



Title	Fast Improvement of TEM Images with Low-Dose Electrons by Deep Learning
Author(s)	Katsuno, Hiroyasu; Kimura, Yuki; Yamazaki, Tomoya et al.
Citation	Microscopy and microanalysis, 28(1), 138-144 https://doi.org/10.1017/S1431927621013799
Issue Date	2021-12-10
Doc URL	https://hdl.handle.net/2115/85768
Rights	This article has been published in a revised form in Microscopy and microanalysis https://doi.org/10.1017/S1431927621013799 . This version is published under a Creative Commons CC-BY-NC-ND. No commercial re-distribution or re-use allowed. Derivative works cannot be distributed. ©The Author(s), 2021.
Rights(URL)	https://creativecommons.org/licenses/by-nc-nd/4.0/
Type	journal article
File Information	Microscopy and microanalysis_28(1)_138-144.pdf



Fast improvement of TEM image with low-dose electrons by deep learning

Hiroyasu Katsuno¹, Yuki Kimura¹, Tomoya Yamazaki¹ and Ichigaku Takigawa^{2,3}

¹*Institute of Low Temperature Science, Hokkaido University, Kita-19, Nishi-8, Kita-ku, Sapporo, Hokkaido, 060-0819, Japan*

²*RIKEN, Center for Advanced Intelligence Project, 1-4-1 Nihonbashi, Chuo-Ku, Tokyo 103-0027, Japan*

³*Institute for Chemical Reaction Design and Discovery (WPI-ICReDD), Hokkaido University, N21 W10, Kita-ku, Sapporo, Hokkaido, 001-0021, Japan*

Corresponding Author: Hiroyasu Katsuno <katsuno@lowtem.hokudai.ac.jp>

Abstract

Low-electron-dose observation is indispensable for observing various samples using a transmission electron microscope; consequently, image processing has been used to improve transmission electron microscopy (TEM) images. To apply such image processing to *in situ* observations, we here apply a convolutional neural network to TEM imaging. Using a dataset that includes short-exposure images and long-exposure images, we develop a pipeline for processed short-exposure images, based on end-to-end training. The quality of images acquired with a total dose of approximately $5 e^-$ per pixel becomes comparable to that of images acquired with a total dose of approximately $1000 e^-$ per pixel. Because the conversion time is approximately 8 ms, *in situ* observation at 125 fps is possible. This imaging technique enables *in situ* observation of electron-beam-sensitive specimens.

Key Words: deep learning, transmission electron microscopy, fast improvement, nanoparticles

(Received XX Y 20ZZ; revised XX Y 20ZZ; accepted XX Y 20ZZ)

Introduction

Transmission electron microscopy (TEM) is a powerful tool for visualization from nano- to atomic-scale. TEM has a very wide range of applications in material sciences, biology, electro-chemistry, fluids, geology, and the environmental sciences. One of the current limitations for the further application of TEM is the influence of electron irradiation on a sample: A beam-sensitive material will often lose its characteristic structure under electron irradiation before an image can be acquired. Cryo-TEM observation is an outstanding technique to reduce the electron-induced damage on a sample (Morán & Dahl, 1952; Chlanda & Sachse, 2014). Nevertheless, an additional technique is required for *in situ* time-resolved imaging using liquid-cell TEM (De Jonge & Ross, 2011). In order to capture an ever-changing phenomenon, it is necessary to follow the phenomenon by looking at the noisy TEM images obtained in a short period of time with a low electron dose. In addition, low electron dose imaging is desirable for liquid-cell TEM to suppress rate-dependent beam effects, such as the generation of radical ions (Schneider et al., 2014). However, it seems that both the improvement of temporal resolution and that of spatial resolution are in conflict (De Jonge et al., 2019).

Over the past decade, image analysis has been dramatically developed by machine learning. Regarding simple image denoising, numerous methods have been proposed for image denoising, including sub-pixel convolutional neural networks (CNNs) (Shi et al., 2016), nuclear norm minimization (Gu et al., 2014), sparse coding (Elad & Aharon, 2006) and domain filtering (Dabov et al., 2007). The machine learning has also developed image analysis methods of TEM imaging. In particular, U-Net architecture (Ronneberger et al., 2015; Falk et al., 2019) has been used for various image processing tasks cell counting, object detection and segmentation in addition to image improvement. Using a deep learning trained on simulated TEM data, the local structures of graphene have been found (Madsen et al., 2018). For high resolution

TEM data, the segmentation between the clean surface and surface contaminant layers has been succeeded (Sadre et al., 2020). Automated analysis has revealed the connectome of the drosophila brain from a huge amount of electron microscope images (Li et al., 2020).

Recently, a new technique of deep learning has been proposed for low-light image enhancement. In this method, an image is acquired with a low-light so that the original image can be used as ground truth for comparison (Lim et al., 2015). A CNN is applied for image improvement, and an image set of short-exposure and long-exposure images are prepared for training (Chen et al., 2018). The object detection in short-exposure images is also succeeded (Loh & Chan, 2019). The corresponding results have shown that not only the machine learning model but also datasets need improvement. Machine learning technique has a large potential for an extension of the application of TEM.

