JOURNAL OF SYNCHROTRON RADIATION

ISSN 1600-5775

Deep-learning-based image registration for nano-resolution tomographic reconstruction

Tianyu Fu,[a,b] Kai Zhang,[a]* Yan Wang,[a] Jizhou Li,[c]* Jin Zhang,[a,b,c] Chunxia Yao,[a,b] Qili He,[a,b] Shanfeng Wang,[a] Wanxia Huang,[a] Qingxi Yuan,[a]* Piero Pianetta[c] and Yijin Liu[c]

[a]Beijing Synchrotron Radiation Facility, X-ray Optics and Technology Laboratory, Institute of High Energy Physics, Chinese Academy of Sciences, Yuquan Road, Shijingshan District, Beijing 100043, People's Republic of China, [b]University of Chinese Academy of Sciences, Yuquan Road, Shijingshan District, Beijing 100043, People's Republic of China, and [c]Stanford Synchrotron Radiation Lightsource, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA. *Correspondence e-mail: zhangk@ihep.ac.cn, lijz@slac.stanford.edu, yuanqx@ihep.ac.cn

Received 6 May 2021
Accepted 15 August 2021

Edited by A. Momose, Tohoku University, Japan

Keywords: full-field transmission X-ray microscopy; nano-tomography; image registration; deep learning; residual neural network.

Nano-resolution full-field transmission X-ray microscopy has been successfully applied to a wide range of research fields thanks to its capability of non-destructively reconstructing the 3D structure with high resolution. Due to constraints in the practical implementations, the nano-tomography data is often associated with a random image jitter, resulting from imperfections in the hardware setup. Without a proper image registration process prior to the reconstruction, the quality of the result will be compromised. Here a deep-learning-based image jitter correction method is presented, which registers the projective images with high efficiency and accuracy, facilitating a high-quality tomographic reconstruction. This development is demonstrated and validated using synthetic and experimental datasets. The method is effective and readily applicable to a broad range of applications. Together with this paper, the source code is published and adoptions and improvements from our colleagues in this field are welcomed.

1. Introduction

X-ray microscopy is a powerful tool for research and industry (Jiang *et al.*, 2020; Lee *et al.*, 2021; Mu *et al.*, 2018). It comes in many different modalities and offers a wide range of resolving power. Compared with visible light and electron beams, X-rays feature a short wavelength and a large penetration depth. X-rays interact with materials in different ways, providing a number of different contrast mechanisms that are sensitive to different material properties. Examples of X-ray imaging with different modalities include, but are not limited to, absorption, fluorescence, valence state, spin, scattering and phase shift (Donoghue *et al.*, 2006; Cagno *et al.*, 2017; Tian *et al.*, 2018; Liu *et al.*, 2013; Gonzalez-Jimenez *et al.*, 2012). One of the most significant advancements in X-ray imaging is tomography, which reconstructs the 3D structure noninvasively. In this work, we tackle the image jittering problem, a technical challenge facing tomography at high resolution.

Specifically, we focus on the nano-resolution full-field transmission X-ray microscopy (TXM) (Andrews *et al.*, 2009), which utilizes a Fresnel zone plate as the objective lens for magnifying the X-ray image and, thus, operates at an effective spatial resolution of ~15–30 nm (Chao *et al.*, 2005; Liu *et al.*, 2011). Compositional and chemical sensitivities have been demonstrated by conducting X-ray microscopy measurements with varying X-ray energy (Liu *et al.*, 2012; Nazaretski *et al.*, 2018). Nano-resolution tomography and, in particular, nano-resolution spectro-tomography are regarded as a powerful

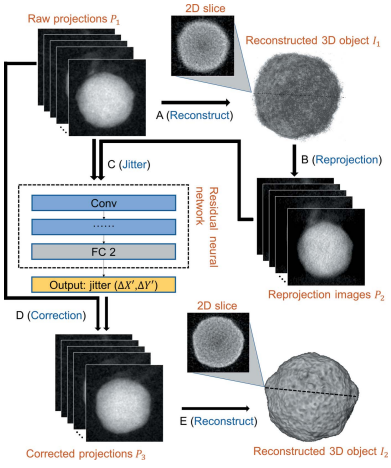

modality of TXM with successful case studies spanning energy materials, environmental science, geoscience, *etc*. With the ongoing efforts in the developments of X-ray optics (Chang & Sakdinawat, 2014) and the next-generation X-ray facilities (Aljdaimi *et al.*, 2018; Dong *et al.*, 2019; Schneider, 1998; Zaman *et al.*, 2019), further improvements in spatial, temporal resolution, and chemical sensitivity can be anticipated.

When conducting nano-resolution tomography with a TXM setup, the sample is physically rotated and imaged in different viewing angles. Although the mechanical stability and accuracy have been greatly improved in state-of-the-art rotation stages, the image misalignment, involving a systematic error component and a random error component, can still jeopardize the quality of the reconstruction. For example, for the nano-resolution full-field transmission X-ray microscope at the 4W1A beamline of BSRF (Yuan *et al.*, 2012), the effective image pixel size is about 15 nm × 15 nm. To achieve a 30 nm-resolution reconstruction, the amount of random image jitter should be smaller than 3 pixels (with the horizontal and the vertical jitter, ΔX and ΔY, combined). This is a rather stringent requirement and often cannot be satisfied in practice.

