Tomography of Materials and Structures 7 (2025) 100046

Contents lists available at ScienceDirect

Tomography of Materials and Structures

journal homepage: www.journals.elsevier.com/tomography-of-materials-and-structures

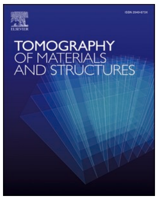

Self-supervised resolution enhancement for anisotropic volumes in edge illumination X-ray phase contrast micro-computed tomography

Jiayang Shi [a,*], Louisa Brown [b], Amir R. Zekavat [b], Daniël M. Pelt [a,1], Charlotte K. Hagen [b,1]

[a] *Leiden Institute of Advanced Computer Science, Leiden University, Einsteinweg 55, Leiden 2333 CC, the Netherlands*
[b] *Department of Medical Physics and Biomedical Engineering, University College London, Malet Place, London WC1E 6BT, United Kingdom*

ARTICLE INFO

Keywords:
Self-supervised learning
Resolution enhancement
Super-resolution
Micro-computed tomography
X-ray phase-contrast imaging

ABSTRACT

X-ray phase contrast micro-computed tomography (micro-CT) can achieve higher contrast than conventional absorption-based X-ray micro-CT by utilizing refraction in addition to attenuation. In this work, we focus on a specific X-ray phase contrast technique, edge illumination (EI) micro-CT. EI uses a sample mask with transmitting apertures that split the X-ray beam into narrow beamlets, enabling detection of refraction-included intensity variations. Between the typical mask designs (circular and slit-shaped apertures), slit-shaped apertures offer practical advantages over circular ones, as they only require sample stepping in one direction, thereby reducing scanning time. However, this leads to anisotropic resolution, as the slit-shaped apertures enhances resolution only along the direction orthogonal to the slits. To address this limitation, we propose a self-supervised method that trains on high-resolution in-plane images to enhance resolution for out-of-plane images, effectively mitigating anisotropy. Our results on both simulated and real EI micro-CT datasets demonstrate the effectiveness of the proposed method.

1. Introduction

X-ray micro-computed tomography (micro-CT) allows analyzing the interior structure of centimeter scale samples in a non-destructive manner [1]. Images are formed by measuring the attenuation of X-ray photons by matter, which leads to high contrast for highly attenuating materials but produces poor contrast for samples composed of materials with weak or similar attenuation properties, which are therefore challenging to visualize [2]. X-ray phase contrast micro-CT describes a class of methods by which images are formed from X-ray refraction alongside attenuation, enhancing contrast and bringing significant benefit for those challenging samples [3]. One of them is the edge illumination (EI) method [4], which uses a sample mask with transmitting apertures to split the X-ray beam into fine beamlets and a matching detector mask to create an array of edges in front of the pixels as shown in Fig. 1. When positioning the sample mask such that half of each beamlet impinges on a pixel while the other half impinges on the detector mask, refraction can be detected, as this causes the beamlets to slightly change direction, leading to an increase or decrease in the intensity measured in each pixel.

When working with the EI method, an important consideration is the design of the sample mask, as this has implications on the spatial resolution in the reconstructed tomographic volumes [5]. In conventional micro-CT, spatial resolution is governed by the source and detector, each introducing a blurring to the images that render features below a certain scale unresolvable [6]. By contrast, in EI, resolution is governed by the width of the sample mask apertures, provided that the spacing between beamlets is sufficiently large to minimize any overlap between them, and that they are not blurred as part of the detection process. This then allows enhancing resolution beyond the limit imposed by the source/detector blurring by using sufficiently small apertures[7]. However, this comes at the cost of extended scan times, given that the apertures restrict the x-ray flux, requiring longer exposures. Furthermore, since only the sample areas traversed by beamlets can contribute to the image, the sample must be stepped across the beamlet array in small increments at each rotation angle to obtain fully sampled datasets [8]. Common sample mask designs feature circular or slit-shaped apertures (Fig. 1). Circular apertures enhance resolution in an isotropic fashion but also imply a greater flux reduction, and require that the sample is stepped along two orthogonal directions at each rotation angle, increasing scan times. Their usage may therefore be considered impractical in certain applications. Slit-shaped apertures transmit more flux and reduce the

---

* Corresponding author.
E-mail address: j.shi@liacs.leidenuniv.nl (J. Shi).
[1] These authors have contributed equally.

