Unsupervised Deep Learning for Laboratory-Based Diffraction Contrast Tomography

Hovad, Emil; Fang, Haixing; Zhang, Yubin; Clemmensen, Line Katrine Harder; Ersbøll, Bjarne Kjær; Juul Jensen, Dorte

Published in:
Integrating Materials and Manufacturing Innovation

Link to article, DOI:
[10.1007/s40192-020-00189-x](https://doi.org/10.1007/s40192-020-00189-x)

Publication date:
2020

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Hovad, E., Fang, H., Zhang, Y., Clemmensen, L. K. H., Ersbøll, B. K., & Juul Jensen, D. (2020). Unsupervised Deep Learning for Laboratory-Based Diffraction Contrast Tomography. *Integrating Materials and Manufacturing Innovation*, 9, 315–321. <https://doi.org/10.1007/s40192-020-00189-x>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Unsupervised deep learning for laboratory-based diffraction contrast tomography

Emil Hovad* · Haixing Fang ·
Yubin Zhang · Line Katrine Harder
Clemmensen · Bjarne Kjær Ersbøll ·
Dorte Juul Jensen

Received: date / Accepted: date

Abstract An important leap forward for the 3D community is the possibility to perform non-destructive 3D microstructural imaging in the home laboratories. This possibility is profiled by a recently developed technique - laboratory X-ray diffraction contrast tomography (LabDCT). As diffraction spots in LabDCT images are the basis for 3D reconstruction of microstructures, it is critical to get their identification as precise as possible. In the present work we use a deep learning (DL) routine to optimize the identification of the spots. It is shown, that by adding an artificial simple constant background noise to a se-

Emil Hovad

DTU COMPUTE, Department of Applied Mathematics and Computer Science. Technical University of Denmark. Richard Petersens Plads. Building 324. 2800 Kgs. Lyngby. *Corresponding author E-mail: emilh@dtu.dk

Haixing Fang

DTU MECHANICAL ENGINEERING, Department of Mechanical Engineering. Technical University of Denmark. Produktionstorvet, Building 425, 2800 Kgs. Lyngby. E-mail: hfang@mek.dtu.dk

Yubin Zhang

DTU MECHANICAL ENGINEERING, Department of Mechanical Engineering. Technical University of Denmark. Produktionstorvet, Building 425, 2800 Kgs. Lyngby. E-mail: yubz@mek.dtu.dk

Line Katrine Harder Clemmensen

DTU COMPUTE, Department of Applied Mathematics and Computer Science. Technical University of Denmark. Richard Petersens Plads. Building 324. 2800 Kgs. Lyngby. E-mail: lkhc@dtu.dk

Bjarne Kjær Ersbøll

DTU COMPUTE, Department of Applied Mathematics and Computer Science. Technical University of Denmark. Richard Petersens Plads. Building 324. 2800 Kgs. Lyngby. E-mail: bker@dtu.dk

Dorte Juul Jensen

DTU MECHANICAL ENGINEERING, Department of Mechanical Engineering. Technical University of Denmark. Produktionstorvet, Building 425, 2800 Kgs. Lyngby. E-mail: doje@mek.dtu.dk

ries of forward simulated LabDCT diffraction images, DL can be trained and then learn to remove high frequency noise and low frequency radial gradients in brightness in the real experimental LabDCT images. The training of the DL routine is unsupervised in the sense that no human intervention is needed for labelling the data. The reduction in high frequency noise and low frequency radial gradients in brightness reduction is demonstrated by comparing line profile scans through the experimental and the DL output images. Finally, the implications of this reduction procedure on the spot identification are analyzed and possible improvements are discussed.

Keywords LabDCT · Deep learning · Noise reduction · Diffraction images · 3D Microstructures

1 Introduction

Non-destructive 3D characterization of materials microstructures using methods at large international synchrotron facilities offers substantial advantages compared to conventional 2D methods. For metals and alloys in particular, the possibility to follow the microstructural evolution in the bulk has led to several major scientific breakthroughs [1–5]. However, to broaden the use of this novel and outstanding characterization possibility, there is a need for techniques that can operate in the home laboratories. Laboratory X-ray diffraction contrast tomography (LabDCT) offers such opportunities [6, 7]. A cornerstone in the LabDCT method is precise identification of the diffraction spots.

