

Online Thickness Determination with Position Averaged Convergent Beam Electron Diffraction using Convolutional Neural Networks

Michael Oberaigner^{1,*} , Alexander Clausen² , Dieter Weber² , Gerald Kothleitner^{1,3} , Rafal E. Dunin-Borkowski² , and Daniel Knez^{1,*} 

¹Institute of Electron Microscopy and Nanoanalysis, Graz University of Technology, Steyrergasse 17, Graz 8010, Austria

²Ernst Ruska-Centre for Microscopy and Spectroscopy with Electrons, Forschungszentrum Jülich, Jülich 52425, Germany

³Graz Centre for Electron Microscopy, Graz 8010, Austria


*Corresponding authors: Michael Oberaigner, E-mail: michael.oberaigner@felmi-zfe.at; Daniel Knez, E-mail: daniel.knez@felmi-zfe.at

Abstract

Position averaged convergent beam electron diffraction (PACBED) is one of the most convenient and precise thickness determination techniques available in a scanning transmission electron microscope. The thickness is determined by finding the best match of the recorded PACBED pattern with a series of simulated diffraction patterns by visual inspection. The automatization of this process can be enhanced by convolutional neural networks (CNNs), making the method fast and easy to apply. However, the simulation of a synthetic dataset and the training of the CNNs carry a high computational cost. With the aim to simplify this process, we propose to build a server-based database of pretrained CNN models that is accessed by the user via a web service directly from the data acquisition and analysis software. We demonstrate a working prototype comprised of a shared CNN database containing three material systems. By this, the microscope operator can determine the specimen thickness by PACBED within a few seconds in a reproducible way during a microscope session, without any prior knowledge about machine learning or multislice modeling. Furthermore, the service is integrated into other software and workflows through the API.

Key words: automatic thickness determination, convolutional neural network, integrated GUI, PACBED, STEM, web service

Graphical Abstract



Introduction

Thickness determination of a sample in the transmission electron microscope (TEM) in the dimension parallel to the beam is essential for many TEM characterization techniques, such as quantitative electron energy loss spectrometry, absorption correction in energy dispersive X-ray spectrometry, interpretation of coherent image contrast, obtaining volumetric information from 2D projected features, and for the direct relation of experimental data with elastic or inelastic multislice calculations (Williams & Carter, 2009; Kirkland, 2010; Pennycook & Nellist, 2011; Knez et al., 2020; Lammer et al., 2022). Several techniques have been developed, based, for instance, on the energy loss of inelastically scattered electrons (Egerton, 2011), sample tilting with surface markers

(Williams & Carter, 2009), electron tomography (Midgley & Dunin-Borkowski, 2009), double cross section preparation of focused ion beam (FIB) lamellae (Boxleitner et al., 2001), or linking TEM analytical spectroscopies (Kothleitner et al., 2014), to name a few. In scanning transmission electron microscope (STEM) mode, in particular, assessing contrast variations within overlapping convergent beam electron diffraction (CBED) disks has proven to provide accurate thickness information for the crystalline specimen (Williams & Carter, 2009). For an Ångström-sized electron probe, as routinely obtained in aberration-corrected STEM, the CBED pattern is not only very sensitive to the thickness of the specimen but also to the position of the probe within the unit cell. To be able to measure the actual thickness, LeBeau et al. introduced

Received: July 28, 2022. Revised: November 10, 2022. Accepted: December 12, 2022

© The Author(s) 2023. Published by Oxford University Press on behalf of the Microscopy Society of America.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

the position averaged convergent beam electron diffraction (PACBED) method, where CBED patterns are averaged over different positions of the beam and the obtained diffraction pattern is compared with multislice simulations via visual inspection (LeBeau et al., 2009, 2010). By averaging, the PACBED patterns are not affected by spherical aberration, defocus, and spatial incoherence (LeBeau et al., 2010). Nevertheless, PACBED features are still influenced by sample tilt, polarization effects, composition, lattice distortions/strain, inelastic phonon scattering, and amorphous surface layers. Dynamical diffraction effects, furthermore, render the contrast variations strongly non-linear with increasing thickness (Ophus, 2019). However, despite these difficulties, PACBED has proven to be a very accurate technique, especially for thin samples (Xu & LeBeau, 2018). The accuracy of this method can be further improved by zero-loss energy filtering (Pollock et al., 2017).

To reduce the influence of human perception and thereby to increase the reliability and reproducibility of the method, efforts have been made to automatize this step using least squares fitting or machine learning (Gao et al., 2010; Ophus et al., 2017; Pollock et al., 2017; Zhang et al., 2020; Shen et al., 2021). In particular, in a seminal paper, Xu and LeBeau demonstrated the use of convolutional neural networks (CNNs) that are trained on multislice simulation data to extract additional information including pixel size in milliradian per pixel, center, and rotation as well as sample mistilt (Xu & LeBeau, 2018).

Based on this approach, we aim to further speed up and simplify the thickness determination process and make it easily accessible even for users without knowledge of multislice simulations or neural networks. We developed a server-based database containing pretrained networks for different materials, experimental conditions, and crystallographic orientations that can be accessed from a local machine through a web service.

To demonstrate the capability and flexibility of an automated thickness determination, we developed a hybrid Python and DM Script application with a graphical user interface (GUI) for the Gatan Microscopy Suite (GMS), a software that is widely used in the community for data acquisition and analysis. Thereby, we are able to obtain the specimen thickness and mistilt during a microscopy session directly from the data acquisition software with a few mouse clicks only. We test the performance of our software on three experimental examples TiO_2 (rutile), SrTiO_3 (STO), and beryl ($\text{Be}_3\text{Al}_2\text{Si}_6\text{O}_{18}$) with two different electron energies, 80 and 300 keV. Our set of CNNs automatically analyzes and corrects for pixel size, center, rotation, small specimen tilt (<10 mrad), and variations of convergence angle. The reliability of the CNN is further tested with artificial and experimental PACBED data containing different levels of Gaussian noise, as well as significant surface amorphization.

We propose a material database collecting different PACBED trained networks that are trained centrally or collected from users, and providing a web service interface to this information to the STEM community to allow an easy, instant, and validated assessment of sample thickness for many users.

Materials and Methods

All code, instructions, and data are available online, see the section “Availability of data and materials”.

PACBED Acquisition

PACBED patterns for the case studies were acquired on a probe-corrected FEI Titan³ 60–300 G2 STEM operated at either 80 or 300 kV. The convergence angle was set for rutile and STO between 16 and 22 mrad. For beryl, a smaller convergence angle of 7 mrad is used. The patterns were recorded with an UltraScan CCD camera with a resolution of 2048×2048 pixels. Data acquisition and analysis were performed with the GMS (version 3.4). Note that for accessing the database from within GMS, its Python capability [available only in newer program versions (≥ 3.4)] is used.