https://doi.org/10.1016/j.tmater.2024.100046
Received 14 November 2024; Received in revised form 12 December 2024; Accepted 20 December 2024

need for sample stepping to only one direction, resulting in faster scans and making this design more practical. As a rough estimate, scans with slit-shaped apertures reduce scan times by at least a factor of *t* compared to circular apertures, where *t* is the number of steps required to step the sample across the beamlet array along one direction. However, resolution is only enhanced along the direction orthogonal to the slits, resulting in tomographic volumes with anisotropic resolution, featuring a higher in-plane (axial) resolution and a lower out-of-plane (sagittal, coronal) resolution. Another distinction between circular and slit-shaped apertures is that the latter are sensitive to refraction along one direction (orthogonal to the slits), whereas circular apertures are omnidirectionally sensitive, which could justify their use despite the longer scan times involved. To overcome the resolution anisotropy in volumes acquired with slit-shaped apertures, deconvolution may be applied to reduce the blur along the out-of-plane direction but this is prone to introducing artifacts [9,10], particularly near edges. More recently, supervised deep learning approaches have shown promise in improving resolution [11–15], but they require high-quality, isotropic reference images for training. This poses a challenge in EI micro-CT, as samples are not usually scanned in sufficiently large numbers to build up training sets. Furthermore, samples would need to be scanned with slit-shaped apertures and circular apertures to obtain matching anisotropic and isotropic volumes, which is extremely time-consuming and may even be infeasible for samples that cannot be kept stable over long time periods.

We present a self-supervised deep learning framework for enhancing the out-of-plane resolution anisotropic EI micro-CT volumes acquired with slit-shaped apertures as shown in Fig. 2. Our approach is based on the idea of training a deep neural network on the higher-resolution in-plane images and applying it to the lower-resolution out-of-plane images [16–18]. This method builds on our previous work for medical helical CT [18] by accurately modelling resolution differences to simulate training images that reflect the anisotropic nature of the data, incorporating the system's point spread functions (PSF) in both the in-plane and out-of-plane directions. Using simulated datasets, we show that the proposed method outperforms conventional deconvolution approaches and other deep learning methods both quantitatively and visually, providing reliable resolution enhancement without hallucinating extraneous features. Furthermore, we demonstrate our method's effectiveness by applying it to experimentally acquired EI micro-CT datasets. Our results demonstrate that isotropic tomographic volumes can be obtained in EI micro-CT while also retaining the faster scan times enabled by using slit-shaped apertures in the sample mask.

2. Method

2.1. Problem statement

We denote the reconstructed tomographic volume obtained with EI micro-CT as $\mathbf{I}(x,y,z) \in \mathbb{R}^{X\times Y\times Z}$, where X,Y,Z represent the number of pixels along the x, y, z axes. The volume can be visualized in three primary orientations: axial, coronal, and sagittal. We denote the axial, coronal, and sagittal slices at positions z', y', and x' as $\mathbf{a}_{z'} = \mathbf{I}(:,:,z')$, $\mathbf{c}_{y'} = \mathbf{I}(:,y',:)$, and $\mathbf{s}_{x'} = \mathbf{I}(x',:,:)$. The voxel size of the volume is represented as $r^x \times r^y \times r^z$, where, for EI micro-CT with slit-shaped apertures in the sample mask: $r^x = r^y = r^{xy} < r^z$. That is, there is anisotropy in the voxel dimensions, with a lower voxel resolution in the z axis (out-of-plane) compared to the xy plane (in-plane). Additionally, we define the PSFs along the in-plane and out-of-plane directions as 1D normalized zero-mean Gaussian functions, where $\mathscr{P}_{xy} = \frac{1}{\sqrt{2\pi\sigma_{xy}^2}}\mathscr{N}(0;\sigma_{xy}^2)$ is applied along both the x and y directions in the xy-plane, while $\mathscr{P}_z = \frac{1}{\sqrt{2\pi\sigma_z^2}}\mathscr{N}(0;\sigma_z^2)$ is applied along the z direction. We choose $\sigma_z > \sigma_{xy}$ to model the more extensive out-of-plane blurring due to the source and detector and the less extensive in-plane blurring due to the sample mask apertures. The choice of Gaussian models for the PSFs has arisen from experimental observations made for EI micro-CT; edge-based measurements of the horizontal and vertical system PSF have both produced approximately Gaussian-shaped curves, albeit with a different standard deviation.

Starting from a high-resolution "clean" volume $\mathbf{I}$, blurring can be introduced by the sequential application of the 1D PSFs along each axis.