To address this technical challenge, efforts have been devoted to developing experimental and computational approaches for correction of the undesired image jitter, which is key to achieving a high-resolution and high-fidelity reconstruction. Various methods have been demonstrated to be effective for alignment of projections in X-ray and electron tomography, but each of them works to a certain degree. A more comprehensive review of recent works can be found in the literature (Odstrčil *et al.*, 2019). Here we discuss three typical approaches as examples. The first method is to use fiducial markers or recognizable features of the sample for registration (Cheng *et al.*, 2014; Fung *et al.*, 1996; Han *et al.*, 2015; Mastronarde & Held, 2017; Olins *et al.*, 1983; Ress *et al.*, 1999). For example, gold particles with a micrometre-level diameter are often used as fiducial markers and a manual or semi-auto correction process can be used to trace the marker positions in different viewing angles. This approach is rather tedious and is not always applicable because such a marker or a recognizable sample feature is not always available. The second method utilizes capacitance sensors or laser interferometers to experimentally measure the runout of the rotation stage and uses the measured values to offset the projective images (Xu *et al.*, 2014). For this experimental approach, the implementation can be expensive, and the measurement accuracy can be affected by the temperature and humidity, decreasing the robustness of the system. The third method is the projective images registration correction method (Gürsoy *et al.*, 2017; Odstrčil *et al.*, 2019; Yu *et al.*, 2018) based on the concept of 'tomographic consistency' (Gürsoy *et al.*, 2014). This method first conducts the reconstruction disregarding the sample jitter to obtain a low-resolution 3D volume, then numerically calculates the projective images in the corresponding viewing angles through a 'reprojection' process. The sample jitter is then corrected by registering the image pairs, each containing an experimental image and its corresponding calculated projective image in the same angle. This process is often carried out iteratively for better results. Our group has implemented this method for the beamline production at SSRL and BSRF. Although the performance of this method is fairly good, the efficiency is far from satisfactory. In addition to the above discussed approaches, we acknowledge that efforts have been put into dealing with more complicated scenarios that involve non-rigid sample transformations (De Andrade *et al.*, 2021; Nikitin *et al.*, 2021). This is a very interesting research direction and is beyond the scope of this paper.

Herein, we present a deep-learning-based method for correcting the undesired image jitters for nano-resolution X-ray tomography. Our method can quickly and precisely calculate the amount of image offsets and correct such image random jitters through a trained residual neural network. We demonstrate that, through our approach, the fidelity and resolution of the tomographic reconstruction can be significantly improved. Compared with the traditional methods, our approach is simple, fast, accurate, and robust. With the source code of our development freely available on our GitHub repository, our method is readily available for deployment and further improvement.

2. Method

2.1. Deep-learning-based image jitter correction

A schematic diagram of our deep-learning-based image jitter correction method for nano-resolution X-ray tomography is shown in Fig. 1. First, the raw projective images (P_1) of the sample were reconstructed using the maximum-likelihood expectation maximization algorithm (MLEM) (Gürsoy *et al.*, 2014; Dempster *et al.*, 1977), resulting in a low-resolution 3D representation of the sample (I_1, see step A). Second, the numerical reprojection process was applied to I_1 in order to obtain the reprojection images (P_2, see step B). Third, P_1 and P_2 are assembled into a two-channel picture and fed into a residual neural network (He *et al.*, 2016). The residual neural network performs feature extraction and matches the image pairs, with the amount of image offsets, $\Delta X'$ and $\Delta Y'$, quantified and recorded (see step C). Fourth, the recorded image jitter, $\Delta X'$ and $\Delta Y'$, were applied to translate P_1, resulting in the corrected projective images (P_3, see step D). Finally, P_3 is fed into the tomographic reconstruction routine and the final result (I_2) can be generated using conventional algorithms, *e.g.* filtered back-projection (FBP) (step E). Conceptually, this process is similar to the 'tomographic consistency' approach. The major development reported herein is the application of a specifically designed residual neural network for the projection image registration without going through an iterative, computationally expensive, and time-consuming process. We will elaborate on our network design in the next section.

2.2. Design of the residual neural network

While the fundamental concept of our approach is similar to 'tomographic consistency', the goal of our development

