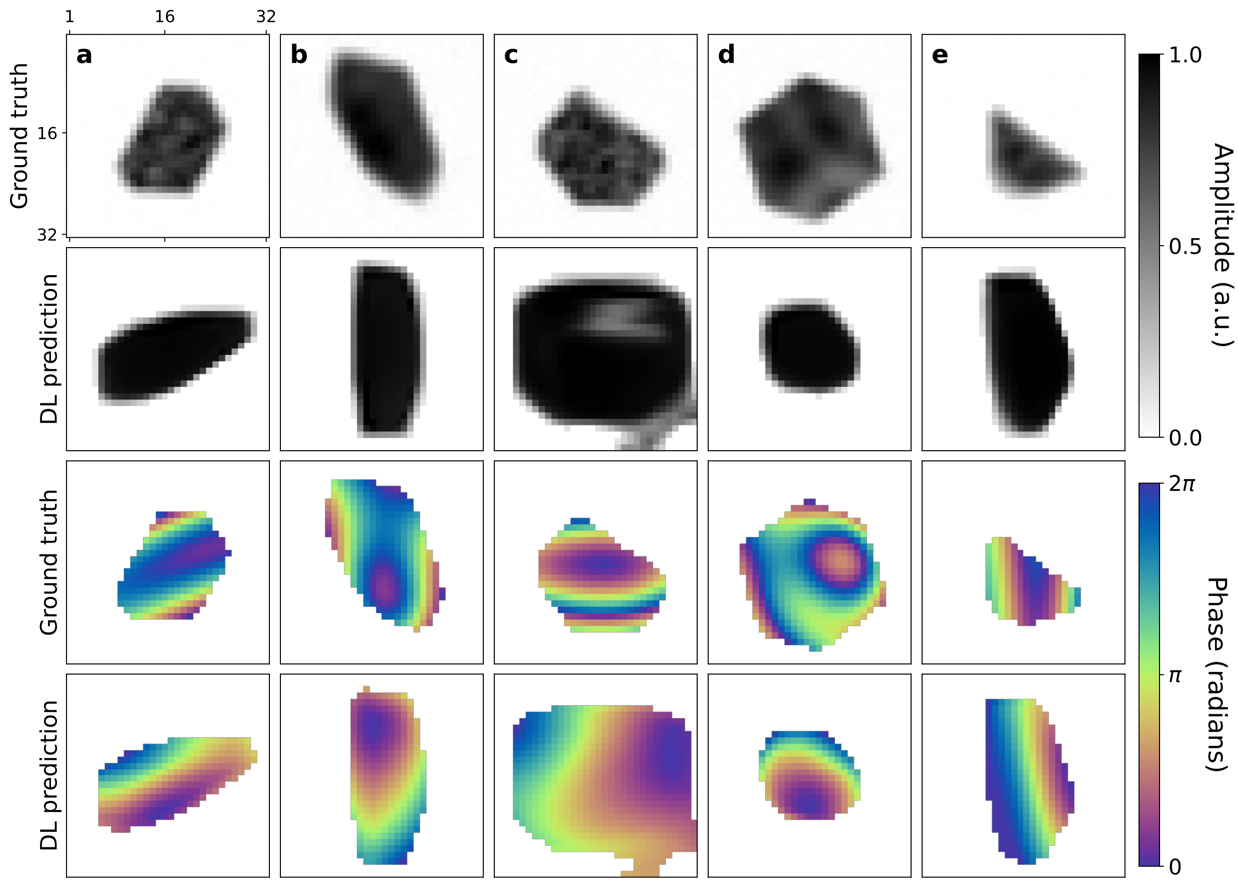

Supplementary Figure 3: Results of the AutoPhaseNN

S3. FURTHER EXAMPLES OF DL ASSISTED EXPERIMENTAL RECONSTRUCTIONS

Here we show two other examples of experimental BCDI patterns acquired at the ID01 beamline of the ESRF and their corresponding reconstructions obtained with standard iterative algorithms (PyNX) as well as the ones obtained with our DL model and with the combination of DL and PyNX. The first is from a Palladium nano-particle measured during acetylene hydrogenation[2] while the second is a nickel particle on sapphire substrate[3].