Sample Preparation

The STO sample was prepared from a single crystalline sample via FIB milling using Ga ions with an energy of 30 keV and a final polishing step at 5 kV. For preparing the rutile sample, wedge polishing and ultra-low-energy milling at 900 V were used, utilizing a Fischione Nanomill. The Beryl sample has been prepared from a natural crystal from Minas Gerais/Brazil by wedge polishing followed by a cleaning step with Ar ions with 1 keV energy using a Gatan PIPS II (10 min@4° angle).

Automatic PACBED Analysis

For the PACBED analysis of a specific material system and the set of experimental conditions (crystal structure, convergence angle range, orientation, and acceleration voltage), three CNN models, each within Fig. 1 depicted architecture, were trained to predict the scale, the sample thickness, and mistilt. The scale is defined as the pixel size ratio of the recorded PACBED to the corresponding simulated PACBED. The CNNs are implemented in a Python environment (v3.10.) by using the Tensorflow (v2.8.) libraries and the included Keras API. Each model consists of the Xception architecture (Chollet, 2017), which is pretrained with the ImageNet dataset (Deng et al., 2009), with an input dimension of 200×200 pixels for the PACBED pattern, two fully connected layers with 1,024 nodes each including dropout layers (Srivastava et al., 2014) with a ratio of 0.2 and a Softmax classifier (Bishop, 2006). Additionally, to increase the available information for the network, the convergence angle is provided to the fully connected layers. This angle is commonly known by the user or given by the microscope.

The CNN models generate a probability figure of merit for each of a range of discrete values of scale, thickness, or mistilt that the network was trained for. A high predicted probability




Fig. 1. CNN architecture used for predicting thickness, mistilt, and scale.

for one of these trained values, while the probability of all other values remains low, indicates a clear prediction by the network.

The training of the CNN relies on a massive data set, which is generated via PACBED multislice simulations with the abTEM code (Madsen & Susi, 2021). The PACBED patterns are simulated for each material system with varying experimental parameters. Here, we demonstrate and apply our program to rutile, STO, and beryl, all in $\langle 001 \rangle$ zone axis with a maximum sample thickness of 100 nm in steps of 1 nm. The high tension was set to 80 kV for rutile and 300 kV for STO and beryl. The convergence angle was varied for beryl between 5 and 15 mrad and for rutile and STO between 15 and 25 mrad in steps of 0.5 mrad. Additionally, the sample mistilt was varied from 0 to 10 mrad in steps of 1 mrad with an azimuth (direction of mistilt) varying from 0 to $\pi/2$ rad in steps of $\pi/10$. This results in nearly 130,000 individual simulated patterns with a dimension of 200×200 px² for each material system. The total time of generating the dataset is around 7 h on a desktop PC with an AMD Ryzen™ 7 3700X and an NVIDIA GeForce RTX 2060 Super GPU, which results in around 2 s per pattern. The dataset is split into a training set with 95% of the patterns and a validation set with 5% of the patterns. Although a large number of training images is available, additional data augmentation is implemented to avoid overfitting and to increase the robustness of the CNNs. Table 1 lists all applied transformations with the corresponding ranges. To the normalized images, random noise is added, which follows a Gaussian distribution with a randomly selected value for the standard deviation in a range between 0 and 0.1. It should be noted that in experimental data, noise is comprised of both a Poisson and a Gaussian contribution dependent on the experimental settings. In the Supplementary Information S1, we demonstrate that the CNNs trained with Gaussian noise are also robust against pure Poisson noise, which justifies our simplified approach. Further, the modulation transfer function of a typical recording device (e.g., Ultrascan CCD camera) does not affect the accuracy of the thickness determination (Pollock et al., 2017) and is therefore not considered. To the normalized convergence angle, Gaussian noise with a standard deviation of 0.05 is also added for the training.

Training takes around 50 min per epoch on a desktop PC with an NVIDIA GeForce RTX 2060 Super GPU. The number of epochs for training depends on the kind of prediction (for thickness 40 epochs, for mistilt 20 epochs, and for scale 5 epochs) and is assessed based on a categorical cross-entropy loss function. The models with the lowest validation loss are saved and used for the PACBED analysis. Early stopping is implemented by terminating the training if the validation loss could not be increased in the last 10 epochs.

Before a recorded PACBED can be analyzed, it has to be preprocessed. If the dimension of the PACBED image exceeds

Table 1. Applied transformations and value ranges for training data augmentation.

Operation	Range
Shearing	0.05° (x and y directions)
Rotation	$\pm 45^\circ$
Shifting	$\pm 10\%$ (x and y directions)
Flipping	Up/down and left/right
Scaling	$\pm 20\%$ (only for thickness and mistilt)
Noise	Gaussian noise (σ from 0 to 0.1)

680×680 px², it is downsized to this dimension to speed up subsequent operations. After preprocessing, the image is resized to 200×200 px² for the CNN network. In case, the PACBED image dimension is lower, and it is upsampled to 200×200 px² already before the preprocessing step. After the PACBED is centered by its center of mass, a background signal is subtracted, dependent on the corresponding CNN applied thereafter. While for the two CNNs predicting the thickness and mistilt a uniform background is subtracted, for the scaling CNN a Gaussian blurred image with a standard deviation of 90% of the image dimension is used. All images are normalized between -1 and 1 . Before thickness and mistilt prediction, we have to ensure that the PACBED pattern is within the size of the training PACBEDs ($\pm 20\%$ of the simulated PACBEDs due to data augmentation). To this end, the pixel size of the PACBED is corrected iteratively with the predicted value of the scale model until it converges or exceeds the maximum runs. To avoid oscillation, the iteratively predicted values are damped before rescaling the PACBED [see eq. (1)].

$$s_k = \frac{k + (n - k) * s_{\text{pred}}}{n} \quad (1)$$

with s_k —damped predicted scale value, s_{pred} —predicted scale value, n —maximum number of iterations, and k —current iteration.

After these processing steps, the PACBED is passed to the other two CNN models for thickness and mistilt.

The functionality of the look-up service is demonstrated with experimental PACBEDs below (Figs. 2, 3), starting from a given crystal structure, to generate a training dataset of PACBED patterns. Since the potential parameter space of material, zone axis, and acceleration voltage is very large, a production deployment should only pretrain common combinations of material, zone axis, and high tension. Additional combinations can be trained on demand.

