



Influence of loss function and electron dose on ptychography of 2D materials using the Wirtinger flow

Max Leo Leidl^{a,b,c}, Benedikt Diederichs^{a,d}, Carsten Sachse^{c,e}, Knut Müller-Caspary^{a,b,*}

^a Department of Chemistry and Centre for NanoScience, Ludwig-Maximilians-University Munich, Butenandtstrasse 11, Munich 81377, Germany

^b Ernst Ruska-Centre for Microscopy and Spectroscopy with Electrons (ER-C-1), Physics of Nanoscale Systems, Forschungszentrum Jülich, Wilhelm-Johnen-Straße, Jülich 52425, Germany

^c Ernst Ruska-Centre for Microscopy and Spectroscopy with Electrons (ER-C-3), Structural Biology, Forschungszentrum Jülich, Wilhelm-Johnen-Straße, Jülich 52425, Germany

^d Institute of Biological and Medical Imaging, Helmholtz Zentrum München, Ingolstädter Landstr. 1, Neuherberg 85764, Germany

^e Department Biology, Heinrich-Heine-Universität Düsseldorf, Universitätsstr. 1, Düsseldorf 40225, Germany

ARTICLE INFO

Keywords:

Transmission electron microscopy
4D-STEM
Phase retrieval
Low-dose TEM

ABSTRACT

Iterative phase retrieval is based on minimising a loss function as a measure of the consistency of an initial guess and underlying experimental data. Under ideal experimental conditions, real data contains Poissonian noise due to counting statistics. In this work, we use the Wirtinger Flow concept in combination with four common loss functions, being the \mathcal{L}_1 loss, the mean-squared error (MSE), the amplitude loss and the Poisson loss. Since only the latter reflects the counting statistics as an asymmetric Poisson distribution correctly, our simulation study focuses on two main cases. Firstly, high-dose momentum-resolved scanning transmission electron microscopy (STEM) of an MoS₂ monolayer is considered for phase retrieval. In this case, it is found that the four losses perform differently with respect to chemical sensitivity and frequency transfer, which we interpret in terms of the substantially different signal level in the bright and dark field part of diffraction patterns. Remedies are discussed using further simulations, addressing the use of virtual ring detectors for the dark field, or restricting loss calculation to the bright field. Secondly, a dose series is presented down to 100 electrons per diffraction pattern. It is found that all losses yield qualitatively reasonable structural data in the phase, whereas only MSE and Poisson loss range at the correct amplitude level. Chemical contrast is, in general, reliably obtained using the Poisson concept, which also provides the most continuous spatial frequency transfer as to the reconstructed object transmission function.

1. Introduction

Contemporary electron ptychography involves a large and still growing spectrum of methods to solve the phase retrieval problem using measured diffraction patterns in scanning transmission electron microscopy (STEM). Recent examples are the spatial resolution record in the field of TEM where reconstructed phase gratings exhibit real-space details at the scale of lattice vibrational amplitudes (Chen et al., 2021), the correction of probe positions (Thibault et al., 2009; Hurst et al., 2010), the recovery of the electron probe (Maiden and Maiden, 2009) or including spatial coherence (Thibault and Menzel, 2013) and inverse frozen phonon multislice, which allows the tracking of atomic positions with picometer resolution in a ferroelectric (Diederichs et al., 2024). Whereas inverse multislice approaches have extended the applicability

of ptychography to thick specimens, also the quantitative imaging of the atomic structure of 2D materials remains challenging. In this respect, both first moment STEM (Waddell and Chapman, 1979; Müller et al., 2014) and ptychography based on the projection assumption (Rodenburg et al., 1993; Rodenburg and Bates, 1655; McCallum and Rodenburg, 1992; Rodenburg and Faulkner, 2004; Maiden and Maiden, 2009) are applicable (Hofer and Pennycook, 2023; Müller-Caspary et al., 2018; Hofer et al., 2023), and both methods achieved a precision and accuracy such that the view into bonding effects becomes feasible (Martinez et al., 2023). In this respect, ptychography is becoming particularly attractive due to its capability of deconvolving probe and object transmission function. On the other hand, 2D materials can sometimes be very dose-sensitive. Since the propagation of noise through ptychographic reconstructions is complex, elucidating the dependence of reconstructed

* Corresponding author.

E-mail address: k.mueller-caspary@cup.lmu.de (K. Müller-Caspary).

<https://doi.org/10.1016/j.micron.2024.103688>

object transmission functions on the invested dose in a 4D-STEM experiment and on the ptychographic reconstruction settings is of central importance.

The noise model for an ideal diffraction pattern recording is known as Poisson shot noise due to the counting process during the detection of single electrons. In practice, further noise contributions are potentially present, for example from the electron source or electronics, which will be ignored in this study, where the ideal experimental setup is considered to explore the ultimate limits set solely by Poissonian noise. Many ptychography algorithms assume no specific noise model such as the single-side-band (SSB) reconstruction (Rodenburg et al., 1993) or Wigner distribution deconvolution (WDD) (Rodenburg and Bates, 1992; McCallum and Rodenburg, 1992). Other approaches feature a flexible choice of the noise statistics (Candes et al., 2015), or presume Gaussian white noise implicitly or explicitly, as, e.g., many iterative algorithms (Zhou et al., 2020; Lee et al., 2023). In particular, for measurements with more than five electrons per single pixel the Gaussian noise becomes a good approximation to the Poisson noise because the root-mean-square difference between both starts to become flat (Thompson, 2001). Therefore, the difference between both noise models is especially interesting for low-dose experiments, not only being relevant for the field of 2D materials, but also for imaging weakly scattering specimens such as covalent/metal-organic frameworks (COFs/MOFs) and proteins. Apart from that, even high-dose experiments may lead to diffraction patterns where the detection of more than five electrons per pixel is unlikely in the dark field. Because high-angle scattering is chemically sensitive, including respective solid angles in the reconstruction is favourable, such that the choice of the noise model is expected to make a difference in the reconstructed phases.

Importantly, gradient-based algorithms employing Wirtinger gradients such as the Wirtinger flow (WF) (Candes et al., 2015) are based on minimising a loss function, which can be chosen freely and, therefore, allow flexibility as to the noise distribution. It was found that the normalised root mean square error was around 11 % for the Gaussian loss function and around 8 % for the Poisson loss using WF in a computational study with a Gaussian sensing matrix (Li et al., 2022). The same work also showed that the Poisson loss function needs fewer measurements to reach the same or better performance than the Gaussian loss function. Moreover, WF using a Poissonian loss converges faster, as has been demonstrated for simulations in light optics using Fourier ptychography (Bian et al., 2016; Yeh et al., 2015).