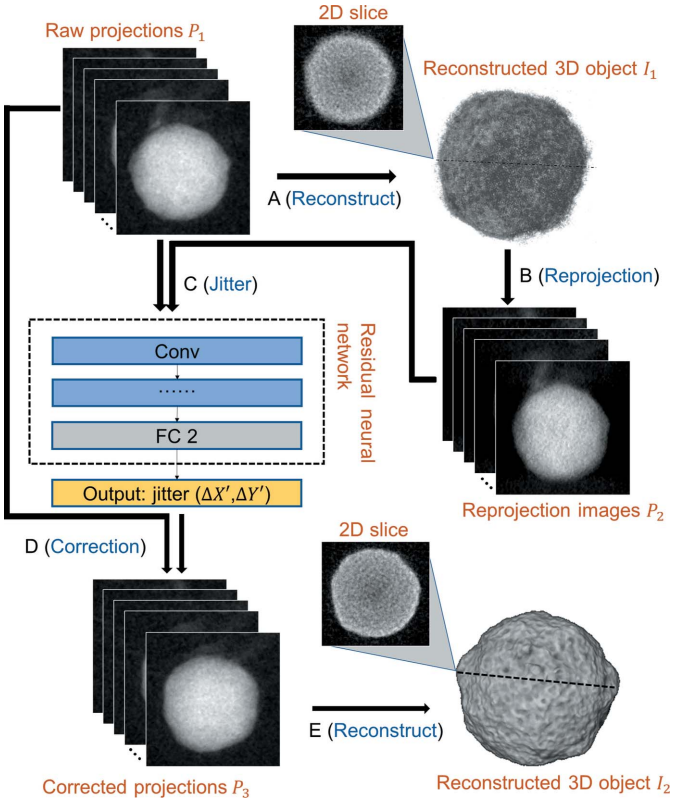

Figure 1
The workflow of the deep-learning-based image jitter correction method for synchrotron nano-resolution tomographic. Step A is the reconstruction of the raw projective images. Step B is reprojection of the sample from different perspectives. Step C acquires image jitter by the residual neural network. Step D corrects the projective images. Step E is to reconstruct the 3D image after correction.

is to accomplish the projective image registration without performing many iterations. This is a nontrivial task as the images generated by reprojection are quite different from the raw experimental images. This is the root cause for the need of multiple iterations in the conventional approach, which improves the quality of the 3D volume progressively. In this work, we demonstrate that, with a properly trained machine learning model, the projective image registration can be achieved without multiple iterations. Our deep-learning-based method does not require iterations because it is not trying to register the images by matching features one-to-one. The network first learns how the features in the raw projection image transform into the characteristics (including nonlocal artifacts) in the reprojection. Therefore, when it is properly trained, this model demonstrates significant improvements over the conventional approach.

Deep neural networks, especially convolutional neural networks, have been a very popular approach to perform feature extraction and subsequently for image registration (Simonyan & Zisserman, 2015; Ronneberger *et al.*, 2015; Balakrishnan *et al.*, 2018; Hoopes *et al.*, 2021). With the increase of network layers, more complex features can be learned progressively (Krizhevsky *et al.*, 2017; LeCun *et al.*, 1989) but the networks are more difficult to train. The residual neural network architecture is used to solve this issue (Sandler *et al.*, 2018; Balduzzi *et al.*, 2017), resulting in the establishment of the high-precision deep network.

We illustrate in Fig. 2 the structure of the residual neural network, which consists of 13 convolutional layers, two pooling layers, and a fully connected layer. In the 13 convolution layers, there is one 7×7 convolution kernel and twelve 3×3 convolution kernels, and these twelve 3×3 convolution kernels constitute three jump connected residual blocks. In the above-mentioned network layer, the layer with a strides factor of 2 can down-sample the image to increase the calculation speed. Besides, the ReLU activation function (Xu *et al.*, 2015) and layer normalization (Ioffe & Szegedy, 2015) are used to connect the convolutional layers, which help to speed up the convergence and improve the final results' accuracy in the tests. The residual neural network takes the image pairs – each has one raw projective image and its corresponding reprojection image – as input and evaluates the amount of image jitter.

A training process is required before the residual neural network can be utilized in real-world applications. Due to the lack of ground truth in the experimental data, we construct synthetic datasets with artificially induced image offsets for the training purpose. In our work, the training procedure involves the forward and backward propagations as illustrated in Fig. 2. For the forward propagation (black arrow in Fig. 2), the jitter value and loss value are output by using the image pair with known image jitter as the input of forwarding propagation. For the backward propagation (orange arrow in Fig. 2), the loss value acquired by forward propagation is used as input, and the parameter of the network is calculated and updated by using the Adam optimization algorithm (Kingma & Ba, 2015) for minimizing the loss function. The backward propagation updates the network and suppresses the loss

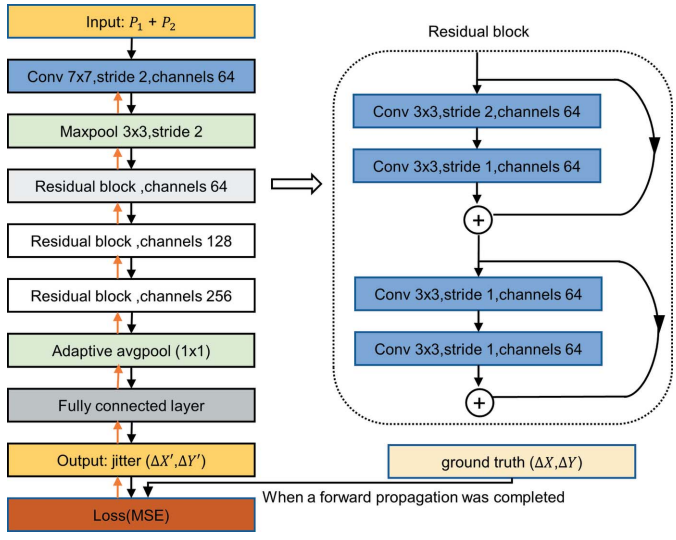

Figure 2
The structure of the residual neural network. The black arrow process represents the forward propagation of the network, and the orange arrow process represents the backward propagation of the network.

value, which ultimately improves the quality of the network for our application.