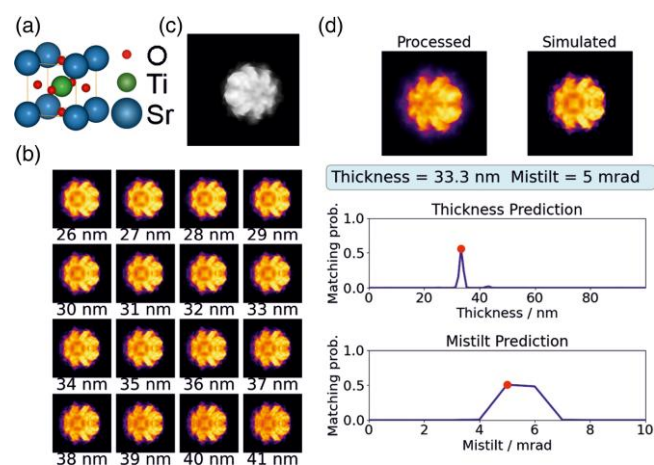


Fig. 2. Functionality of the look-up service in the example of STO. (a) Crystal structure of STO, which is used to simulate PACBEDs. (b) Set of simulated PACBEDs from STO at 300 kV high tension, 19.5 mrad convergence angle, and 5 mrad mistilt with varying specimen thickness. (c) Experimental PACBED of STO at 300 kV and 19.6 mrad convergence angle (thickness estimated by visual inspection: 29–35 nm). (d) Feedback given to the user from the service by submitting the experimental PACBED from (c), showing the processed PACBED, the predicted PACBED, and the outputs of the CNNs to validate the results.

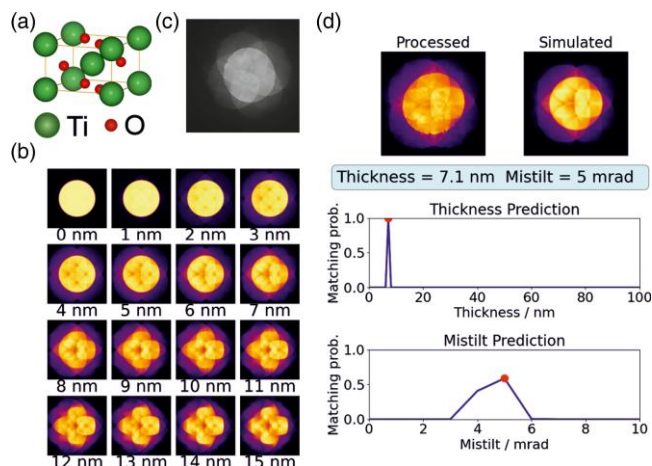


Fig. 3. Functionality of the look-up service in the example of rutile. (a) Crystal structure of rutile, which is used to simulate PACBEDs. (b) Set of simulated PACBEDs from rutile at 80 kV high tension, 20 mrad convergence angle, and 4 mrad mistilt with varying specimen thickness. (c) Experimental PACBED of rutile at 80 kV and 20.6 mrad convergence angle (thickness estimated by visual inspection: 6–7 nm). (d) Feedback given to the user from the service by submitting the experimental PACBED from (c), showing the processed PACBED, the predicted PACBED, and the outputs of the CNNs to validate the results.

Web Service

An overview and the currently implemented components of the service are shown in Fig. 4. The CNN look-up is wrapped into a simple web service. The user submits parameters (crystal, material orientation, high tension, and convergence angle) and a PACBED pattern and receives thickness, mistilt, scale, and validation information. The request and return type definitions are documented in Fig. 5. Examples and source code are available online, see the section “Availability of data and materials”. For the first iteration, the material identifier can be the string “Rutile”, “Strontium Titanate”, or “Beryl”. Validation information consists of a plot with side-by-side

comparison of the submitted PACBED pattern with the best-matching simulation in combination with a graph of the matching probability of the CNN look-up for each possible return value (Fig. 2d). The plot is returned as a base64-encoded PNG image. In the future, the material identifier should match identifiers from a widely accepted structure database such as the ICDD Powder Diffraction File, and the returned information can be extended with information in numerical form to match the requirements of the clients.

Currently, an error is returned if no pretrained model with matching parameters is available. In the future, an option to train the requested model could be provided to users. This would queue the training job, create a ticket for users to check the status and progress, and possibly notify users as soon as the training is complete so that the requested look-up can be performed.

Clients

The database can be accessed from three different client applications. First, a web-based form (Fig. 6a) is integrated with the web service and where a PACBED pattern can be uploaded. Second, a GUI is running within the GMS environment (Fig. 6b). And third, a generic Python-based form can be used from within any Python-based program. The GMS client combines data acquisition, submission to the web service, and display of the result into a convenient, streamlined workflow, as shown in Fig. 6b. For this, the PACBED pattern in question is required to be the foremost active image. Clicking on the “Get PACBED pattern” button initializes the process with this image, and its name is then shown below the pattern. The GUI allows the user to define the crystal structure and orientation of the sample. For now, the drop-down list only allows the selection between the three pretrained test structures STO $\langle 0\ 0\ 1 \rangle$, rutile $\langle 0\ 0\ 1 \rangle$, and beryl $\langle 0\ 0\ 1 \rangle$. The high tension and convergence angle used to acquire the PACBED are also required and are filled automatically if the information is

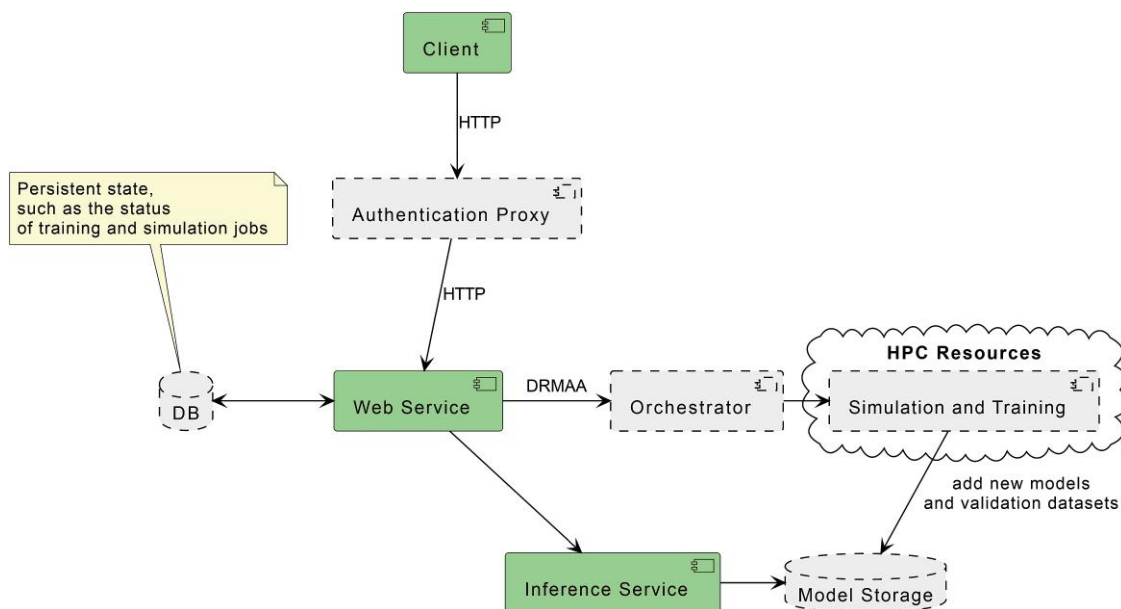


Fig. 4. Service architecture. The currently implemented components are denoted with a solid border, and the planned components with a dashed border. DRMAA, distributed resource management API; HTTP, Hypertext transfer protocol; DB, database; HPC, high-performance computing.