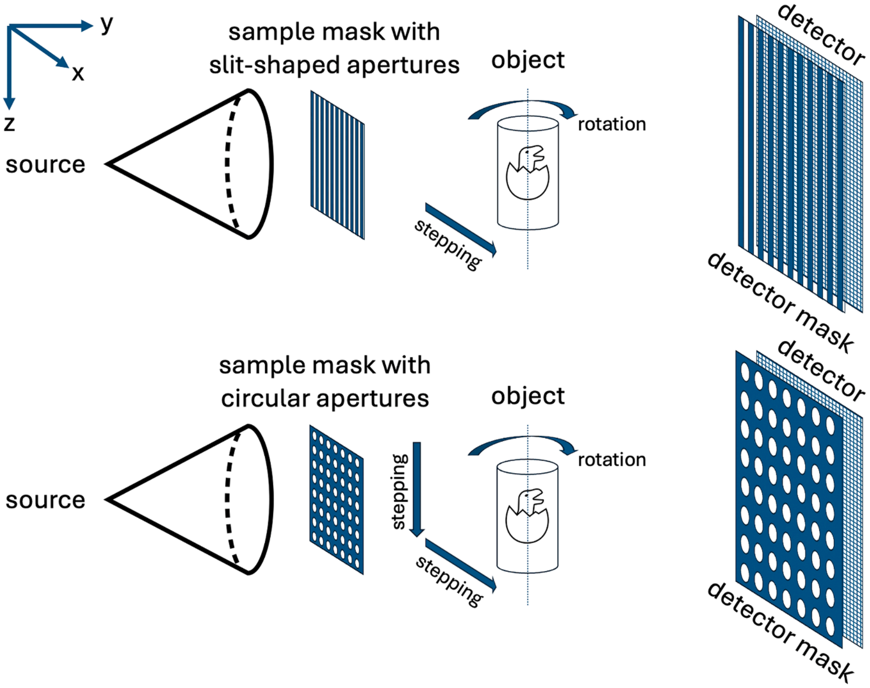

Fig. 1. Schematic of the EI system setup. The sample mask can have either circular or slit-shaped apertures. Circular apertures provide isotropic resolution but require two-dimensional stepping, resulting in long scanning times, which may be impractical for some applications. In contrast, slit-shaped apertures require translation in only one direction, significantly reducing scanning time. However, this approach yields anisotropic resolution.

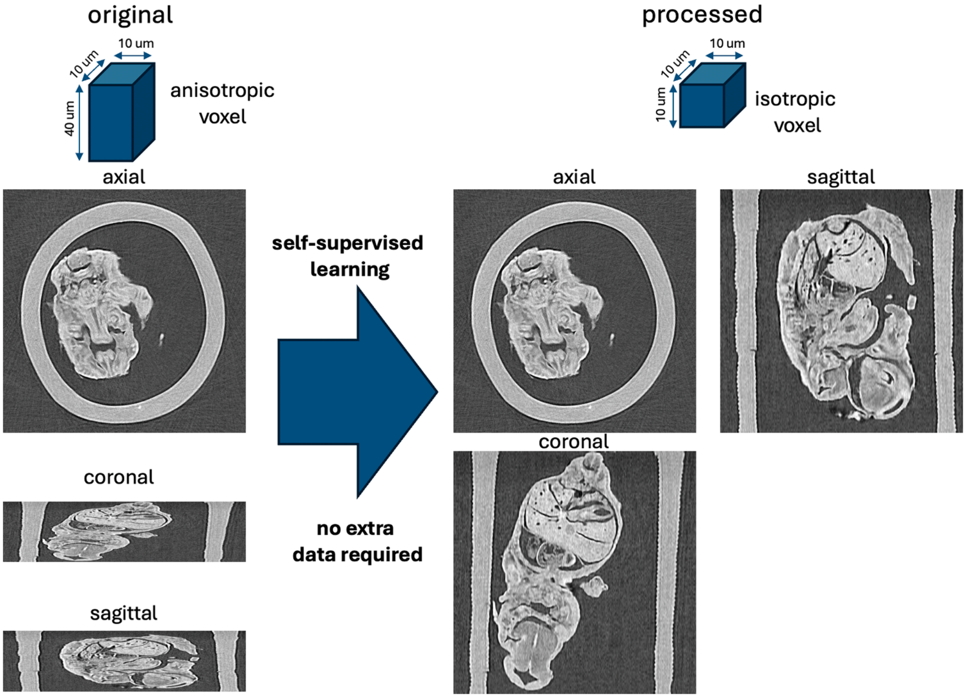

Fig. 2. Problem statement. When EI micro-CT is performed with a sample mask with slit-shaped apertures, the reconstructed volumes exhibit anisotropic resolution, with different voxel sizes and point spread functions (PSFs) (varying blur levels) between the in-plane and out-of-plane directions. Our self-supervised learning approach leverages the high-resolution in-plane images to generate training data, enabling improved out-of-plane resolution. The processed images feature isotropic resolution without the need to acquire data with a sample mask with circular apertures, which is time-consuming.

In the xy-plane, blurring is applied by convolving both rows and columns of each slice with $\mathscr{P}_{xy}$, while in the z-direction, each column along the z-axis is convolved with $\mathscr{P}_z$. This leads to the blurred volume $\widehat{\mathbf{I}}$, where $\widehat{\mathbf{I}} = \mathbf{I} * \mathscr{P}$, with $\mathscr{P}$ being the 3D anisotropic Gaussian PSFs, representing the anisotropic volumes reconstructed from EI micro-CT data.

