












Weak signal extraction enabled by deep neural network denoising of diffraction data

Received: 26 August 2022

Accepted: 8 January 2024

Published online: 13 February 2024

 Check for updates

Jens Oppliger ¹✉, M. Michael Denner ¹, Julia Küspert ¹, Ruggero Frison ¹, Qisi Wang ^{1,2}, Alexander Morawietz¹, Oleh Ivashko ³, Ann-Christin Dippel ³, Martin von Zimmermann ³, Izabela Biato ^{1,4}, Leonardo Martinelli¹, Benoît Fauqué⁵, Jaewon Choi ⁶, Mirian Garcia-Fernandez⁶, Ke-Jin Zhou⁶, Niels Bech Christensen⁷, Tohru Kurosawa⁸, Naoki Momono^{8,9}, Migaku Oda⁸, Fabian D. Natterer¹, Mark H. Fischer ¹, Titus Neupert¹ & Johan Chang¹✉

The removal or cancellation of noise has wide-spread applications in imaging and acoustics. In applications in everyday life, such as image restoration, denoising may even include generative aspects, which are unfaithful to the ground truth. For scientific use, however, denoising must reproduce the ground truth accurately. Denoising scientific data is further challenged by unknown noise profiles. In fact, such data will often include noise from multiple distinct sources, which substantially reduces the applicability of simulation-based approaches. Here we show how scientific data can be denoised by using a deep convolutional neural network such that weak signals appear with quantitative accuracy. In particular, we study X-ray diffraction and resonant X-ray scattering data recorded on crystalline materials. We demonstrate that weak signals stemming from charge ordering, insignificant in the noisy data, become visible and accurate in the denoised data. This success is enabled by supervised training of a deep neural network with pairs of measured low- and high-noise data. We additionally show that using artificial noise does not yield such quantitatively accurate results. Our approach thus illustrates a practical strategy for noise filtering that can be applied to challenging acquisition problems.

In recent years, remarkable progress has been made in the field of image restoration through the application of deep learning techniques^{1–6}. A central task in image restoration is removing noise from an image^{7–11}, where pixel j is composed of the intrinsic signal s_j and noise n_j , $x_j = s_j + n_j$. A typical benchmark problem has correlated signal between neighbouring pixels whereas the noise is uncorrelated and white. Such denoising problems have been the subject of both supervised^{2,3} and

unsupervised^{4,12,13} machine-learning approaches. Supervised algorithms rely on either ground-truth (x_j, s_j) or noise-2-noise training pairs (x_j, x'_j) . In the latter case, the image pairs have different (or equal) noise levels. In both cases, deep convolutional neural networks (CNNs) have been successfully applied to images with Gaussian noise^{1–5}. Unsupervised approaches, sometimes dubbed noise-2-self, noise-2-void¹² or noise-as-clean¹⁴, have also been employed. Their realization relies on

¹Physik-Institut, Universität Zürich, Zurich, Switzerland. ²Department of Physics, The Chinese University of Hong Kong, Hong Kong, China. ³Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany. ⁴Faculty of Physics and Applied Computer Science, AGH University of Krakow, Krakow, Poland. ⁵JEIP, USR 3573 CNRS, Collège de France, PSL University, Paris, France. ⁶Diamond Light Source, Didcot, UK. ⁷Department of Physics, Technical University of Denmark, Kongens Lyngby, Denmark. ⁸Department of Physics, Hokkaido University, Sapporo, Japan. ⁹Department of Applied Sciences, Muroran Institute of Technology, Muroran, Japan. ✉e-mail: jens.oppliger@physik.uzh.ch; johan.chang@physik.uzh.ch