As considering a new application to *in situ* imaging using liquid-cell TEM imaging, low-dose observation is necessary. Moreover, it is better to realize the high temporal resolution. In the present work, in order to solve both matters, we introduce a CNN model for TEM imaging and evaluate the image quality and the speed for image improvement. Low-dose observation TEM images are enhanced by using the CNN model trained with the dataset, which contains high-dose-electron (HDE) and low-dose-electron (LDE) TEM images with the same field of view. In general, image conversion is fast, although it takes a long time to train a CNN model. Using a CNN help increase the speed of image improvement, enabling *in situ* observation.

Methods

We used a transmission electron microscope equipped with a field-emission gun (JEM-2100F, JEOL, Tokyo) operated at an acceleration voltage of 200 kV and a CMOS camera, OneView IS (Gatan, Inc., Pleasanton, CA, USA). For training, we prepare sets of high-dose-electron and low-dose-electron TEM images with the same field of view. All images were acquired with drift correction using the function incorporated into the software (Digital Micrograph, Gatan. Inc., Pleasanton, CA, USA) used to operate the CMOS camera. The HDE image resolution

Table 1. The number of images for training and validation.

Set No.	Material	Training	Validation
1	Ni	224	36
2	FeNi	260	40
3	SiC	140	20
4	Silicate	176	24
5	Alumina	200	40

was 4096×4096 pixels, and its exposure time was 5 s. After an HDE image was acquired, the corresponding LDE image is taken at a resolution of 512×512 pixels and with an exposure time of 3.3 ms. The typical total doses in each view were $10^{10} e^-$ and $10^6 e^-$ in HDE and LDE images, respectively, on the camera. The dose rate was calibrated using a Faraday cage (JEOL Ltd., Tokyo, Japan). The numbers of electrons on each pixel were $\sim 1000 e^-$ and $\sim 5 - 10 e^-$ in the HDE and LDE images, respectively. Typical magnifications used in the present study were $25,000\times$ and $30,000\times$; the dose rate on the samples was therefore $\sim 10^2 e^- \text{nm}^{-2} \text{s}^{-1}$ in case of LDE observation. The number of images for training data is summarized in Table 1. The typical samples were particles of Ni (metal), FeNi (alloy), SiC (semiconductor), silicate (mineral), and alumina (ceramics) with diameters of 30–200 nm.

The original binary data of the digital micrograph were converted to grayscale images in tiff format with the dark current subtracted and including the intensity with 16-bit expression in each pixel. The obtained images were preprocessed with intensity rescaling. Before training, the position of each HDE image was adjusted so that its position fit the corresponding LDE image, thereby compensating for sample drift in the sequences.

Our deep learning model has the U-Net architecture (Ronneberger et al., 2015) with the ResNet encoder/decoder (He et al., 2016) using the segmentation package in PyTorch (Yakubovskiy, 2020). The schematic of our image improvement is shown in Fig. 1. The model parameters were fine-tuned against the pre-trained Resnet-18 model from the ImageNet data

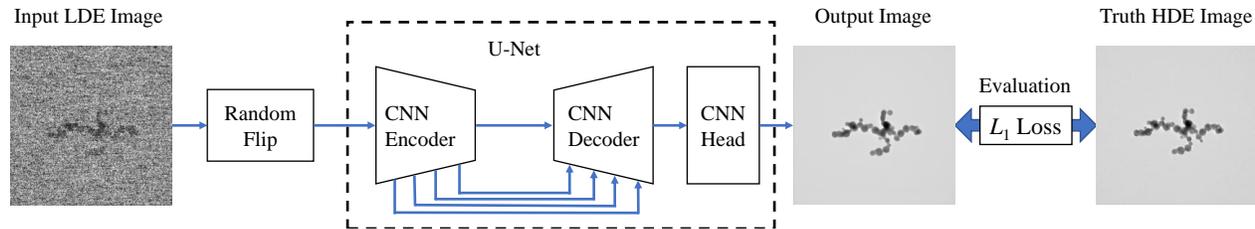


Fig. 1. Schematic of image improvement using our deep learning model.

using the L_1 loss function and the Adam optimizer (Kingma & Ba, 2015) with the learning rate of 10^{-4} . The concrete form of L_1 function for two images is written as

$$L_1 = \frac{1}{N} \sum_x |I_A(x) - I_B(x)|, \quad (1)$$

where x represents the position of a pixel. $I_A(x)$ and $I_B(x)$ are the normalized intensity of two images at the position x , and N is the total number of pixels. In each iteration, an LDE image was randomly flipped horizontally and/or vertically for data augmentation. We did not apply rotations and random crops since they were not effective in our study because pixelwise lateral noise in LDE images is broken. It suggests that our current model does not make any improvement for other cameras. All training was conducted on a Linux machine with a 10-core Intel i9-9900X 3.5 GHz CPU and an NVIDIA Quadro RTX 8000 graphics card (see Machine No. 1 in Table 3).