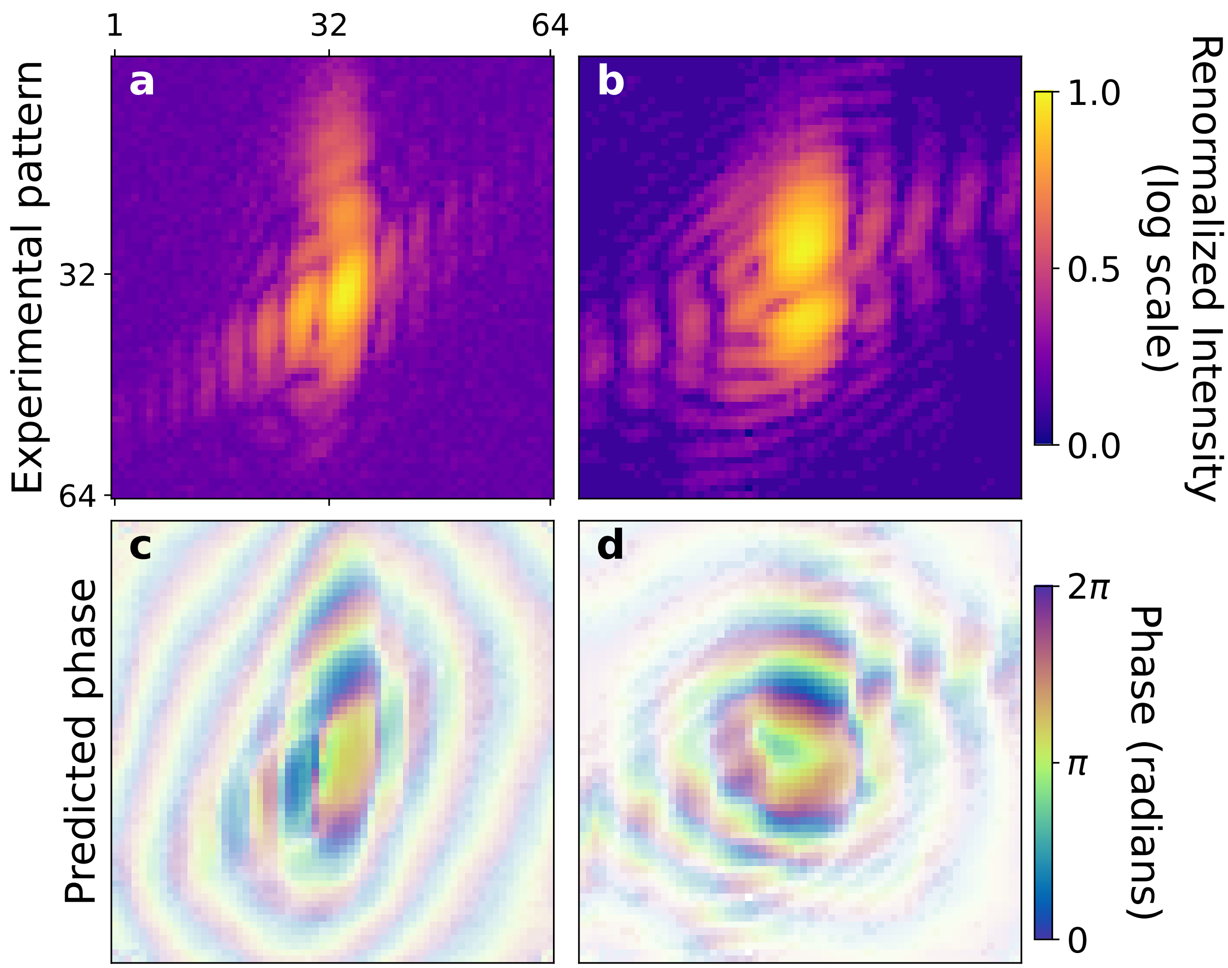

Supplementary Figure 4: Central slice of two examples of experimental BCDI data.a - b) XY Central slices of two experimental BCDI patterns. **c-d)** Corresponding slices of the DL model phase predictions.