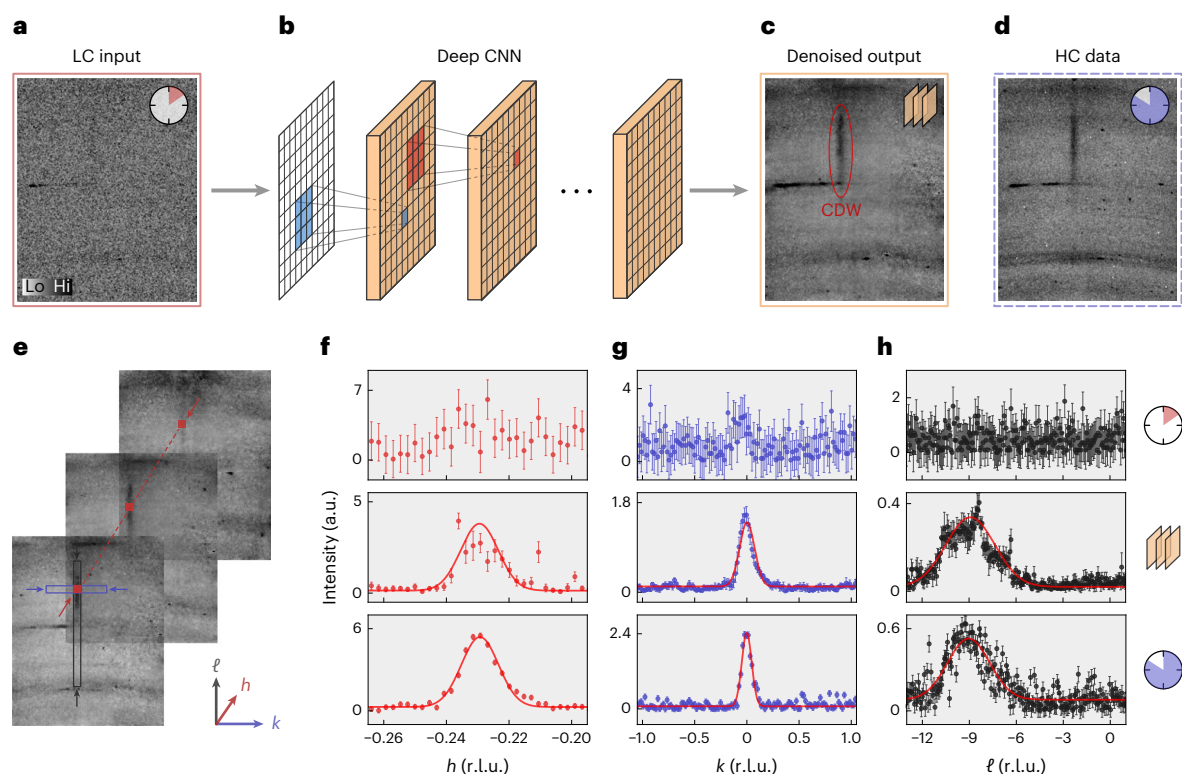


Fig. 1 | An example of denoising X-ray diffraction data using a deep CNN.

a, b, A real experimental LC frame (exposure time 1 s) (**a**) is used as an input to a deep CNN (**b**) trained to remove the noise. **c**, The denoised output reveals a CDW signal (red), barely visible in the raw LC data. **d**, The real experimental HC frame (exposure time 20 s) for comparison. **e**, A stack of denoised X-ray intensity frames as in **c**. Arrows indicate the projected reciprocal coordinates $Q = (h, k, \ell)$. **f–h**, One-dimensional projected scans through $Q \approx (0.23, 0, 8.5)$ along the h (**f**), k (**g**) and ℓ (**h**) reciprocal space axes, in units of r.l.u. For each projected scan, a

background subtraction has been performed (see main text). Gaussian fits for HC and denoised output profiles are indicated by solid red lines. The data points depicted in the denoised output profile are computed as the mean value over five training runs of the IRUNet neural network with different initial conditions. Error bars for LC and HC are shown under the assumption of counting statistics. Error bars for the denoised output are shown as the s.d. over the mentioned training runs. The clock symbols indicate relative counting time, and the network symbol indicates the denoised LC produced by the neural network.

less training information as a ground truth is absent. Unsupervised approaches therefore deliver (slightly) inferior performance compared with supervised algorithms.

Many scientific disciplines utilize digital data recording. One-, two- or three-dimensional data structures can always be transformed into a pixel-based picture format. Two-dimensional detectors are common across experimental fields such as astronomy, materials science and medical imaging. Counting of events with time-independent probability is expected to follow Poisson statistics. As such, virus–cell infection, radioactivity and particle scattering are events following a Poisson distribution. That is, the signal s_j and the noise n_j are no longer independent as the s.d. is given by $\sigma_j = \sqrt{s_j}$. Poisson noise can generally be reduced by using a sufficient acquisition time. However, long exposure times are not always possible. For example, for radiation of molecules, proteins, or human tissue, low exposure times are required to avoid beam damage¹⁵. Diffraction experiments in pulsed magnetic fields, by construction, have limited counting times and hence suffer from low-count (LC) statistics^{16,17}. Finally, many experiments explore multi-dimensional parameter spaces that are almost impossible to cover completely with sufficient statistics. Thus, there is clear potential in developing robust methods to denoise LC statistics data to produce results of comparable quality to what would be obtained from high-count (HC) statistics data. By extension, noise filtering can speed up exploratory approaches by orders of magnitudes.

However, the removal of noise from experimental data is challenging. This can be attributed to the fact that experimental noise is the sum of multiple noise sources, such as Poisson and read-out noise, with

their respective statistical properties. It is therefore difficult or often impossible to simulate experimental noise accurately. The common approach to analyse artificially added noise n_j to a ground-truth signal s_j is not directly applicable. Limited experimental training data however prevents further progress on this important problem^{18,19}.

In this Article, we present experimental training data recorded by X-ray diffraction, $(x_j, x'_j) = (x_j^{\text{LC}}, x_j^{\text{HC}})$, where LC and HC refer to low- and high-count statistics, respectively. Two deep CNNs were trained on such pairs to remove noise from the LC data. In a further step, the performance of the neural networks trained on experimental data was compared with the same networks trained on artificial training pairs where the HC data were corrupted with synthetic Poisson noise. We found that noise filtering of experimental noise—the ultimately relevant task—is significantly improved by training the neural networks on experimental data. This fact is particularly evident when analysing physical length scales associated with weak signals. As such, we provide a noise filtering approach for scientific data with challenging signal-to-noise features.