Results & Discussion

Overview of the Image Improvement

Figure 2 shows the epoch dependence of the loss of the training dataset and the validation dataset. As the training proceeded, the value of the loss decreased until 100 epochs. Whereas the loss value for the training data decreased after 100 epochs, that for the validation data was steady. We also checked the loss in more than 1000 epochs, and observed a slight increase in the loss value of the validation data because of overfitting. In the present study, we used

: Fast improvement by deep learning

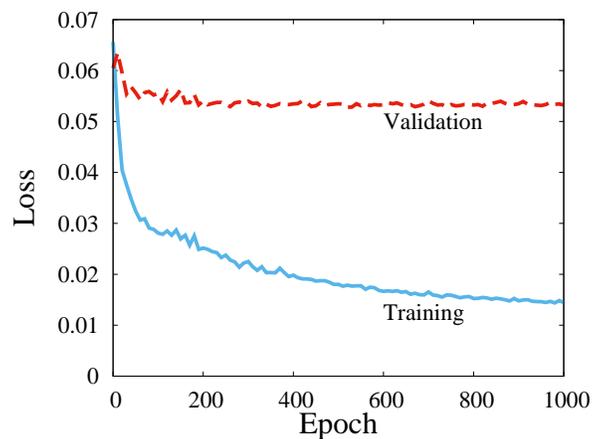


Fig. 2. Training loss and validation loss.

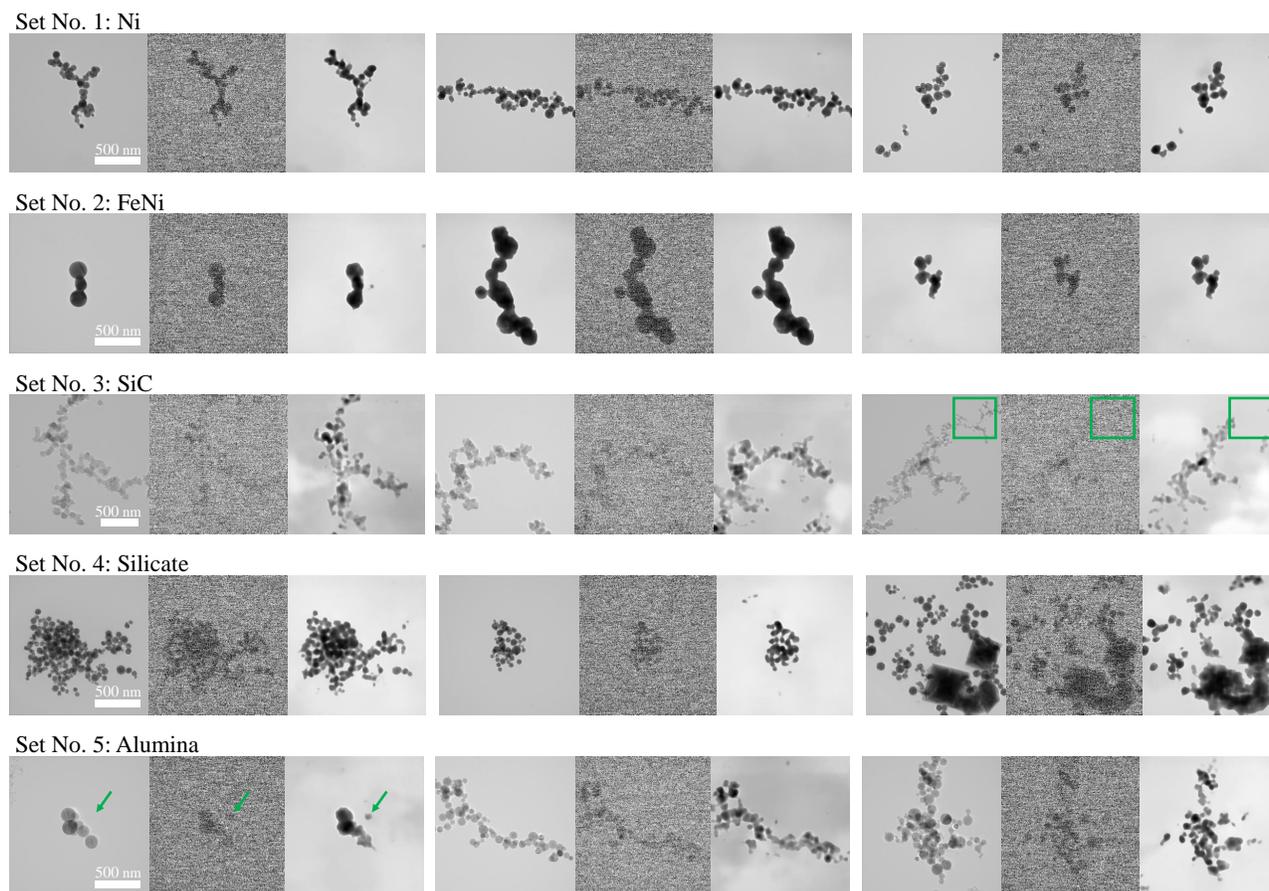


Fig. 3. Examples of HDE, LDE and output images (left to right); images of Ni, FeNi, SiC, silicate, and alumina displayed top to bottom. Arrows and boxes indicate examples of the failure results of the improvement. The presented images were used for validation, not training.

