

Hybrid physics-machine learning models for quantitative electron diffraction refinements

Received: 6 August 2025

Accepted: 27 March 2026

Cite this article as: Malik, S.A., Doherty, T.A., Colmey, B. *et al.* Hybrid physics-machine learning models for quantitative electron diffraction refinements. *Nat Commun* (2026). <https://doi.org/10.1038/s41467-026-71673-9>

Shreshth A. Malik, Tiarnan A. S. Doherty, Benjamin Colmey, Stephen J. Roberts, Yarin Gal & Paul A. Midgley

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Hybrid Physics-Machine Learning Models for Quantitative Electron Diffraction Refinements

Shreshth A. Malik^{1,†} Tiarnan A.S. Doherty^{1,2,†,*} Benjamin Colmeley²
 Stephen J. Roberts³ Yarin Gal^{1,*} Paul A. Midgley^{2,*}

¹OATML, Department of Computer Science, University of Oxford, Wolfson Building, Parks Rd, Oxford OX1 3QG, United Kingdom

²Department of Materials Science and Metallurgy, University of Cambridge, 27 Charles Babbage Rd, Cambridge CB3 0FS, United Kingdom

³Machine Learning Research Group, Department of Engineering Science, University of Oxford, Eagle House, Walton Well Road, Oxford OX2 6ED, United Kingdom

[†]These authors contributed equally to this work.

*Corresponding authors: td404@cam.ac.uk, yarín.gal@cs.ox.ac.uk, pam33@cam.ac.uk

Abstract

High accuracy electron microscopy simulations required for quantitative crystal structure refinements face a fundamental challenge: while physical interactions are well-described theoretically, real-world experimental effects are challenging to model analytically. To address this gap, we present a hybrid physics-machine learning framework that integrates differentiable physical simulations with neural networks. By leveraging automatic differentiation throughout the simulation pipeline, our method enables gradient-based joint optimization of physical parameters and neural network components representing experimental variables, offering superior scalability compared to traditional second-order methods. We demonstrate this framework through application to three-dimensional electron diffraction (3D-ED) structure refinement, where our approach learns complex thickness distributions directly from diffraction data rather than relying on simplified geometric models. This method achieves state-of-the-art refinement performance across synthetic and experimental datasets, recovering atomic positions, thermal displacements, and thickness profiles with high fidelity. The modular architecture proposed can naturally be extended to accommodate additional physical phenomena and extended to other electron microscopy techniques. This establishes differentiable hybrid modeling as a powerful paradigm for quantitative electron microscopy, where experimental complexities have historically limited analysis.

1 Introduction

Recent advances in scientific machine learning have demonstrated remarkable potential to complement traditional physics-based simulations across diverse domains of computational science. Hybrid physics–machine learning (ML) approaches, which combine quantitative physical models with the expressive capacity of neural networks [1–5], offer a principled framework for addressing complex phenomena that remain challenging for purely analytical methods.

In these approaches, well-established physical theories that govern the forward model of a simulation are explicitly used to maintain theoretical rigor and interpretability, while neural networks are used as universal function approximators [6] to parameterize complex, system-specific effects that are difficult to model explicitly. Thus these hybrid methods bridge the gap between first-principles theory and experimental reality. In recent years, this approach has shown substantial promise in fields ranging from dynamical systems [7–9] and atmospheric physics [10] to quantum physics [11, 12], where the learned components augment traditional simulators to improve realism, accuracy, and/or computational efficiency.

A key enabler of these innovations is differentiable physics, in which physical simulators are constructed to support gradient evaluation through every computational step. In traditional physics-based parameter refinement, the application of gradient-based optimization requires explicitly deriving and implementing analytical expressions for the derivatives of the forward model. This process naturally becomes prohibitively complex for large-scale systems and more complex forward model theories. Automatic differentiation [13, 15–17] overcomes this limitation by systematically applying the chain rule to the elementary operations comprising the simulation to ‘backpropagate’ gradients to the parameters of interest with a computational cost comparable to that of the forward model, without the need for analytic derivative computation. Thus by implementing the forward model within modern automatic differentiation frameworks such as PyTorch [27] or JAX [37], the backpropagated gradients can be used to efficiently refine both physical parameters and neural network components jointly within a unified, end-to-end optimization framework.

Electron microscopy represents an ideal setting for the use of a hybrid modeling approach because while the elastic interaction of the electron beam with matter is well described by established physical theory [14, 38], real experiments involve sample morphologies, inelastic scattering, electron beam damage, atomic scale defects and other sample heterogeneities that are notoriously difficult to parameterize. Omitting these contributions in simulations imposes a systematic ceiling on the precision and accuracy with which structural and compositional information can be recovered from experimental data. However, to the best of our knowledge, there has not yet been a demonstration of the utility of hybrid modeling approaches when applied to any electron microscopy technique.

One such electron microscopy modality that could benefit substantially from a hybrid modeling approach is the field of three dimensional electron diffraction (3D-ED), where a crystal is rotated relative to an electron beam and a diffraction pattern is obtained at every orientation (Figure 1). 3D-ED has emerged as a compelling quantitative technique for structure determination in molecular and materials science [18], particularly for nanocrystalline samples unsuitable for traditional X-ray analysis. The interaction of electrons with matter, which is Coulombic in nature, is orders of magnitude stronger than that of X-rays [19]; this

allows high-resolution diffraction patterns to be obtained from μm or even nm scale crystals [20, 39] and drives pronounced ‘dynamical’ interactions between the incident and diffracted beams, so that the experimentally recorded intensities become exquisitely sensitive to subtle variations in the underlying structure factors [14, 40]. Practically, this extreme sensitivity enables the precise location of weakly scattering hydrogen atoms [18], robust determination of the absolute structure of chiral systems [22], unambiguous discrimination between atomic species possessing nearly identical scattering factors [21] and even information on charge density deformation due to bonding to be obtained [24], provided that dynamical calculations are performed and incorporated into the refinement process in order to accurately interpret the experimental data.

For dynamical refinements of 3D-ED data, the Bloch wave formalism [14, 41] is typically employed to simulate diffracted intensities from a structural starting model, compare these intensities to those observed experimentally, update the model and repeat the process until convergence is achieved [23]. While the Bloch wave approach provides a physically accurate description of multiple elastic scattering, each evaluation of the scattering matrix scales as $\mathcal{O}(N^3)$ and a full 3D-ED refinement, which typically involves 10^3 – 10^5 evaluations, can quickly scale to hours or even days. As dynamical diffracted intensities are highly sensitive to thickness variations [25], accurate modeling of how thickness varies as a function of sample orientation throughout the dataset is essential for quantitative refinement [24]. Real crystals possess irregular shapes, bending, surface roughness, and defects that induce complex, thickness profiles that vary as a function of crystal orientation.

In this work, we introduce a hybrid physics–machine learning approach that integrates a neural-network that models these complex thickness profiles with a fully differentiable implementation of the Bloch-wave formalism [13, 26]. While the Bloch-wave component governs the rigorous calculation of dynamical diffracted intensities given a structural model, the neural network augments this by learning orientation-dependent thickness parameters directly from data. Crucially, the differentiability of the Bloch-wave solver enables gradients of the simulated intensities to be backpropagated automatically through both the scattering calculation and the learned thickness model. This facilitates the use of gradient-based optimization to jointly refine structural parameters, thermal displacements, and thickness variations within a single, end-to-end framework.

We demonstrate that this hybrid differentiable framework matches or exceeds the performance of traditional refinements on both synthetic and experimental 3D-ED datasets. By leveraging first-order optimizers, which incur only linear storage and compute costs, we can co-refine thousands of thickness-model parameters alongside structural degrees of freedom in a single pass. This scalability, which is infeasible with traditional Hessian-based, second-order schemes, is a direct benefit of the differentiable formulation. While our focus here is on 3D-ED, the same differentiable-physics plus learned-submodel strategy can be applied to other electron microscopy modalities (e.g., 4D-STEM, convergent-beam diffraction, scanning-precession ED) where complex, hard-to-model experimental effects likewise limit quantitative analysis.