Results

X-ray diffraction data

An example of X-ray diffraction intensities recorded on the high-temperature superconductor $\text{La}_{1.88}\text{Sr}_{0.12}\text{CuO}_4$ is shown in Figs. **1a, d** and **2b, c** with LC and HC frames, respectively. The experimental set-up is schematically depicted in Fig. **2a** and further described in Methods. Although the data cover volumes of reciprocal space, the training is carried out on two-dimensional slices (so-called frames).

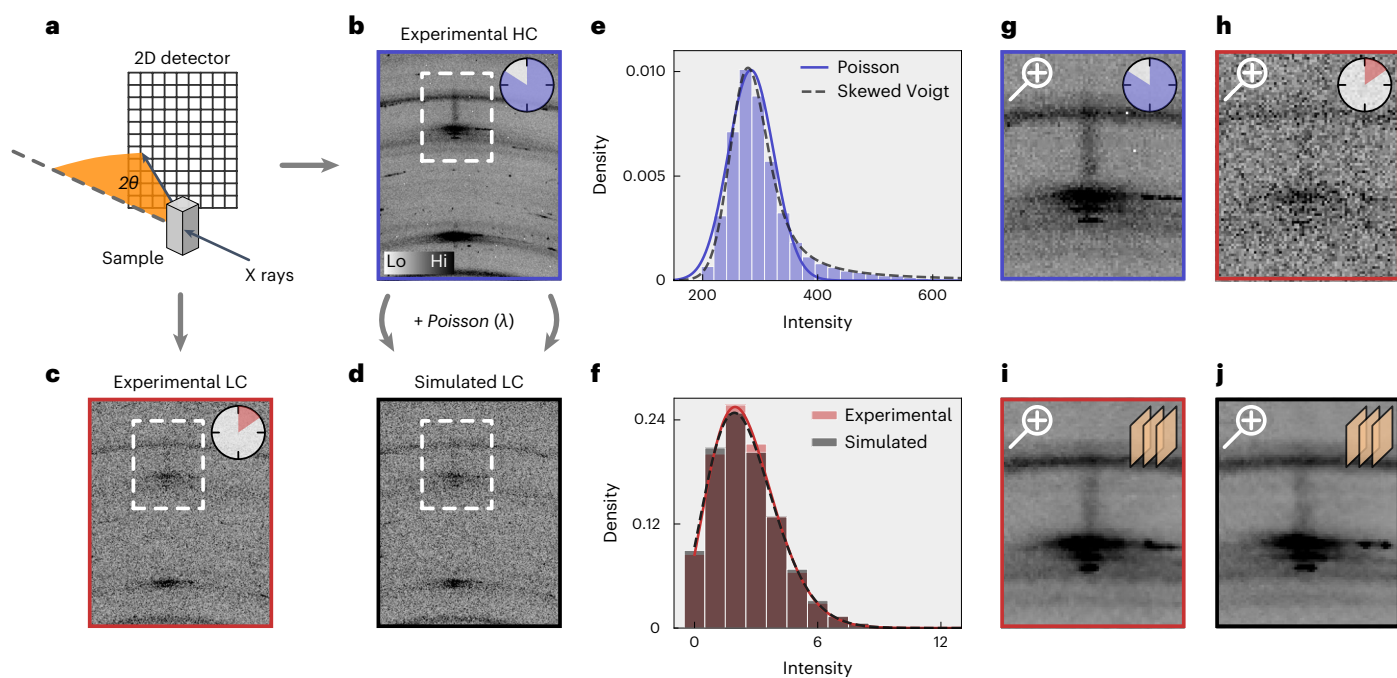


Fig. 2 | A comparison of experimental and simulated noise statistics.

a, A schematic of the experimental X-ray diffraction set-up with scattering angle 2θ . **b, c**, Long exposure time leads to an HC frame (**b**), while short exposure time leads to an LC frame (**c**). **d**, A simulated LC frame obtained when adding Poisson noise to the experimental HC frame in **b**. **e**, The intensity distribution of the HC frame in

b with fitted Poisson and skewed Voigt profiles. **f**, The intensity distribution of the experimental and simulated LC frame in **c** and **d** with a fitted Poisson profile.

g, h, A zoom of the white dashed rectangular region in **b** (**g**) and **c** (**h**) encircling the CDW reflection. **i, j**, Zooms of the white dashed rectangular region in **c** (**i**) and **d** (**j**) after denoising using the IRUNet network trained on the respective noise distributions.

Therefore, the neural networks do not have access to the three-dimensionality of the data set but rely on the two-dimensional correlation of pixels in individual frames. The LC (HC) data are recorded typically for 1 s (20 s). Such experimental pairs were recorded successively with all experimental parameters fixed. The entire data set contains 7,134 frame pairs (194×242 pixels each) and includes signals with intensities varying over six orders of magnitude. Weak two-dimensional charge density wave (CDW) order^{20–22} manifests by vertical rod-like shapes. In cuprates, the exact nature of CDW ordering is still being debated. On an atomistic level, the CDW in $\text{La}_{1.88}\text{Sr}_{0.12}\text{CuO}_4$ represents monoclinic distortions of the fundamental orthorhombic crystal structure²³. Fundamental Bragg peaks (not shown) are more intense and distributed circularly over much fewer pixels. The data also contain Debye–Scherrer (powder) rings originating from the polycrystalline sample environment. Finally, the data include spurions (unidentified signal) and dead pixels. Bragg scattering implies a direct connection between scattering angle (that is, position on the detector) and incident photon energy/wavelength. As such, conclusions drawn here are invariant under different scattering angles defined by incident photon energy or sample (lattice parameter). Furthermore, a different doping concentration in our $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$ show case would change the charge order incommensurability²⁴ but not the overall data content.