Table 2. Qualitative comparison of output images for validation using the mean absolute error (MAE) and the peak signal-to-noise ratio (PSNR).

Set no.	Material	MAE	PSNR
1	Ni	0.075 ± 0.026	20.82 ± 2.13
2	FeNi	0.034 ± 0.011	25.10 ± 1.86
3	SiC	0.074 ± 0.020	20.80 ± 1.60
4	Silicate	0.062 ± 0.025	21.60 ± 2.82
5	Alumina	0.075 ± 0.026	20.69 ± 2.06
	(Average)	0.063 ± 0.028	21.97 ± 2.79

the finetuned parameter at 900 epochs.

Figure 3 shows some examples of HDE images, LDE images, and output images obtained using our model from corresponding LDE images. These HDE and LDE image sets were not used for training. Images of Ni, FeNi, SiC, silicate, and alumina are displayed from the top row to the bottom row in Fig. 3. Various sizes and contrasts of nanoparticles are observed in HDE images. The preprocessed LDE images are similar to the images displayed by the camera software (Gatan Digital Micrograph) during low-dose-rate observations. All of the images show noise resembling a sandstorm. A common feature of all of the output images is that the noise is removed, although a haze remained in the background of the images in the third and fifth rows of Fig. 3. The overlapping nanoparticles, such as those in the fourth row and first column image of Fig. 3, were reproduced in the output images. Thus, we obtained an image comparable to an HDE image from a sandstorm image, although the improvement failed in some cases: the appearance of imaginary nanoparticles indicated by the arrow in the image in the fifth row and first column of Fig. 3, and the disappearance of nanoparticles, such as in the boxed area in the image in the third row and third column of Fig. 3.

The quality of the improvement appears to depend on the sample. For a qualitative comparison of the quality of the output images, the mean absolute error (MAE) and the peak

signal-to-noise ratio (PSNR) are listed in Table 2. The MAE is the difference between images at the pixel level (see details in the next subsection) and the large MAE value indicates a large difference between images. The PSNR is a logarithm of the inverse of the mean square error where a large value of PSNR indicates that two images are similar. Both indicators show that our model is more effective for set No. 2 (FeNi) than the other sets. In set No. 2, the magnitude of the electron count on nanoparticles is less than 50% of that on the background in HDE images. However, in sets No. 3 (SiC) and No. 5 (alumina), the magnitude of the electron count on the nanoparticles is almost 90% of that on the background in HDE images; i.e., most of the SiC and alumina nanoparticles exhibit relatively weaker contrast than FeNi particles. The small difference in contrast between the background and the nanoparticles, which is only 10%, may prevent the image improvement. In sets No. 1 (Ni) and No. 4 (silicate), the HDE images tend to aggregate nanoparticles with weak and strong contrast. In particular, the large standard deviation of the PSNR of set No. 4 might be reflected the variety of images.

Since all datasets are acquired under similar experimental conditions [focus, magnification, lens conditions, apertures and substrate (amorphous carbon thin film)], we speculate the origin of the difference in our improvement as follows: When the size of a nanoparticle is small, the amount of electron beam transmission at the nanoparticle is close to that at the background and it is difficult to improve the image. This effect leads to the size dependence of nanoparticles on image improvement. The degree of image improvement differs between nanoparticles with the same size but different electron beam transparency. It leads to the material dependence on image improvement. At least, the statistical data indicates that our model parameters in this study provides the same level of image quality for 30–200 nm nanoparticles of various materials with the magnification of $25,000\times$ – $30,000\times$.

Example of Visualization of LDE Image

In a microscope image, indicators of the performance include the accuracy of the size of an object and the accuracy of the separation of two adjacent objects. We studied the improvement

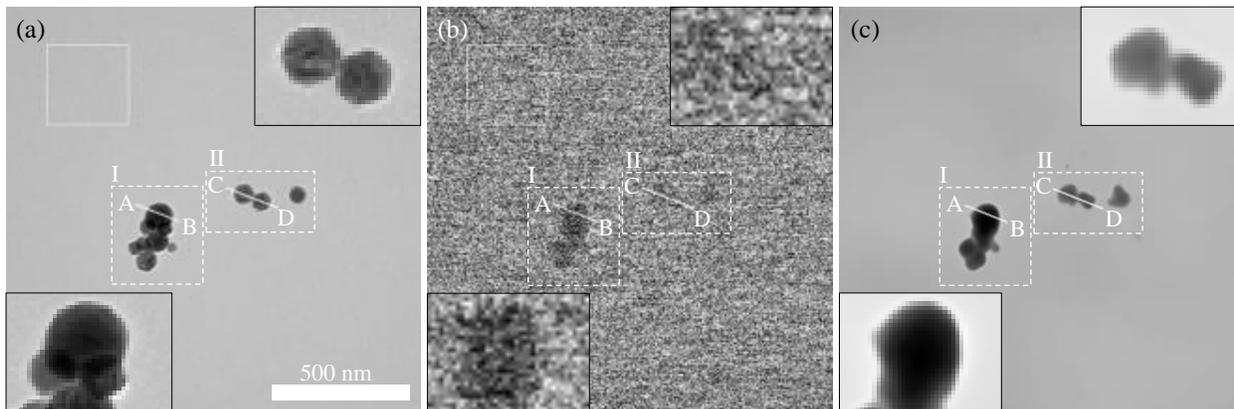


Fig. 4. Examples of (a) an HDE image, (b) a LDE image and (c) an output image of set No. 4. The average dose is about $1200 e^-$ per pixel and $4.7 e^-$ per pixel in a box. The insets in the figures show enlarged images of parts of region I and II. The presented images were used for validation, not training.

rate and these two indicators in detail as an example in set No. 4.