Here, we present a systematic study of the impact of counting noise in 4D-STEM diffraction patterns on ptychographic reconstructions of a 2D MoS₂ monolayer employing the WF concept. We evaluate four different loss functions, namely the \mathcal{L}_1 loss, the \mathcal{L}_2 loss (also known as mean square error, MSE), the amplitude and Poissonian loss for simulated data using the phase gratings of simulations performed with electron energies of 300 keV and 60 keV as known ground truths. The focus lies on the 300 keV case because the knock-on damage, and thus dose sensitivity, becomes more significant than at 60 keV. Besides drawing general conclusions about using different loss functions, this work is thus relevant in practice when 2D materials are studied with comparably high incident electron energies, e.g., due to the currently better performing electron optics. In these cases, the dose should be as low as possible. Characteristics of the reconstructed amplitudes and phases of the object transmission function are worked out in both real and diffraction space in dependence on electron dose and focus. Special attention is drawn to cases which evaluate the Ronchigram region only, as well as both the Ronchigram and the dark field, the latter challenging reconstruction algorithms by involving signals with high and low counts simultaneously. In that respect, concepts such as partitioned and virtual dark field detectors are investigated as a remedy to improve the stability of the inversion and to exploit not only structural but also chemical contrast. To begin with, Sec. 2 introduces the theoretical background of WF-based object retrieval and the various loss functions. Section 3 summarises the used methods including simulation parameters, before

results are presented in Sec. 4. The paper closes with a discussion and central conclusions.

2. Theory

Ptychography aims to solve the phase retrieval problem and therefore tries to recover the object transmission function (OTF) $o(\mathbf{r})$ and the probe $a(\mathbf{r}, \mathbf{R})$ from a set of measured intensities in diffraction space. \mathbf{r} is the real space coordinate and \mathbf{R} the coordinate of the scan position. $\Psi(\mathbf{r}, \mathbf{R})$ is the specimen exit wave function for scan position \mathbf{R} . For STEM, the measured far-field diffraction pattern recorded at probe position \mathbf{R} reads

$$I(\mathbf{k}, \mathbf{R}) = |\mathcal{F}_{\mathbf{r}}[\Psi(\mathbf{r}, \mathbf{R})]|^2 = |\mathcal{F}_{\mathbf{r}}[a(\mathbf{r} - \mathbf{R})o(\mathbf{r})]|^2, \quad (1)$$

with the assumption that the exit wave is given by a multiplication of $o(\mathbf{r})$ and $a(\mathbf{r}, \mathbf{R})$ in real space. This means that multiple scattering and propagation inside the specimen can be neglected.

Different loss functions are used to compare the forward simulated intensities \hat{I}_i with the measured intensities I_i , where i is the index for the current scan position. According to the negative log-likelihood function, a measurement corrupted by Poisson noise shall be optimised by the loss function (Bian et al., 2016)

$$\mathcal{L}_p = \sum_i \hat{I}_i - I_i \ln(\hat{I}_i + \epsilon), \quad (2)$$

where $\epsilon \in \mathbb{R}^{>0}$ is a small positive number to avoid divergence of the natural logarithm. Depending on the starting value, the logarithm can diverge, and the Poisson loss can be difficult to handle for experimental data. Therefore, it is often assumed that the guessed OTF is a good approximation of the underlying true OTF ($\hat{o} \approx o$). In this case, a Taylor expansion of Equation 2 can be performed for the variable $\sqrt{\hat{I}_i}$ around $\sqrt{I_i}$ for each i . Ignoring all constant terms and scaling constants, the so-called amplitude loss function can be derived (Thibault and Guizar-Sicairos, 2012; Konijnenberg et al., 2018),

$$\mathcal{L}_A = \sum_i \left(\sqrt{\hat{I}_i} - \sqrt{I_i} \right)^2. \quad (3)$$

Recently, Melnyk has discussed in detail that the (extended) ptychographic iterative engine (PIE/ePIE) (Rodenburg and Faulkner, 2004; Maiden and Maiden, 2009) can be seen as a particular case of stochastic gradient descent using the \mathcal{L}_A loss instead of intensities (Melnyk, 2023).

Using the negative log-likelihood function (Thibault and Guizar-Sicairos, 2012) and assuming that the noise model can be approximated by Gaussian white noise, the loss function for the mean squared error (MSE) can be derived

$$\mathcal{L}_{\text{MSE}} = \sum_i (\hat{I}_i - I_i)^2, \quad (4)$$

which is also known as intensity or \mathcal{L}_2 loss. The assumption of Gaussian white noise means that the standard deviation is constant over all pixels. As a consequence, the Gaussian loss has the same weight for the whole dynamic range, which is different from the amplitude or the Poisson loss functions, whose weights scale with the square root of the measured intensities because the standard deviation $\sigma_i = \sqrt{I_i}$ is not constant (Yeh et al., 2015).

In literature, the \mathcal{L}_1 loss

$$\mathcal{L}_1 = \sum_i |\hat{I}_i - I_i| \quad (5)$$

is frequently used (Zhao et al., 2023; Zhou et al., 2023; Jiang et al., 2018; Zhang et al., 2019, 2021). Therefore, the \mathcal{L}_1 loss is included as the fourth criterion in our subsequent comparison. The difference to the MSE loss is the linear instead of the quadratic consideration of the difference, which

mainly affects outliers in a different manner.

The challenge is to determine $o(\mathbf{r})$ using a set of measurements as in eq. (1). This problem can be solved, for example, by iteratively using the gradient descent method with the introduced loss functions to optimise the initial guess of the OTF \hat{o} (Cauchy et al., 1847; Courant, 1943)

$$\hat{o}^{(t+1)}(\mathbf{r}) = \hat{o}^{(t)}(\mathbf{r}) - \mu \nabla_{\hat{o}} \mathcal{L}(\hat{I}(\mathbf{k}, \mathbf{R}), I(\mathbf{k}, \mathbf{R})) \quad (6)$$

Here, μ is the step size, $\nabla_{\hat{o}}$ the gradient with respect to the complex conjugate of \hat{o} , \mathcal{L} the loss function, \hat{I} the intensities calculated from the specimen guess $\hat{o}^{(t)}$ and t the index of the current iteration. One iteration stands for an update employing all scan positions.