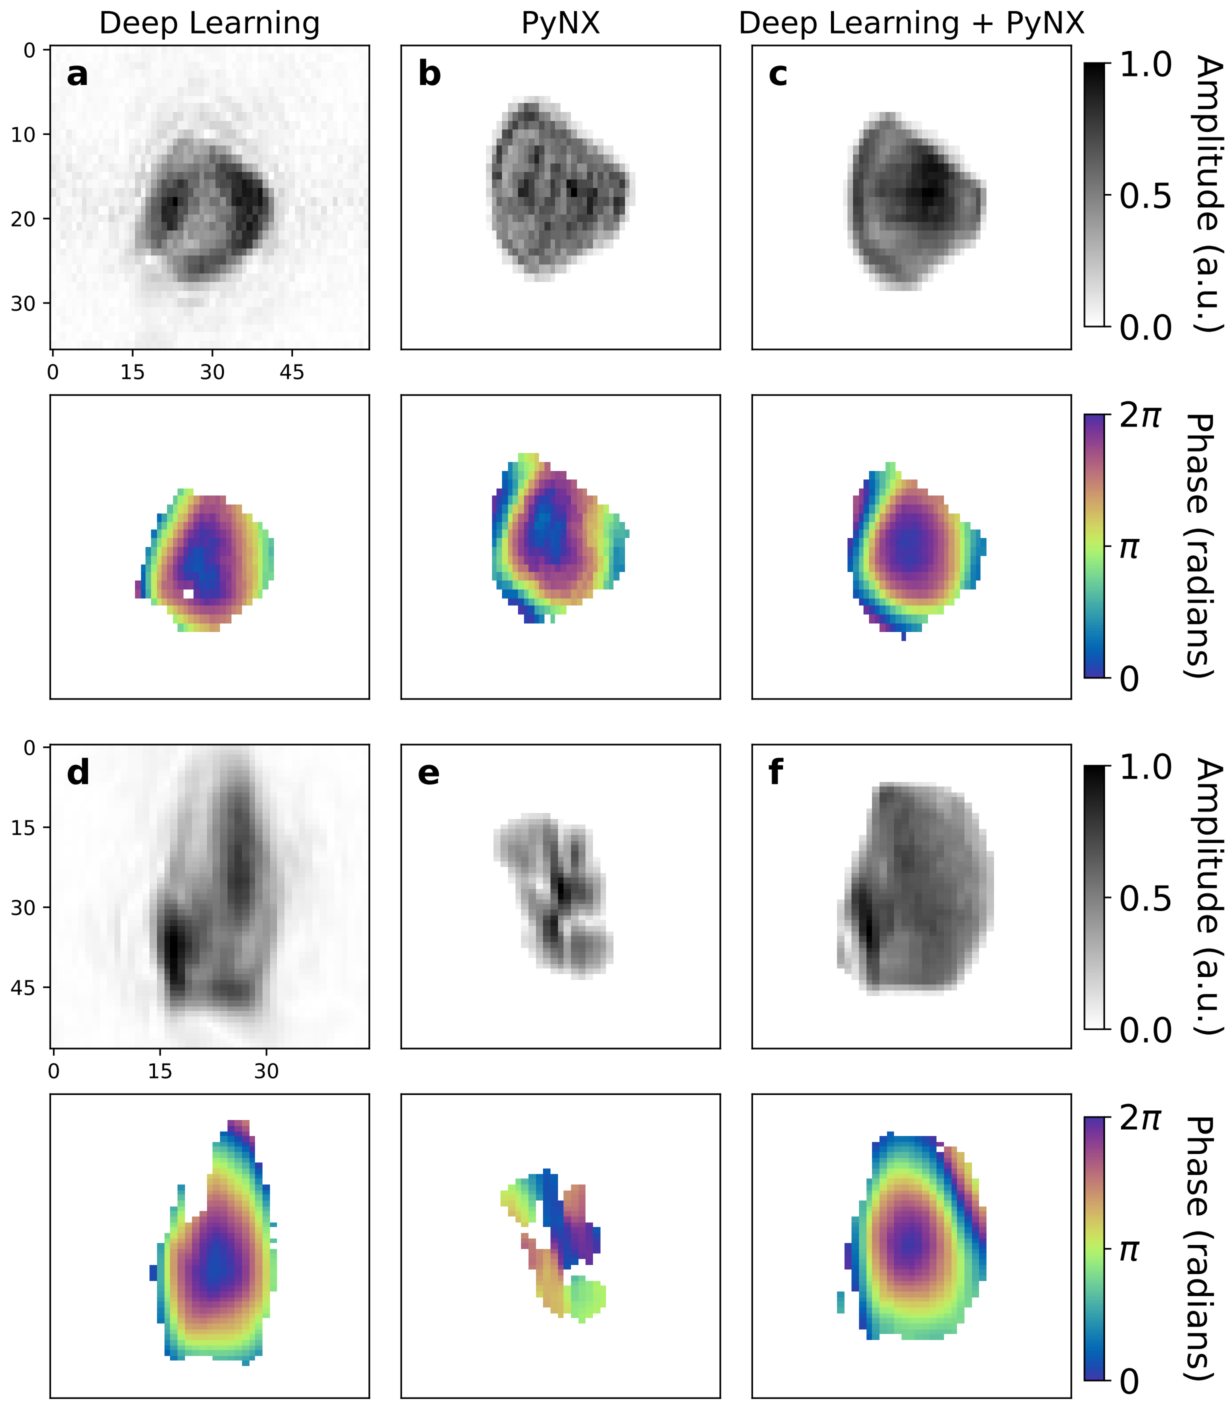

Supplementary Figure 5: Reconstructed objects from the diffraction patterns in Supplementary Figure 4 a-d) Object obtained from the DL prediction of the reciprocal space phase. **b-e)** Reconstructed object using 400 HIO + 1000 RAAR + 300 ER iterations (best of 60 different reconstructions). **c-f** Reconstructed object using 400 ER iterations starting from the object obtained from the DL phase prediction.

S4. INCREASING ITERATIVE PHASING ROBUSTNESS

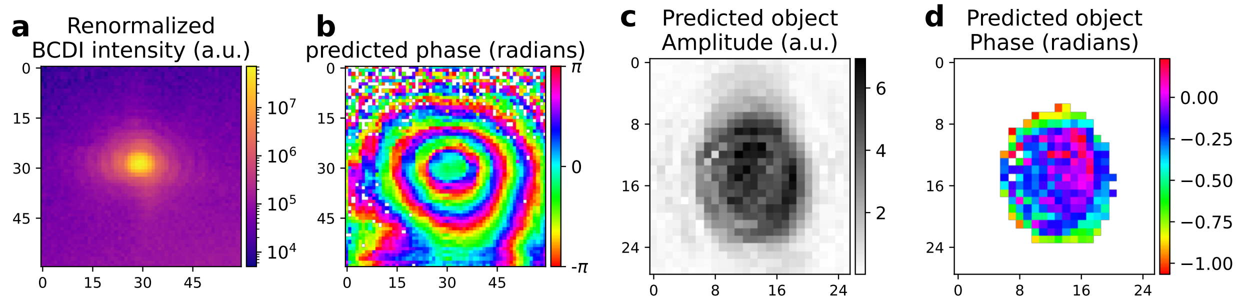

Supplementary Figure 6: a) BCDI data of Pd particle. **b)** Deep learning model prediction of the reciprocal space phase. **c-d)** Corresponding predicted object amplitude and phase.

Fig 6**a** shows experimental BCDI data from a palladium (Pd) nano-particle. Despite a strong noise background from both air scattering and a nearby aluminum powder ring from a furnace dome, our model prediction of the reciprocal space phase shown in **b** is still relatively accurate. The corresponding reconstructed object amplitude and phase are shown in **c** and **d** respectively. Despite the reconstruction artefacts induced by the noise background on our data, the object has the correct expected shape and phase field and is a good starting point for iterative phase retrieval refinement as shown below.