In the above-described training process for the residual neural network, a carefully defined loss function is essential. For the forward propagation, the loss function determines the functionality of the model. And for the backward propagation, the loss function guides optimization of the network. In this paper, the mean square error (MSE) of the known image offsets and the corresponding values calculated by the residual neural network is chosen as the loss function, which can be expressed as

$$\mathrm{MSE} = \frac{1}{n}\sum_{n}(\Delta X_n - \Delta X_n')^2 + (\Delta Y_n - \Delta Y_n')^2, \qquad (1)$$

where ΔX_n, ΔY_n, $\Delta X_n'$, $\Delta Y_n'$ represent the horizontal and vertical image jitters in the nth projective viewing angle, respectively.

When the MSE value reaches the minimum, the residual neural network training will be stopped. The trained residual neural network is ready for being applied to correcting the image jitter in real-world tomographic data and only the forward propagation of the residual neural network will be processed.

3. Experiments and discussion

3.1. Evaluation of the deep-learning-based image registration using data with synthetic jitters

To verify the effectiveness of our deep-learning-based image jitter correction approach, we first conduct nano-tomography on an isolated gold particle as our test sample. This experiment was carried out at the 4W1A beamline of BSRF. A total of 181 projective images were recorded over an angular range of 0° to 180°. In our experimental configuration, the effective pixel size is 15 nm. As described above, a ground truth is needed for the training process. Therefore, we conducted an iterative image registration process (Yu *et al.*, 2018) to get this dataset aligned and used this result as our ground truth. We then purposely induced image offsets by a random amount (from −20 pixels to +20 pixels) in both horizontal and vertical directions to mimic the random image jittering. With known offset values, this dataset was fed into our model for the training process.

We show in Fig. 3 a comparison of the results of the gold particles. Figs. 3(*a*)–3(*d*) are the ground-truth slice, the slice with synthetic random jitters, the slice with manual alignment, and the slice with our model-based registration, respectively. Although the manual alignment process can effectively supress the blurred particle edge, the lack of recognizable features in the sample makes it very difficult to precisely register the projection images in different angles, resulting in visible artifacts near the edge of the particle. On the other hand, our deep-learning model successfully reconstructed the tomographic data and the presented central slice faithfully restored the features of the ground-truth image. To present the effectiveness of the image alignment from a different

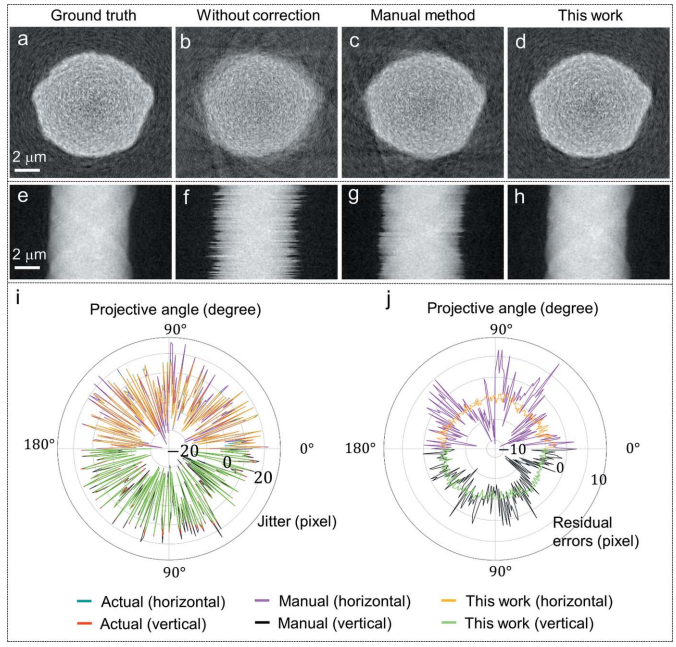

Figure 3
The reconstructed slice images and sinogram images of the gold particle sample. Reconstruction slice images results of projective images data with ground truth (*a*), without correction (*b*), with manual correction method (*c*), and with deep-learning-based correction method (*d*). (*e*)–(*h*) Sinogram images corresponding to the slice images (*a*)–(*b*). (*i*) Radar chart of images jitter distribution. (*j*) Radar chart of sample residual errors distribution.

perspective, we show the sinograms in Figs. 3(*e*) to 3(*h*). The image jittering is clearly visible on the edges of the sinogram, which is clearly supressed after our model is applied for correction. We quantify the residual error by calculating the mean square error [MSE, as shown in equation (1)] between the model-estimated amount of offset and the ground truth. This calculation yields a value at 0.87 pixels, indicating a rather successful correction using our model. The absolute amounts of image offset and the residual errors in different angles are plotted in Figs. 3(*i*) and 3(*j*). Although both the manual correction and the deep-learning correction could follow the trend of the artificially induced random jitter, the residual errors are significantly suppressed in the deep-learning-corrected dataset to a single-pixel level.