The data set is separated into a training, validation and test set. All frames containing obvious CDW signals (our main feature of interest) are excluded from the training and validation set. These frames are instead moved to the test set, which is used for performance evaluation. Overall, the size of the training set is 3,280 pairs while the size of the validation set is 820 pairs.

Artificial noise generation

As shown in Fig. 2f, the experimental LC data follow an approximately Poisson distribution. Therefore, to complement the experimental LC data, we artificially create LC data by adding Poisson noise to the experimental HC data. X-ray diffraction data are governed by counting

statistics, where the probability of a single photon hitting pixel j is theoretically given by the Poisson probability distribution for large total count N (ideally $N \rightarrow \infty$). For a fair comparison, artificial and experimental LC data should be statistically similar. To achieve this, we define λ_j as the ratio between the frame-integrated LC M_f^{LC} and HC M_f^{HC} (where f is a frame index) and $\lambda = \text{median}(\lambda_j)$. Each HC frame is then normalized with λ and LC frames being generated by adding the associated Poisson noise, resulting in simulated LC frames (Fig. 2d). Notice that signal intensities may vary by many orders of magnitude across the detector pixels, therefore the HC data typically display an asymmetric probability distribution (Fig. 2e).

Deep neural network architectures and training

We implement two neural network architectures, referred to as VDSR²⁵ and IRUNet²⁶ (see Fig. 1a–c for a schematic illustration). The networks learn the intrinsic features of the LC input frames and produce a denoised output using the HC frames as reference. The VDSR architecture relies on stacking many convolutional layers and uses a residual learning approach to extract the noise-free data from its noisy variant^{2,27}. The IRUNet architecture combines convolutional layers with an encoder/decoder framework, utilizing skip connections to reduce the vanishing gradient problem and increase accuracy. An Adam optimizer²⁸ with the AMSGrad variant²⁹ is used to improve convergence. All frames are normalized by their total intensity, ensuring equal scaling between LC and HC frames. During training, we apply data augmentation in the form of mirroring the frames along the ℓ and k direction and randomly adjusting the global brightness of the frames. Additional information can be found in Methods.

Analysis

The performance evaluation of the trained neural networks (on test data) is illustrated by one-dimensional line cuts (along the reciprocal h , k and ℓ directions) through the CDW ordering vector (Fig. 1e–h). This involves the summation of pixel intensities within a region of interest




Fig. 3 | The enhancement of the SRBR when using CNN denoising via the IRUNet network trained on experimental data. Multiple frames containing CDW signals analysed along the reciprocal directions h (left), k (centre) and l (right) in a similar fashion to Fig. 1f–h. The SRBR of the CDW reflection in the

HC frame is plotted against the SRBR of the LC frame and its denoised version. We observe that the denoising of the LC frames improves the SRBR and, in many cases, even leads to better results than the HC data (scattered data points above the dashed diagonal line).

Table 1 | Average Gaussian fitting results of different training and evaluation protocols using multiple frames from the test set containing CDW signals

	$\mu_h (\times 10^2)$	$\mu_k (\times 10^2)$	$\mu_l (\times 10)$	$\sigma_h (\times 10^2)$	$\sigma_k (\times 10^2)$	$\sigma_l (\times 10)$	SRBR _{h}	SRBR _{k}	SRBR _{l}
LC	0.66 (05)	1.94 (13)	2.48 (37)	0.18 (04)	1.17 (13)	2.85 (54)	0.39 (07)	0.41 (04)	0.31 (07)
IRUNet									
Poisson \rightarrow Poisson	0.29 (03)	0.56 (09)	1.08 (12)	0.14 (02)	0.75 (10)	1.03 (15)	0.90 (10)	0.96 (03)	1.13 (04)
Poisson \rightarrow Exp.	0.75 (04)	0.87 (11)	4.17 (14)	0.31 (04)	1.65 (13)	1.37 (17)	1.02 (25)	0.95 (03)	0.97 (04)
Exp. \rightarrow Exp.	0.19 (03)	0.65 (07)	1.47 (12)	0.31 (03)	0.69 (08)	1.50 (15)	1.41 (24)	1.00 (02)	1.21 (05)
VDSR									
Poisson \rightarrow Poisson	0.38 (02)	0.58 (08)	1.01 (11)	0.14 (02)	0.78 (09)	1.23 (14)	0.95 (08)	1.05 (02)	1.22 (05)
Poisson \rightarrow Exp.	0.46 (03)	0.72 (14)	2.62 (11)	0.19 (02)	1.28 (18)	1.18 (14)	0.94 (07)	1.15 (03)	1.24 (05)
Exp. \rightarrow Exp.	0.32 (02)	0.63 (08)	1.11 (11)	0.16 (02)	0.73 (08)	0.95 (14)	0.97 (09)	1.09 (02)	1.25 (04)