2 Results

2.1 Refinement workflow

Figure 2 is a graphical overview of our differentiable refinement workflow. Beginning with the asymmetric unit (ASU), we apply space group symmetry operations to expand the ASU into the unit cell. The unit cell is then rotated to an orientation that matches the experimentally observed diffraction pattern. Following this, we perform a Bloch wave calculation to obtain the simulated intensities at this crystal orientation, where the sample thickness used in the simulation is determined by a lightweight neural network (ThicknessNN, Section 2.2). A loss is then computed – in this case the crystallographic R-factor – between the integrated observed and simulated intensities. Using backpropagation, we evaluate the gradients of this loss with respect to the ASU parameters, including atomic positions and isotropic or anisotropic displacement parameters, and the ThicknessNN. Finally, these parameters are updated via gradient descent to minimize the loss.

The refinement is carried out against experimental integrated intensities, obtained after the full reduction of the raw diffraction patterns. This involves indexing and extracting per-hkl intensities with uncertainties informed by detector gain and noise characteristics, and the variance obtained from the background pixels surrounding each Bragg disc [29]. The data are then organised into virtual frames by segmenting the tilt series into angular ranges, and each hkl is allocated to the appropriate frame to allow direct matching of simulated and experimental hkl, with refinements carried out on a frame-by-frame basis, as described in [23, 29–31].

All datasets contained in this paper are continuous rotation type datasets, with the exception of the Paracetamol experimental dataset presented in section 2.3, which is a precession assisted rotation dataset.

2.2 Thickness modelling

Diffraction intensities are highly sensitive to crystal thickness. Since a typical 3D-ED experiment illuminates a large volume of an irregular crystal simultaneously, the recorded diffraction pattern inevitably integrates signals arising from many variations in crystal thickness across the probed volume. An additional complication is that as the crystal rotates, the projected thicknesses of the sample relative to the diffracted beam change as a function of rotation angle. Previous work has addressed this complexity with a dual pronged approach. First, by invoking the column approximation [42], the crystal can be simplified into a collection of independent, parallel scattering columns (Figure 3(a)). The diffracted intensity can then be treated as an incoherent sum of diffraction patterns emerging from these small columns. To make this calculation tractable, the unknown morphology of the real crystal is approximated with an idealized geometric model such as a wedge, lens, or cylinder from which the required thickness probability density functions can be derived analytically [23]. Secondly, to account for sample rotation, these models are then adjusted by applying empirical corrections that interpolate between two limiting cases of sample geometry: an infinite flat plane, where, for tilt angle θ , the apparent thickness varies as $t(\theta) = t_0 / \cos \theta$, and a cylindrical or spherical object, where the thickness remains constant, $t(\theta) = t_0$ [24].

A limitation of this approach is that it relies on rigid geometric priors, which may lack the flexibility to accurately describe the complex, irregular morphologies encountered in real experimental samples. Moreover, current implementations allow the thickness probability distribution to become decoupled from the tilt correction. This can create physical inconsistencies in the refinement if not done with care. For example, using a probability density function derived from a wedge with a tilt correction parameter for a sphere or cylinder is contradictory. In reality, these properties are intrinsically coupled: the underlying 3D shape of the crystal determines both the statistical distribution of thicknesses and how that distribution evolves under rotation. Finally, as these geometric models are rarely optimized jointly with atomic parameters, significant computational effort is often required to identify the correct shape approximation through manual tuning or grid searches outside the main refinement loop.

Here, we instead propose a flexible, data-driven alternative that replaces fixed geometric priors with a lightweight neural network. Rather than assuming a specific crystal shape (e.g., a wedge), our model learns the statistics of the thickness distribution directly from the diffraction data. We assume a Gaussian form for the thickness distribution at each orientation (tilt angle θ) as shown in Figure 3(b), whose mean $\mu(\theta)$ and standard deviation $\sigma(\theta)$ are predicted by the neural network. Representative thickness values are then sampled from this distribution and used as input to the dynamical diffraction simulator. The resulting intensities from each different thickness simulation are then averaged to account for the distribution of thicknesses present in the illuminated volume at that orientation. The network operates jointly with the atomic parameter refinement: as the crystallographic refinement progresses, the network weights are updated via backpropagation to minimize the weighted R-factor (wR_{all}). This allows the model to predict the optimal $\mu(\theta)$ and $\sigma(\theta)$ for every orientation dynamically, adapting readily to complex, irregular morphologies (Figure 3(c)). In practice, we find that explicitly sampling from the predicted thickness distribution and averaging the resulting intensities produces refinement R -values comparable to those obtained by evaluating the intensity at the mean thickness alone (see Supplementary Note 3 for discussion). We therefore use only the predicted mean thickness in all refinements reported here for computational efficiency. Further implementation details and comparisons with prior approaches are provided in the Methods section.

2.3 Refinement on Synthetic Data

To evaluate the ability of our refinement workflow to jointly recover both atomic parameters and sample thickness profiles, we first tested our methodology on synthetic 3D-ED datasets. We generated synthetic datasets for three representative structures: paracetamol, CsPbBr₃, and quartz (see Methods). Each 3D-ED dataset was corrupted with Poisson noise and was assigned a randomly generated thickness profile that varied smoothly as a function of rotation angle θ , mimicking realistic variations in apparent crystal thickness during data acquisition.

Initial models for the refinement were created by adding random atomic displacements drawn from a uniform distribution over the range $[-\Delta, \Delta]$ Å, with maximum displacement magnitude of $\Delta = 0.2$ Å. Displacements were generated in a symmetry-aware fashion in that atomic positions were perturbed randomly while preserving crystallographic symmetry operations and special Wyckoff positions. Thermal displacement parameters for the start-

ing models were created by scaling the ground truth thermal displacement parameters in proportion to the magnitude of the atomic displacements. Additionally, for CsPbBr₃, which has anisotropic thermal displacement parameters, the starting models were initialized with simplified isotropic values to increase the difficulty of the refinement task.

Figure 4 shows representative refinements for CsPbBr₃ and paracetamol. We observe that wR_{all} , the root mean squared distance (RMSD) to ground truth atomic positions, and the Frobenius norm (see Methods 4.5) of the difference in thermal displacement parameters all converge effectively to zero within 200 epochs. These results are promising for several reasons. Firstly, a 0.2 Å displacement represents a relatively poor initial model for a crystallographic refinement and we do not employ any stereochemical restraints, yet convergence is still achieved. Secondly, in the cases of CsPbBr₃, the refinement successfully recovers anisotropic thermal displacements even when initialized with isotropic ones with a large displacement. Thirdly, the ground truth thickness distribution is successfully recovered alongside the atomic parameters and this remarkable performance continues even when the thickness profile has a complex dependence on θ (Supplementary Note 1). All of this is achieved using the Adam optimizer [28] with default settings and without any hyperparameter tuning. Similar results are observed for the quartz datasets (Table 1 and Supplementary Figure S5), and for models initialised with displacements of 0.1 Å (Supplementary Figure S4) and 0.3 Å (Supplementary Figure S6).

To understand the importance of utilising a fully dynamical forward model in the refinement of 3D-ED data, we compare in Figure 5 our results to those obtained with a kinematical model which we approximate by fixing the sample thickness to 10 Å such that multiple scattering is negligible. As expected, a degradation in refinement quality is observed when using a kinematical model and it is most pronounced when strong dynamical scattering is present in the dataset. For instance, in CsPbBr₃, where substantial dynamical effects are expected due to the combination of the presence of heavy elements (Cs and Pb) and a large simulated crystal thickness (Figure 4), the final RMSD in atomic positions and the Frobenius norm (Methods 4.5) of the thermal displacements remain substantially large compared to the dynamical models. In contrast, the performance gap is smaller for paracetamol, which consists of lighter atoms and was simulated with a thinner crystal.

Summary metrics for the converged refinements across all three structures are shown in Table 1. Here we also ablate performance against a dynamical refinement with an imperfect thickness model, where we assume a constant effective thickness for the entire tilt series rather than a detailed thickness profile. The recovered single thickness values along with the ground truth thickness distributions are shown in Supplementary Figure S3. We find that the single-thickness model also consistently under-performs compared to the hybrid approach. This is especially pronounced in quartz where strong variation with θ is present. We observe similar trends for various levels of structural perturbation. Full refinement curves and results across datasets are provided in Supplementary Figures S4, S5 and S6. We further note runtime comparisons to second order methods in Supplementary Note 2.