Figure 4 shows an example of (a) an HDE image, (b) an LDE image, and (c) a corresponding output image from (b). In Fig. 4(a), there are nine silicate nanoparticles in region I indicated by the dashed-line box and three nanoparticles in region II. The nanoparticles are typically 100 nm in diameter. Figure 4(b) shows a preprocessed LDE image. Sandstorm-like noise is present in the whole image. Only two particles are recognized in region I, and it is difficult to recognize a particle in region II. Figure 4(c) is the output image generated from Fig. 4(b) by our model. Particles are recognized in both regions I and II, and no particles are observed elsewhere, as in Fig. 4(a).

All images have been normalized and have intensity values from 0 to 1 in each pixel, and the L_1 function indicates the difference in the intensity at pixels (see Eq. 1). The maximum value of L_1 is 1, where an image is composed of only black and white pixels and another image is composed of the opposite color, e.g., a set of all black and all white images. On the contrary, the minimum value is 0; where two images are in complete agreement. The value of the L_1 loss in Fig. 4(a) and Fig. 4(b) is 0.34. After the training, the value decreases to 0.02 in the

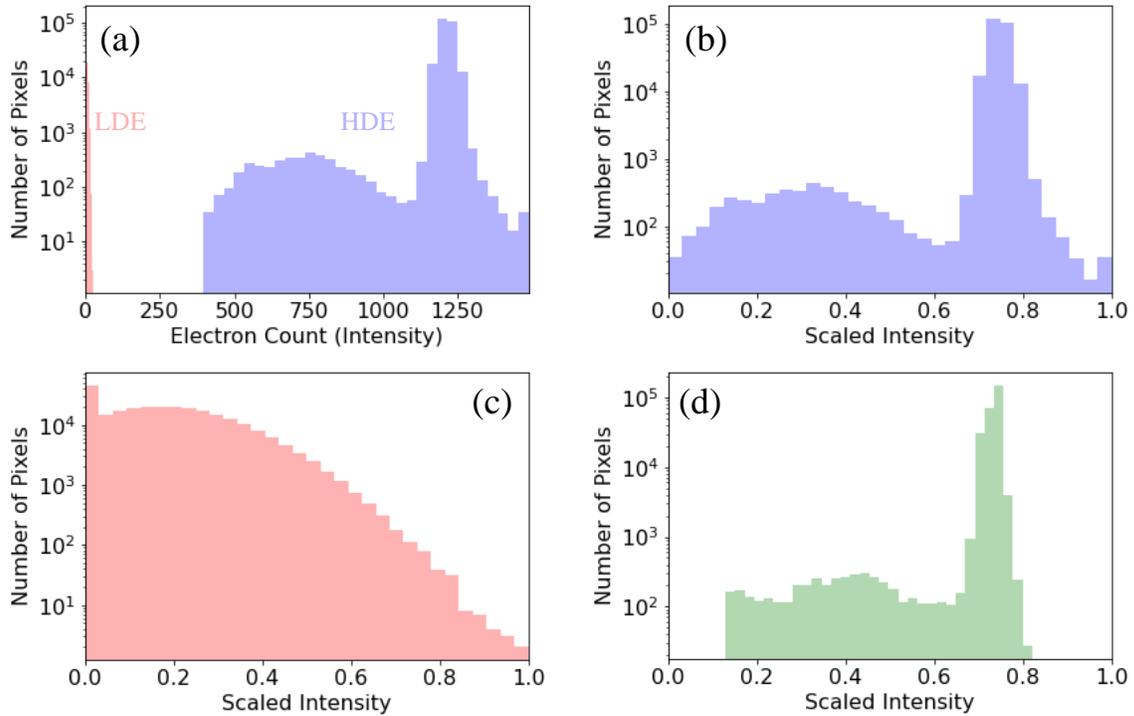


Fig. 5. (a) Raw intensity histograms of the HDE and LDE images shown in Fig. 4. Scaled intensity histograms of (b) the HDE image, (c) the corresponding LDE image, and (d) the output image in Figs. 4(a)–(c), respectively. The abscissa is the electron count and the ordinate is the number of pixels. The data is plotted as semi-log plots.