The first column indicates the used training and evaluation methodology, for example, training on artificial Poisson noise and evaluation on experimental noise (Poisson \rightarrow Exp.). Values given for training and evaluation on experimental noise are additionally highlighted in bold for visual guidance. The Gaussian peak position μ_α and s.d. σ_α with reciprocal space direction $\alpha=(h, k, l)$ are given as the mean absolute error between the Gaussian parameter obtained from the denoised and the one obtained from the HC signal (the lower the better). Values for the SRBR _{α} are given as the absolute ratio of the Gaussian parameter obtained from the denoised and the one obtained from the HC signal (the higher the better). Values for μ_α as well as values for σ_α are scaled as indicated. Because of the broader peak in the l direction, a scaling of 10 for μ_l and σ_l has been chosen over a scaling of 100.

(ROI) and subtraction of neighbouring ROIs (Supplementary Fig. 6). This subtraction is applied to eliminate the background surrounding the CDW signal, such as powder rings. As the background subtraction is not always perfect, the one-dimensional line cuts are composed of signal and a small residual 'background'. To avoid negative residual background intensities, a small constant shift was applied. In Fig. 1f–h, we analyse the line cuts by fitting a Gaussian model. The resulting parameters are the amplitude A , the peak position μ , the s.d. σ and the constant residual background C . We furthermore define the signal to residual background ratio (SRBR) as A/C . Figure 3 shows the SRBR for 50 different examples of CDW order from the test set. Denoising using a neural network significantly improves the SRBR of the CDW order, oftentimes surpassing the results obtained from the HC data. Owing to the random nature of noise, the network is not able to learn the small but finite noise component of the HC data, resulting in efficient noise removal. These results are summarized in Table 1 in comparison with values extracted from the unfiltered LC data. We also compare the training with experimental and artificial noise of similar noise levels. Additionally to the SRBR, we calculate the mean absolute error between the denoised peak position $\mu_{h,k,l}$ and the s.d. $\sigma_{h,k,l}$ with the HC values. From these results, we conclude that training on experimental data greatly improves the noise filtering. This conclusion holds even in the case when the amount of artificial training data is larger than the amount of experimental training data. A considerable improvement can be achieved by employing a multiscale training procedure where the artificial training data cover a wide range of statistics (Supplementary

Table I). A table containing standard image quality metrics describing the denoising performance can be found in Supplementary Table III. Finally, we observe that both the IRUNet and VDSR networks, on average, achieve comparable results, despite their different architectures.

Discussion

The removal of experimental noise is the ultimate goal of noise filtering algorithms. Many studies have focused on filtering artificial noise from photographs^{1–6}. The artificial noise typically has a single statistical distribution (Gaussian, Poisson or Bernoulli), and the photographs are represented by red–green–blue colour scales from 0 to 255. Experimental noise poses a much harder problem as multiple noise sources are present and the signal can vary by many orders of magnitude. This suggests that denoising algorithms should (also) be benchmarked on more challenging, experimental data. Here, we provide an X-ray diffraction data set where the signal intensity varies by six orders of magnitude. Not surprisingly, we find that networks trained to remove artificial noise perform well on exactly this task. However, this high performance, unfortunately, does not carry over to the filtering of experimental noise. Our results suggest that neural networks filter experimental noise better when trained on experimental noise rather than artificial noise profiles. This statement is especially true when the noise levels of the experimental and artificial noise are comparable.

To illustrate the generality and robustness of this work, we apply the trained network to resonant inelastic X-ray scattering (RIXS) data.

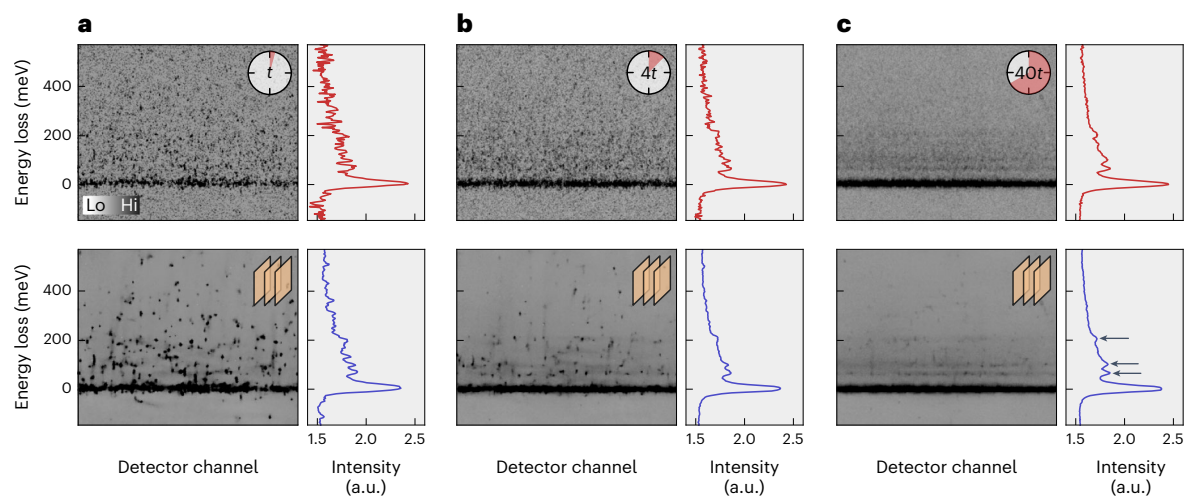


Fig. 4 | RIXS spectra recorded on SrTiO₃. **a–c**, Top row: RIXS spectra with counting statistics of 1 **(a)**, 4 **(b)** and 40 **(c)** times 3 min (t). Left: counting intensities as detector channel versus energy loss. Right: horizontally projected RIXS spectra. Bottom row: corresponding denoised neural network outputs. The arrows in **c** highlight three inelastic peaks.