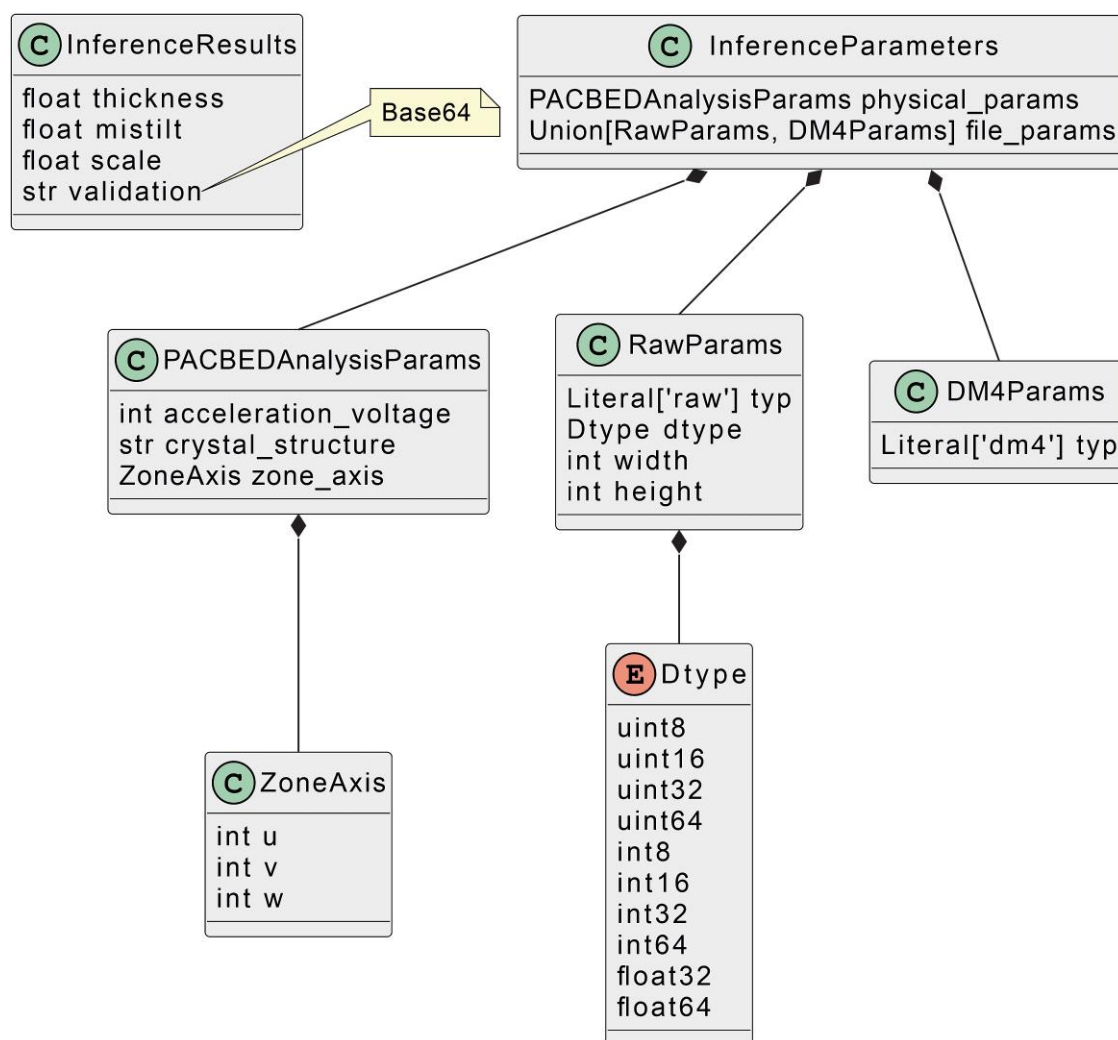


Fig. 5. Request and response schema, showing the type definition of each variable. C, class; E, enumerated type.

found in the image tags. The look-up request is then started via a click on the “Determine Thickness” button, which then uploads the image and metadata to the service, invokes the CNN PACBED analysis, and returns the thickness and the mistilt of the specimen. The result can then be further evaluated by the user via the feedback sheet (Figs. 2d, 3d) that is also generated and shown in a new image window in GMS.

Both the generic Python client and the GMS GUI are available online (<https://github.com/MichaelO1993/PACBED-CNN/blob/main/client/examples/>).

Results

The thickness determination should be fast enough in order to provide on-the-fly information directly at the microscope during an experiment. Therefore, the round-trip time for a request has been tested and optimized. We found that the look-up takes less than 3 s over a GBit/s scale network connection. The request size can be large, in the region of 50 MB, if data are acquired using a high-resolution camera and sent in its native size, since the PACBED pattern is transmitted uncompressed in numerical form. Uploading this amount of data can dominate the round-trip time over a slower network.

High-resolution patterns can be scaled down before submitting to improve performance since the inference will be performed at a fixed resolution. See “Automatic PACBED analysis” in the Materials and Methods section for details.

A trained model has a size of 275 MB, and the accompanying validation dataset takes about 5 GB. However, this validation data, which is only needed to allow the user to check the correct prediction, could also be compressed or even generated by an “on-the-fly” calculation using fast simulation codes such as PRISM (Ophus, 2017) or its implementation in abTEM (Madsen & Susi, 2021).