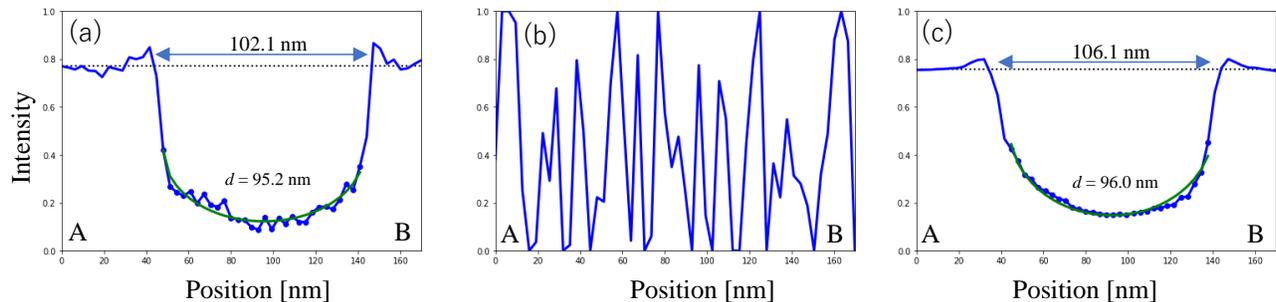


Fig. 6. Line profile of A–B shown in Fig. 4: (a) the HDE image, (b) the LDE image, and (c) the output image. Circles are the data used to calculate the diameter of the nanoparticle. The dotted line shows the average of the background.

case of Fig. 4(a) and Fig. 4(c). The image is improved by our model.

Figure 5(a) shows the raw intensity histograms of the HDE and LDE images shown in Fig. 4 as semi-log plots. As the minimum electron count of the HDE image is 396 and the maximum value of the LDE image is 24, the histogram of the LDE image is barely visible in this plot. The histograms of the HDE and the LDE images are scaled as shown in Figs. 5(b) and 5(c), respectively. On the histogram corresponding to the HDE image [Fig. 5(b)] shows two peaks. The sharp peak arises from the background, whose value is about 0.7–0.8. Another broad peak at 0.3 originates from the nanoparticles. Two peaks are clearly separated on the intensity histogram. However, the intensity histogram corresponding to the LDE image [Fig. 5(c)] shows a single peak. The intensity of the nanoparticles is not substantially different from that of the background. Figure 5(d) shows the histogram of the output image. Despite the conversion of the LDE image [Fig. 4(b)], whose histogram shows a broad single peak [Fig. 5(c)], a sharp peak at 0.7 and a broad peak at 0.4 appear. The shape of the histogram is similar to that of the histogram of the HDE image although there is no intensity greater than 0.8 or smaller than 0.15.

We here focus on the discrimination of nanoparticles from viewpoint of image improvement. We first investigated whether the size of nanoparticles could be accurately reproduced in the output images. A magnified view of location A–B in the LDE image is shown in the bottom-left

inset in Fig. 4. Although all nanoparticles are easily found in the HDE and output images, all of the nanoparticles are difficult to observe in the LDE image. The line profiles corresponding to A–B in Fig. 4 are shown in Fig. 6. In the HDE image [Fig. 6(a)] and the output image [Fig. 6(c)], a concave curve appears at the center because of the nanoparticle. In the LDE image shown in Fig. 6(b), random noise is present and the concave curve does not appear. When the shape of the nanoparticle is assumed to be spherical, the diameter of the nanoparticle is found to be 95.2 nm in the HDE image and 96.0 nm in the output image on the basis of the data corresponding to the center of the concave curve. The output image generated by our model reproduces the size of the nanoparticle with an accuracy within 1%.

In addition, we investigated the edge width of a nanoparticle. The edge width was estimated by comparing a region of strong contrast and the estimated diameter of a particle with an assumed spherical shape. The distance between crosspoints composed of the line profile (blue solid line) and the background (horizontal dotted line) in Figs. 6(a) and 6(c) was assumed to represent the pseudo-diameter of the nanoparticle indicated by A–B in Fig. 4. By subtracting the estimated particle diameter from the pseudo-diameter of the nanoparticle, we obtained the value of twice the edge width. The obtained edge widths were 6.9 nm and 10.1 nm for the nanoparticle in the HDE image and that in the output image, respectively. Thus, the magnitude of the edge width in the output image is comparable to that in the HDE image.

We also investigated whether two adjacent nanoparticles could be distinguished. The line profiles corresponding to C–D in Fig. 4 are shown in Fig. 7. The line profile of the LDE image [Fig. 7(b)] shows random noise, as does that in Fig. 6(b). At location C–D, nanoparticles are difficult to recognize in the LDE image. By contrast, nanoparticles are recognized in the HDE and output images. From the line profile of the HDE image [Fig. 7(a)], we obtained diameters of $d_1^{\text{HDE}} = 58.3$ nm and $d_2^{\text{HDE}} = 73.0$ nm for the nanoparticles in the HDE image. From the nanoparticles' center position indicated by dotted vertical lines shown in Fig. 7(a), the distance between two nanoparticles is obtained as $l^{\text{HDE}} = 78.4$ nm. The condition $(d_1^{\text{HDE}} + d_2^{\text{HDE}})/2 < l^{\text{HDE}}$ is satisfied; that is, the system has sufficient resolution to distinguish two nanoparticles.