The gradients of \mathcal{L}_1 (Jiang et al., 2018), \mathcal{L}_{MSE} (Candes et al., 2015), \mathcal{L}_A (Thibault and Guizar-Sicairos, 2012) and \mathcal{L}_P (Thibault and Guizar-Sicairos, 2012) are given by:

$$\nabla_{\hat{o}} \mathcal{L}_1 = \sum_i \nabla_{\hat{o}} \hat{I}_i \left(\frac{1}{2} \frac{\hat{I}_i - I_i}{|\hat{I}_i - I_i|} \right), \quad (7)$$

$$\nabla_{\hat{o}} \mathcal{L}_{\text{MSE}} = 2 \sum_i \nabla_{\hat{o}} \hat{I}_i (\hat{I}_i - I_i), \quad (8)$$

$$\nabla_{\hat{o}} \mathcal{L}_A = \sum_i \nabla_{\hat{o}} \hat{I}_i \left(1 - \frac{\sqrt{\hat{I}_i}}{\sqrt{I_i}} \right), \quad (9)$$

$$\nabla_{\hat{o}} \mathcal{L}_P = \sum_i \nabla_{\hat{o}} \hat{I}_i \left(1 - \frac{I_i}{\hat{I}_i + \epsilon} \right). \quad (10)$$

Note again the divergence for $\hat{I}_i \rightarrow 0$ and $\epsilon = 0$. Finally, the gradient of the inner differentiation, being the same for all four loss functions, reads

$$\nabla_{\hat{o}} \hat{I}_i(x) = \alpha^* (\mathbf{r} - \mathbf{R}) \cdot \mathcal{F}^{-1}(\mathcal{F}_{\mathbf{r}}[\Psi(\mathbf{r}, \mathbf{R})] \cdot x). \quad (11)$$

3. Methods

An MoS₂ monolayer is considered as a test sample for the simulations. A supercell of 18×10 unit cells (Fig. 1 a) was generated, each with a dimension of $3.19 \text{ \AA} \times 5.53 \text{ \AA}$. The supercell was centred in a larger, square simulation box which had a size of $72.6 \text{ \AA} \times 72.6 \text{ \AA}$ sampled with 2048×2048 pixels. In the forward direction, the multi-slice approach (Cowley and Moodie, 1957) was used with one slice per atomic layer, amounting to three slices in total.

The scan consists of 128×128 scan positions with a STEM pixel size of $0.25 \lambda \alpha^{-1} = 27.3 \text{ pm}$, where λ is the electron wavelength and α the convergence semi-angle in rad. This STEM pixel size fulfils the Nyquist theorem for frequencies up to twice the bright field (BF) aperture radius at in-focus conditions without additional aberrations (Rodenburg et al., 1993). The STEM simulations were performed with an acceleration voltage of 60 kV and 300 kV, respectively, at in-focus conditions as well




Fig. 1. MoS₂ structure and dose effects on diffraction patterns. a) Projection of one unit cell of MoS₂. b) - d) Full diffraction patterns that are used for the reconstructions. The cases are given in the labels, and the maximal probability/count number for the shown pattern is given in each quadrant. In the mathematical positive sense from the top left given in electrons per diffraction pattern: infinite, high 10^6 , medium 10^4 and low 10^2 .

as with an under-focus of 25 nm. The convergence semi-angle was chosen to 18 mrad and 44.5 mrad for 300 kV and 60 kV, respectively. Thus, the maximum spatial frequency transfers were identical in both settings.

As a basis for the reconstruction, the 4D data set in eq. (1) was created by the central patch of 256×256 pixels only which has been cut from the full simulated diffraction patterns of the forward simulation with size 2048×2048 . Note that the total probability that an electron hits the central, cropped part is then smaller than one which has been taken into account properly when finite doses have been added to the data. This mimics both the sampling and array size of contemporary pixelated diffraction cameras, and the fact that real recordings are incomplete in that electrons scattered beyond the solid angle covered by the detector are lost in experiments. Here, this corresponds to a maximum spatial frequency of 1.749 \AA^{-1} which equals cutoff angles of 34.4 and 85.2 mrad for 300 and 60 keV electrons, respectively. To account for the Poisson shot noise, 10^6 , 10^4 and 10^2 electrons were distributed on each diffraction pattern for the high-, medium- and low-dose case using the Poisson distribution, which results in $1.34 \times 10^7 e^-/\text{\AA}^2$, $1.34 \times 10^5 e^-/\text{\AA}^2$ and $1.34 \times 10^3 e^-/\text{\AA}^2$ for 300 kV. The diffraction pattern of the first scan point is exemplarily shown in Fig. 1 b-d for the different forward simulations. No additional noise sources were assumed so as to deal with an ideal experiment being only limited by counting statistics.

A single-slice model was assumed for the reconstructions, and the Wirtinger Flow (WF) (Candes et al., 2015) was used with different loss functions. In experiments, the probe can, in principle, be retrieved at high dose even in a damaged specimen area. For this case, we found that the defocus converges from an initialised value of 0 \AA to $-100 \pm 5 \text{ \AA}$ for all loss functions in our simulation study where the true defocus of -100 \AA was known. A parameterised probe has been used (Diederichs et al., 2024) for this reconstruction, as well as a specimen thickness of 5.88 \AA . The use of a parameterised probe avoids scaling ambiguities between the reconstructed specimen and the reconstructed probe as described in refs. (Fannjiang and Chen, 2020; Li et al., 2016), because the amplitude of the probe is fixed. Due to the accurate probe retrieval, e.g., from high-dose data, we assume the probe being known in the following and focus on the effect of the dose and the loss concepts on the reconstruction of the specimen. Vacuum was used as the initial guess for the initialisation and $\epsilon = 10^{-4}$ for the Poisson loss function. The losses have been calculated by using the full available diffraction patterns of 256×256 pixels, however, the reciprocal space of the reconstructions extended up to eight times larger spatial frequencies to allow for super-resolution. The update step size was constant and set such that the amplitude change of the specimen in the first epoch was between 1×10^{-3} and 5×10^{-2} . We made this choice to avoid the additional complexity of variable step size, which allows for faster convergence. To ensure convergence and to avoid overshoots we monitored the loss for every epoch. For all cases, 100 iterations were performed. The results are shown for the last iteration.

A normalisation is needed to compare the forward simulation, which calculates the probability distribution on the detector, with the (simulated) measurements, which count the electrons on the detector. The only tested loss function not invariant under scaling is the Poisson loss in eq. (2). To obtain a diffraction pattern \hat{I} being quantitatively comparable to the measurement I , each diffraction pattern obtained during the reconstruction process via the forward simulation employing the current object guess \hat{o} is first cut to the size of the detector to match the dimensions of the recordings, and then multiplied by the number of electrons in the recording I . To assure conservation of the number of electrons in a setup where only a part \hat{I} of a model diffraction pattern is used, the number of electrons is multiplied by the total probability divided by the probability landing on the detector.