The performance of the automatic thickness determination by the CNNs is compared with the least square fitting (LSF) method and visual inspection on the example of rutile at a high tension of 80 kV. For this purpose, 20 PACBED patterns are simulated with random parameters of thickness, mistilt, azimuthal direction, and convergence angle within the range of the training dataset of the CNNs. These PACBED patterns were never seen by the CNNs. The percentage deviation of the predicted thickness from the true thickness is compared between the different methods with increasing noise levels (Fig. 7). The random noise follows a Gaussian distribution with different standard deviations, which is added to the normalized image. The performance of the CNNs with Poisson

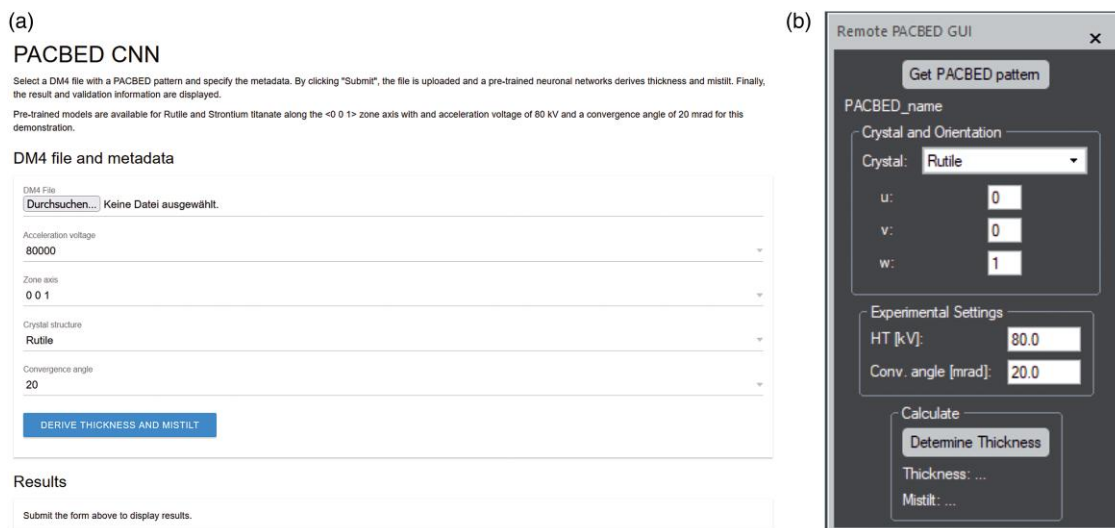


Fig. 6. Demonstration of the client forms. The screenshots are taken from the web-based form (a) and the GUI in GMS (b) for the PACBED thickness determination service.

noise is similar as demonstrated in [Supplementary Information S1](#). The thickness prediction by visual inspection is done by comparing the PACBEDs manually with the training dataset to find the best match, starting with the noisiest PACBEDs. Note that this process is subjective and can contain human errors. For the LSF, a pixelwise L2 norm is calculated between the PACBED and a part of the training dataset with the closest convergence angle and 0 mrad mistilt. The thickness is taken from the PACBED pattern with the lowest L2 norm. [Figure 7](#) demonstrates that the automatic thickness determination by CNN is less sensitive to noise and has better

accuracy compared to the other methods. The low accuracy of LSF could be caused by the low dimension of the images with 200×200 px² and the mistilt. To improve the results, the resolution of the PACBEDs should be increased and the database should include mistilts with different azimuthal directions. However, this demands higher computational costs and memory. Furthermore, the position, rotation, and pixel size of experimental PACBEDs have to be precisely aligned to match the simulated PACBEDs. Klicken oder tippen Sie hier, um Text einzugeben.

Although the CNN models have good accuracy in the prediction of simulated PACBEDs, they have to perform also on real-world data. Therefore, the same performance test is applied to five experimental PACBEDs from rutile ([Fig. 8](#)). Reference thicknesses are estimated by visual inspection from the PACBEDs with the lowest noise. The automatic thickness determination by CNN is again robust against noise. Although the percentage deviation reaches up to 20%, the maximum absolute difference is only 2 nm from the visual thickness estimation, which can contain human errors ([Fig. 8c](#)).

To keep the number of required models low, each model for a particular set of materials, orientation, and high tension is trained for a large convergence angle range. [Figure 9](#) shows the performance at different convergence angles and degrees of surface amorphization on the example of rutile (0 0 1). The PACBEDs are recorded with two different convergence angles (16.4 and 23 mrad) at the same spot position. Even though the thickness prediction becomes more challenging with increasing convergence angles, thicknesses, and levels of surface amorphization, the model performs well in [Figure 9a](#). In [Figure 9b](#), a mismatch between the thickness prediction is visible. The CNN predicts a larger thickness for the PACBED, which is recorded with a larger convergence angle. However, a second peak is visible, which matches with the PACBED with the smaller convergence angle. This thickness can also be confirmed by visual inspection. This demonstrates the importance of a feedback, providing the user the information about how reliable the prediction is, and which other thicknesses have a high matching probability. Moreover, the

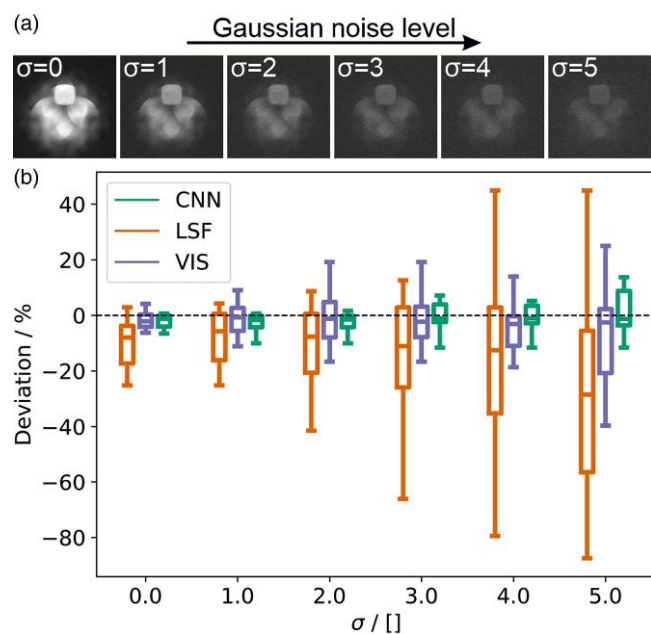


Fig. 7. Performance of the different methods with simulated PACBED patterns at different levels of artificial Gaussian noise. (a) Example of a PACBED pattern with different levels of noise. (b) Percentage deviation of the predicted thicknesses from the true thickness. CNN, automatic thickness determination by CNN; LSF, least square fitting; VIS, visual inspection.