2.4 Refinement on Experimental Data

In the refinement of synthetic datasets, we assume purely elastic scattering with added Poisson noise, and the experimental orientation is considered to be perfectly known. In contrast,

real experimental data are affected by additional complexities, including inelastic scattering processes—such as plasmon losses and thermal diffuse scattering [33, 34]—as well as detector noise [29]. Accurate frame-by-frame orientation determination is also essential to match simulations to measured intensities [29]. To evaluate the robustness of our proposed hybrid physics–ML refinement approach on real data, we reproduced dynamical refinements for three previously published 3D electron diffraction (3D-ED) datasets: quartz [31], CsPbBr₃ [24], and paracetamol [32]. In each case, the original analyses involved data processing using the PETS2 software package [29], followed by dynamical refinements with DYNAGO and JANA [43].

For our refinements, we used the same PETS2 output files reported in the literature but replaced the DYNAGO-JANA pipeline with our own hybrid physics-ML refinement code. To the best of our knowledge, this is the first independent demonstration and benchmarking of an alternative codebase to JANA and DYNAGO for dynamical refinement of 3D-ED datasets. Other dynamical refinement codebases [25, 44–46] in the past have focused on the refinement of non-integrated intensities extracted from individual diffraction patterns (often convergent beam electron diffraction; CBED) collected near major crystallographic zone axes, which is a fundamentally different problem to the one presented here. Table 2 summarizes our results. Across all datasets, the number of N_{obs} ($I > 3\sigma$) and N_{all} intensities compares very closely between our refinements and the original studies. The small differences in N_{obs} and N_{all} between our results and those in the literature likely arise from variations in orientation search protocols and parameterizations. For continuous rotation datasets (quartz and CsPbBr₃), we perform a three-parameter simplex search involving rotations about the crystal’s x , y , and z axes (see Methods 4.3). In contrast, published studies appear to predominantly use a two-parameter Euler angle approach — ϕ (rotation about the laboratory z -axis) and θ (tilt about the new x -axis) — for both continuous rotation and precession datasets [30, 31]. For precession electron diffraction data (Paracetamol), we follow this convention, where ϕ defines the tilt direction and θ its amplitude.

For the quartz dataset, the hybrid physics-ML method achieves a lower R_{obs} value (4.2 %) and wR_{all} value (4.3 %) compared to the DYNAGO-JANA values of (5.7 %) and (6.6 %), indicating improved fit to the observed intensities. This is likely due to the thickness modeling employed in the hybrid physics-ML refinement; in Ref. [31], from which the quartz dataset is obtained, there is no explicit mention of modeling the thickness of the crystal as a function of θ beyond a single refined thickness value of ~ 44 nm. By comparison, the ThicknessNN utilised in the hybrid physics-ML method recovers a more complex relationship of apparent thickness as a function of θ in addition to an overall larger crystal thickness at each orientation (~ 85 nm average). The discrepancy between these thickness values is difficult to explain. Notably though, if we turn off the ThicknessNN and use a single thickness value in our simulator, we observe a minimum in wR_{all} at 87.2 nm (Supplementary Figure S8) which compares very closely with the mean thickness recovered by the ThicknessNN. Additionally, interrogating images of the crystal from which the 3D-ED dataset is recorded – though a qualitative analysis at best – suggests that the true mean crystal thickness is likely closer to 85 nm than 44 nm (See Supplementary Figure S9 and Supplementary Note 4). The recovered thickness distribution (Figure 6) is plausible given the complex geometry of the quartz crystal and possible specimen movement during data acquisition (See Supplementary Video 1 and Supplementary Note 5). However, we note that it is difficult

to derive accurate 3D specimen geometries from experimental electron microscopy images without a full tomographic analysis.

For CsPbBr₃, the recovered R_{obs} and wR_{all} values are comparable between our hybrid Physics-ML model (6.4% R_{obs} , 6.7% wR_{all}) and those reported in the original study (5.5% R_{obs} , 6.8% wR_{all}) [24]. These similar residuals are expected given that both approaches incorporate an explicit thickness model. However, the underlying recovered thickness distributions differ. The original study reports a constant thickness of 69.1 nm across the dataset and noted that introducing a tilt correction parameter—intended to vary thickness as a function of θ —did not improve the residuals. This suggests a sample geometry approximating a uniform cylindrical rod, such that $t(\theta) = t_0$. In contrast, our model recovers a thickness distribution with a mean of 44.0 nm and a mild dependence on θ (see Figure 6). An advantage of our approach is that we do not require an explicit tilt correction parameter, which in the original study may have necessitated multiple refinement cycles to identify an optimal form. In this case, both models yield similar residuals because the thickness distribution is relatively simple. However, in more complex scenarios—such as the quartz dataset discussed earlier—a manually defined tilt correction function would likely fail to capture the true variation, highlighting the flexibility and generality of our approach.

For the paracetamol dataset, our R_{obs} is higher than reported (10.15% vs 9.2%), but our wR_{all} is lower (8.4% vs 10.34%), which is atypical, but has been observed in other dynamical refinements that utilise the same weighting in the calculation of wR_{all} (see Methods) [30]. The recovered thickness distribution with the hybrid physics-ML model is asymmetric, which is roughly consistent with the rotation range covered by the goniometer during the experiments (-52° to $+30^\circ$), or alternatively, could be explained by the electron beam intentionally being moved across the sample during acquisition to mitigate beam damage as previously reported [18]. Interestingly, the mean value of thickness distribution recovered from the hybrid physics-ML model (47.5 nm) compares closely to the single thickness value refined with the DYNGO-JANA model (51.1 nm [32], Figure 6).

One of the key advantages of electron diffraction is its high sensitivity to the chirality (handedness) of non-centrosymmetric crystal structures. In such systems, chirality is encoded in the violation of Friedel’s Law, which results in intensity differences between symmetry-related reflections \mathbf{g} and $-\mathbf{g}$. This asymmetry arises from multiple scattering effects and cannot be interpreted using the kinematical approximation alone. Instead, full dynamical simulations are required, involving refinements of both enantiomorphic models to determine the correct absolute structure [22, 31]. To verify that our hybrid physics-ML model can recover these differences we perform independent refinements of both enantiomers ($P3_221$ and $P3_121$) of quartz to replicate the results presented in [31]. We obtain substantially lower residual factors with $P3_221$ enantiomer (Table 3), which is consistent with previous results [31] and demonstrates that the hybrid physics-ML model is capable of accurately determining the absolute structure of chiral systems.

3 Discussion

Our results demonstrate that hybrid physics–machine learning models, built on a foundation of differentiable physics, can offer a powerful framework for the dynamical refinement of 3D

electron diffraction data. As a proof of principle, we have focused here on the application of this framework to modeling crystal thickness variations. Our results on both synthetic and experimental data are encouraging. In the case of synthetic data, recovery of ground truth thickness is almost flawless suggesting that it could someday be possible to eliminate thickness variations as an uncertainty in experimental refinements. While a complete evaluation of this hypothesis was not possible with the experimental data utilised in this study, future work that is performed on data extracted from experimental specimens of known geometry, such as those obtained from focused ion beam milling, will enable a rigorous assessment of how closely our framework can recover ground-truth thickness distributions.

Future work can also explore the applicability of this framework to other sources of error within existing refinement protocols. For example, models – neural network based or otherwise – of beam damage, inelastic scattering or detector noise that are physically constrained to avoid overfitting could be developed. Incorporating such physical phenomena into the refinement process would further lower residual R -factors while simultaneously yielding dataset-specific estimates of, for example, beam damage and inelastic scattering.

An even more compelling prospect is the potential for this framework to serve as a development base for robust quantum electron crystallography. In our Bloch wave forward model, as supposedly in greater than 99% of crystallographic refinements [47], we implicitly assume that each atom is described by a neutral, spherically symmetric, charge density centered on the atomic nucleus. This approximation is known as the independent atom model. In real materials and molecules, chemical bonding between atoms leads to deformations of charge density – and thus electrostatic potential – that are directly measurable with electron scattering [24, 45, 46]. However, quantifying these features remains non-trivial. Approaches such as multipole modeling and Hirshfeld atom refinement, originally developed for X-ray crystallography, have shown promising potential for 3D-ED refinements [24, 48, 49] but further development of these frameworks, and others, will be required to determine which physical descriptions are most appropriate for dynamical electron scattering.

Regardless of the specific physical model chosen, the differentiable nature of the framework presented here significantly lowers the barrier to entry for future developments, as it eliminates the need to analytically derive gradients for complex descriptions of the scattering potential, and naturally accommodates future adoption of hybrid physics-machine learned representations of the electrostatic potential.