Four virtual detector geometries are tested for an acceleration voltage of 300 kV and the high-dose case. First, the whole pixelated

detector is used to calculate the loss. Second, the BF disk is used in combination with 64 virtual ring detectors for the dark field (DF). The aim is to avoid the sparsity of the electron counts in the dark field and to reduce the dynamic range by reducing the difference between the BF and the DF due to the integration. The number of iterations was increased to 600 because the reconstructions converge more slowly and for the first 100 iterations only the BF was used. Third, only the pixels inside the BF disk are used as input for the loss functions to test whether the dynamic range between BF and DF is responsible for the differences in the reconstructions, or if the differences can also be seen for the pixels inside the BF for the high-dose case. As has been shown by Yang et al. (Yang et al., 2015), substantial binning of the detector does not affect the quality of the reconstruction for in-focus measurements significantly. Therefore, we tested a binning of the detector from 256×256 down to 32×32 for the reconstruction with the \mathcal{L}_P loss for the high- and low-dose case, as a fourth detector geometry.

The projected potential scaled by the relativistic constant is used as the ground truth. The global phase to the ground truth is minimised before the comparisons. All computations were performed with an in-house written software package (Diederichs et al., 2024) that is based on the PyTorch library (Paszke et al., 2019).

4. Results

4.1. Gradient characteristics of different losses

Before considering explicit reconstructions for different loss functions, we shed light on their update characteristics during the optimisation according to eq. (6). In particular, the Fourier-transform with respect to the specimen coordinate \mathbf{r} reads

$$\begin{aligned} \mathcal{F}_{\mathbf{r}}[\hat{o}^{(t+1)}(\mathbf{r})] &= \mathcal{F}_{\mathbf{r}}[\hat{o}^{(t)}(\mathbf{r})] \\ &- \mu \mathcal{F}_{\mathbf{r}}[\nabla_{\hat{o}} \mathcal{L}(\hat{I}(\mathbf{k}, \mathbf{R}), I(\mathbf{k}, \mathbf{R}))] \end{aligned} \quad (12)$$

and shows that the reconstructed spatial frequencies directly depend on the spatial frequencies of the gradient. Consequently, one expects that the different loss functions from sec. 2 update the Fourier coefficients of the object function \hat{o} in a significantly different manner. The effect of the chosen loss function on the gradient can be seen in Figs. 2 and 3, corresponding to a scan position on an Mo atom and for an interstitial scan point in the centre of the hexagon, respectively. For illustration purposes, only the high-dose case is considered here, and the gradient after 10 iterations (Yeh et al., 2015) has exemplarily been chosen. Note the linear scaling in the left half and the logarithmic plot in the right half. The circles represent single and twice the Ronchigram radius. All data is normalised for visualisation, with the lower left quadrant normalised such that only spatial frequencies above $2 \cdot \alpha$ are taken into account.

Considering the gradients for the probe position on Mo in Fig. 2 first, the \mathcal{L}_1 and \mathcal{L}_{MSE} losses lead to similar updates of spatial frequencies, concentrated to the BF disc and faintly extending beyond $2 \cdot \alpha$. In contrast, \mathcal{L}_A and \mathcal{L}_P promote a broad spatial frequency band that is maximum at the BF radius and extends to twice the BF radius and slightly beyond as well. Whereas all gradients have spatial frequencies above twice the BF radius, as shown in the lower left quadrant of each panel, the high frequencies are strongly suppressed in the case of \mathcal{L}_{MSE} and are negligible compared to the frequencies within the BF disc, as seen in panel (b). A quantitative evaluation shows that the suppression of high spatial frequencies is lower for \mathcal{L}_1 . Considering \mathcal{L}_A and \mathcal{L}_P , the gradients extend significantly further in Fourier space than the \mathcal{L}_1 or \mathcal{L}_{MSE} counterpart, again looking at the lower left quadrants of Fig. 2 (a-d).

Although the power spectra of the gradients naturally look different for the STEM probe positioned in the centre of a hexagon of the MoS₂ lattice in Fig. 3, some of the conclusions from Fig. 2 still hold. The \mathcal{L}_{MSE} loss is the most compact one in Fourier space not reaching beyond $2 \cdot \alpha$, the \mathcal{L}_1 , \mathcal{L}_A and \mathcal{L}_P losses update spatial frequencies partly




Fig. 2. Spatial frequency characteristics of the gradient for an on-atom probe. Power spectra of the gradient after 10 iterations corresponding to a probe on an Mo atom for the different loss functions as noted in the caption. The left half is scaled linearly, the right adopts a logarithmic scale. The inner ring indicates the bright field (BF) radius α , the outer ring twice the BF radius. Because the absolute value of the gradient only becomes meaningful in combination with the learning rate, the shown gradients are separately normalised for the left and right half to show values between 0 and 1. The lower left corner shows an enhancement of frequencies outside twice the bright field (BF) radius α and is normalised separately.

significantly beyond $3 \cdot \alpha$. Among these, \mathcal{L}_P implies the most homogeneous characteristics in Fourier space and best reflects the sixfold symmetry of the scan site. A faint three-fold symmetry can be observed in Fig. 2, too, recognizable mostly for the gradient of the Poisson loss in (d). Of course, exploring the tenth iteration and considering only two scan points is a case study of limited generalisability. Nevertheless, the respective guesses of the OTFs, which are used to calculate the gradients, vary significantly between the loss functions. Up to here, studies took place at the high-dose limit, such that the observed different characteristics are rather fundamental and not related to dose effects, which will be elucidated next.

4.2. Reconstruction overview

In the following subsections, we study the impacts of various parameters and settings on reconstructions of an MoS₂ monolayer from 4D-STEM data in a recurrent systematic manner. We show the reconstructed amplitudes and phases for the four losses in figure parts (a) and (b). Figure parts (c) always depict the power spectra of the reconstructed OTFs, and (d) shows line profile averages across Mo-S₂ pairs obtained in the reconstructed phases so as to explore the presence of chemical

sensitivity. We first address the pixel-wise calculation of the losses employing the full diffraction patterns. Even for the high-dose case considered here, this includes both the high intensities in the BF and rather low ones in the DF. The sparsity of the latter motivate the second and third studies, in which the DF is treated as annular rings or fully neglected, respectively. Different focus settings are discussed. Finally, reconstructions for a low dose case are assessed.