We point out that the synthetic data are originated from the experimentally measured images, which already contain experimental noise associated with the actual synchrotron measurements. Nevertheless, to further evaluate the robustness of our method, we conduct a testing with purposely induced salt-and-pepper noise (Zhang *et al.*, 2021) at different levels [ranging from 0% to 60%, shown in Figs. 4(*a*) to 4(*d*)]. The salt-and-pepper noise can be represented as randomly occurring white and black pixels in the image. The presence of this noise not only makes the structural details contained in the images submerged, but also introduces a lot of interference information, which will harm the reconstruction quality of the conventional registration methods due to their limited adaptability to noise (Yüksel & Başţürk, 2003; Singh *et al.*, 2018). We plot the MSE values of the projection alignment

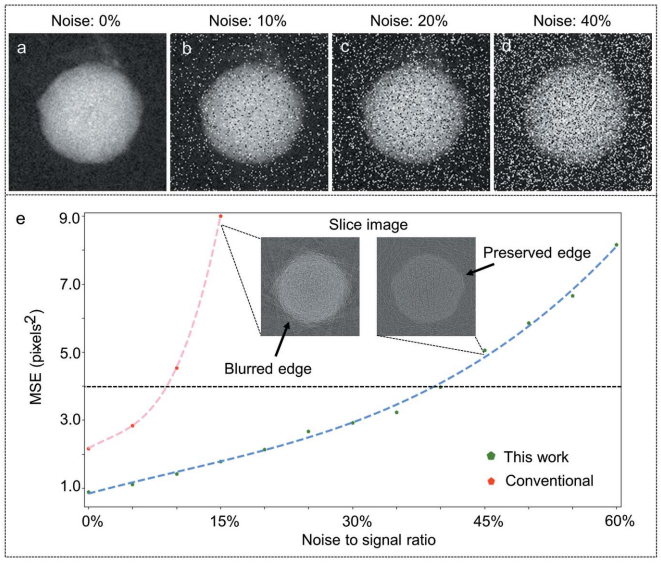

Figure 4
(*a*)–(*d*) The original projection image of the gold particle and its noise corrupted images at various levels of salt-and-pepper noise: (*a*) noise ratio at 0%; (*b*) noise ratio at 10%; (*c*) noise ratio at 20%; (*d*) noise ratio at 40%. (*e*) The MSE values of the conventional and the proposed methods (red and blue, respectively) are plotted with respect to the noise ratio. The reconstructed image slices are shown in the inset.

result of the conventional phase correlation algorithm and that of the deep-learning-based method against the levels of noise in Fig. 4(*e*). In general, the features of the gold particle in the projection images are overwhelmed and, therefore, the registration accuracy is gradually reduced as the salt-and-pepper noise level increases. When the noise level reaches above 40%, the conventional method failed while the herein developed model can still lead to an MSE value below 4, which means that the residual jitter errors are smaller than 2 pixels. The corresponding reconstructed slices [insets in Fig. 4(*e*)] demonstrate a reasonable quality. This shows that the proposed method can robustly correct the undesired image jitters for nano-resolution X-ray tomography, even with high noise conditions.

3.2. Application of the deep-learning-model to realistic nano-tomographic data on a battery cathode particle

Nano-tomographic data on battery cathode particles ($LiNi_{0.6}Mn_{0.2}Co_{0.2}O_2$) with complicated internal structure are used to further test the developed deep-learning-based image jitter correction approach. In a typical dataset, a total of 180 projective images are recorded over an angular range of 1° to 180° (Xia *et al.*, 2018). Similarly, the training data were needed for training the network. Generally, it is well appreciated that the performance of supervised deep neural network relies heavily on the size and properties of the training data. However, in many cases, it is difficult to create training datasets that are large enough. To solve this problem, we train the network with different types of the synthetic datasets with artificially induced jitters and evaluate their performance on the test set. The results show that, if the training data and the test data have a high shape similarity, the well trained network can also be obtained even with a fairly small training set. On the contrary, the proposed method may work poorly for other types of data and a retraining process will be necessary. In this paper, as a demonstration example, we have trained our model on external datasets (Andrews *et al.*, 2008; and data from BSRF), which has similar object shape with the battery cathode particle. The results demonstrate its effective jitter correction performance. Importantly, our framework is very flexible, and the open-sourced simulation and training codes allow the proposed approach to be quickly adapted to new datasets with different properties. In addition, the proposed method can also serve as the initial step within the popular re-projection framework (Yu *et al.*, 2018) to significantly accelerate the convergence of the iterative process.

We display the reconstructed 3D volume and selected slices of the imaged battery particle in Fig. 5. Figures 5(*a*) and 5(*b*) show the results with and without the image registration using our model. The presented deep-learning-based correction significantly enhances the image quality and contrast, which is clear to a visual assessment of the images shown in Figs. 5(*a*) and 5(*b*). The grayscale histogram plot in Fig. 5(*c*) also shows a better separation of the two peaks that are associated with the background/crack/void phase and the solid phase, respectively.