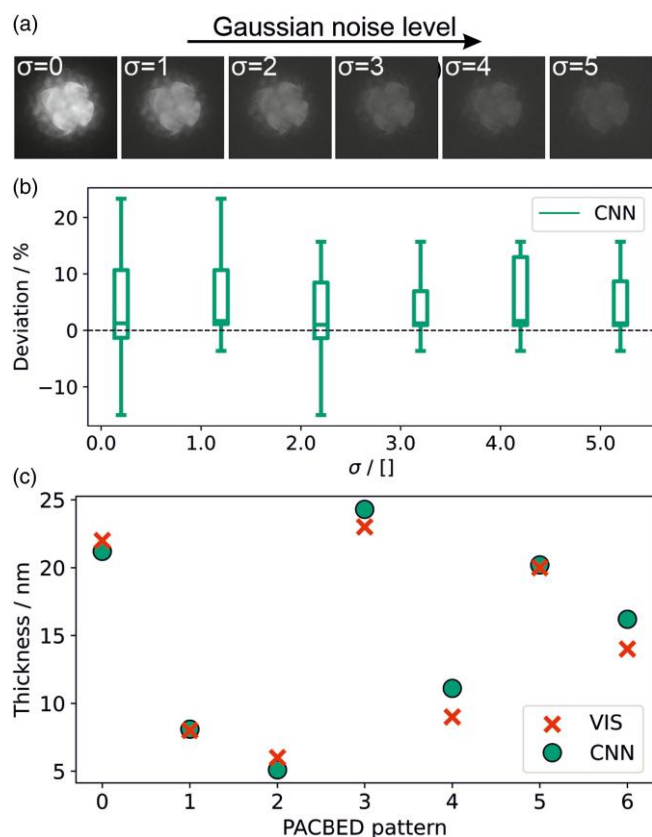


Fig. 8. Performance of the CNN method with experimental PACBEDs at different levels of artificial Gaussian noise. (a) Example of an experimental PACBED pattern with different levels of noise. (b) Percentage deviation of the CNN prediction from the visual inspection. (c) Predicted thicknesses of the PACBED pattern with the lowest noise by CNN and visual inspection.

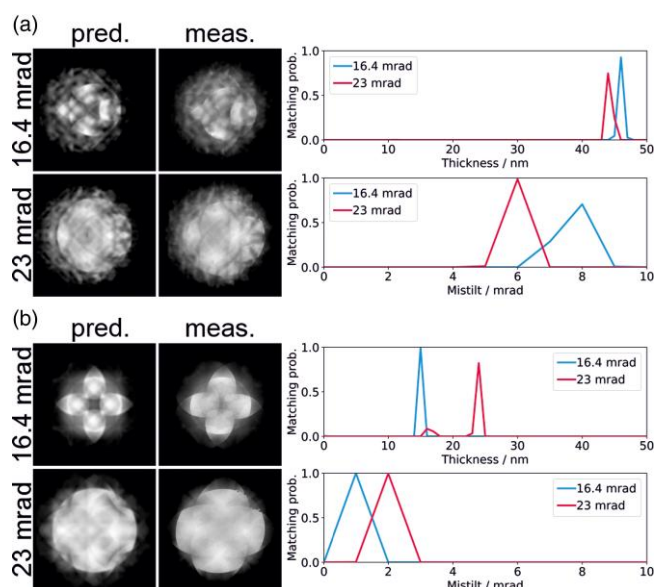


Fig. 9. Performance of the CNN for rutile at experimental PACBEDs with different convergence angles. (a) Predictions of both convergence angles are overlapping (thickness estimated by visual inspection: 44–48 nm). (b) Model fails at the larger convergence angle, which is recognizable by comparing the recorded PACBED with the predicted PACBED. A second peak is visible, which matches with visual inspection (thickness estimated by visual inspection: 15–17 nm).

structure mismatch can be used to identify incorrect input parameters by comparing the submitted pattern with the best match of the CNN (see [Supplementary Information S2](#)).

The reliability and accuracy of the thickness determination of the PACBED method depend on the applied acceleration voltage, convergence angle, and the thickness itself ([Supplementary Information S3](#)). This statement is valid both for a classic visual determination and for the CNN prediction. This behavior can already be observed in the analysis of STO ([Fig. 2b](#)) and of rutile ([Fig. 3b](#)). The PACBED patterns of rutile, simulated at 80 kV, exhibit more distinct contrast changes with thickness compared to the PACBEDs from STO, simulated at 300 kV. This evolves at the broadening and lowering of the thickness prediction peak of the CNN ([Fig. 2d](#)) compared to the thickness prediction at rutile ([Fig. 3d](#)). The advantage of the CNN prediction is that it might give a confidence interval of the estimation. However, it should be kept in mind that the PACBED method is not suitable for arbitrary large acceleration voltages, convergence angles, and thicknesses for precise thickness determination. For example, a maximum convergence angle of about 20 mrad is suggested for STO at 300 kV ([LeBeau et al., 2010](#)). For other materials such as beryl and other experimental conditions, different convergence angles are necessary in order to obtain meaningful results. Moreover, the less unique the patterns are, the less robust the prediction of the models will be against noise and contamination, which, in principle, imposes the same limitations as with visual inspection of PACBED data ([LeBeau et al., 2010](#)), see [Supplementary Information S3](#). It shall be noted that the enhanced accuracy by zero-loss energy filtering, as shown by [Pollock et al. \(2017\)](#), is lower than the thickness resolution of approximately 1 nm that is obtained in our case and given by the simulation settings.

The performance of the service also allows creating thickness maps, as demonstrated by [Xu & LeBeau \(2018\)](#). A 4D STEM dataset containing PACBED patterns is generated by subpixel scanning over several unit cells. [Figure 10](#) demonstrates a thickness map of beryl, recorded at a high tension of 300 kV, and a convergence angle of 7 mrad. The map contains 400 analyzed PACBED patterns. Each pixel can be further investigated by the corresponding validation plot.

Discussion

The service is currently developed as a demonstrator for three material systems. As such, it is already useful for a local deployment if the thickness of many specimens with the same material and imaging conditions are to be determined. A much larger number of materials and imaging conditions could be supported for a deployment that is available for a wider circle of users. This is a promising candidate for data science repositories and services such as European Open Science Cloud (EOSC) or Novel Materials Discovery (NOMAD) ([Draxl & Scheffler, 2019](#)). This could be facilitated by the following technical improvements in the implementation.

Authentication

Since using the service will bind significant resources compared to a simple web page, access should be restricted and per-user limits should be implemented. Authentication should be integrated with the existing educational identity ecosystem,