2.2. Training and inference

Our proposed method simulates training images that resemble the lower voxel resolution and blurrier out-of-plane images by generating degraded versions of the high-resolution in-plane images. These simulated images are then used to train a neural network, which can subsequently be applied to enhance the resolution of coronal and sagittal images, as illustrated in Fig. 3.

To achieve a blur level in the in-plane images that matches the out-of-plane blur, the in-plane images are convolved with a 1D Gaussian PSF, defined as $\mathscr{P}_{dif} = \frac{1}{\sqrt{2\pi\sigma_{dif}^2}}\mathscr{N}(0;\sigma_{dif}^2)$, where $\sigma_{dif} = \sqrt{\sigma_z^2 - \sigma_{xy}^2}$. For each original in-plane image $\mathbf{a}$:

- Convolving every row with $\mathscr{P}_{dif}$ produces the horizontally blurred image $\widehat{\mathbf{a}}^{hor}$;
- Convolving every column with $\mathscr{P}_{dif}$ produces the vertically blurred image $\widehat{\mathbf{a}}^{ver}$.

To simulate the effect of reduced voxel resolution, we define downsampling and upsampling functions that mimic the degradation caused by lower resolution along the off-slice direction. The row-wise downsampling function, $\mathscr{F}^{hor}_{\downarrow}$ intakes an image with resolution $r^{xy} \times r^{xy}$ and downscales it along rows to $r^{xy} \times r^{z}$, effectively reducing the pixel count in the horizontal direction. Similarly, the upsampling function $\mathscr{F}^{hor}_{\uparrow}$ rescales the image back to its original resolution by linearly upsampling the rows. This down- and upsampling process simulates the degradation effect of reduced resolution while maintaining the original image size. Analogously, the column-wise downsampling and upsampling functions, $\mathscr{F}^{\mathrm{ver}}_{\downarrow}$ and $\mathscr{F}^{\mathrm{ver}}\uparrow$, operate along the vertical direction. For each original axial image $\mathbf{a}$, we obtain:

- A horizontally degraded image $\mathscr{F}^{hor}_{\uparrow}(\mathscr{F}^{hor}_{\downarrow}(\widehat{\mathbf{a}}^{hor}))$;
- A vertically degraded image: $\mathscr{F}^{ver}_{\uparrow}(\mathscr{F}^{ver}_{\downarrow}(\widehat{\mathbf{a}}^{ver}))$.

During training, these horizontally and vertically degraded images, generated from all Z axial slices along the z axis, are used to train a neural network f_θ to map the degraded images back to the original axial images. Since resolution enhancement always takes place along the z axis (the vertical axis for coronal and sagittal images), we define a rotation function $\mathscr{R}$ to rotate the horizontally degraded image by 90°, ensuring that degradation consistently aligns with the vertical direction. The weights θ of the neural network are optimized by minimizing the following objective function:

$$\theta* = \min_{\theta}\sum_{i=1}^{Z} L[f_\theta(\mathscr{F}^{ver}_{\uparrow}(\mathscr{F}^{ver}_{\downarrow}(\widehat{\mathbf{a}}^{ver})),\mathbf{a}] + L[\mathscr{R}(f_\theta(\mathscr{F}^{hor}_{\uparrow}(\mathscr{F}^{hor}_{\downarrow}(\widehat{\mathbf{a}}^{hor}))),\mathscr{R}(\mathbf{a})], \tag{1}$$

where L is a loss function that computes the difference between the network output to the target image.

After training, the network has learned the mapping from a degraded image with voxel resolution r^z and PSF $\mathscr{P}_z$ along z axis to an image with r^{xy} and PSF $\mathscr{P}_{xy}$. With the assumption that tomographic volumes exhibit similar features across orientations, we apply this learned mapping to enhance the resolution of coronal and sagittal images. For inference, the coronal and sagittal images are first linearly upsampled with $\mathscr{F}^{ver}_{\uparrow}$, then the trained network f_θ is applied to obtain images with enhanced resolution and reduced blur. Empirically, we find that combining the enhanced coronal and sagittal volumes yields no additional improvement in image quality. Consequently, we use only the improved coronal images as the final output.