In the commercially available LabDCT software, the images of LabDCT projections are normally first pre-processed with different filters, e.g. rolling median filter followed by Laplacians of Gaussians, and then binarized to separate the diffraction spots from the background by applying thresholds for spot size and grey value. While this method works well in general, it shows deficiencies in identifying precisely weak spots surrounded locally by a large amount of noise, i.e. in the situations with variations of the background intensity, which may be hard to overcome using conventional image segmentation methods.

In recent years, new image processing methods have been developed using advanced deep learning (DL) algorithms for image restoration and segmentation [8]. The aim of the present work is to establish a new framework using DL for processing LabDCT datasets. As a first attempt, a DL routine consisting of two pretrained sub-networks, as proposed in [9] for a super resolution problem, is used. To ease the training of the DL routine, we have here chosen to use noise-free diffraction images generated by forward simulation as input for the training procedure. The advantage is that with such images, the ground-truth is known thus avoiding time-consuming human labeling (i.e. the training is unsupervised). In this first attempt, a very simple constant background noise was added to the simulated LabDCT projections for the training. After the DL network was trained, we applied the DL routines for real experimental LabDCT images to understand how it removes high frequency noise and low frequency

Fig. 1 Training data: (a) Simulated LabDCT diffraction image (content image y_c), (b) the synthetic input image x where the darkest pixel values are brightened Δ , (c) the prediction image $\hat{y}$ (output image).

radial gradients in brightness in realistic situations. It was found that in spite the network was trained removing some super simple and unrealistic simple constant background noise from forward simulated images, it manages quite well also to remove a significant fraction of the much more complex high frequency noise and to some extent low frequency radial gradients in brightness in real diffraction images. The results from this study is thus promising and provide guidelines for developing more sophisticated training processing that is essential for future restoration of noise and artifact free LabDCT datasets with DL.

2 Data

A dataset of LabDCT projections was generated with a forward simulation model to train the deep convolutional neural network to remove high frequency noise and low frequency radial gradients in brightness in the images for an subsequent enhanced identification of the diffraction spots. These simulated projections are referred to as content images (y_c) in the following. An example of a forward simulated LabDCT projection (content image) is shown in Fig. 1a.

The forward simulated projections contain all the possible diffraction spots from grains in a virtual Al sample. The sample was created as a cylinder with a diameter of $400 \mu\text{m}$ and a height of $600 \mu\text{m}$. There were 144 grains with an average grain size of $100 \mu\text{m}$ and a standard deviation of $10 \mu\text{m}$ in the sample and the grain orientations were generated randomly. Here we concentrated on diffraction images from a sample with relatively large grains. But, the forward simulation model can handle any grain size. Since, LabDCT currently is limited to strain-free samples with zero-mosaicity within the grains, the grain in the virtual input were also assigned zero-mosaicity. Each grain was meshed into polyhedrons with a mesh size of about $12 \mu\text{m}$, which has been proved to be a suitable size for balancing the simulation accuracy and the computational speed [11,12]. Diffraction events were considered individually

for each polyhedron. An X-ray spectrum of a typical X-ray tube with tungsten target operating at an acceleration voltage of 140 kV was used as the photon source [12,13]. For each pixel on the detector, the intensity was calculated by summing the intensities from all the possible diffraction signals and the background. In the present simulation the background intensity is described in section 3.1. More details about the forward simulation algorithm can be found in [11,12,14]. The simulation was performed for typical LabDCT experimental conditions: The detector, recording the diffraction patterns, has an area of 1015×1015 pixels with an effective pixel size of $6.72 \mu\text{m}$. The sample-to-source distance (L_{ss}) and the sample-to-detector distance (L_{sd}) were set to be the same and equals 11 mm. This symmetric geometry results in the so-called Laue focusing of the diffraction spots. A virtual beam stop was used to block the direct beam. Therefore, the middle area of the simulated projections records no signals from the sample and is thus black (intensity is zero) in the simulations (see Fig. 1a). A total of 181 projections were computed for a full sample rotation of 360° with a step size of 2° . There are 625 ± 22 spots present in each simulated projection.

Experimental LabDCT diffraction images were used to test the trained DL network capability of removing of high frequency noise and image artifacts for an enhanced identification of real diffraction spots. For this, a silicon sample was characterized using a Zeiss Xradia 520 Versa X-ray tomography instrument equipped with the LabDCT module. The measurement was performed in the Laue focusing condition ($L_{ss} = L_{sd} = 13 \text{ mm}$) with an acceleration voltage of 80 kV and an exposure time of 30 s. The grain reconstruction shows that the average grain size is about $125 \mu\text{m}$.