Finally, while this work focused on 3D-ED using the Bloch wave formalism, the paradigm of differentiable hybrid modeling is broadly applicable. Similar hybrid physics-ML frameworks could readily be adapted to other simulation methods, such as the multislice algorithm, and extended to diverse modalities like 4D scanning transmission electron microscopy (4D-STEM), ultimately providing an updated standard for quantitative analysis across the electron microscopy community.

4 Methods

4.1 Hybrid Physics-ML Model

Our forward diffraction model, shown in Figure 2, consists of a hybrid neural-network based thickness model and a differentiable physics simulator, combining the physical accuracy of the Bloch wave formalism and the flexibility of neural network function approximation.

4.1.1 Differentiable Bloch Wave Propagation

The Bloch wave method, developed by Bethe [41] and later expanded by others [14, 40, 42], models elastic electron scattering in perfect crystals across arbitrary orientations and thicknesses, yielding diffracted intensities that can be directly compared to 3D-ED data. Its matrix formulation is naturally differentiable, making it well adapted to gradient-based optimization in dynamical refinement [30]. The Bloch wave formalism expands the wavefunction of the incident electron $\psi(\mathbf{r})$ into a finite set of plane-wave states permitted by the periodic crystal potential $V(\mathbf{r})$. For a complete theoretical treatment, see Spence and Zuo [50]. The periodic crystal potential is expressed as $V(\mathbf{r}) = \sum_{\mathbf{g}} V_{\mathbf{g}} e^{2\pi i \mathbf{g} \cdot \mathbf{r}}$, where each reciprocal lattice vector \mathbf{g} corresponds to a Fourier component $V_{\mathbf{g}}$, determined by the underlying atomic model.

To compute $V_{\mathbf{g}}$, the ASU is expanded to the full unit cell using the crystal's space group operations. Each atom i contributes an electron scattering factor f_i^e , evaluated at scattering vector $s = |\frac{\mathbf{g}}{2}|$, modulated by the Debye–Waller factor B_i that accounts for atomic thermal motion. $V_{\mathbf{g}}$ then takes the form:

$$V_{\mathbf{g}} = \frac{1}{\Omega} \sum_i f_i^e(\mathbf{s}) \exp(-B_i s^2) \exp(-2\pi i \mathbf{g} \cdot \mathbf{r}_i), \quad (1)$$

where Ω is the unit cell volume and \mathbf{r}_i is the atomic position. Thermal vibrations reduce the sharpness of the periodic potential, leading to overall damping of diffracted intensities, with stronger attenuation at higher scattering angles. For isotropic displacements, the Debye–Waller factor is given by $B_i = 8\pi^2 \langle u_i^2 \rangle$, where $\langle u_i^2 \rangle$ is the Cartesian atom-specific mean-square displacement. In the anisotropic case, thermal motion is described by a symmetric second-rank tensor U_{ij} , defined as the covariance $U_{ij} = \langle u_i u_j \rangle$ of atomic displacements along Cartesian directions i and j [51]. To express U_{ij} in crystallographic fractional coordinates, it is transformed using the orthogonalization matrix A :

$$U_{ij}^* = A^{-1} U_{ij}^{\text{cart}} A^{-\text{T}} \quad (2)$$

This transformation ensures compatibility with space group symmetry operations. To maintain physical validity and enable stable refinement, the tensor is parameterized via its Cholesky decomposition $U = LL^{\text{T}}$, where L is a lower triangular matrix. This guarantees positive-definiteness and enables smooth optimization. During asymmetric unit expansion, each U_{ij}^* is rotated under the space group operations and then converted back to Cartesian form for use in scattering calculations.

4.1.2 Apparent Thickness Neural Network

To model orientation-dependent variations in apparent crystal thickness, we implemented a lightweight neural network module (`ThicknessNN`) in PyTorch [27] that maps a rotation angle θ to a predicted mean thickness $\mu(\theta)$ and associated standard deviation $\sigma(\theta)$. The model consists of a feedforward architecture with two hidden layers of 64 units each and `Tanh` activation functions (to promote smoothness in θ), followed by a two-node output layer that returns μ and $\log \sigma^2$. The predicted standard deviation is scaled to lie within a user-defined range to ensure numerical stability.

Users may often have prior knowledge of both the approximate thickness range (e.g. 10–200 nm) and the qualitative shape of its variation across orientations. Such priors can be incorporated to warm-start the optimization. For example, we incorporate an optional quadratic prior by expressing $\mu(\theta)$ as a weighted combination of a learnable quadratic function $a\theta^2 + b\theta + c$ (with $a > 0$ enforced via exponential parameterization) and a residual term learned from the network. This hybrid formulation allows the model to learn deviations from simple geometric thickness profiles while retaining regularizing inductive structure. In synthetic data experiments, we utilized this prior for quartz and paracetamol, however we soon realised that it was not required for reliable thickness distribution recovery. For subsequent experiments including synthetic CsPbBr₃ and experimental data, we did not add this prior, though we note in particular high-noise or data scarce environments (such as refining on partial datasets), these priors may prove useful.

General Framework. For a reflection h , the diffracted intensity at a given crystal thickness t is denoted $I_h(t)$. If the illuminated region contains a distribution of thicknesses, the observed intensity corresponds to the expectation of $I_h(t)$ over the thickness probability distribution function (PDF) $f(t)$:

$$I_h^{\text{obs}} = \int I_h(t) f(t) dt, \quad \int f(t) dt = 1. \quad (3)$$

Both the present method and earlier approaches [23] are based on this same principle, differing primarily in the parameterisation of $f(t)$. Previous approaches typically derived $f(t)$ from assumed *geometric priors* (e.g. wedge, cylindrical, ribbon, or convex-lens morphologies), from which closed-form cumulative thickness distributions could be obtained. Tilt dependence was often introduced through empirical corrections that adjusted the mean thickness as a function of orientation.

Neural Net Parameterisation. In contrast, our approach replaces the fixed geometric parameterisation with a data-driven parameterisation of the thickness probability density function. At each orientation (tilt angle θ), we assume a Gaussian form for the thickness distribution $f(t | \theta)$, whose mean $\mu(\theta)$ and standard deviation $\sigma(\theta)$ are predicted by a neural network:

$$f(t | \theta) = \frac{1}{\sigma(\theta)\sqrt{2\pi}} \exp\left[-\frac{(t - \mu(\theta))^2}{2\sigma(\theta)^2}\right]. \quad (4)$$

This formulation enables the thickness distribution to adapt smoothly to complex, non-geometric variations in the specimen while remaining differentiable with respect to the dis-

tribution parameters $\mu(\theta)$ and $\sigma(\theta)$. No manual tuning of empirical parameters is required, as the network learns the distribution directly from the experimental data.

Computation. To compute the expected intensity, we use the reparameterization trick to generate differentiable samples analogously to variational autoencoders [35]:

$$t = \mu(\theta) + \sigma(\theta) \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1), \quad (5)$$

and approximate the expectation as

$$I_h^{\text{obs}}(\theta) \approx \frac{1}{N} \sum_{j=1}^N I_h(\mu(\theta) + \sigma(\theta) \epsilon_j). \quad (6)$$

This enables gradients to propagate through both $\mu(\theta)$ and $\sigma(\theta)$ during refinement, allowing the model to learn orientation-dependent thickness statistics directly from data.

Simplified Evaluation. In practice, we found that sampling from the predicted Gaussian and averaging the resulting intensities yielded similar refinement R -values to simply evaluating the intensity at the mean thickness $\mu(\theta)$ (Supplementary Note 3). For computational efficiency, all refinements reported in this work therefore used

$$I_h^{\text{obs}}(\theta) \approx I_h(\mu(\theta)). \quad (7)$$

Nevertheless, the framework remains fully compatible with explicit distributional averaging and could be extended to more flexible, non-Gaussian formulations if required in future applications.