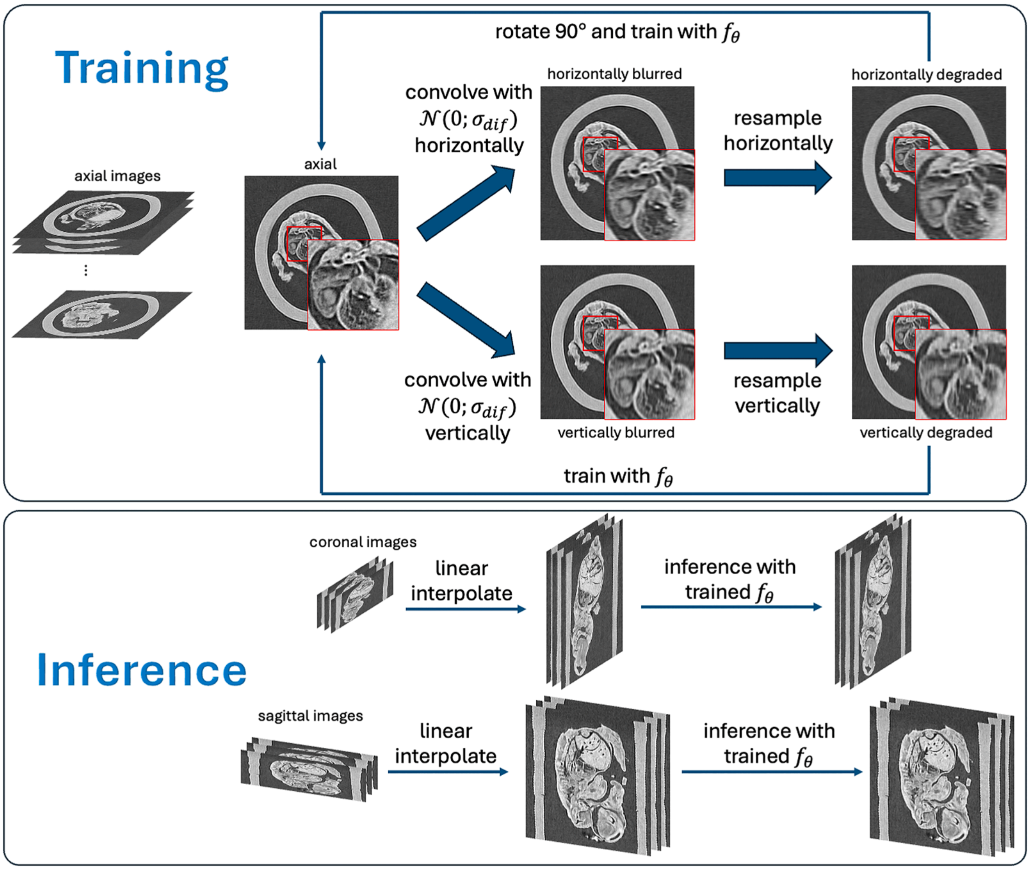

Fig. 3. Training and inference in our method. A neural network is trained on purposefully degraded in-plane images to enhance resolution the out-of-plane images. To degrade the in-plane images, they are convolving with a PSF representing the difference between the in-plane and out-of-plane blur, followed by resampling to simulate the reduced voxel resolution. For inference, the trained network is applied to linearly upsampled coronal and sagittal images, producing final outputs with improved voxel resolution and reduced blur.

Any image-to-image neural network can be used as f_θ. In our experiments, we employ a 2D Mixed Scale Dense (MS-D) network [19] with 100 layers, training it for 200 epochs using an L2 loss function and the ADAM optimizer [20].

3. Results and discussion

3.1. Datasets

We evaluate our method on three simulated datasets and three real EI micro-CT datasets acquired with slit-shaped apertures in the sample mask. We initially use isotropic tomographic volumes to simulate EI micro-CT characteristics by applying anisotropic PSFs ($\sigma_{xy} = 0.25$, $\sigma_z = 2.5$) along different orientations, as described in Section 2.1. The volumes are then resampled along the z axis, yielding an anisotropic voxel size where the voxel resolution along the z axis is four times worse than in the xy plane. We compare the resolution-enhanced results against the original isotropic high-resolution volumes.

- **Foam Phantom:** Using the `foam_ct_phantom` package [21], we simulate a 512 × 512 × 512 voxels cylindrical foam phantom with 100,000 non-overlapping bubbles of various sizes.
- **LoDoInd:** From the LoDoInd dataset [22], we use a reference tube sample comprising 15 materials (e.g., coriander, pine nuts), providing complex, heterogeneous content. We crop and downsample the middle portion of the tube to 512 × 512 × 512 voxels.
- **Lung:** We select a human lung volume from *Task06 Lung* in the Medical Segmentation Decathlon dataset [23], resampling it to isotropic resolution with a shape of 404 × 512 × 512 voxels.

Real data were acquired using a custom EI micro-CT scanner [5]. The sample mask had slit-shaped apertures (10 μm wide), while the pixel size of the detector was 50 μm (40 μm when demagnified to the sample plane). The reconstructed volumes therefore had voxel sizes of 10 × 10 × 40 μm^3. The PSFs describing the blur were modeled as Gaussian functions with $\sigma_{xy} = 0.39$ and $\sigma_z = 1.21$, based on edge response function measurements, resulting in a calculated $\sigma_{dif} = 1.14$. However, we found empirically that a slightly larger σ_{dif} improved results, likely due to minor inaccuracies in the edge response function measurements such as edge imperfections, inclinations, or noise. Therefore, we used $\sigma_{dif} = 1.5$ for our experiments.