3 Method

The DL method learns from examples of input images x and transforms the images into desired output images $\hat{y}$. As mentioned above, every simulated LabDCT image serve as a content image y_c and the input image x is generated by adding simple constant background noise into the simulated image to vaguely mimick the background noise of the experimental LabDCT images. After training, the DL network is then used to remove the noise from experimental images.

3.1 Synthetic input data

The synthetic input data x (see Fig. 1b) is made by increasing the brightness values of any dark pixels with values less than a certain value, α , to a brighter value of Δ for all the 181 content images y_c . Namely, each pixel value in the input image $x_{(i,j)}$ is created in the following way from the corresponding pixel

Fig. 2 An illustration of the transformation network on the left and an example of the loss network $\phi(z)$ of the type VGG16 on the right. The j different layers used in the loss network is illustrated. The figure is adapted from [9].

value in the content image $y_{c(i,j)}$:

$$x_{(i,j)} = \begin{cases} \Delta, & \text{if } y_{c(i,j)} < \alpha, \Delta \sim U[25, 225] \\ y_{c(i,j)}, & \text{if } y_{c(i,j)} \geq \alpha \end{cases}, \quad (1)$$

where $\alpha = 5$, i denotes the row and j the column. The pixel value of Δ is sampled from a discrete uniform distribution (denoted as U) in the range from 25 to 225.

3.2 DL algorithm

The pipeline of the DL routine is illustrated in Fig. 2. The routine consists of two components: an image transformation network $f_w(x) = \hat{y}$ and a loss network $\phi(z)$ where $z = \{\hat{y}, y_c\}$, i.e. two networks as in [9]. The transformation network is trained to generate desired output images by mapping $f_w(x) = \hat{y}$ using synthetic input images (x). Here the U-net is selected for the transformation network (the first network) [8].

The loss network VGG16 [10] (the second network) is used in the loss function and this network remains unchanged during the training process. The loss network gets inputs from both the transformation network output (predicted $\hat{y}$) and the content image (y_c). The loss is calculated as differences between the outputs at different layers. Three outputs are used in this study ($j = \{0, 1, 2\}$) where j denotes the number of the output in the loss function and not the respective layer position in the loss network, see Fig. 2(right). The comparison of the layers in the loss consists of both the L1-norm and the Gram loss (denoted as the function g). A description of the Gram loss can be

found in [9].

$$loss(\hat{y}, y_c) = \frac{|\hat{y} - y_c|}{CHW} + \sum_{j=0}^2 \left(\frac{1}{C_j H_j W_j} |\phi_j(\hat{y}) - \phi_j(y_c)| + g(\phi_j(\hat{y}), \phi_j(y_c)) \right), \quad (2)$$

where j is the j th output, C is the number of channels, H and W is the height and width of the image in pixels, respectively.

The implementation of this DL routine is based on the fast.ai library's Jupyter notebook in the Python programming language [15] and [16]. The training time is roughly 5-6 hours on a moderately powerful desktop pc with a NVIDIA GeForce RTX 2080 with 8 GB of RAM using the CUDA software library. After the model is trained, it takes a few seconds to remove the noise and artifacts for a regular LabDCT data set consisting of 181 images.

3.3 DL network training

Due to a limited available capacity of GPU RAM, the network input and output are both decreased from $W \times H = 1015 \times 1015$ to $W \times H = 500 \times 500$ with a training batch size of 1.

The optimal learning rate is first found with a learning rate finder $lr = 0.01$. Then the network is first trained with frozen layers in the down-sampling (untrained) and the parameters in the up-sampling are changed during training of the network for the first 10 epochs.

Then the whole network is un-frozen and both the down-sampling and up-sampling parameters are updated during training. The network is trained until the validation loss stops decreasing. The training loss is the value of the loss function for images that have been used for training. Here back-propagation is used on the training data, to update the parameters. Validation loss is calculated using images that the network has not been trained on (unseen images). The data split consist of the training data using 90 % of the images (163 images) versus validation data using 10 % of the images (18 images). The more detailed training steps are described in [17].