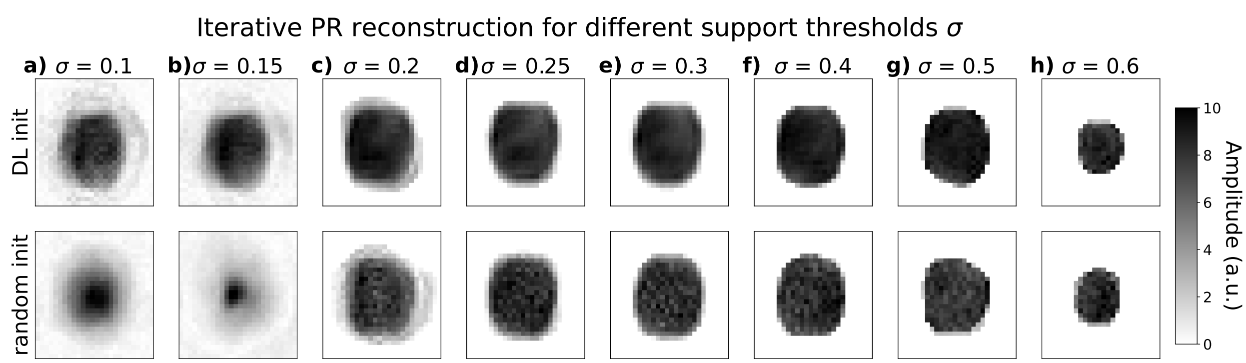

Supplementary Figure 7: Iterative Phase Retrieval (PR) reconstructions for different support thresholds σ starting from the DL model prediction (1^{st} row) and with a random start (2^{nd} row).

In Supplementary Figure 7, we compare standard iterative phase retrieval algorithms starting either from our model prediction (1^{st} row) or a random initialization (2^{nd} row) and this for several values of the support threshold parameter σ. We used a short phase retrieval recipe as combination of Hybrid-Input-Output (HIO) and Error-Reduction (ER) as 200 HIO + 200 ER + 200 HIO + 600

ER. Furthermore, for each value of the σ parameter, we made 20 reconstructions and combined the best 3 using mode decomposition.

Comparing the 1^{st} (DL model initialization) and 2^{nd} rows (random initialization) in Supplementary Figure 7, we can conclude that our model initialization induces a better and more robust final result. When the support threshold is too low (σ=0.1-0.2), our model prediction is already close enough to the expected result that the final object still has the correct shape, despite some noise. On the other hand, a random initialization doesn't give a proper object shape. Finally, for correct support threshold values 0.25<σ<0.4, the reconstructed object with DL model initialization shows a more homogeneous amplitude, indicating a better and more reproducible reconstruction.