4.3. Reconstructions on a pixelated basis

The comparison of reconstructions in which the loss has been calculated pixel-wise from the full diffraction pattern of 256×256 pixel size is shown in Fig. 4. An acceleration voltage of 300 kV was used with probe settings at in-focus conditions, and the high-dose case was applied to generate Poisson noise within the 4D-STEM data. The amplitudes obtained from the \mathcal{L}_1 and \mathcal{L}_{MSE} loss functions on the one hand, and from the \mathcal{L}_A and \mathcal{L}_P losses on the other hand are similar, as seen in panel (a). Fig. 4 b demonstrates that all loss concepts yield reliable phases that represent the structure excellently. It is to note that the dynamic range is nearly a factor of two larger for the \mathcal{L}_P loss. This is also obvious in the power spectra in Fig. 4 c, where the Poisson-based




Fig. 3. Spatial frequency characteristics of the gradient for an interstitial probe. Power spectra of the gradient after 10 iterations corresponding to a probe at an interstitial position in the centre of the hexagon for the different loss functions as noted in the caption. The left half is scaled linearly, the right adopts a logarithmic scale. The inner ring indicates the bright field (BF) radius α , the outer ring twice the BF radius. Because the absolute value of the gradient only becomes meaningful in combination with the learning rate, the shown gradients are separately normalised for the left and right half to show values between 0 and 1. The lower left corner shows an enhancement of frequencies outside twice the bright field (BF) radius α and is normalised separately.

solution yields the highest spatial frequencies. Moreover, \mathcal{L}_P is the only loss which does not pronounce frequencies within a disc with twice the BF radius. The power spectrum of the \mathcal{L}_{MSE} decays the fastest with increasing frequencies, followed by \mathcal{L}_1 . The disk marking twice the BF radius can be seen for \mathcal{L}_{MSE} . The amplitude and Poisson loss functions show reflections up to twice the frequency recorded by the detector.

Chemical contrast between molybdenum and the two sulphur atoms is only visible in the amplitude for all loss functions and in the phase for \mathcal{L}_A and \mathcal{L}_P , which is most obvious in the line profiles of the phases in Fig. 4 (d). The peaks of the Poisson loss function are around 1/3 and for the other loss functions around 1/4 of the peak height of the ground truth.

Except for minor differences, the results for an acceleration voltage of 60 kV, in-focus conditions and the high dose setting in Fig. 5 are similar to the 300 kV reconstructions. A faint difference in the phases of \mathcal{L}_1 and \mathcal{L}_{MSE} can be noticed. In general, the Poisson loss yields the clearest fingerprint of chemical contrast in the phase as in Fig. 4, whereas the phase difference between the Mo and S₂ site is higher for 60 keV electrons.

In summary, the Poisson loss performs best for a pixel-wise loss calculation using both bright and dark field data, followed by the

amplitude loss. The \mathcal{L}_1 criterion still yields significant spatial frequencies beyond twice the convergence semi-angle, which appears as a limit for the \mathcal{L}_{MSE} in the present case.

4.4. Reconstructions augmented with annular dark field detectors

From counting statistics, one expects the Poisson loss to make a difference when signals are low. Although the preceding section considered the high-dose case, the DF exhibits a rather sparse distribution of single electron events according to Fig. 1. This could be one explanation for the better performance of \mathcal{L}_P above. On the other hand, the intensity distribution in the DF is usually much smoother compared to the Ronchigram, such that the dominant information is rather gained from the modulus of the momentum than from its direction. Consequently, it is now studied how well a hybrid model performs in which the BF contributes to the losses pixel-wise, the DF in ring-wise integrated manner.

Assuming an acceleration voltage of 300 kV and using the pixelated BF disc in combination with 64 virtual ring detectors for the dark field, the reconstructions in Fig. 6 are obtained. Compared to the full pixelated detector, the differences between the loss functions become smaller. The




Fig. 4. Reconstruction on pixelated basis at high dose (300 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^6 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra the minimum is zero in all cases and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.

line profile of the phase in panel (d) shows that the peak heights of \mathcal{L}_1 , \mathcal{L}_{MSE} and \mathcal{L}_P are similar and all reconstructions show the chemical contrast faintly. In contrast to the full pixelated case, the amplitude loss function shows significantly higher peaks and nearly no chemical contrast.

Figure 6 c shows the similarities also in the power spectra with a slightly faster decay for the MSE loss for the high frequencies. Interestingly, the power spectra of the reconstructions performed with the \mathcal{L}_1 and \mathcal{L}_{MSE} losses now contain spatial frequencies substantially above twice the BF radius while the power spectra obtained with \mathcal{L}_A and \mathcal{L}_P decay slightly faster. A reason for this observation might be the artificial introduction of isotropy in the DF via rings compared to the anisotropic frequencies measured by a square detector. In fact, the Poisson result got even slightly worse in terms of the maximum phase difference, being 0.84 rad in Fig. 4 b and 0.8 rad in Fig. 6 b. This indicates that the Poisson loss is not reliant on reducing the sparsity of counts via integrating over solid angles in which count rates are low, whereas \mathcal{L}_1 and \mathcal{L}_{MSE} profit

significantly.

4.5. Reconstructions using the Ronchigram only

The investigations above point towards a dependence of reconstructions using different loss concepts on how the DF intensity is treated. We thus turn towards the high-count regime and study to which extent reconstructions differ among the different losses calculated pixel-wise from the BF region only. The results are shown in Fig. 7 for identical settings as before, i.e., 300 keV electrons and in-focus conditions. Indeed, the reconstructions for the different loss functions become nearly identical as to amplitudes, phases and spatial frequency characteristics. The chemical contrast can only be seen in the amplitude. The differences in the power spectra also become negligible, however, all reconstructions show only frequencies up to twice the BF radius.

We interpreted the different outcomes of the loss functions in the previous sections as originating from the high dynamic range in the




Fig. 5. Reconstruction on pixelated basis at high dose (60 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^6 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra the minimum is zero in all cases and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.

experimental diffraction patterns containing BF and DF. The BF exhibits only weak intensity modulations for thin monolayer specimens considered here, which could explain the identical outcomes among all four losses. Therefore, it is interesting to analyse BF data with more structure in the Ronchigram, which we introduced by an under-focus of 25 nm as exemplarily shown in Fig. 1 c. Defocusing increases the measured dynamic range as seen from the broadening of the second peak in Fig. 8 and is a frequently used experimental setting for phase retrieval. Figure 9 shows the reconstructions for the defocused scan. The results are comparable to the case without a defocused probe with small differences in the amplitude, which demonstrates that, for high-dose settings, all four loss functions yield the same OTFs if only the Ronchigram is evaluated.