- **Sponge:** A piece of household sponge in a 3 mm diameter plastic straw was scanned with 450 projections, 8 steps per projection, and a 1.5 s exposure per step. Data were flat and dark field corrected, underwent single-image phase retrieval [24], and were subsequently reconstructed into tomographic images via Filtered Back Projection (FBP). The reconstructed volume contained 225 × 400 × 400 voxels.
- **Mouse Embryo:** A deceased and ethanol dehydrated mouse embryo (wild type, 14.5-day gestation) generated as surplus during research [25] performed under the regulation of the UK Animals (Scientific Procedures) Act 1986 and the National Centre for the 3Rs' Responsibility in the Use of Animals for Medical Research (2019) was scanned with 1200 projections, 8 steps per projection, and a 1.2 s exposure per step. The tomographic volume, reconstructed in the same way as the sponge data, was cropped to a size of 128 × 512 × 512 voxels. In addition, we scanned the same mouse embryo using a sample mask with finer apertures (5 μm wide), with 2400 projections, 16 steps per projection, and a 2 s exposure per step. The

tomographic volume, again reconstructed as described above, had a voxel size of 5 × 5 × 40 μm^3 and was cropped to a size of 143 × 1024 × 1024 voxels.

3.2. Comparison with other methods

To validate the effectiveness of the proposed method, we compare it with the following classical and deep learning-based methods: Linear Upsampling (out-of-plane images are linearly upsampled along the z axis), Wiener Deconvolution using PSF with $\sigma_{dif} = 2.49$ for simulated data and $\sigma_{dif} = 1.5$ for real data (following linear upsampling, Wiener deconvolution [26] is applied along the z axis to address differences in blur levels), Lucy-Richardson Deconvolution using PSF with $\sigma_{dif} = 2.49$ for simulated data and $\sigma_{dif} = 1.5$ for real data (similarly, Lucy-Richardson deconvolution [10] is applied along the z axis), SAINT (convolutional neural network-based method for CT image super-resolution with a fixed upsampling factor of 4 [11]), and PLHR (a vision transformer-based approach for CT image super-resolution with a fixed upsampling factor of 5 [13]; to enable comparison, the outputs of RPLHR were downsampled by a factor of 1.25). As discussed in Section 1, acquiring isotropic high-resolution EI micro-CT volumes, which are required as reference data for training supervised learning models, is often challenging and time-consuming. Consequently, we did not retrain these methods and instead used the authors' pre-trained weights, originally trained on human lung CT images. This may introduce domain-shift effects when applied to other datasets in this work. Additionally, the traditional methods were applied with default parameters, which could potentially introduce a bias favoring the more optimized deep learning-based approaches.

3.3. Results

Simulated datasets. The results of using our proposed method, as well as the various methods of comparison, to increase the out-of-plane resolution in the degraded simulated datasets are shown in Fig. 4. It can be seen that linear upsampling results in images that almost entirely retain their blurriness, which is expected. While Wiener and Lucy-Richardson deconvolution appear effective in deblurring, they also introduce artifacts which is undesirable. SAINT appears to fail to adequately deblur the image, and RPLHR introduces visible patch-based artifacts. These arise because RPLHR processes images in patches due to the high computational demands of the transformer model, limiting its ability to handle entire images seamlessly. In comparison, our proposed method produces deblurred images without introducing hallucinated features or other unwanted artifacts. We also quantify the performance of the various methods in Table 1, showing that our method consistently yields the highest Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) [27] values across all datasets. The sub-optimal performance of SAINT and RPLHR may be explained by the domain shift between training and testing data.

Real EI micro-CT datasets. Results obtained for the real EI micro-CT data are shown in Fig. 5. Due to the difficulty in acquiring isotropic EI micro-CT images (requiring the use of a sample mask with circular apertures and involving excessive scan times), ground truth images are not available in this case. This also means that quantitative metrics (PSNR, SSIM) cannot be extracted. However, visually the results are consistent with those observed on the simulated datasets. Linear upsampling results in blurry images, deconvolution methods introduce new artifacts likely due to their sensitivity to slightly erroneous PSF estimates, and SAINT outputs noticeably blurred images. RPLHR produces better visual results than on the simulated datasets but the results also exhibits a degree of hallucination, with generated features that visibly diverge from the true image structures. By contrast, our method again delivers high-resolution out-of-plane images with significantly reduced blur while preserving the original image features, even for the challenging case of an initial 8 × resolution mismatch.