S5. HIGH-STRAIN RECONSTRUCTION

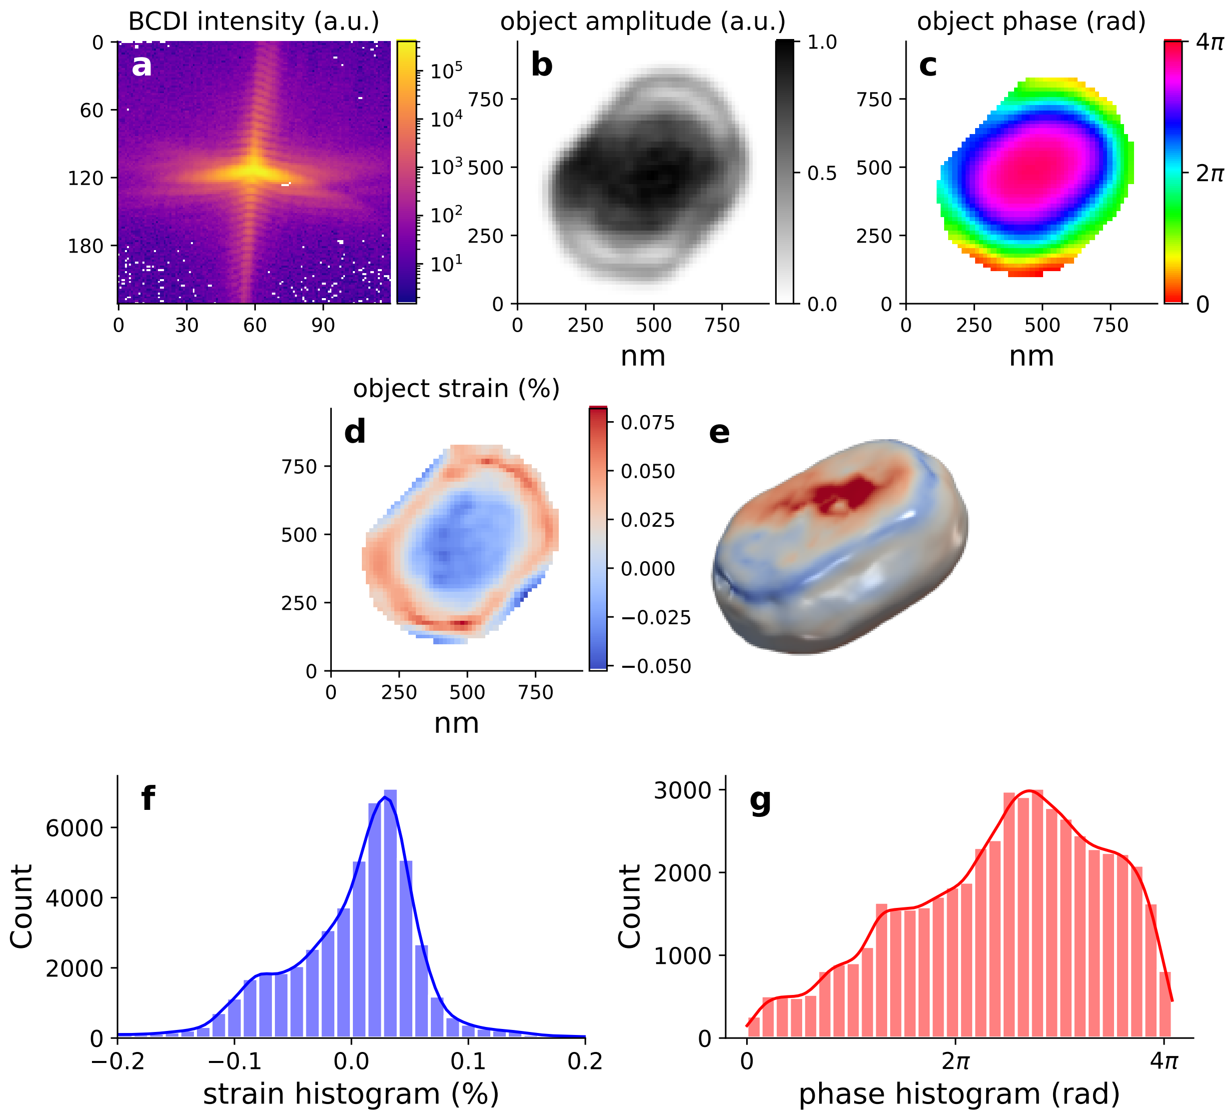

Supplementary Figure 8: Experimental strain reconstructiona Experimental BCDI data. **b-c** reconstructed object amplitude and phase after orthogonalization. **d-e** Strain map and 3D iso-surface. **f** Strain histogram. **g** Object phase histogram.

In Supplementary Figure 8 we show the post-processed particle reconstruction shown in the main text Figure 6f. Using our model prediction, we perform phase retrieval refining using PyNX with 20 different sets of support update parameters. Combining the 3 best reconstructions with mode decomposition we obtain the object shown in Supplementary Figure 8b-c. The corresponding strain map is shown in Supplementary Figure 8d and on a 3D iso-surface in Supplementary

Figure 8e. One can observe from the histograms Supplementary Figure 8f-g a large phase and strain range explaining the difficulty to obtain a correct reconstruction using iterative phase retrieval only.

For completeness, an attempt of iterative phase retrieval of this dataset with genetic selection was performed as well. While this approach yielded higher-quality results than conventional alternating-projections phase retrieval (Fig. 6e of the main text), it remained inferior to the DL-PyNX phasing outcome (Fig. 6f of the main text).

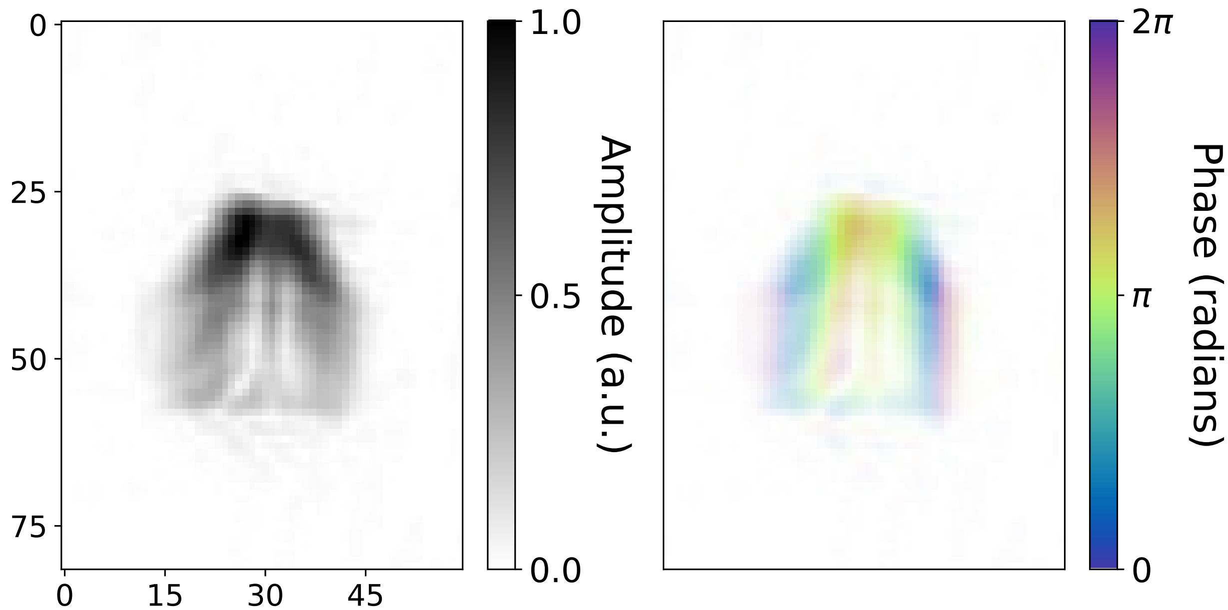

Supplementary Figure 9: Reconstruction with genetic algorithm.