4.2 Refinements, loss functions and reflection filtering

For calculation of residuals wR_{all} and R_{obs} we followed the conventions and definitions outlined in Klar et al. [31].

$$R = \frac{\sum |\sqrt{I_{\text{obs}}} - \sqrt{I_{\text{calc}}}|}{\sum \sqrt{I_{\text{obs}}}} \quad (8)$$

$$wR = \sqrt{\frac{\sum (w |I_{\text{obs}} - I_{\text{calc}}|)^2}{\sum (w I_{\text{obs}})^2}} \quad (9)$$

$$w = \left(\sigma \left(\sqrt{I_{\text{obs}}} \right)^2 + \left(u \sqrt{I_{\text{obs}}} \right)^2 \right)^{-1/2} \quad (10)$$

Where the sums run over all reflections for the calculation of wR_{all} , and only over observed reflections with $I_{\text{obs}} > 3\sigma(I_{\text{obs}})$ for R_{obs} . The instability factor u was set to 0.01. Reflections with either weak or negative intensities, with $I_{\text{obs}} < 0.01\sigma(I_{\text{obs}})$ were set to 0, and their uncertainties were adjusted such that

$$\sigma \left(\sqrt{I_{\text{obs}}} \right) = 5\sqrt{\sigma(I_{\text{obs}})}. \quad (11)$$

All refinements of atomic positions, anisotropic thermal displacement, ThicknessNN weights and biases and orientations proceeded against wR_{all} and did not utilise any restraints or constraints beyond those imposed by crystal symmetry for a particular spacegroup. For all experiments, we refined parameters using the Adam optimizer [28] over 200 epochs (passes over the full dataset of rotations) using a full batch gradient descent (we average the loss over all rotations for each gradient step), and a learning rate of 0.001 for all parameters.

For selection of integrated intensities, the methodologies to calculate S_g , S_g^{max} , D_{sg} and R_{sg} described in Ref. [31] were again followed. Values for S_g , S_g^{max} , D_{sg} and R_{sg} utilized in the refinements of experimental data were as previously reported [24, 31, 32].

4.3 Orientation refinement

Orientation refinement for each rotation was performed using either a three-parameter Nelder–Mead simplex search or a modified simplex approach based on the methodology described by Palatinus et al. [30]. In the Nelder–Mead approach, the crystal orientation was parameterized as a product of three small-angle rotations about the crystal’s x , y , and z axes, with the optimization objective being the minimization of the residual wR_{all} . The three rotation angles α , β , and ω were optimized using the `scipy.optimize.minimize` function with `method='Nelder-Mead'`, initialized from a regular simplex [36]. At each optimization step, an updated orientation matrix was constructed and used in the forward simulation of dynamical diffraction intensities. The updated orientation matrix had the form

$$R = R_z(\omega)R_y(\beta)R_x(\alpha) \quad (12)$$

The alternative two parameter orientation search followed a method whereby a discrete hexagonal scan over orientations defined by two angular parameters: (ϕ), a rotation about the laboratory z -axis, and θ , a tilt about the rotated x -axis was performed. Trial orientations were generated by rotating the current orientation matrix as

$$R_{\text{new}} = R_z(\phi)R_x(\theta)R_z(-\phi)R_{\text{current}}, \quad (13)$$

with (ϕ) varied in 60° increments and θ halved iteratively if no improvement in wR_{all} was found. This approach effectively performs a simplex-like local search over a cone of directions around the current orientation.

Both search protocols used the full experimental–simulated diffraction matching pipeline described above, including filtering of reflections and scaling factor optimization, to evaluate the objective function at each orientation step.

4.4 Synthetic Data Generation and Noise Addition

The synthetic data utilised in this study was generated from atomic co-ordinates obtained from successful experimental dynamical refinements of quartz [18, 24].

These three structures were chosen as they span many of the current application areas of dynamical refinements of 3D-ED data (inorganic materials science and small molecule crystallography) and will exhibit varying degrees of dynamical effects. For example. we expect CsPbBr₃ to have more dynamical scattering than paracetamol for similar thicknesses.

To model experimental noise in simulated diffraction data, we applied Poisson-distributed noise across the full dataset. All simulated intensities were first linearly scaled such that the brightest reflection across the dataset matched a predefined dynamic range maximum (e.g. $I_{\max} = 10^6$ or $I_{\max} = 10^2$ etc). This ensures control over the signal-to-noise ratio (SNR), with lower dynamic range values producing proportionally higher noise relative to signal, thereby emulating low-dose data acquisition conditions. After scaling, reflections were removed from each individual pattern if their intensities were more than a factor of the dynamic range below the global maximum, i.e., if their scaled intensity fell below one count. This step mimics a practical detector noise floor and enforces the selected dynamic range as a hard cutoff on observable intensities. Poisson noise was then applied to the remaining scaled intensities using the `torch.poisson` function from the PyTorch library, which samples from a Poisson distribution with mean equal to the scaled intensity. This implementation captures counting statistics typical of electron diffraction experiments while preserving the statistical structure of the dataset. The synthetic data presented in the main text was generated using a dynamic range of 10^6 , consistent with the capabilities of state-of-the-art photon-counting detectors. The methodology was also evaluated under high-noise conditions using a reduced dynamic range of 10^2 to assess robustness in low-SNR regimes (Supplementary Figure S2). Uncertainties for each reflection were estimated as the square root of the noisy intensities, i.e., $\sigma_i = \sqrt{I_i}$, consistent with Poisson counting statistics.

4.5 Frobenius Norm

For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, the Frobenius norm is defined as

$$\|\mathbf{A}\|_F = \left(\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2 \right)^{1/2}. \quad (14)$$

In this work, the Frobenius norm is used to quantify the distance of the refined thermal displacement parameter matrices ($\mathbf{U}_{\text{refined}} \in \mathbb{R}^{3 \times 3}$) from their ground-truth values ($\mathbf{U}_{\text{ground truth}} \in \mathbb{R}^{3 \times 3}$) during crystallographic refinements. We compute and report the average $\|\mathbf{U}_{\text{deviation}}\|_F = \|\mathbf{U}_{\text{ground truth}} - \mathbf{U}_{\text{refined}}\|_F$ across atoms in all tables and plots.

4.6 Code Development

The code is developed in `python` and is built on Bloch Wave code originally developed by the creators of `abTEM` [52].

Data Availability Relevant data generated during and supporting the findings of this study are available at the following Zenodo link: [10.5281/zenodo.18281349](https://zenodo.org/record/18281349).

Code availability

A version of the codebase suitable for reproducing the results in this paper has been deposited at the following Zenodo link: [10.5281/zenodo.18281349](https://zenodo.org/record/18281349). The authors are collaborating with software developers to make a more stable and mature version of the codebase available for general release on github.

References

- [1] Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT press, <http://www.deeplearningbook.org>, (2016)
- [2] Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Bartlett, P.L., Pereira, F.C.N., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a Meeting Held December 3-6, 2012, Lake Tahoe, Nevada, United States, pp. 1106–1114 (2012). <https://proceedings.neurips.cc/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html>
- [3] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
- [4] Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al.: Mastering chess and shogi by self-play with a general reinforcement learning algorithm. arXiv preprint arXiv:1712.01815 (2017)
- [5] Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., et al.: Mastering atari, go, chess and shogi by planning with a learned model. Nature **588**(7839), 604–609 (2020)
- [6] Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. Neural networks **2**(5), 359–366 (1989)
- [7] Avila Belbute-Peres, F., Smith, K., Allen, K., Tenenbaum, J., Kolter, J.Z.: End-to-end differentiable physics for learning and control. Advances in neural information processing systems **31** (2018)
- [8] Allen, K.R., Lopez-Guevara, T., Stachenfeld, K., Sanchez-Gonzalez, A., Battaglia, P., Hamrick, J., Pfaff, T.: Physical design using differentiable learned simulators. arXiv preprint arXiv:2202.00728 (2022)
- [9] Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.W.: Learning mesh-based simulation with graph networks. arXiv preprint arXiv:2010.03409 (2020)
- [10] Arcomano, T., Szunyogh, I., Wikner, A., Pathak, J., Hunt, B.R., Ott, E.: A hybrid approach to atmospheric modeling that combines machine learning with a physics-based numerical model. Journal of Advances in Modeling Earth Systems **14**(3), 2021–002712 (2022)
- [11] Deringer, V.L., Caro, M.A., Csányi, G.: Machine learning interatomic potentials as emerging tools for materials science. Advanced Materials **31**(46), 1902765 (2019)