4 Results

4.1 Training the network

The loss should be investigated while training and it is desired that the loss function for the validation loss becomes as small as possible. The loss function values are shown in Fig. 3 for the "un-frozen" training where all the loss values decrease from epoch 0 to epoch 9. At the 10th epoch, the validation loss stops decreasing and starts to increase, therefore the training is stopped here. The loss curve is typically not perfectly smooth due to the complex optimization

Fig. 3 The loss values as a function of the epoch number.

problem of the many parameters changed during training in the convolutional neural network.

An example of the removal of the constant background noise is shown in Fig. 1. The noise in the synthetic input image (x) Fig. 1b is nicely cleaned in the transformed image ($\hat{y}$) Fig. 1c, which is thus very similar to the content image (y_c) (i.e. the ground-truth) shown in Fig. 1a.

4.2 Model application

After the DL network is trained, the critical question is how well it performs on real experimental data. An example of the results is shown in Fig. 4. The figure shows that close to the beam stop, the high frequency noise in the real LabDCT image (Fig. 4a) is mostly removed together with the ‘block’ (which is due to the beam stop) in the middle of the DL reconstructed image (Fig. 4b). Some high frequency noise (speckles) still remain at the four corners in Fig. 4b. This is likely due to the low frequency radial gradients in brightness towards the four corners of the image, this type of low frequency radial gradients in brightness is not used in the training data. Since, the constant “noise” (Δ) used for the input image (x) during the training does not contain low frequency radial gradients in brightness or high frequency noise as well as texts remain (see the bottom left corner for the red text in Fig. 4a). As a consequence, this low frequency radial gradients in brightness together with the red text is not learnt to be removed by the DL algorithm. Some weak diffraction spots become clearly visible in the reconstructed image. The subtraction of the two images by the arithmetic $x - \hat{y}$ is shown in Fig. 4c, which illustrates that the reconstructed image has preserved most of the spots, especially the relatively bright ones. A few medium bright spots located close to the beam stop can still be completely or partially seen in Fig. 4c.

Fig. 4 Application example: (a) An experimental LabDCT diffraction image (x) (b) the model output image ($\hat{y}$) and (c) the image is obtained by $(x - \hat{y})$ showing the absolute differences between the two.

To evaluate more quantitatively the efficiency of the DL routines to remove noise from experimental labDCT images, three line scans of pixel values in a selected regions of interest (ROI) of the input and the output images are performed. The positions of the three line scans can be seen in Fig. 5a and b). It should be noted that in order to enhance the visibility of the output image the brightness and contrast are "artificially enhanced" in Fig. 5b as compared to the original output image in Fig. 4b. However, the values reported in the lines scans in Fig. 5c), d) and e) are the real pixel values originating from Fig. 4b.

Line A, going from the top left corner towards the center, shows that the low frequency radial gradients in brightness at the corner is decreased in the DL output image although it still remains high (see Fig. 5c). Notably, the pixel values are getting smaller "radially" along line A in both the input and output image. A horizontal line scan (line B) going through regions with small low frequency radial gradients in brightness, is shown in Fig. 5d. Here it is clearly seen that the high frequency noise is almost cleaned to zero while the sharp peaks of diffraction spots are preserved. More strikingly, the weak spots, among which one example is marked by the green arrow in Fig. 5a and d, are also nicely retrieved and the peak is distinctly visible. This shows that the trained DL network performs very well for regions with small radial gradients in brightness. A third scan (line C) going through some bright spots (located close to the outer edge) as well as several very weak spots (located close to the beam stop) is performed. The spot marked by the blue arrow in Fig. 5a and e is retrieved less well, although the spot is distinctly visible in Fig. 5b. All together this illustrates that after DL cleaning most of the weak spots are reasonable well preserved.

From the quantitative line scan comparison, we learn that the DL network performs very well with respect to removing high frequency noise around spots located in regions with relatively small low frequency radial gradients in brightness. Concerning removing high frequency noise around the spots located in regions with a large low frequency radial gradients in brightness, it performs

Fig. 5 Quantitative line scan profiles showing the performance of the DL network for removing noise. (a) a region of interest (ROI) cropped from Fig. 4a and (b) the corresponding ROI of the output image cropped from Fig. 4b and both are visualized in the grey value range of [0-255]. (c-e) show grey value profiles of lines A, B and C in (a) and (b), respectively. Spatial positions have been normalized. Arrows in (a) and (e) highlight a weak spot close to the beam stop.

less well. Since we did not apply low frequency radial gradients in brightness of the pixel values when synthesising the input image (x) provided for the training, our trained DL network naturally performs less well for such regions with image artifacts. This, however, points out the direction to further improve the DL network - to include low frequency radial gradients in brightness as close to that in the real experimental LabDCT images as possible for the training. Nevertheless, the results from the current DL network have shown great potential for removing high frequency noise around the spots in LabDCT images and have set the future direction.