Supplementary Figure 9 shows a reconstruction performed using a genetic algorithm approach, similar to the method described in Ref.[4]. Starting from 15 randomly initialized guesses, the following iterative chain was applied:

- 400 iterations of HIO
- 300 iterations of RAAR
- 100 iterations of ER

One reconstruction was selected from the population based on a sharpness criterion, defined as the sum over all voxels of the absolute value of the reconstruction raised to the fourth power (see Ref.[4]).

The amplitude of this selected reconstruction was then used as a seed for the next 10 generations of the base chain, with the phase kept unchanged. This update followed Wilkin's approach[5]:

$$\mathscr{A} = \sqrt{\mathscr{A} \times \mathscr{A}_{\text{best}}} \tag{1}$$

where $\mathscr{A}$ is the amplitude of an individual in the population being updated, and $\mathscr{A}_{\text{best}}$ is the amplitude with the lowest sharpness from the previous generation.

S6. ORTHOGONALIZED PARTICLE 1

Here we show instead Particle 1 of the main text after orthogonalization, to better assess the size of the particle and the Winterbottom shape.

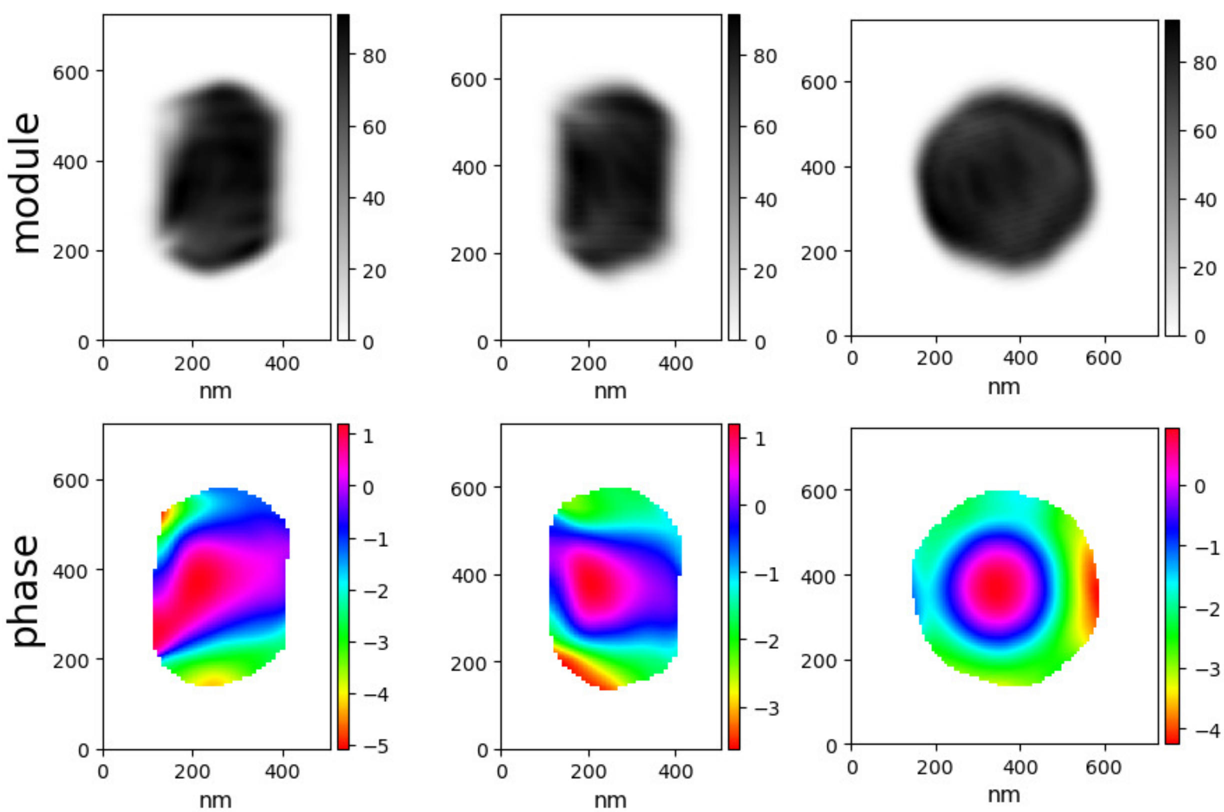

Supplementary Figure 10: Three central slices of Particle 1 of the main text reconstructed with the DL + PyNX method and orthogonalized.

S7. NOISE ROBUSTNESS

For completeness we have tested the DL model on simulated data for different noise conditions. Supplementary Figure 11 depicts the same diffraction pattern with an increasing amount of photons (from 3, to 10^6) and the respective ground truth and predictions. The model is not severely affected by the noise.