- [12] Schoenholz, S., Cubuk, E.D.: Jax md: a framework for differentiable physics. *Advances in Neural Information Processing Systems* **33**, 11428–11441 (2020)
- [13] Baydin, A.G., Pearlmutter, B.A., Radul, A.A., Siskind, J.M.: Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research* **18**, 1–43 (2018)
- [14] Humphreys, C.J.: The scattering of fast electrons by crystals. *Reports on Progress in Physics* **42**(11), 1825 (1979) <https://doi.org/10.1088/0034-4885/42/11/002> . Accessed 2023-09-28
- [15] Pennington, R.S., Broek, W., Koch, C.T.: Third-dimension information retrieval from a single convergent-beam transmission electron diffraction pattern using an artificial neural network. *Physical Review B* **89**(20), 205409 (2014) <https://doi.org/10.1103/PhysRevB.89.205409> . Publisher: American Physical Society. Accessed 2025-04-30
- [16] Van Den Broek, W., Koch, C.T.: Method for Retrieval of the Three-Dimensional Object Potential by Inversion of Dynamical Electron Scattering. *Physical Review Letters* **109**(24), 245502 (2012) <https://doi.org/10.1103/PhysRevLett.109.245502> . Accessed 2024-01-12
- [17] Diederichs, B., Herdegen, Z., Strauch, A., Filbir, F., Müller-Caspary, K.: Exact inversion of partially coherent dynamical electron scattering for picometric structure retrieval. *Nature Communications* **15**(1), 101 (2024) <https://doi.org/10.1038/s41467-023-44268-x> . Publisher: Nature Publishing Group. Accessed 2025-04-30
- [18] Palatinus, L., Brázda, P., Boullay, P., Perez, O., Klementová, M., Petit, S., Eigner, V., Zaarour, M., Mintova, S.: Hydrogen positions in single nanocrystals revealed by electron diffraction. *Science* **355**(6321), 166–169 (2017) <https://doi.org/10.1126/science.aak9652> . Publisher: American Association for the Advancement of Science. Accessed 2024-01-10
- [19] Henderson, R.: The potential and limitations of neutrons, electrons and x-rays for atomic resolution microscopy of unstained biological molecules. *Quarterly reviews of biophysics* **28**(2), 171–193 (1995)
- [20] Nannenga, B.L., Shi, D., Leslie, A.G., Gonen, T.: High-resolution structure determination by continuous-rotation data collection in microed. *Nature methods* **11**(9), 927–930 (2014)
- [21] Gemmi, M., Mugnaioli, E., Gorelik, T.E., Kolb, U., Palatinus, L., Boullay, P., Hovmöller, S., Abrahams, J.P.: 3D Electron Diffraction: The Nanocrystallography Revolution. *ACS Central Science* **5**(8), 1315–1329 (2019) <https://doi.org/10.1021/acscentsci.9b00394> . Publisher: American Chemical Society. Accessed 2024-01-12
- [22] Brázda, P., Palatinus, L., Babor, M.: Electron diffraction determines molecular absolute configuration in a pharmaceutical nanocrystal. *Science* **364**(6441), 667–669 (2019) <https://doi.org/10.1126/science.aaw2560> . Publisher: American Association for the Advancement of Science. Accessed 2024-01-10

- [23] Palatinus, L., Petříček, V., Corrêa, C.A.: Structure refinement using precession electron diffraction tomography and dynamical diffraction: theory and implementation. *Acta Crystallographica Section A: Foundations and Advances* **71**(2), 235–244 (2015) <https://doi.org/10.1107/S2053273315001266> . Number: 2 Publisher: International Union of Crystallography. Accessed 2023-09-28
- [24] Suresh, A., Yörük, E., Cabaj, M.K., Brázda, P., Výborný, K., Sedláček, O., Müller, C., Chintakindi, H., Eigner, V., Palatinus, L.: Ionisation of atoms determined by kappa refinement against 3D electron diffraction data. *Nature Communications* **15**(1), 9066 (2024) <https://doi.org/10.1038/s41467-024-53448-2> . Publisher: Nature Publishing Group. Accessed 2025-05-01
- [25] Zuo, J.M., Spence, J.C.H.: Automated structure factor refinement from convergent-beam patterns. *Ultramicroscopy* **35**(3), 185–196 (1991) [https://doi.org/10.1016/0304-3991\(91\)90071-D](https://doi.org/10.1016/0304-3991(91)90071-D) . Accessed 2025-05-01
- [26] Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. *Nature* **323**(6088), 533–536 (1986) <https://doi.org/10.1038/323533a0> . Publisher: Nature Publishing Group. Accessed 2025-04-30
- [27] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E.Z., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 8024–8035 (2019). <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>
- [28] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015). <http://arxiv.org/abs/1412.6980>
- [29] Palatinus, L., Brázda, P., Jelínek, M., Hrdá, J., Steciuk, G., Klementová, M.: Specifics of the data processing of precession electron diffraction tomography data and their implementation in the program PETS2.0. *Acta Crystallographica Section B: Structural Science, Crystal Engineering and Materials* **75**(4), 512–522 (2019) <https://doi.org/10.1107/S2052520619007534> . Publisher: International Union of Crystallography. Accessed 2024-01-14
- [30] Palatinus, L., Jacob, D., Cuvillier, P., Klementová, M., Sinkler, W., Marks, L.D.: Structure refinement from precession electron diffraction data. *Acta Crystallographica Section A: Foundations of Crystallography* **69**(2), 171–188 (2013) <https://doi.org/10.1107/S010876731204946X> . Publisher: International Union of Crystallography. Accessed 2025-05-06

- [31] Klar, P.B., Krysiak, Y., Xu, H., Steciuk, G., Cho, J., Zou, X., Palatinus, L.: Accurate structure models and absolute configuration determination using dynamical effects in continuous-rotation 3D electron diffraction data. *Nature Chemistry* **15**(6), 848–855 (2023) <https://doi.org/10.1038/s41557-023-01186-1> . Number: 6 Publisher: Nature Publishing Group. Accessed 2024-01-10
- [32] PETS. <http://pets.fzu.cz/> Accessed 2025-05-06
- [33] Mendis, B.: Modelling dynamical 3D electron diffraction intensities. II. The role of inelastic scattering. *Acta Crystallographica Section A* **80**(2) (2024) <https://doi.org/10.1107/s2053273323010690> . Publisher: International Union of Crystallography. Accessed 2024-08-27
- [34] Mendis, B.G.: A “Phase Scrambling” Algorithm for Parallel Multislice Simulation of Multiple Phonon and Plasmon Scattering Configurations. *Microscopy and Microanalysis* **29**(3) (2023) <https://doi.org/10.1093/micmic/ozad052> . Publisher: Cambridge University Press. Accessed 2024-08-22
- [35] Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: Bengio, Y., LeCun, Y. (eds.) 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings (2014). <http://arxiv.org/abs/1312.6114>
- [36] Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020) <https://doi.org/10.1038/s41592-019-0686-2>
- [37] Bradbury, J., Frostig, R., Hawkins, P., Johnson, M.J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., Zhang, Q.: JAX: composable transformations of Python+NumPy programs. <http://github.com/google/jax> (2018)
- [38] Kirkland, E.J.: *Advanced Computing in Electron Microscopy*. Springer US, Boston, MA (2010) <https://doi.org/10.1007/978-1-4419-6533-2>
- [39] Oyonarte, E.C., Rebecchi, L., Gholam, S., Diouf, M.D.F., Bigard, E., Pralong, V., Prestipino, C., Kriegel, I., Castellanos-Aliaga, A., Hadermann, J., *et al.*: 3D Electron Diffraction on Nanoparticles: Minimal Size and Associated Dynamical Effects. *ACS nano* **19**(22), 20599 (2025)
- [40] Cowley, J.M., Moodie, A.F.: The scattering of electrons by atoms and crystals. I. A new theoretical approach. *Acta Crystallographica* **10**(10), 609–619 (1957) <https://doi.org/10.1107/S0365110X57002194>