Computation cost. The computational efficiency of the proposed method was assessed using the mouse embryo dataset with the size of 128 × 512 × 512 voxels. Our approach simulates 256 pairs of training

Table 1
Comparison of various resolution enhancement methods, including our proposed approach, applied to the described simulated volumes. Each entry shows the Peak Signal to Noise Ratio (PSNR)/Structural Similarity Index Measure (SSIM) [27] metrics calculated relative to the ground truth volumes, with the highest values highlighted in bold.

Dataset	Linear	Wiener	Lucy-Richardson	SAINT	RPLHR	Ours
Lung	33.59/ 0.90	35.80/ 0.92	36.77/0.94	33.53/ 0.88	32.10/ 0.87	**40.20/ 0.97**
LoDoInd	32.20/ 0.83	33.37/ 0.86	33.86/0.88	31.88/ 0.82	27.00/ 0.70	**35.20/ 0.91**
Foam	18.50/ 0.67	19.23/ 0.66	20.47/0.77	18.18/ 0.64	17.73/ 0.61	**24.84/ 0.82**

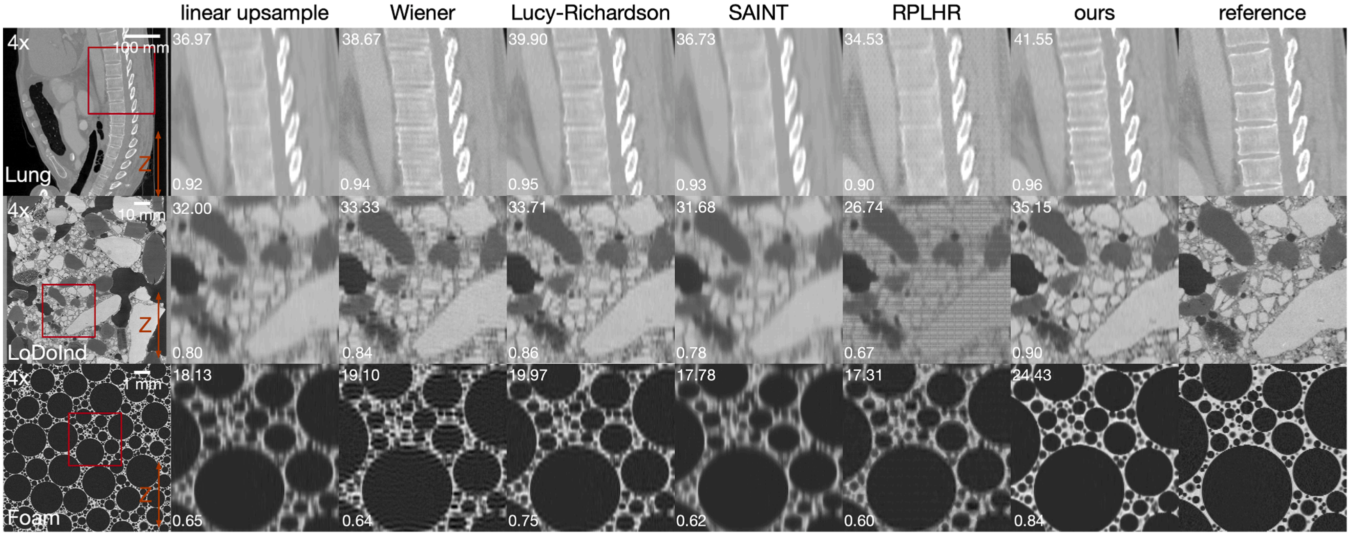

Fig. 4. Visual comparison of various resolution enhancement methods, including our proposed approach, on tomographic volumes with simulated non-isotropic resolution. PSNR and SSIM values for each cropped image are displayed in the upper-left and lower-left corners, respectively. The locations of the selected patches are indicated in the first column from the left. The labels 4 × in the upper-left corners denote the resolution enhancement factor along the off-plane direction.