4.6. Reconstructions at medium and low doses

After working out the inherent and loss function characteristics irrespective of the recorded dose, we now shed light on the practical aspects. In particular, the strategy of solely pixel-wise loss calculation

using the whole diffraction space from sec. 4.3 is applied while the dose is lowered.

The medium dose level of 10^4 electrons per diffraction pattern leads to the reconstructions shown in Fig. 10. Compared to the high-dose case in Fig. 4, fewer electron counts affect all reconstructions significantly, except the Poisson result. The \mathcal{L}_1 reacts on counting noise by over-estimating the amplitude in Fig. 10 a, \mathcal{L}_P underestimates it and \mathcal{L}_{MSE} yields a noisy amplitude distribution, however, being one on average. Moreover, the atomic structure becomes less pronounced, although still resolved qualitatively in all cases with regards to the phase in Fig. 10 b. Chemical contrast appears only in the result based on the Poisson loss, as is obvious from the line profiles in Fig. 10 d. Note that the dynamic range of the phase is identical to the high dose case for \mathcal{L}_P although the dose was reduced by two orders of magnitude.

In the power spectra in Fig. 10 c, the \mathcal{L}_1 and \mathcal{L}_{MSE} loss only show meaningful frequencies up to twice the BF radius, meaning that the resolution obtainable by \mathcal{L}_1 decreased with the dose. The amplitude and Poisson losses maintain their broad frequency transfer band far




Fig. 6. Reconstruction based on pixelated BF and 64 DF rings at high dose (300 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^6 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra the minimum is zero in all cases and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.

beyond twice the BF radius, whereas the amplitude loss tends to pronounce noise within the $2 \cdot \alpha$ frequency band as a limited dose effect. Of course also the \mathcal{L}_P based power spectrum gets noisier, but the noise distribution is continuous and thus more realistic up to the highest reconstructed spatial frequencies.

Decreasing the dose further by another factor of 100 results in the reconstructions compiled in Fig. 11. This low-dose case with 10^2 electrons per diffraction pattern shows a dramatic underestimation of the mean amplitude for the \mathcal{L}_1 and \mathcal{L}_A losses. Meanwhile, the MSE and Poisson loss function's mean amplitude stays around one. However, no meaningful structural contrast is remaining in any of the amplitude distributions in Fig. 11 a. Nevertheless, slight phase contrast remains visible for \mathcal{L}_1 and \mathcal{L}_A in Fig. 11 b. For the MSE and Poisson loss, the phase contrast is highest but lower than for the high-dose cases. None of the loss functions shows reliable chemical contrast between Mo and S in the phase anymore, as seen in Fig. 11 d.

The \mathcal{L}_1 and amplitude loss functions show only the first reflections

in the power spectra in Fig. 11 c. In contrast, \mathcal{L}_{MSE} and \mathcal{L}_P show reflections up to twice the BF radius. However, the reflections start to fade in the noisy background. The Poisson loss function shows the highest values above twice the BF radius, but no further structural reflections there. Thus, only high-frequency noise is reconstructed for frequencies larger than twice the BF radius, and also the \mathcal{L}_P norm tends to pronounce noise within the $2 \cdot \alpha$ radius now. Comparing with the medium dose in in Fig. 10 c, the Poisson-based reconstruction becomes the most promising one. Although it picks the noise with frequencies well beyond $2 \cdot \alpha$ at low doses, it will be the first approach among those considered here, which will be able to reconstruct true spatial frequencies of structural origin above the noise level there.

Figure 11 e-h shows the reconstruction using the Poisson loss \mathcal{L}_P for the high- and low-dose case, whereas the detector has been binned down to 32×32 pixels. Compared to the 256×256 detector, the amplitude contrast appears weaker in panel e with random contrast between the atom types. On the contrary, the reconstructed phase in panel f depicts




Fig. 7. Reconstruction restricted to the pixelated BF at high dose (300 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^6 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra the minimum is zero in all cases and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.




Fig. 8. Effect of focus on dynamic range. Distribution of electron counts for the full simulated 4D-STEM scan for 300 keV using an in-focus setting and an under-focus of 25 nm. The first peak of both distributions shows the dark field counts and the second, broader peak the bright field counts.

the atomic structure clearly. Compared to the low-dose case in Fig. 11 b, the atomic peaks in Fig. 11 f are sharper, although the atoms are partly elongated in random directions. In the power spectrum in Fig. 11 g, the high-frequency peaks have a slightly better signal-to-noise ratio than in panel c, and the disc edge at twice the maximum bright field radius is no longer visible for the low-dose case. The line profile in panel h shows slightly broader peaks for the high-dose case and slightly higher peaks for the low-dose case than for the 256×256 detector. Thus, at in-focus conditions, a small number of detector pixels is indeed sufficient (Yang et al., 2015) and can partly even improve the reconstruction.

In Fig. 12, the low-dose case with 10^2 electrons per diffraction pattern is shown for an acceleration voltage of 60 kV. Observations similar to the 300 keV case in Fig. 11 are made. The reconstructions with \mathcal{L}_1 and \mathcal{L}_A show an underestimation of the mean amplitude while the reconstructions with \mathcal{L}_{MSE} and the \mathcal{L}_P show slightly lower values than in the high-dose case. The phases in panel b are rather comparable with higher peaks for the reconstruction with \mathcal{L}_A and \mathcal{L}_P . In the power




Fig. 9. Reconstruction using BF data with 25 nm defocus at high dose (300 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^6 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra the minimum is zero in all cases and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.

spectrum, the disk containing the noisy background has a larger radius than for 300 keV. An exception is found for the reconstruction with \mathcal{L}_P , where a smooth decay of the Fourier coefficients with increasing spatial frequencies can be observed. The reflections corresponding to the crystal lattice show a better signal-to-noise ratio than in the 300 keV case, especially at higher frequencies. The average line profile of the phase shows an inverted contrast between the atom types for the reconstructions with \mathcal{L}_{MSE} and \mathcal{L}_P . However, the difference is statistically not significant.

5. Discussion

The amplitude and Poisson losses lead to a more homogeneous frequency update over the whole frequency space than \mathcal{L}_1 and \mathcal{L}_{MSE} as to the gradient according to Figs. 2 and Figs. 3. Concerning reconstructions of the atomic lattice, the intensity difference between the bright and dark field matters significantly concerning the decision which loss

function to take. Whereas a pixel-wise comparison of the full diffraction pattern including bright and dark field intensity works reliably for the Poisson loss, the results of the other loss functions can be stabilised by using virtual ring detectors for the dark field. For the MoS_2 monolayer studied here, this resulted in a mean electron count of 102 electrons for the inner and 26 for the outermost ring, such that, e.g., Gaussian statistics applies. The difference between the loss functions disappears besides minor differences in the amplitude and loss of chemical sensitivity if the reconstruction is limited to the pixels of the BF disc. Introducing more features into the BF disc by defocusing the electron probe, the variation of the number of electrons per pixel is too small to cause significant differences between the loss functions, which is beneficial for applications. Note that constant differences between the loss functions are compensated by the learning rate.