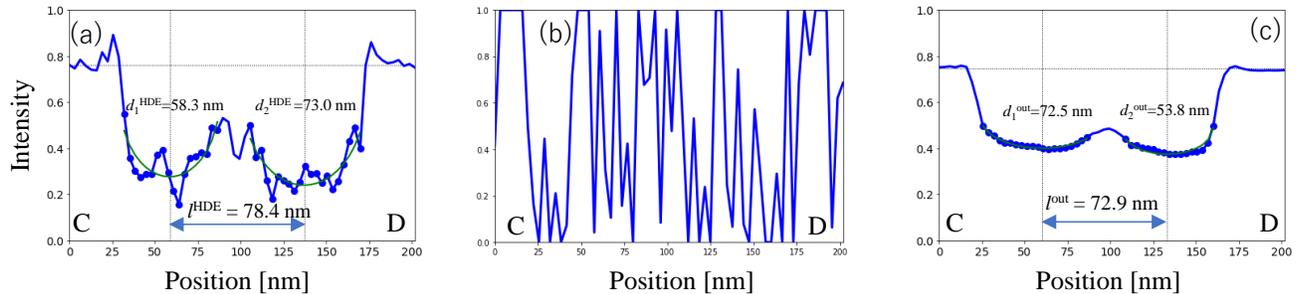


Fig. 7. Line profile of C–D shown in Fig. 4: (a) the HDE image, (b) the LDE image, and (c) the output image. Circles are the data used to calculate the diameter of the nanoparticle. The dotted line shows the average of the background.

In the case of the output image in Fig. 7(c), the sizes of nanoparticles are $d_1^{\text{out}} = 72.5$ nm and $d_2^{\text{out}} = 53.8$ nm. The difference in nanoparticle size determined from the HDE and output images is more than 20%. The distance between nanoparticles is $l^{\text{out}} = 72.9$ nm, which is close to the l^{HDE} . The condition $(d_1^{\text{out}} + d_2^{\text{out}})/2 < l^{\text{out}}$ is also satisfied for output image. After applying our CNN model, we could distinguish between two adjacent nanoparticles at this size scale.

Our model removes noise from LDE images, clearly revealing the presence of nanoparticles. The center position and the distance between adjacent nanoparticles can be reproduced at this size scale. Although the size of the nanoparticles in the output image matches that in the HDE image, the nanoparticles' shape is uncertain (see Fig. 4). Moreover, the intensity of the line profile due to the internal structure of the nanoparticles can disappear during the image processing because our model is a CNN that appears to spatially average the intensity. Therefore, restoring the image of agglomerated nanoparticles like those in region I of Fig. 4 tends to be difficult, although isolated nanoparticles are correctly reproduced after our image processing. For instance, we can count only five or seven nanoparticles in region I of the output image [Fig.4(c)] against nine nanoparticles in the corresponding HDE image [Fig. 4(a)].

In our dataset, there is no significant difference in the choice of architecture. Therefore, we show the results using U-Net architecture, which is the simple and conventional architecture.

Table 3. Waiting time for an output image in various processors. The unit is in milliseconds. The values were obtained as an average from 100 samples.

Machine No.	Processor	Model number	Calc.	Calc. including data transfer
1	CPU	Intel Core i9-9900X 3.50 GHz	106 ± 0.7	–
	GPU	NVIDIA Quadro RTX 8000	4.4 ± 0.0	8.0 ± 0.1
2	CPU	Intel Core i7-9700 3.00 GHz	300 ± 10	–
	GPU	NVIDIA Geforce GTX 1650	5.3 ± 0.6	25 ± 0.5

It is also works well for *in situ* observation where fast image improvement is necessary (see the next subsection). Further improvement in prediction accuracy may be expected by using other architecture for other datasets.

Waiting Time for Improvement

For *in situ* TEM observations with low electron doses, improving the LDE image on a timescale that approaches the camera speed is important. In maintaining the high performance of the TEM observation, the output image should be generated from the LDE image faster than the frame rate of the camera. The framerates of the CMOS camera (OneView IS) are 25 and 300 frames per second in normal mode (output size: 4096×4096 pixels) and binning 8 mode (output size: 512×512 pixels), respectively. Our goal was to convert to one image within 40 ms in the first step and then every 3.3 ms thereafter. Table 3 shows the waiting time for converting an LDE image using our model and two types of machines. The first configuration (Machine No. 1) is a calculation machine for numerical calculation, where the training in this study is performed. The second configuration (Machine No. 2) is a personal computer for normal use. In both cases, the CPU calculation requires hundreds of milliseconds. The GPU calculation is fast, and the performance is more than 40 times greater than that of CPU calculation; that

is, the conversion time is several milliseconds. The time of conversion using the Geforce GTX 1650 graphics card is approximately the same as that using the Quadro RTX 8000 graphics card. However, GPU calculation requires the data transfer from the CPU to the GPU and from the GPU to the CPU. When the data transfer is taken into account, the total waiting time is 8 ms with Machine No. 1 and 25 ms with Machine No. 2. Although the waiting time of 8 ms is more than two times longer than the 3.3 ms maximum temporal resolution of the Gatan OneView camera operating in binning 8 mode, it is substantially shorter than the 40 ms estimated from the frame rate of the OneView camera in normal use (25 fps).