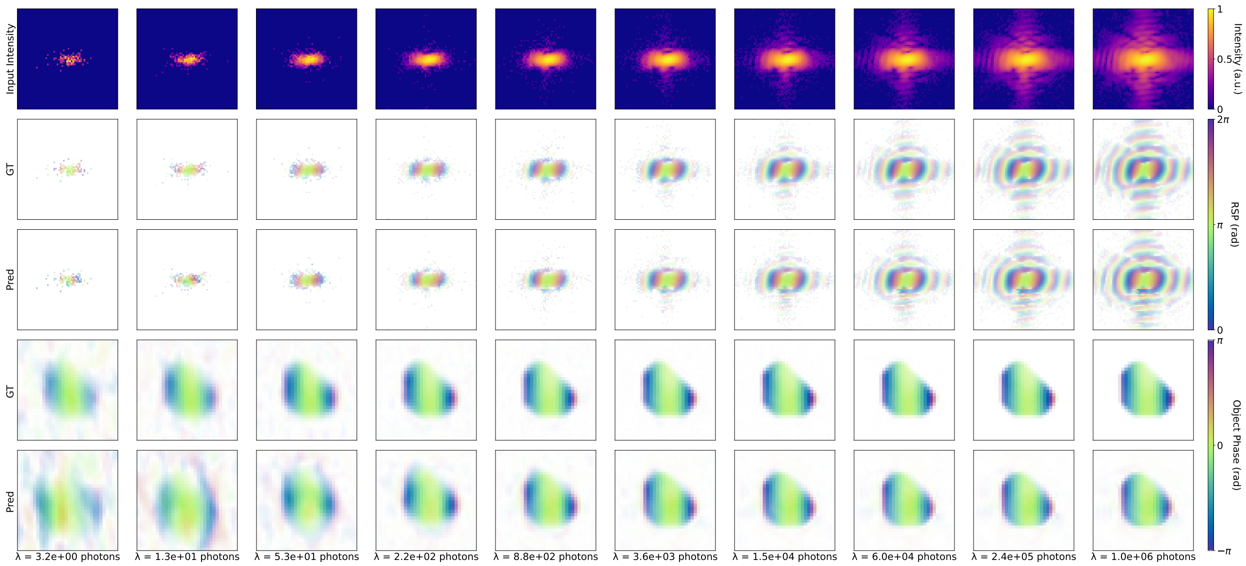

Supplementary Figure 11: DL model predictions for a particle with a Gaussian phase field simulated for different noise levels.

S8. SIZE ADJUSTMENT

The model accepts and predicts cubic arrays of the size of $64 \times 64 \times 64$ but experimentally, the size of each dataset can be different and usually larger than the volume handled by the model. For this reason, the combination of DL-model and ER refinement requires fundamental cropping/padding and resizing of the data. In particular, the procedure is here summarized in the Supplementary Figure 12.

It follows that, though the DL model works on a low-resolution version of the dataset, the ER refinement is operated using the original full-size data, thus overcoming the limitations caused by the initial data resizing to a $64 \times 64 \times 64$ pixel-size volume.