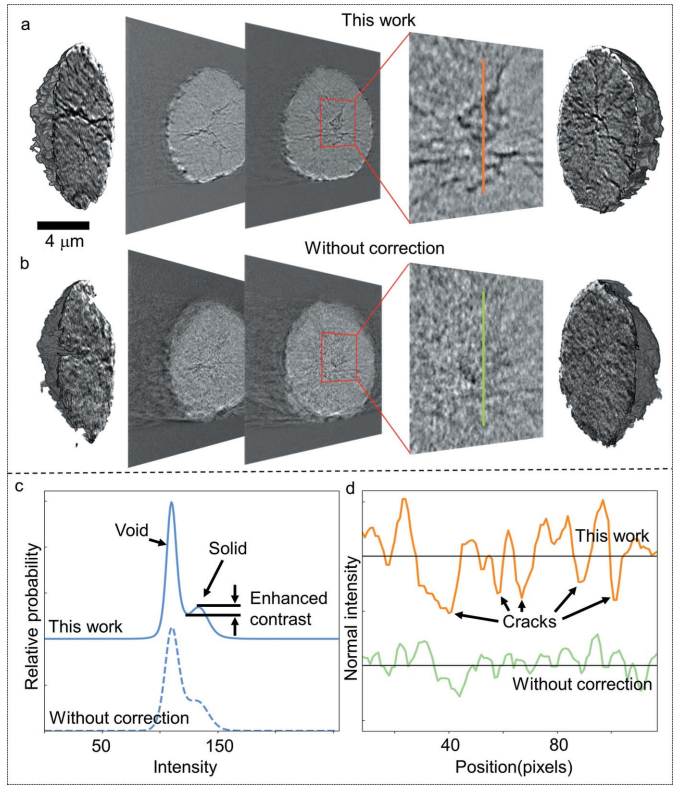

Figure 5
The reconstruction 3D volume and slice images result of projective images with deep-learning-based correction method (*a*), and without correction (*b*). The right-had images show enlargement views of the interested regions marked as red squares in the respective slice image. (*c*) Intensity histogram plot of the tomographic reconstruction. (*d*) Intensity over the lines highlighted (crack) in (*a*)–(*b*).

In Figs. 5(*d*) and 5(*e*), we compare the intensity profiles over the highlighted lines in Figs. 5(*a*) and 5(*b*), respectively. The enhanced contrast is supported by the clearly visible intensity drops at the crack positions. Accurate visualization and quantification of the spatial arrangement of the cracks and pores is important to understanding the chemomechanical processes in battery cathodes (Yang *et al.*, 2019), which largely determines the battery performance and lifespan. The herein presented development is clearly valuable to researching battery materials using TXM.

3.3. Runtime of the proposed method

The herein reported deep-learning-based image jitter correction method is developed using *PyTorch* (Paszke *et al.*, 2019), which is a deep-learning framework for GPU-accelerated tensor calculation. Our package was operated on NVIDIA Quadro P6000, a 25 GB memory graphics card. The pre-trained model can quickly correct a set of nano-resolution projective data composed of 181 sequential images in less than one second.

4. Conclusion

In this paper, a deep-learning-based image jitter correction method for synchrotron nano-resolution tomographic reconstruction is presented. Compared with the traditional correction methods, our approach is efficient, effective, and does not require any expensive experimental developments. After correction, the residual jitter errors are at a level of 0–2 pixels, which greatly improves the quality of the 3D reconstruction. In addition, our method demonstrates a significantly improved computing speed. The robustness of our approach is also reflected in the fact that our model can be easily retrained for applications for different samples and for different beamlines. We publish our source code and welcome adoptions and improvements from our colleagues in this field.

5. Data availability

The data supporting the findings of the study is available at the GitHub repository: https://github.com/SSRL-LiuGroup/corresnet.

Acknowledgements

We acknowledge the 4W1A beamline of the Beijing Synchrotron Radiation Source and the 6-2c beamline of the Stanford Synchrotron Radiation Light Source for the experimental data and facilities provided. The engineering support from D. Van Campen, D. Day and V. Borzenets for the TXM experiment at beamline 6-2c of Stanford Synchrotron Radiation Lightsource is gratefully acknowledged. Use of Stanford Synchrotron Radiation Lightsource (SSRL) at SLAC National Accelerator Laboratory is supported by the US Department of Energy (DOE), Office of Science, Office of Basic Energy Sciences under contract number DE-AC02-76SF00515.

Funding information

The following funding is acknowledged: National Key Research and Development Program of China (award No. 2016YFA0400900); National Natural Science Foundation of China (award No. U2032107; award No. U1632110).