X-ray diffraction and RIXS are fundamentally different experimental techniques (Methods). Figure 4 shows a raw RIXS spectrum recorded on SrTiO₃ with different counting statistics as indicated (top panels). The bottom panels show the corresponding denoised output obtained from a CNN trained on experimental X-ray diffraction data exclusively. As the used RIXS detector does not offer single-photon sensitivity, the signal is not expected to follow pure Poisson statistics. Despite the dissimilar experimental technique, different sample and different noise distribution, the trained neural network achieves a visible noise reduction and consequently enhances the SRBR. The successful denoising of RIXS data likely stems from the rich variation of signals (powder rings, charge order and lattice Bragg peaks, spurious and dead pixels) and noise sources in the X-ray diffraction training data.

Our results, therefore, encourage the collection of even more diverse training data with different compositions of noise sources from other scattering, spectroscopy and microscopy techniques. Small-angle neutron and (resonant inelastic) X-ray scattering data³⁰ would be obvious choices to extend the training data. Data obtained from spectroscopies^{7,8} and microscopies^{9,10} such as angle-resolved photoemission electron spectroscopy³¹ and transmission electron microscopy^{32,33} could also help expand the amount and variety of training data. Furthermore, the application of transfer learning³⁴ using a pre-trained model might prove beneficial in reducing the amount of distinct training data needed. By applying our method to future studies, a large amount of beamtime could be saved, or a fixed beamtime budget could be used more efficiently by, for example, being able to probe a larger parameter space.

Methods

X-ray diffraction

The training data were recorded on a La_{1.88}Sr_{0.12}CuO₄ single crystal³⁵ at beamline P21.1 at the PETRA III storage ring at the Deutsches Elektronen-Synchrotron in Hamburg. The scattering intensities were recorded using a DECTRIS PILATUS3 X CdTe 100k detector. This detector provides 195 × 487 pixels per frame and a bit depth of 32. Each pixel is associated with a horizontal and a vertical scattering angle from which reciprocal space coordinates can be reconstructed as described in ref. 23. The CNN training is independent of this reconstruction that is done to extract correlation lengths. The diffractometer was operated with 100 keV photons, and the sample was cooled to around 30 K, where a CDW order is fully developed. The charge order has a short correlation

length along the c -axis directions. Hence, along the reciprocal c -axis (ℓ), the CDW order manifests by a long rod of scattering intensity.

RIXS

The oxygen K-edge RIXS spectra were recorded at the I21 beamline³⁶ at the Diamond Light Source on a SrTiO₃ crystal. Linear vertical light polarization and a photon energy of ~531 eV were used. The sample temperature was 20 K, and the momentum transfer was set to $(h, k, \ell) = (0, 0, 0.245)$ reciprocal lattice units (r.l.u.).

Loss function

During each training epoch, the performance of the neural networks was determined by comparing the denoised output with the HC frame. The used loss function L is given by a combination of the mean absolute error and multiscale structural similarity (MS-SSIM)^{7,37}. We find that this loss function results in a better overall denoising performance when compared with other losses such as the mean squared error (L2 loss) (Supplementary Fig. 1).

CNNs

Although networks for three-dimensional data structures exist³⁸, we employed architectures designed for uncorrelated two-dimensional images. A comprehensive review of deep learning and CNNs applied to noise filtering of images is given in ref. 11. Generally, many networks display comparable performance. In this work, we implemented two different neural network architectures, referred to as VDSR²⁵ and IRUNet²⁶. For the VDSR architecture, we did not include the final addition layer as we do not find a significant performance change (Supplementary Table II). The weights of the convolutional layers are randomly initialized using the He method³⁹. For the VDSR model, we make use of a parametric rectifying linear unit³⁹ after each convolutional layer, while a normal rectifying linear unit is used in the IRUNet architecture. The VDSR network was trained for 150 epochs using a batch size of 8 and an initial learning rate of 5×10^{-4} . The IRUNet network was trained for 200 epochs using a batch size of 16 and an initial learning rate of 5×10^{-4} . The learning rate was decreased after a certain number of epochs to ensure good convergence. For the VDSR model, we multiplied the learning rate by 0.5 after every 50 epochs. For the IRUNet model, the learning rate was multiplied by 0.5 after 150 epochs. The total training duration of the VDSR and IRUNet models was on average around 20 and 10 h, respectively, on an Nvidia Tesla P100 GPU with 10 GB of VRAM

using TensorFlow 2.4.1. A discussion of the receptive field of the neural networks can be found in Supplementary Section G.