5 Conclusion and Outlook

A DL method has been developed for removing high frequency noise from experimental LabDCT diffraction images. Based on training with idealized forward simulated LabDCT images and a simple constant background noise as

input, the removal of high frequency noise in real experimental labDCT images by the trained model is surprisingly efficient. It is shown that the complex and spatially varying background in the experimental images is replaced by an essentially zero background level, which makes reconstruction including even weak spots more straightforward. Some low frequency radial gradients in the pixel brightness values towards the corners of the experimental images still remain after the DL cleaning as the DL routines are not yet trained to identify this type of gradients. The combination of the per pixel loss local focus combined with the network loss contextual understanding of visual features, makes it an appealing method for this type of image reconstruction problem addressing the different types of high frequency noise and low frequency radial gradients in brightness occurring locally in the samples (per pixel loss) and the beam block larger size (network loss). The principle of using forward simulation for creating training data, so no time demanding human labelling is needed, makes this DL routine fully automatic and an appealing framework.

Acknowledgements This work is financially supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (M4D – grant agreement No. 788567).

References

1. H. Poulsen, S. Fæster Nielsen, E. Lauridsen, S. Schmidt, R. Suter, U. Lienert, L. Margulies, T. Lorentzen, D. Juul Jensen, Three-dimensional maps of grain boundaries and the stress state of individual grains in polycrystals and powders, *Journal of Applied Crystallography* 34 (2001) 751–756.
2. J. Tischler, W. Liu, R. Xu. Micro-diffraction in 3D. *Proceedings of 40th Risø International Symposium on Materials Science: Metal Microstructures in 2D, 3D and 4D*, 2019, pp. 329–334.
3. Y. Zhang Quantification of local boundary migration in 2D/3D. *IOP Conference Series: Materials Science and Engineering*, 2019, 580, 012015.
4. A. Bhattacharya, Y. F. Shen, C. M. Hefferan, S. F. Li, J. Lind, R. M. Suter, G. S. Rohrer, Three-dimensional observations of grain volume changes during annealing of polycrystalline ni, *Acta Materialia* 167 (2019) 40–50.
5. C. M. Hefferan, J. Lind, S. F. Li, U. Lienert, A. D. Rollett, R. M. Suter, Observation of recovery and recrystallization in high-purity aluminum measured with forward modeling analysis of high-energy diffraction microscopy, *Acta Materialia* 60 (10) (2012) 4311–4318.
6. F. Bachmann, H. Bale, N. Gueninchault, C. Holzner, E. M. Lauridsen, 3D grain reconstruction from laboratory diffraction contrast tomography, *Journal of Applied Crystallography* 52 (3) (2019) 643–651.
7. J. Oddershede, J. Sun, N. Gueninchault, F. Bachmann, H. Bale, C. Holzner, E. Lauridsen, Non-destructive characterization of polycrystalline materials in 3d by laboratory diffraction contrast tomography, *Integrating Materials and Manufacturing Innovation* 8 (2) (2019) 217–225.
8. O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. 2015.
9. J. Johnson, A. Alahi, and F. F. Li. Perceptual losses for real-time style transfer and super-resolution. *Arxiv*, 2016.
10. K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. 2015.
11. H. Fang, D. Juul Jensen, Y. Zhang, An efficient method to improve the spatial resolution of laboratory x-ray diffraction contrast tomography, *I O P Conference Series: Materials Science and Engineering* 580 (1) (2019) 012030

-
12. <https://github.com/haixingfang/LabDCT-forward-simu-model>
 13. J. M. Boone, J. A. Seibert, An accurate method for computer-generating tungsten anode x-ray spectra from 30 to 140 kv, *Medical Physics* 24 (11) (1997) 1661–1670.
 14. H. Fang, D. Juul Jensen, Y. Zhang, A flexible and standalone forward simulation model for laboratory X-ray diffraction contrast tomography, *Acta Crystallographica, A* 76 (2020).
 15. <https://github.com/fastai/course-v3/blob/master/nbs/dl1/lesson7-superres.ipynb>
 16. <https://course.fast.ai/videos/?lesson=7>
 17. <https://github.com/hiromis/notes/blob/master/Lesson7.md>