The chemical contrast between the two sulphur and the molybdenum atoms is small. As a result, the difference can only be seen in the amplitude for all loss functions in the high-dose case. Here, the




Fig. 10. Reconstruction on pixelated basis at medium dose (300 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^4 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra the minimum is zero in all cases and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.

amplitude of the OTF accounts for the scattering beyond the used detector. Only the amplitude and Poisson loss show the chemical contrast in the phase and only the Poisson loss for the medium dose. Introducing virtual rings in the dark field reduces the difference of the measured electrons between the bright field and the rings in the dark field. As a result, the frequency transfer of the different loss functions becomes similar.

A reduction of the dose reduces the reconstruction quality of the \mathcal{L}_1 and amplitude loss significantly (Figs. 10 to Figs. 12), which is strongly reflected in the deviation of the mean amplitude from one. This problem does not appear for the MSE and Poisson losses or the high-dose case. In general, our observations are consistent with the weightings due to the standard deviation assumed for the different numbers of electrons per pixel. The results are in agreement with the findings of Yeh et al., who show a flatter frequency transfer of the gradients for the amplitude and Poisson loss functions compared to the MSE loss (Yeh et al., 2015). They used Fourier ptychography for light optics and also discussed the

different standard deviations as a cause.

The findings as to the amplitude loss might become relevant to interpret phase retrieval results at low dose obtained with the ePIE algorithm (Maiden and Maiden, 2009) because it has been shown that ePIE basically exploits the amplitude loss function (Godard et al., 2012; Melnyk, 2023; Konijnenberg et al., 2018). Finally, the impact of the choice of a respective loss function can also be reduced by using momentum-based update schemes such as Adam (Kingma and Ba, 2014; Diederichs et al., 2024) instead of the gradient descent applied in the present paper.

6. Summary and Conclusion

For iterative phase retrieval, the Wirtinger Flow approach provides sufficient freedom to choose the loss function as an estimator to minimise the difference between the model and experiment in a flexible manner. The loss function, in combination with the evaluated solid




Fig. 11. Reconstruction on pixelated basis at low dose (300 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^2 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra, the minimum is zero in all cases, and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness. (e)-(h) Results obtained for a binning of the detector from 256×256 pixels down to 32×32 for \mathcal{L}_P . *Left:* High-dose case. *Right:* Low-dose case.




Fig. 12. Reconstruction on pixelated basis at low dose (60 keV). The data shows results for in-focus STEM with a dose per diffraction pattern of 10^2 electrons. The loss function for each quarter starting from the top left in a clockwise direction: \mathcal{L}_1 , \mathcal{L}_{MSE} , \mathcal{L}_A and \mathcal{L}_P . Panels show reconstructed (a) amplitudes, (b) phases, (c) OTF power spectra and (d) average of eight line profiles of the phase with the ground truth as a black dashed line. Numbers indicate minimum and maximum of the colourscale. For the power spectra, the minimum is zero in all cases, and the indicated ring shows frequencies corresponding to α . The line profiles are shifted by 0.5, 1.0, 1.5 and 2.0 for a better distinctness.

angle, can have a decisive impact on the outcome. In particular, using the Poisson loss appears robust against such settings and limited dose effects, in contrast to the amplitude, \mathcal{L}_1 loss and mean-squared error. To a certain extent, integrating the sparse dark field counts within virtual ring detectors can stabilise their results such that the number of counts varies less across the detector.

CRediT authorship contribution statement

Knut Müller-Caspary: Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Carsten Sachse:** Supervision, Resources, Funding acquisition. **Benedikt Diederichs:** Software, Methodology. **Max Leo Leidl:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation.

Declaration of Competing Interest

The authors declare no competing financial interests.

Data Availability

No data was used for the research described in the article.

Acknowledgements

Funding from the European Research Council within the Horizon Europe Innovation Funding Programme under Grant Agreement 101118656 (ERC Synergy project 4D-BioSTEM) is kindly acknowledged. Further funding has been gratefully received from the Helmholtz Association (Germany) under contract VH-NG-1317.