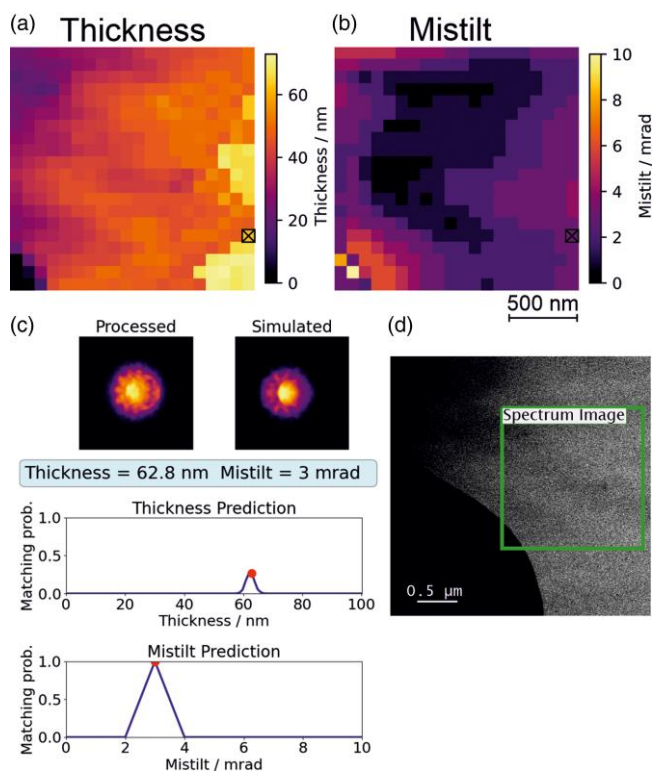


Fig. 10. Thickness (a) and mistilt (b) determination for beryl. The lower left corner of the mistilt map contains artifacts due to a hole in the specimen. Inside the crystal, a mistilt of ~ 6 mrad was found. (c) Example of the corresponding validation plot from the marked pixel in (a). (d) shows the reference HAADF image.

for example, leveraging eduGAIN or related identity provider services (Torroglosa et al., 2018). This provides ease of use and a very low barrier of entry while shielding from malicious usage. Per-user limits can be implemented using standard rate-limiting techniques, constraining the maximum number of requests per time unit and per user. The rate limit could be realized using an in-memory key-value database, like Redis (Carlson, 2013). The concrete rate limit/quota configuration depends on the available resources and the number of concurrent users. If necessary, this can be augmented by a persistent allow/deny list, for example, to persistently block users, or allow more-than-default quota.

Request Size

In the future, compressed formats could be explored alongside the current uncompressed DM4 or numerical raw data, and downscaling could be integrated with the web service client Application Programming Interface (API) for convenience.

Resource Use of Back-End

Currently, the simulated training dataset for each model is stored on the back-end server to look up and display the best-matching simulated PACBED pattern for a request. With 10^6 structures in the ICDD Powder Diffraction File and potentially many combinations of zone axis, acceleration voltage, and convergence angle per material, an absolute upper limit of 10^6 models is assumed for a service that is used ubiquitously for thickness

determination in STEM. Due to the fact that many structures are not suitable for STEM investigation, this can only be a very rough estimate for illustration purposes. Scaling the database to 10^6 models would require 5 PB if the validation dataset is stored, while the models alone would only require 275 TB. The estimated costs of saving a complete database with 10^6 models are 5,600 USD per month for the models and 105,000 USD per month for the validation dataset, assuming a storage price of 0.02 USD per GB for a hot data storage (according to Google, <https://cloud.google.com/storage/pricing>, accessed on 05 October 2022). Measures for size reduction, such as a fast on-the-fly simulation, caching, or compression, could be beneficial for a production deployment for that reason.

Reducing the model size, utilizing crystal symmetries to avoid redundant models, and improving training performance by transfer learning (Pan & Yang, 2009) from similar materials, hyperparameter optimization (Falkner et al., 2018; Yang & Shami, 2020), and further optimizing of the source code could also be explored to reduce the cost of a production deployment. With a simulation time of 7 h on a GPU and 65 training epochs with 50 min per epoch, generating a model requires roughly 61 GPU-hours. At a price of about 3 USD per hour (according to Amazon, <https://aws.amazon.com/ec2/instance-types/p3/>, accessed on 26 July 2022), training a model costs about 184 USD. That means a deployment of this service should focus on materials and imaging conditions that are used often, and it may require financial or in-kind contributions from users for the training of additional models of their choice. Optimization of data generation, training, and model storage as well as sufficient funding respectively a business model will be imperative for a large-scale deployment that approaches 10^6 models to bring down the cost. The cost of inference is low in comparison, which is estimated to be about 0.013 ct per request based on measured 12 CPU-seconds computation at 4 ct per CPU-hour (according to Amazon, <https://aws.amazon.com/ec2/instance-types/t3/>, accessed on 26 July 2022) and 100 kB outbound traffic per request at 9 ct per GB. Nevertheless, rate limits and potentially fees for heavy users should be considered.

Deployment and Scaling

Since microscope control computers are often not connected to the internet, a local proxy or even a local deployment of the look-up engine that pulls trained models from a central database could be considered. Furthermore, such local deployments could train models locally and submit them to a centralized model database, leading to a “sharing economy” between users of the service. This would require quality assurance procedures for a model to be accepted into the database.

Assuming 1,000 concurrent users, a centralized service should serve 16 requests per second if each user performs one thickness determination per minute. Given the performance of the inference, this seems feasible.

Reliability

The quality and robustness test of the look-up is described in the Results section. For a production deployment, one should explore in more detail how errors can be recognized

automatically and/or by users. Blind trust of the results should be avoided since material, specimen preparation, alignment, microscope state and settings, charging, contamination, and camera characteristics open a large parameter space that will probably be explored only gradually by users. A feedback mechanism to report issues and disclaimers or user agreements to manage liability for wrong results would be of importance for a service that is used in production. Manual thickness determination suffers from similar potential impact of experimental conditions on the result. However, it requires more attention from users because it does not provide a result as a “magic number” but requires interpretation by the user.

Results can be easier to reproduce if models are archived and requests are logged. Creating a DOI for a particular look-up with archived request data and model identifier would allow to document and reference results that are used in scientific publications reproducibly.

Conclusion

The web service to determine specimen thickness based on PACBED patterns using pretrained CNN-based models shows a promising use case for the integration of online services, reproducibility, and metadata in practical microscopy workflows. The proof-of-concept client that was integrated with Gatan Digital Micrograph is convenient and fast compared to previous manual thickness determination. Integration into other software, including electronic lab notebooks or data repositories, seems feasible. Our CNN model has been tested with different experimental parameters and acquisition artifacts and already provides thickness predictions with a high level of accuracy and robustness.

In its present form, our program can be used as a local deployment for routine analysis of specimens with uniform material, orientation, and imaging conditions. Further developments to improve performance and scalability, allow easy deployment, implement training on demand, and improve quality assurance will likely allow offering this as a service for a wider audience. Furthermore, a business model to sustain such a service should be found. This could consist of public funding, in-kind provision of resources, or usage fees from users, to name examples.

Availability of data and materials

All code, instructions, and example data are available from the GitHub site of this project: <https://github.com/MichaelO1993/PACBED-CNN>. The code used for this publication is also available at <https://doi.org/10.3217/m55nr-g7v97>.

Supplementary material

To view [Supplementary material](#) for this article, please visit <https://doi.org/10.1093/micmic/ozac050>.

Acknowledgments

The authors thank Martina Dienstleder for sample preparation, Prof. Ferdinand Hofer for providing the beryl sample, and Johannes Biskupek for providing the rutile sample.