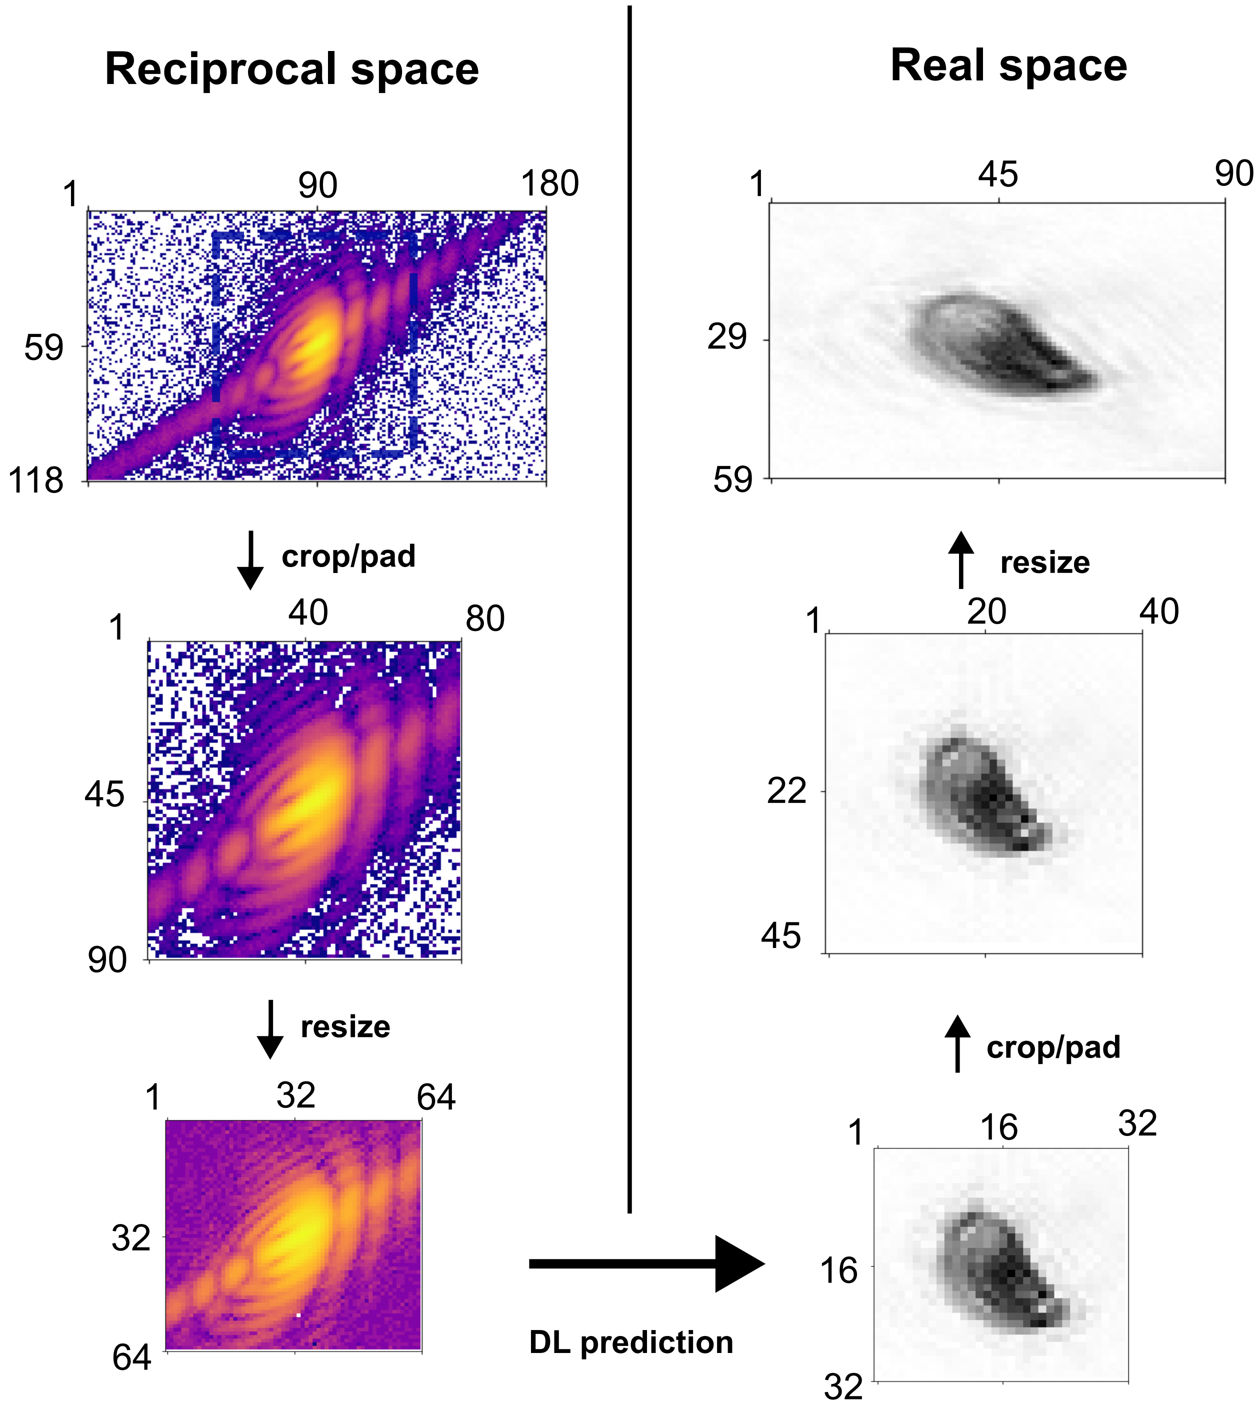

Supplementary Figure 12: Resizing procedure adopted for Particle 1 of the main text. Notice that cropping/padding in reciprocal space corresponds to interpolation in direct space and vice versa.

SUPPLEMENTARY REFERENCES

[1] Xi Yu, Longlong Wu, Yuewei Lin, Jiecheng Diao, Jialun Liu, Jörg Hallmann, Ulrike Boesenberg, Wei Lu, Johannes Möller, Markus Scholz, Alexey Zozulya, Anders Madsen, Tadesse Assefa, Emil S. Bozin, Yue Cao, Hoydoo You, Dina Sheyfer, Stephan Rosenkranz, Samuel D. Marks, Paul G. Evans, David A. Keen, Xi He, Ivan Božović, Mark P.M. Dean, Shinjae Yoo, and Ian K. Robinson. Ultrafast bragg coherent diffraction imaging of epitaxial thin films using deep complex-valued neural networks. *npj Computational Materials*, 10, 12 2024.

[2] M. Grimes, M.-I. Richard, K. Olson, E. Bellec, J. Zhao, and A. Viola. Bragg coherent diffraction imaging of palladium nanoparticles (version 1) [dataset], 2025. Includes data from _0015_scan_3.

[3] E. Bellec, S. Leake, M.-I. Richard, and J. Zhao. Bragg coherent modulation imaging for highly strained nanocrystal [dataset], 2027. [Dataset].

[4] A. Ulvestad, M. J. Welland, W. Cha, Y. Liu, J. W. Kim, R. Harder, E. Maxey, J. N. Clark, M. J. Highland, H. You, P. Zapol, S. O. Hruszkewycz, and G. B. Stephenson. Three-dimensional imaging of dislocation dynamics during the hydriding phase transformation. *Nature Materials*, 16(5):565–571, May 2017.

[5] Matthew J. Wilkin, Siddharth Maddali, Stephan O. Hruszkewycz, Anastasios Pateras, Richard L. Sandberg, Ross Harder, Wonsuk Cha, Robert M. Suter, and Anthony D. Rollett. Experimental demonstration of coupled multi-peak Bragg coherent diffraction imaging with genetic algorithms. *Physical Review B*, 103(21):214103, June 2021.