References

Aljdaimi, A., Devlin, H., Dickinson, M., Burnett, T. & Slater, T. J. A. (2018). *Clin. Oral Invest.* **23**, 2279–2285.
Andrews, J. C., Brennan, S., Liu, Y., Pianetta, P., Almeida, E. A. C., Meulen, M. C. H. V. D., Wu, Z., Mester, Z., Ouerdane, L., Gelb, J., Feser, M., Rudati, J., Tkachuk, A. & Yun, W. (2009). *J. Phys. Conf. Ser.* **186**, 012081.
Andrews, J. C., Brennan, S., Patty, C., Luening, K., Pianetta, P., Almeida, E., van der Meule, M. C. H., Feser, M., Gelb, J. & Rudati, J. (2008). *Synchrotron Radiat. News*, **21**(3), 17–26.
Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J. & Dalca, A. V. (2018). *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18–23 June 2018, Salt Lake City, UT, USA.
Balduzzi, D., Frean, M., Leary, L., Lewis, J., Wan-Duo Ma, K. & McWilliams, B. (2017). *arXiv*:1702.08591.
Cagno, S., Brede, D. A., Nuyts, G., Vanmeert, F., Pacureanu, A., Tucoulou, R., Cloetens, P., Falkenberg, G., Janssens, K., Salbu, B. & Lind, O. C. (2017). *Anal. Chem.* **89**, 11435–11442.
Chang, C. & Sakdinawat, A. (2014). *Nat. Commun.* **5**, 4243.
Chao, W. L., Harteneck, B. D., Liddle, J. A., Anderson, E. H. & Attwood, D. T. (2005). *Nature*, **435**, 1210–1213.
Cheng, C., Chien, C., Chen, H., Hwu, Y. & Ching, Y. (2014). *PLoS One*, **9**, e84675.
De Andrade, V., Nikitin, V., Wojcik, M., Deriy, A., Bean, S., Shu, D., Mooney, T., Peterson, K., Kc, P., Li, K., Ali, S., Fezzaa, K., Gürsoy, D., Arico, C., Ouendi, S., Troadec, D., Simon, P., De Carlo, F. & Lethien, C. (2021). *Adv. Mater.* **33**, 2008653.
Dempster, A. P., Laird, N. M. & Rubin, D. B. (1977). *J. R. Stat. Soc.* **39**, 1–38.
Dong, K., Osenberg, M., Sun, F., Markötter, H., Jafta, C. J., Hilger, A., Arlt, T., Banhart, J. & Manke, I. (2019). *Nano Energy*, **62**, 11–19.
Donoghue, P. C. J., Bengtson, S., Dong, X. P., Gostling, N. J., Huldtgren, T., Cunningham, J. A., Yin, C., Yue, Z., Peng, F. & Stampanoni, M. (2006). *Nature*, **442**, 680–683.
Fung, J. C., Liu, W. P., de Ruijter, W. J., Chen, H., Abbey, C. K., Sedat, J. W. & Agard, D. A. (1996). *J. Struct. Biol.* **116**, 181–189.
Gonzalez-Jimenez, I. D., Cats, K., Davidian, T., Ruitenbeek, M., Meirer, F., Liu, Y., Nelson, J., Andrews, J. C., Pianetta, P., de Groot, F. M. F. & Weckhuysen, B. M. (2012). *Angew. Chem. Int. Ed.* **51**, 11986–11990.
Gürsoy, D., De Carlo, F., Xiao, X. & Jacobsen, C. (2014). *J. Synchrotron Rad.* **21**, 1188–1193.
Gürsoy, D., Hong, Y. P., He, K., Hujsak, K., Yoo, S., Chen, S., Li, Y., Ge, M. Y., Miller, L. M., Chu, Y. S., De Andrade, V., He, K., Cossairt, O., Katsaggelos, A. K. & Jacobsen, C. (2017). *Sci. Rep.* **7**, 11818.
Han, R. M., Wang, L. S., Liu, Z. Y., Sun, F. & Zhang, F. (2015). *J. Struct. Biol.* **192**, 403–417.
He, K., Zhang, X., Ren, S. & Sun, J. (2016). *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 27–30 June 2016, Las Vegas, NV, USA.
Hoopes, A., Hoffmann, M., Fischl, B., Guttag, J. & Dalca, A. V. (2021). *arXiv*:2101.01035.
Ioffe, S. & Szegedy, C. (2015). *arXiv*:1502.03167.
Jiang, Z., Li, J., Yang, Y., Mu, L., Wei, C., Yu, X., Pianetta, P., Zhao, K., Cloetens, P., Lin, F. & Liu, Y. (2020). *Nat. Commun.* **11**, 2310.
Kingma, D. P. & Ba, J. L. (2015). *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*, 7–9 May 2015, San Diego, CA, USA.