Data availability

The experimental data used in this work can be found at <https://doi.org/10.5281/zenodo.8237173> (ref. 40).

Code availability

The code used for the set-up and training of the neural networks is available at <https://doi.org/10.5281/zenodo.10245374> (ref. 41).

References

- Jain, V. & Seung, S. Natural image denoising with convolutional networks. In *Advances in Neural Information Processing Systems* Vol. 21 (eds Koller, D. et al.) (Curran Associates, 2009).
- Zhang, K., Zuo, W., Chen, Y., Meng, D. & Zhang, L. Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. *IEEE Trans. Image Process.* **26**, 3142–3155 (2017).
- Zhang, K., Zuo, W. & Zhang, L. FFDNet: toward a fast and flexible solution for CNN based image denoising. *IEEE Trans. Image Process.* **27**, 4608–4622 (2018).
- Lehtinen, J. et al. Noise2Noise: learning image restoration without clean data. In *Proc. 35th International Conference on Machine Learning* (eds Dy, J. & Krause, K.) 2965–2974 (PMLR, 2018).
- Lefkimiatis, S. Universal denoising networks : a novel CNN architecture for image denoising. In *Proc. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 3204–3213 (IEEE, 2018).
- Tian, C. et al. Deep learning on image denoising: an overview. *Neural Netw.* **131**, 251 (2020).
- Kim, Y. et al. Deep learning-based statistical noise reduction for multidimensional spectral data. *Rev. Sci. Instrum.* **92**, 073901 (2021).
- Schuetzke, J., Szymanski, N. J. & Reischl, M. A universal synthetic dataset for machine learning on spectroscopic data. Preprint at <https://arxiv.org/abs/2206.06031> (2022).
- Zhang, Y. et al. A Poisson-Gaussian denoising dataset with real fluorescence microscopy images. In *Proc. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 11702–11710 (IEEE, 2019).
- Joucken, F. et al. Denoising scanning tunneling microscopy images with machine learning. *Phys. Rev. Mater.* **6**, 123802 (2022).
- Elad, M., Kowar, B. & Vaksman, G. Image denoising: the deep learning revolution and beyond—a survey paper. *SIAM J. Imag. Sci.* **16**, 1594–1654 (2023).
- Krull, A., Buchholz, T.-O. & Jug, F. Noise2Void - learning denoising from single noisy images. In *Proc. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2124–2132 (IEEE, 2019).
- Batson, J. & Royer, L. Noise2Self: blind denoising by self-supervision. In *Proc. 36th International Conference on Machine Learning* (eds Chaudhuri, K. & Salakhutdinov, R) 524–533 (PMLR, 2019).
- Xu, J. et al. Noisy-as-clean: learning self-supervised denoising from corrupted image. *IEEE Trans. Image Process.* **29**, 9316–9329 (2020).
- Leijten, Z. J. W. A., Keizer, A. D. A., de With, G. & Friedrich, H. Quantitative analysis of electron beam damage in organic thin films. *J. Phys. Chem. C* **121**, 10552 (2017).
- Knafo, W. et al. Field-induced spin-density wave beyond hidden order in URu₂Si₂. *Nat. Commun.* **7**, 13075 (2016).
- Ruff, J. P. C. et al. Susceptibility anisotropy in an iron arsenide superconductor revealed by X-ray diffraction in pulsed magnetic fields. *Phys. Rev. Lett.* **109**, 027004 (2012).
- Bansal, M. A., Sharma, D. R. & Kathuria, D. M. A systematic review on data scarcity problem in deep learning: solution and applications. *ACM Comput. Surv.* **54**, 1–29 (2022).
- Wang, Q., Farahat, A., Gupta, C. & Zheng, S. Deep time series models for scarce data. *Neurocomputing* **456**, 504–518 (2021).
- Christensen, N. B. et al. Bulk charge stripe order competing with superconductivity in La_{2-x}Sr_xCuO₄ (x=0.12). Preprint at <https://arxiv.org/abs/1404.3192> (2014).
- Croft, T. P., Lester, C., Senn, M. S., Bombardi, A. & Hayden, S. M. Charge density wave fluctuations in La_{2-x}Sr_xCuO₄ and their competition with superconductivity. *Phys. Rev. B* **89**, 224513 (2014).
- Thampy, V. et al. Rotated stripe order and its competition with superconductivity in La_{1.88}Sr_{0.12}CuO₄. *Phys. Rev. B* **90**, 100510(R) (2014).
- Frison, R. et al. Crystal symmetry of stripe-ordered La_{1.88}Sr_{0.12}CuO₄. *Phys. Rev. B* **105**, 224113 (2022).
- Hücker, M. et al. Competing charge, spin, and superconducting orders in underdoped YBa₂Cu₃O_y. *Phys. Rev. B* **90**, 054514 (2014).
- Kim, J., Lee, J. K. & Lee, K. M. Accurate image super-resolution using very deep convolutional networks. In *Proc. 2016 IEEE Conference on Computer Vision and Pattern Recognition* 1646–1654 (IEEE, 2016).
- Gil Zuluaga, F. H., Bardozzo, F., Rios Patino, J. I. & Tagliaferri, R. Blind microscopy image denoising with a deep residual and multiscale encoder/decoder network. In *Proc. 2021 43rd Annual International Conference of the IEEE Engineering in Medicine Biology Society* 3483–3486 (IEEE, 2021).
- He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proc. 2016 IEEE Conference on Computer Vision and Pattern Recognition* 770–778 (CVPR, 2016).
- Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. In *Proc. 3rd International Conference on Learning Representations* (eds Bengio, Y. & LeCun, Y.) (ICLR, 2015).
- Reddi, S. J., Kale, S. & Kumar, S. On the convergence of Adam and beyond. In *Proc. 6th International Conference on Learning Representations* (ICLR, 2018).
- Mühlbauer, S. et al. Magnetic small-angle neutron scattering. *Rev. Mod. Phys.* **91**, 015004 (2019).
- Sobota, J. A., He, Y. & Shen, Z.-X. Angle-resolved photoemission studies of quantum materials. *Rev. Mod. Phys.* **93**, 025006 (2021).
- Kushwaha, H. S., Tanwar, S., Rathore, K. & Srivastava, S. De-noising filters for TEM (transmission electron microscopy) image of nanomaterials. In *Proc. 2012 Second International Conference on Advanced Computing Communication Technologies* 276–281 (ACCT, 2012).
- Mevenkamp, N. et al. Poisson noise removal from high-resolution STEM images based on periodic block matching. *Adv. Struct. Chem. Imag.* **1**, 1–19 (2015).
- Pan, S. J. & Yang, Q. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **22**, 1345 (2010).
- Chang, J. et al. Tuning competing orders in La_{2-x}Sr_xCuO₄ cuprate superconductors by the application of an external magnetic field. *Phys. Rev. B* **78**, 104525 (2008).
- Zhou, K.-J. et al. I21: an advanced high-resolution resonant inelastic X-ray scattering beamline at Diamond Light Source. *J. Synchrotron Radiat.* **29**, 563 (2022).
- Zhao, H., Gallo, O., Frosio, I. & Kautz, J. Loss functions for image restoration with neural networks. *IEEE Trans. Comput. Imag.* **3**, 47–57 (2017).
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T. & Ronneberger, O. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *Proc. Medical Image Computing and Computer-Assisted Intervention 2016* 424–432 (Springer International, 2016).
- He, K., Zhang, X., Ren, S. & Sun, J. Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In *Proc. 2015 IEEE International Conference on Computer Vision* 1026–1034 (IEEE, 2015).