References

- Bian, L., Suo, J., Chung, J., Ou, X., Yang, C., Chen, F., Dai, Q., 2016. Fourier ptychographic reconstruction using poisson maximum likelihood and truncated wirtinger gradient. *Sci. Rep.* 6 (1), 27384.
- Candes, E.J., Li, X., Soltanolkotabi, M., 2015. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Trans. Inf. Theory* 61 (4), 1985–2007.
- Cauchy, A., et al., 1847. Méthode générale pour la résolution des systèmes d'équations simultanées. *Comp. Rend. Sci. Paris* 25 (1847), 536–538.
- Chen, Z., Jiang, Y., Shao, Y.-T., Holtz, M.E., Odstrčil, M., Guizar-Sicairos, M., Hanke, I., Ganschow, S., Schlom, D.G., Müller, D.A., 2021. Electron ptychography achieves atomic-resolution limits set by lattice vibrations. *Science* 372 (6544), 826–831.
- R. Courant, Variational methods for the solution of problems of equilibrium and vibrations (1943).
- Cowley, J.M., Moodie, A.F., 1957. The scattering of electrons by atoms and crystals. i. a new theoretical approach. *Acta Crystallogr.* 10 (10), 609–619.
- Diederichs, B., Herdegen, Z., Strauch, A., Filbir, F., Müller-Caspary, K., 2024. Exact inversion of partially coherent dynamical electron scattering for picometric structure retrieval. *Nat. Commun.* 15 (1), 101.
- Fannjiang, A., Chen, P., 2020. Blind ptychography: uniqueness and ambiguities. *Inverse Probl.* 36 (4), 045005.
- Godard, P., Allain, M., Chamard, V., Rodenburg, J., 2012. Noise models for low counting rate coherent diffraction imaging. *Opt. Express* 20 (23), 25914–25934.
- C. Hofer, J. Madsen, T. Susi, T.J. Pennycook, Detecting charge transfer at defects in 2d materials with electron ptychography, arXiv preprint arXiv:2301.04469 (2023).
- Hofer, C., Pennycook, T.J., 2023. Reliable phase quantification in focused probe electron ptychography of thin materials. *Ultramicroscopy* 254, 113829. <https://doi.org/10.1016/j.ultramic.2023.113829>.
- Hurst, A., Edo, T., Walther, T., Sweeney, F., Rodenburg, J., 2010. Probe position recovery for ptychographical imaging. *J. Phys.: Conf. Ser.* 241 (1), 012004.
- Jiang, S., Guo, K., Liao, J., Zheng, G., 2018. Solving fourier ptychographic imaging problems via neural network modeling and tensorflow. *Biomed. Opt. Express* 9 (7), 3306–3319.
- D.P. Kingma, J. Ba, A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- Konijnenberg, A., Coene, W., Urbach, H., 2018. Model-independent noise-robust extension of ptychography. *Opt. Express* 26 (5), 5857–5874.
- Lee, J., Lee, M., Park, Y., Ophus, C., Yang, Y., 2023. Multislice electron tomography using four-dimensional scanning transmission electron microscopy. *Phys. Rev. Appl.* 19 (5), 054062.
- Li, P., Edo, T., Batey, D., Rodenburg, J., Maiden, A., 2016. Breaking ambiguities in mixed state ptychography. *Opt. Express* 24 (8), 9038–9052.
- Li, Z., Lange, K., Fessler, J.A., 2022. Poisson phase retrieval in very low-count regimes. *IEEE Trans. Comput. Imaging* 8, 838–850.
- Maiden, A.M., Rodenburg, J.M., 2009. An improved ptychographical phase retrieval algorithm for diffractive imaging. *Ultramicroscopy* 109 (10), 1256–1262.
- G.T. Martinez, B.X. Shi, T.C. Naginey, L. Jones, C.M. O'Leary, T.J. Pennycook, R.J. Nicholls, J.R. Yates, P.D. Nellist, Direct imaging of charge redistribution due to bonding at atomic resolution via electron ptychography (2023). arXiv:1907.12974.
- McCallum, B., Rodenburg, J., 1992. Two-dimensional demonstration of wigner phase-retrieval microscopy in the stem configuration. *Ultramicroscopy* 45 (3-4), 371–380.
- O. Melnyk, Convergence properties of gradient methods for blind ptychography, arXiv preprint arXiv:2306.08750 (2023).
- Müller, K., Krause, F.F., Beche, A., Schowalter, M., Galioit, V., Löffler, S., Verbeeck, J., Zweck, J., Schattschneider, P., Rosenauer, A., 2014. Atomic electric fields revealed by a quantum mechanical approach to electron picodiffraction, 5653 *Nat. Commun.* 5, 1–8. <https://doi.org/10.1038/ncomms6653>.
- Müller-Caspary, K., Duchamp, M., Rösner, M., Migunov, V., Winkler, F., Yang, H., Huth, M., Ritz, R., Simson, M., Ihle, S., Soltau, H., Wehling, T., Dunin-Borkowski, R. E., Van Aert, S., Rosenauer, A., 2018. Atomic-scale quantification of charge densities in 2D materials. *Phys. Rev. B* 98, 121408. <https://doi.org/10.1103/PhysRevB.98.121408>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* 32.
- Rodenburg, J., Bates, R., 1992. The theory of super-resolution electron microscopy via wigner-distribution deconvolution, *Philosophical Transactions of the Royal Society of London. Ser. A: Phys. Eng. Sci.* 339 (1655), 521–553.
- Rodenburg, J.M., Faulkner, H.M., 2004. A phase retrieval algorithm for shifting illumination. *Appl. Phys. Lett.* 85 (20), 4795–4797.
- Rodenburg, J., McCallum, B., Nellist, P., 1993. Experimental tests on double-resolution coherent imaging via stem. *Ultramicroscopy* 48 (3), 304–314.
- Rodenburg, J., McCallum, B., Nellist, P., 1993. Experimental tests on double-resolution coherent imaging via stem. *Ultramicroscopy* 48 (3), 304–314.
- Thibault, P., Dierolf, M., Bunk, O., Menzel, A., Pfeiffer, F., 2009. Probe retrieval in ptychographic coherent diffractive imaging. *Ultramicroscopy* 109 (4), 338–343.
- Thibault, P., Guizar-Sicairos, M., 2012. Maximum-likelihood refinement for coherent diffractive imaging. *N. J. Phys.* 14 (6), 063004.
- Thibault, P., Guizar-Sicairos, M., 2012. Maximum-likelihood refinement for coherent diffractive imaging. *N. J. Phys.* 14 (6), 063004.
- Thibault, P., Menzel, A., 2013. Reconstructing state mixtures from diffraction measurements. *Nature* 494 (7435), 68–71.
- Thompson, W.J., 2001. Poisson distributions. *Comput. Sci. Eng.* 3 (3), 78–82.
- Waddell, E.M., Chapman, J.N., 1979. Linear imaging of strong phase objects using asymmetrical detectors in STEM. *Optik* 54, 83–96.
- Yang, H., Pennycook, T.J., Nellist, P.D., 2015. Efficient phase contrast imaging in stem using a pixelated detector. part ii: optimisation of imaging conditions. *Ultramicroscopy* 151, 232–239.
- Yeh, L.-H., Dong, J., Zhong, J., Tian, L., Chen, M., Tang, G., Soltanolkotabi, M., Waller, L., 2015. Experimental robustness of fourier ptychography phase retrieval algorithms. *Opt. Express* 23 (26), 33214–33240.
- Zhang, J., Xu, T., Shen, Z., Qiao, Y., Zhang, Y., 2019. Fourier ptychographic microscopy reconstruction with multiscale deep residual network. *Opt. Express* 27 (6), 8612–8625.
- Zhang, J., Xu, T., Zhang, J., Chen, Y., Li, J., 2021. Cross-level channel attention network for fourier ptychographic microscopy reconstruction. *IEEE Photonics J.* 14 (1), 1–8.
- Zhao, L., Zhou, X., Lu, X., Tong, H., Fang, H., 2023. Transformer-based reconstruction for fourier ptychographic microscopy. *IEEE Access*.
- Zhou, H., Feng, B.Y., Guo, H., Lin, S.S., Liang, M., Metzler, C.A., Yang, C., 2023. Fourier ptychographic microscopy image stack reconstruction using implicit neural representations. *Optica* 10 (12), 1679–1687.
- Zhou, Y., Hua, X., Zhang, Z., Hu, X., Dixit, K., Zhong, J., Zheng, G., Cao, X., 2020. Wirtinger gradient descent optimization for reducing gaussian noise in lensless microscopy. *Opt. Lasers Eng.* 134, 106131.