Krizhevsky, A., Sutskever, I. & Hinton, G. E. (2017). *Commun. ACM*, **60**, 84–90.
LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. & Jackel, L. D. (1989). *Neural Comput.* **1**, 541–551.
Lee, H. R., Liao, L., Xiao, W., Vailionis, A., Ricco, A. J., White, R., Nishi, Y., Chiu, W., Chu, S. & Cui, Y. (2021). *Nano Lett.* **21**, 651–657.
Liu, Y., Andrews, J. C., Meirer, F., Mehta, A., Gil, S. C., Sciau, P., Mester, Z. & Pianetta, P. (2011). *AIP Conf. Proc.* **1365**, 357–360.
Liu, Y., Meirer, F., Wang, J., Requena, G., Williams, P., Nelson, J., Mehta, A., Andrews, J. C. & Pianetta, P. (2012). *Anal. Bioanal. Chem.* **404**, 1297–1301.
Liu, Y., Nelson, J., Holzner, C., Andrews, J. C. & Pianetta, P. (2013). *J. Phys. D Appl. Phys.* **46**, 4001.
Mastronarde, D. N. & Held, S. R. (2017). *J. Struct. Biol.* **197**, 102–113.
Mu, L. Q., Lin, R. L., Xu, R., Han, L. L., Xia, S. H., Sokaras, D., Steiner, J. D., Weng, T. C., Nordlund, D., Doeff, M. M., Liu, Y. J., Zhao, K. J., Xin, H. L. L. & Lin, F. (2018). *Nano Lett.* **18**, 3241–3249.
Nazaretski, E., Xu, W., Yan, H., Huang, X., Coburn, D. S., Ge, M., Lee, W. K., Gao, Y., Xu, W. & Fuchs, M. R. (2018). *Synchrotron Radiat. News*, **31**(5), 3–8.
Nikitin, V., De Andrade, V., Slyamov, A., Gould, B. J., Zhang, Y., Sampathkumar, V., Kasthuri, N., Gursoy, D. & De Carlo, F. (2021). *IEEE Trans. Comput. Imaging*, **7**, 272–287.
Odstrčil, M., Holler, M., Raabe, J. & Guizar-Sicairos, M. (2019). *Opt. Express*, **27**, 36637–36652.
Olins, D. E., Olins, A. L., Levy, H. A., Durfee, R. C., Margle, S. M., Tinnel, E. P. & Dover, S. D. (1983). *Science*, **220**, 498–500.
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., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. & Chintala, S. (2019). *NeurIPS*. *arXiv*:1912.01703.
Ress, D., Harlow, M. L., Schwarz, M., Marshall, R. M. & McMahan, U. J. (1999). *J. Electron Microsc.* **48**, 277–287.
Ronneberger, O., Fischer, P. & Brox, T. (2015). In *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)* Vol. 9351 of *Lecture Notes in Computer Science*. Cham: Springer.
Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. & Chen, L.-C. (2018). *IEEE Conference on Computer Vision and Pattern Recognition, 2018*, pp. 4510–4520. IEEE.
Schneider, G. (1998). *Ultramicroscopy*, **75**, 85–104.
Simonyan, K. & Zisserman, A. (2015). *ICLR. arXiv*:1409.1556.
Singh, V., Dev, R., Dhar, N. K., Agrawal, P. & Verma, N. K. (2018). *IEEE Trans. Fuzzy Syst.* **26**, 3170–3176.
Tian, C., Xu, Y., Nordlund, D., Lin, F., Liu, J., Sun, Z., Liu, Y. & Doeff, M. (2018). *Joule*, **2**, 464–477.
Xia, S., Mu, L., Xu, Z., Wang, J., Wei, C., Liu, L., Pianetta, P., Zhao, K., Yu, X. & Lin, F. (2018). *Nano Energy*, **53**, 753–762.
Xu, B., Wang, N., Chen, T. & Li, M. (2015). *arXiv*:1505.00853.
Xu, W., Lauer, K., Chu, Y. & Nazaretski, E. (2014). *J. Synchrotron Rad.* **21**, 1367–1369.
Yang, Y., Xu, R., Zhang, K., Lee, S. J., Mu, L. Q., Liu, P. F., Waters, C. K., Spence, S., Xu, Z. R., Wei, C. X., Kautz, D. J., Yuan, Q. X., Dong, Y. H., Yu, Y. S., Xiao, X. H., Lee, H. K., Pianetta, P., Cloetens, P., Lee, J. S., Zhao, K. J., Lin, F. & Liu, Y. J. (2019). *Adv. Energy Mater.* **9**, 1900674.
Yu, H., Xia, S., Wei, C., Mao, Y., Larsson, D., Xiao, X., Pianetta, P., Yu, Y.-S. & Liu, Y. (2018). *J. Synchrotron Rad.* **25**, 1819–1826.
Yuan, Q., Zhang, K., Hong, Y., Huang, W., Gao, K., Wang, Z., Zhu, P., Gelb, J., Tkachuk, A., Hornberger, B., Feser, M., Yun, W. & Wu, Z. (2012). *J. Synchrotron Rad.* **19**, 1021–1028.
Yüksel, M. E. & Başẗurk, A. (2003). *Int. J. Electron. Commun.* **57**, 214–219.
Zaman, W., Hortance, N., Dixit, M. B., De Andrade, V. & Hatzell, K. B. (2019). *J. Mater. Chem. A*, **7**, 23914–23921.
Zhang, J., Hu, J., Jiang, Z., Zhang, K., Liu, P., Wang, C., Yuan, Q., Pianetta, P. & Liu, Y. (2021). *J. Synchrotron Rad.* **28**, 278–282.