- [41] Bethe, H.: Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie. *Annalen der Physik* **392**, 55–129 (1928) <https://doi.org/10.1002/andp.19283920202>
- [42] Hirsch, P.B., Howie, A., Nicholson, R.B., Pashley, D.W., Whelan, M.J.: *Electron Microscopy of Thin Crystals*. Butterworths, London (1965)
- [43] Petříček, V., Palatinus, L., Plášil, J., Dušek, M.: Jana2020 – a new version of the crystallographic computing system Jana. *Zeitschrift für Kristallographie - Crystalline Materials* **238**(7-8), 271–282 (2023) <https://doi.org/10.1515/zkri-2023-0005>
- [44] Jansen, J., Tang, D., Zandbergen, H.W., Schenk, H.: MSLS, a Least-Squares Procedure for Accurate Crystal Structure Refinement from Dynamical Electron Diffraction Patterns. *Acta Crystallographica Section A: Foundations of Crystallography* **54**(1), 91–101 (1998) <https://doi.org/10.1107/S0108767397010489>
- [45] Zuo, J.M., Spence, J.C.H., O’Keeffe, M.: Bonding in GaAs. *Physical Review Letters* **61**(3), 353–356 (1988) <https://doi.org/10.1103/PhysRevLett.61.353>
- [46] Midgley, P.A., Saunders, M.: Quantitative electron diffraction: From atoms to bonds. *Contemporary Physics* **37**(6), 441–456 (1996) <https://doi.org/10.1080/00107519608217535>
- [47] Dominiak, P.M., Martín Pendás, Á., Woźniak, K.: Focus on Quantum Crystallography. *IUCrJ* **12**(6), 610–613 (2025) <https://doi.org/10.1107/S2052252525008759>
- [48] Yörük, E., Brázda, P., Palatinus, L.: Hirshfeld atom refinement on 3D electron diffraction data: the path to a free refinement of hydrogen atoms. *Journal of Molecular Structure* **1343**, 142798 (2025) <https://doi.org/10.1016/j.molstruc.2025.142798>
- [49] Chodkiewicz, M.L., Olech, B., Jha, K.K., Dominiak, P.M., Woźniak, K.: Hirshfeld atom refinement and dynamical refinement of hexagonal ice structure from electron diffraction data. *IUCrJ* **11**(5), 730–736 (2024) <https://doi.org/10.1107/S2052252524006808>
- [50] Spence, J.M., Zuo, J.-M.: *Electron Microdiffraction*. Springer US, New York (2020) <https://doi.org/10.1007/978-1-4899-2353-0>
- [51] Grosse-Kunstleve, R.W., Adams, P.D.: On the handling of atomic anisotropic displacement parameters. *Journal of Applied Crystallography* **35**(4), 477–480 (2002) <https://doi.org/10.1107/S0021889802008580>
- [52] Madsen, J., Susi, T.: The abTEM code: transmission electron microscopy from first principles. *Open Research Europe* **1**, 24 (2021) <https://doi.org/10.12688/openreseurope.13015.2>

Acknowledgments

TASD acknowledges the support of a Schmidt Science Fellowship and an Oppenheimer Research Fellowship. SM acknowledges funding from EPSRC Centre for Doctoral Training in

Autonomous Intelligent Machines and Systems (Grant No: EP/S024050/1). BC acknowledges funding from Queens' College Cambridge and the Stamps Scholars Program. YG is supported by a Turing AI Fellowship financed by the UK government's Office for Artificial Intelligence, through UK Research and Innovation (grant reference EP/V030302/1) and delivered by the Alan Turing Institute. PAM acknowledges funding from the Engineering and Physical Sciences Research Council (Nos. EP/V007750/1 and EP/R008779/1). The authors would like to thank Lukas Palatinus for making experimental data from their work on dynamical refinements of 3D-ED freely available and for useful discussions.

Author contributions

Conceptualisation: T ASD, PAM, YG. **Methodology:** T ASD, SM, PAM, YG **Investigation:** T ASD, SM, BC. **Data Curation:** SM, T ASD. **Formal Analysis:** SM, T ASD **Software:** SM, T ASD, BC. **Writing - original draft:** SM, T ASD, BC. **Writing - review and editing:** All authors. **Supervision:** T ASD, PAM, YG, SR. **Funding Acquisition:** T ASD, PAM, YG.

Competing interests

The authors declare no competing interests.

Table 1: Results comparing refinements with different modeling assumptions (kinematical, constant thickness dynamical, and the ThicknessNN hybrid model) on synthetic 3D-ED datasets. We report the weighted R-factor (wR_{all}), the root mean squared distance (RMSD) to ground truth ASU, and the Frobenius norm of the difference in thermal displacements for converged models after refining for 200 epochs from an initial 0.2 Å randomly displaced starting model.

Synthetic Sample	Metric	Kinematical	Dynamical (Single Thickness)	Hybrid Physics-ML (ThicknessNN)
Quartz	wR_{all}	0.480	0.121	0.013
	ASU RMSD (Å)	0.0384	0.00035	0.00067
	Thermal Disp. (Å)	0.0139	0.00045	0.00010
CsPbBr ₃	wR_{all}	0.557	0.126	0.0217
	ASU RMSD (Å)	0.0482	0.0069	0.00040
	Thermal Disp. (Å)	0.0664	0.0145	0.0010
Paracetamol	wR_{all}	0.327	0.0265	0.0112
	ASU RMSD (Å)	0.0163	0.00070	0.00048
	Thermal Disp. (Å)	0.0101	0.00105	0.00025

Table 2: Results of dynamical refinements performed on experimental data. * $N_{\text{obs}}, N_{\text{all}}$ for the paracetamol dataset were not explicitly reported at 1.2 (\AA^{-1}) resolution. These values were estimated by examining the number of reflections that passed R_{obs} and wR_{all} filters in the experimental dataset provided in Ref. [32]

Sample	Metric	DYNGO- JANA	Hybrid Physics- ML
Quartz	R_{obs}	5.7 % [31]	4.2 %
	(Continuous) wR_{all}	6.6 % [31]	4.3 %
Rotation)	$N_{\text{obs}}, N_{\text{all}}$	994, 1710	957, 1734
	$g_{\text{max}}(\text{\AA}^{-1})$	1.6	1.6
CsPbBr₃	R_{obs}	5.5 % [24]	6.4 %
	(Continuous) wR_{all}	6.8 % [24]	6.7 %
Rotation)	$N_{\text{obs}}, N_{\text{all}}$	2709, 16345	2714, 17095
	$g_{\text{max}}(\text{\AA}^{-1})$	2.0	2.0
Paracetamol	R_{obs}	9.2 % [32]	10.15 %
	(Precession) wR_{all}	10.34 % [32]	8.4 %
	$N_{\text{obs}}, N_{\text{all}}$	2738*, 6669*	2697, 6621
	$g_{\text{max}}(\text{\AA}^{-1})$	1.2	1.2

Table 3: Absolute structure determination from the 3D-ED quartz dataset.

Sample	Enantiomorph 1	Enantiomorph 2
Quartz		
Spacegroup	$P3_221$	$P3_121$
$g_{\text{max}}(\text{\AA}^{-1})$	1.6	1.6
$N_{\text{obs}}, N_{\text{all}}$	957, 1734	957, 1734
wR_{all}	4.3 %	7.1 %
R_{obs}	4.2 %	6.2 %

Figure 1: Schematic of the 3D-ED method. An irregularly shaped crystal is continuously rotated under the electron beam while diffraction patterns are recorded. Three representative orientations of a quartz crystal with corresponding simulated dynamical diffraction patterns are shown.


Figure 2: Schematic diagram for the proposed hybrid physics-machine learning model for dynamical electron diffraction. Spatial (x, y, z) and thermal displacement (U_{ij}) parameters of the atoms in the asymmetric unit (ASU), along with neural network parameters of the thickness neural network (ThicknessNN) are jointly optimized via backpropagation to minimize the diffraction loss wR_{all} .