40. Oppliger, J. et al. X-ray diffraction dataset for experimental noise filtering. *Zenodo* <https://doi.org/10.5281/zenodo.8237173> (2022).
41. Oppliger, J. X-ray diffraction denoising using deep convolutional neural networks. *Zenodo* <https://doi.org/10.5281/zenodo.10245374> (2022).

Acknowledgements

We thank N. Serra for GPU time used for training the neural networks and A. Steppke, N. Dennler and Y. Kim for insightful discussions. We acknowledge DESY (Hamburg, Germany), a member of the Helmholtz Association HGF, for the provision of experimental facilities. Parts of this research were carried out at beamline P21.1 at PETRA III. J.O. acknowledges support from a Candoc grant (no. K-72334-06-01) of the University of Zurich. J.K., Q.W. and J. Chang. express gratitude to the Swiss National Science foundation for funding under grant no. 200021_188564. Q.W. is supported by the Research Grants Council of Hong Kong (ECS no. 24306223). J.O. and F.D.N. thank the SNF (PPOOP2_176866) and ONR (N00014-20-1-2352) for generous support. I.B. and L.M. acknowledge support from the Swiss Government Excellence Scholarship under project nos. ESKAS-Nr: 2022.0001 and ESKAS-Nr: 2023.0052. M.M.D. and T.N. acknowledge support from the European Union's Horizon 2020 research and innovation programme (ERC-StG-Neupert-757867-PARATOP). M.M.D. acknowledges support from a Forschungskredit of the University of Zurich (grant no. FK-22-085). J.K., R.F. and A.M. were supported by the project CALIPSOplus under grant agreement 730872 from the EU Horizon 2020 framework programme for research and innovation. J.K. acknowledges funding from the German Academic Scholarship Foundation. N.B.C. thanks the Danish Agency for Science, Technology and Innovation for funding the instrument centre DanScatt and acknowledges support from the Q-MAT ESS Lighthouse initiative. Finally, we acknowledge Diamond Light Source for providing beamtime at I21 under proposal MM31819.

Author contributions

T.K., N.M. and M.O. grew and characterized the LSCO single crystals. The SrTiO₃ samples were obtained from commercial source. J.K., R.F., Q.W., A.M., N.B.C., O.I., A.-C.D. and M.v.Z carried out the X-ray diffraction data collection. I.B., L.M., B.F., J. Choi, M.G-F. and K.-J.Z. recorded the RIXS data. J.O. did the machine learning training and data analysis with supervision from M.M.D., F.D.N., M.H.F., T.N. and J. Chang. All authors contributed to the manuscript text.

Funding

Open access funding provided by University of Zurich.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42256-024-00790-1>.

Correspondence and requests for materials should be addressed to Jens Oppliger or Johan Chang.

Peer review information *Nature Machine Intelligence* thanks Ayana Ghosh, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024