Summary

We improved TEM images acquired using LDE by applying a simple CNN. Our model is based on the U-Net architecture with the ResNet encoder. We demonstrated that enabling the observation of objects that are difficult to visualize in the LDE image because our model can reproduce objects from the noise in addition to removing noise. The position of nanoparticles in the HDE images was reproduced in the corresponding output images, and the size and the edge width of nanoparticles were similar in the HDE and output images. In contrast, their shape reproduction requires further improvement. The time necessary for the image conversion is approximately 8 ms, making the method applicable for *in situ* observation at a frame rate of 125 fps or lower. Our model is effective for investigating fast dynamic processes such as nucleation from a solution or tracking the motion of nanoparticles via LDE TEM observation.

Acknowledgments This work is supported by JSPS KAKENHI Grant Numbers 20H05657 and 21K03379.

References

Chen, C., Chen, Q., Xu, J. & Koltun, V. (2018). Learning to See in the Dark. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3291–3300.

- Chlanda, P. & Sachse, M.** (2014). Cryo-electron microscopy of vitreous sections. *Methods Mol Biol* **1117**, 193–214.
- Dabov, K. Foi, A., Katkovnik, V. & Egiazarian, K.** (2007). Image denoising by sparse 3-D transform-domain. *2007 IEEE Trans on Image Process* **16**, 2080–2095.
- De Jonge, N., Houben, L., D-Borkowski, R.E. & Ross, F.M.** (2019). Resolution and aberration correction in liquid cell transmission electron microscopy. *Nat. Rev. Mater.* **4**, 61–78.
- De Jonge, N. & Ross, F.M.** (2011). Electron microscopy of specimens in liquid. *Nat. Nanotechnol.* **6**, 695–704.
- Elad, M. & Aharon, M.** (2006). Image denoising via sparse and redundant. *2006 IEEE trans Image Process* **15**, 3736–3745.
- Falk, T., Mai, D., Bensch, R., Özgün, Ç., Abdulkadir., A., Marrakchi, Y. Böhm A., Deubner, J., Jäckel, Z., Seiwald, K., Dovzhenko., A., Tietz, O., Bosco, C.D., Walsh, S., Saltukoglu, D.I., Tay, T.L., Prinz, M., Palme, K., Simons, M., Diester, I., Brox, T. & Ronneberger, O.** (2019). U-Net:deep learning for cell counting, detection, and morphometry. *nature methods* **16**, 67–70.
- Gu, S., Zhang, L., Zou, W. & Feng, X.** (2014). Weighted nuclear norm minimization with application to image denoising. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2862–2869.
- He, K., Zhang, X., Ren, S., & Sun, J.** (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778.
- Kingma, D.P. & Ba, J.** (2015). Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning*; *arXiv* 1412.6980v5.
- Li, P.H., Lindsey, L.F., Januszewski, M., Zheng, Z., Bates, A.S., Taisz, I., Tyka, M., Nichols, M., Li, F., Perlman, E., Maitin-Shepard, J., Blakely, T., Leavitt, L., Jefferis, G.S.X.E, Bock, D. & Jain. V.** (2020). Automated reconstruction of a serial-Section EM drosophila brain with flood-filling networks and local realignment. *bioRxiv*:605634

Lim, J., Kim, J.-H., Sim, J.-Y., & Kim, C.-S. (2015). Robust contrast enhancement of noisy low-light images: denoising-enhancement completion. *2015 IEEE International Conference on Image Processing (ICIP)*, pp.4131–4135.

Loh, Y.P. & Chan, C.S. (2019). Getting to know low-light images with the exclusively dark dataset. *Comput Vision Image Understanding* **178**, 30–42.

Madsen, J., Liu, P., Kling, J., Wagner, J.B., Hansen, T.W., Winther, O. & Schiøtz, J. (2018). A Deep Learning Approach to Identify Local Structures in Atomic-Resolution Transmission Electron Microscopy Images. *Adv. theory simul.* **1**, 1800037.

Fernández-Morán, H. & Dahl, A.O. (1952). Electron microscopy of ultrathin frozen sections of pollen grains. *Science* **116**, 465–467.

Ronneberger, O., Fischer P. & Thomas, B. (2015). U-Net: convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015.*, Lecture Notes in Computer Science , Vol. 9351. Navab, N., Hornegger, J., Wells, W., Frangi, A. (Eds.), pp.234–241. Springer, Cham.

Sadre, R., Ophus, C., Butko, A. & Weber, G.H. Deep learning segmentation of complex features in atomic-resolution phase contrast transmission electron microscopy images. *arXiv* 2012.05322v1.

Schneider, N.M., Norton, M.M, Mendel. B.J., Gorgan, J.M., Ross, F.M. & Bau H.H. (2014). Electron-water interactions and implications for liquid cell electron microscopy. *J Phys Chem C* **118**, 22373–22382.

Shi, W., Caballero, J., Hanzar, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D. & Wang, Z. (2016). Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1874–1883.

Yakubovskiy, P. (2020). Segmentation Models: Python library with Neu-

: *Fast improvement by deep learning*

ral Networks for Image Segmentation based on PyTorch. GitHub repository,
https://github.com/qubvel/segmentation_models.pytorch.