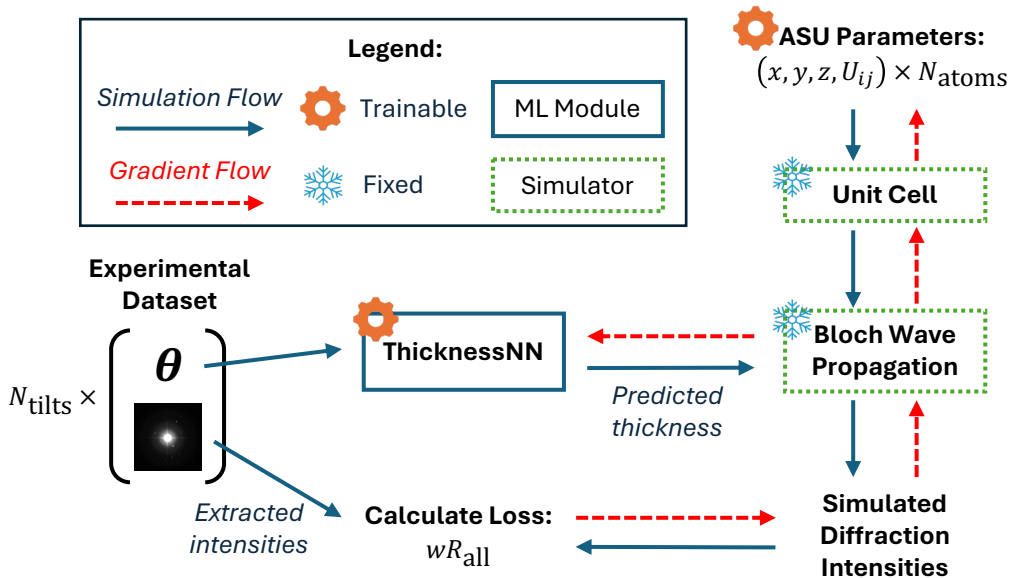
Figure 3: Neural network learning of orientation-dependent crystal thickness distributions. (a) Schematic showing a crystal roughly approximated by a collection of small, parallel slabs with varying thickness at different orientations. (b) Gaussian distributions fit over the distribution of thicknesses for each crystal orientation shown in (a). A neural network learns to predict $\mu(\theta)$ (the mean) and $\sigma(\theta)$ (the standard deviation) for each orientation during crystallographic refinement by minimizing wR_{all} . (c) Plot of representative predicted $\mu(\theta)$ and $\sigma(\theta)$ for a dataset spanning $\pm 60^\circ$. Stars represent the location of Gaussians plotted in (b).

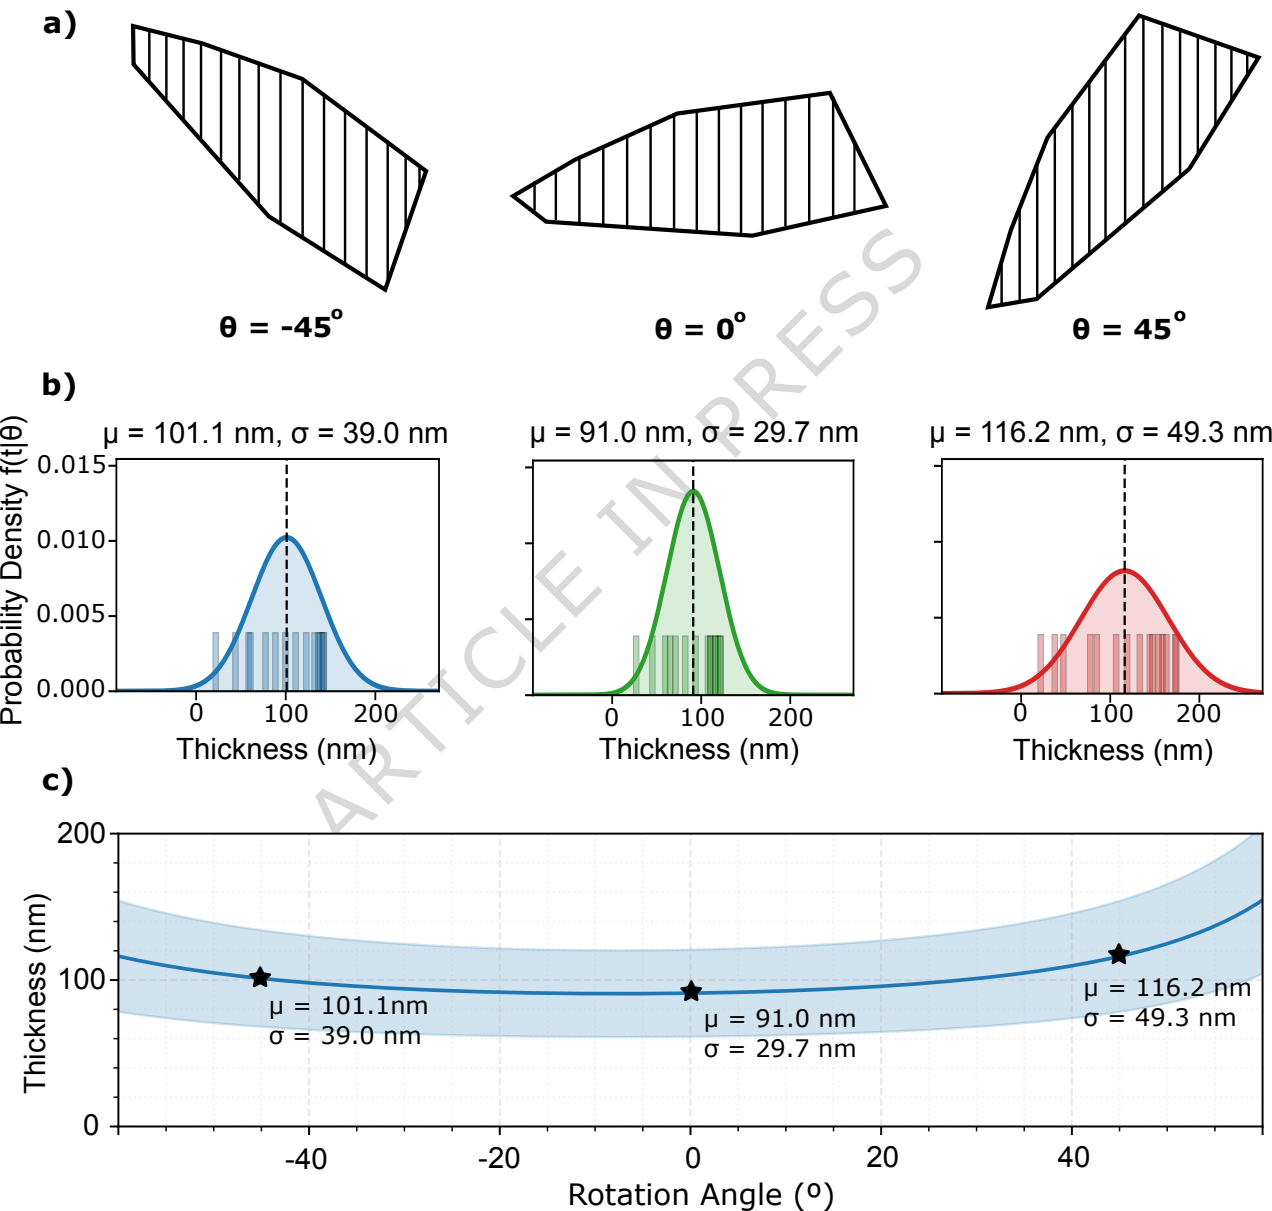
Figure 4: Hybrid Physics-ML model (ThicknessNN) refinements on synthetic data. Refinement results on (a) Paracetamol and (b) CsPbBr₃ synthetic 3D-ED datasets. The weighted R-factor (wR_{all}), the root mean squared distance (RMSD) to ground truth ASU, the Frobenius norm of the difference in thermal displacements, and the thickness distributions are shown as the refinement progresses from initial 0.2 Å randomly displaced starting models. The unit cells, ground truth, starting, and refined models of the asymmetric unit (ASU) are also shown. Thermal displacements are represented by spheres (isotropic) or ellipsoids (anisotropic) centered on atomic sites. The size of the visualized thermal displacements is proportional to their magnitude.

Figure 5: Kinematical refinements on synthetic data. Kinematical refinement results on (a) Paracetamol and (b) CsPbBr₃ synthetic 3D-ED datasets. The weighted R-factor (wR_{all}), the root mean squared distance (RMSD) to ground truth asymmetric unit (ASU), and the Frobenius norm of the difference in thermal displacements are shown alongside the refined ASU and ground truth overlap from initial 0.2 Å randomly displaced starting models. Refinement curves showing the refinement using the proposed hybrid model (ThicknessNN) are also shown for comparison.



Figure 6: Applying the ThicknessNN to experimental datasets. Experimental thickness curves recovered from the hybrid physics-ML model and DYNGO-JANA (dashed lines) for Quartz, CsPbBr₃, and Paracetamol datasets













a) Paracetamol

ASU
Ground Truth


Displaced



Post Refinement

b) CsPbBr₃ASU
Ground Truth

Displaced



Post Refinement