Financial support

This project has received funding from the European Union’s Horizon 2020 Research and Innovation Programme (Grant No. 823717, project “ESTEEM3”) and the Austrian Science Fund (FWF) under grant number I4309-N36 “Orbitalkartierung an Grenzflächen”. For the purpose of open access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission. The authors also acknowledge fundings by the Zukunftsfond Steiermark for purchasing the K2 camera.

Conflict of interest

The authors declare that they have no competing interests.

References

- Bishop CM (2006). *Pattern Recognition and Machine Learning*, 1st ed. New York: Springer.
- Boxleitner W, Hobler G, Klüppel V & Cerva H (2001). Simulation of topography evolution and damage formation during TEM sample preparation using focused ion beams. *Nucl Instrum Meth Phys Res Sect B: Beam Interact Mater Atoms* 175–177, 102–107.
- Carlson J (2013). *Redis in Action*, 1st ed. USA: Manning Publications Co.
- Chollet F (2017). Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, July 21–26, 2017, pp. 1800–1807. USA: IEEE.
- Deng J, Dong W, Socher R, Li L-J, Li K & Fei-Fei L (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, June 20–25, 2009, pp. 248–255. USA: IEEE.
- Draxl C & Scheffler M (2019). The NOMAD laboratory: From data sharing to artificial intelligence. *J Phys Mater* 2, 036001. <https://doi.org/10.1088/2515-7639/ab13bb>.
- Egerton RF (2011). *Electron Energy-Loss Spectroscopy in the Electron Microscope*. New York: Springer.
- Falkner S, Klein A & Hutter F (2018). BOHB: Robust and efficient hyperparameter optimization at scale. In *Proceedings of the 35th International Conference on Machine Learning*, Dy J, Krause A (Eds.), vol. 80. pp. 1437–1446. USA: PMLR. Available at <https://proceedings.mlr.press/v80/falkner18a.html>.
- Gao W, Addiego C, Wang H, Yan X, Hou Y, Ji D, Heikes C, Zhang Y, Li L, Huyan H, Blum T, Aoki T, Nie Y, Schlom DG, Wu R & Pan X (2019). Real-space charge-density imaging with sub-ångström resolution by four-dimensional electron microscopy. *Nature* 575, 480–484.
- Kirkland EJ (2010). *Advanced Computing in Electron Microscopy*, 2nd ed. New York: Springer.
- Knez D, Dražić G, Chaluvadi SK, Orgiani P, Fabris S, Panaccione G, Rossi G & Ciancio R (2020). Unveiling oxygen vacancy superstructures in reduced anatase thin films. *Nano Lett* 20, 6444–6451.
- Kothleitner G, Grogger W, Dienstleder M & Hofer F (2014). Linking TEM analytical spectroscopies for an assumptionless compositional analysis. *Microsc Microanal* 20, 678–686.
- Lammer J, Berger C, Löffler S, Knez D, Longo P, Kothleitner G, Hofer F, Haberehner G, Bucher E, Sitte W & Grogger W (2022). A method for a column-by-column EELS quantification of barium lanthanum ferrate. *Ultramicroscopy* 234, 113477. <https://doi.org/10.1016/j.ultramic.2022.113477>.
- LeBeau JM, Findlay SD, Allen LJ & Stemmer S (2010). Position averaged convergent beam electron diffraction: Theory and applications. *Ultramicroscopy* 110, 118–125.

- LeBeau JM, Findlay SD, Wang X, Jacobson AJ, Allen LJ & Stemmer S (2009). High-angle scattering of fast electrons from crystals containing heavy elements: Simulation and experiment. *Phys Rev B* **79**, 214110. <https://doi.org/10.1103/PhysRevB.79.214110>.
- Madsen J & Susi T (2021). The abTEM code: Transmission electron microscopy from first principles. *Open Res Eur* **1**, 13015. <https://doi.org/10.12688/openreseurope.13015.1>.
- Midgley PA & Dunin-Borkowski RE (2009). Electron tomography and holography in materials science. *Nat Mater* **8**, 271–280.
- Ophus C (2017). A fast image simulation algorithm for scanning transmission electron microscopy. *Adv Struct Chem Imaging* **3**, 13. <https://doi.org/10.1186/s40679-017-0046-1>.
- Ophus C (2019). Four-dimensional scanning transmission electron microscopy (4D-STEM): From scanning nanodiffraction to ptychography and beyond. *Microsc Microanal* **25**, 563–582.
- Ophus C, Ercius P, Huijben M & Ciston J (2017). Non-spectroscopic composition measurements of SrTiO₃-La_{0.7}Sr_{0.3}MnO₃ multilayers using scanning convergent beam electron diffraction. *Appl Phys Lett* **110**, 063102. <https://doi.org/10.1063/1.4975932>.
- Pan SJ & Yang Q (2009). A survey on transfer learning. *IEEE Trans Knowl Data Eng* **22**, 1345–1359.
- Pennycook SJ & Nellist PD (2011). *Scanning Transmission Electron Microscopy: Imaging and Analysis*, 1st ed. New York: Springer.
- Pollock JA, Weyland M, Taplin DJ, Allen LJ & Findlay SD (2017). Accuracy and precision of thickness determination from position-averaged convergent beam electron diffraction patterns using a single-parameter metric. *Ultramicroscopy* **181**, 86–96.
- Shen B, Chen X, Wang H, Xiong H, Bosch EGT, Lazić I, Cai D, Qian W, Jin S, Liu X, Han Y & Wei F (2021). A single-molecule van der Waals compass. *Nature* **592**, 541–544.
- Srivastava N, Hinton G, Krizhevsky A, Sutskever I & Salakhutdinov R (2014). Dropout: A simple way to prevent neural networks from overfitting. *J Mach Learn Res* **15**, 1929–1958.
- Torroglosa E, Ortiz J & Skarmeta A (2018). Matching federation identities, the eduGAIN and STORK approach. *Future Gener Comput Syst* **80**, 126–138.
- Williams DB & Carter CB (2009) *Transmission Electron Microscopy: A Textbook for Materials Science*, 2nd ed. New York: Springer.
- Xu W & LeBeau JM (2018). A deep convolutional neural network to analyze position averaged convergent beam electron diffraction patterns. *Ultramicroscopy* **188**, 59–69.
- Yang L & Shami A (2020). On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing* **415**, 295–316. Available at <https://www.sciencedirect.com/science/article/pii/S0925231220311693>
- Zhang C, Feng J, DaCosta LR & Voyles Paul M (2020). Atomic resolution convergent beam electron diffraction analysis using convolutional neural networks. *Ultramicroscopy* **210**, 112921. <https://doi.org/10.1016/j.ultramic.2019.112921>.