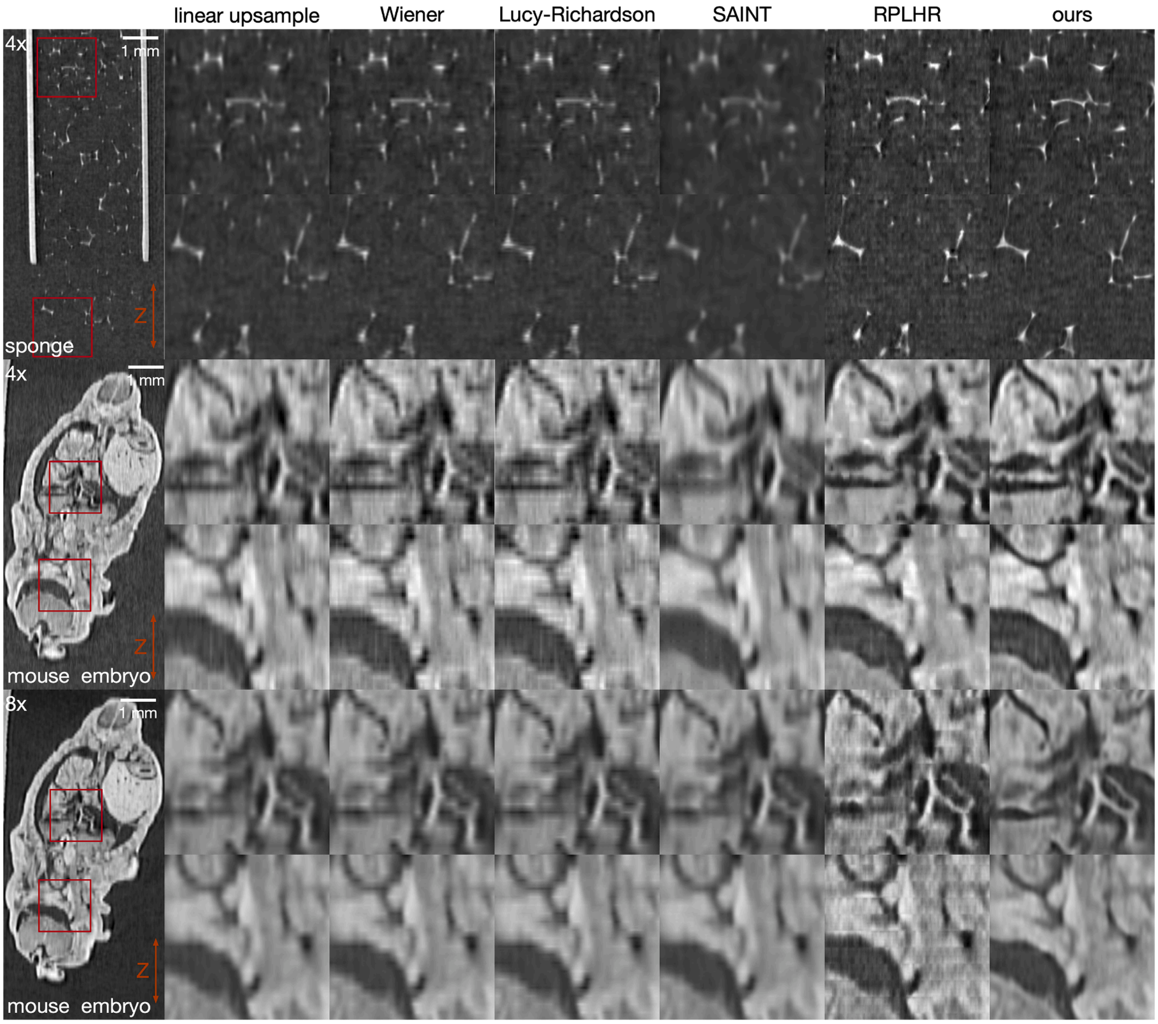

Fig. 5. Visual comparison of various resolution enhancement methods, including our proposed approach, on real EI micro-CT images. The red insets indicate the locations of the enlarged patches. The labels 4 × and 8 × in the upper-left corners denote the resolution enhancement factor along the off-plane direction. Note that, although the same mouse embryo was scanned at different resolutions, slight positional shifts of the sample components can occur between scans. As a result, precise registration is challenging, so the images only display patches from approximately similar locations, with some differences in image features.

and target images. Training with the depth-100 MSD network requires approximately 39 s per epoch. For the default 200 training epochs, the total training time amounts to approximately 130 min. Once the network is trained, the inference process for generating the final resolution-enhanced CT volume takes approximately 20 s. All computations were performed on a workstation equipped with an Intel i7-11700KF CPU and an Nvidia RTX 4080 GPU.

4. Conclusion

We have proposed a self-supervised framework for enhancing EI micro-CT data by leveraging the high-resolution, in-plane images to improve the lower-resolution, out-of-plane images. Our approach simulates training images by accurately modelling differences in voxel resolution and blurring through orientation-specific PSFs, allowing the network to learn and correct anisotropic effects inherent to EI micro-CT data acquired with slit-shaped apertures in the sample mask. Our method's performance was evaluated on simulated data and real EI micro-CT datasets, and compared to other methods for resolution enhancement. Our approach consistently produced the highest-quality results, reducing blur while retaining image features without introducing artifacts.

However, our method relies on the assumption that images across different orientations share similar features in EI micro-CT. When significant differences exist between orientations, the performance of our approach may degrade. For instance, orientation-specific artifacts, such as ring artifacts in in-plane images, could be inadvertently introduced into out-of-plane images. Additionally, while our self-supervised approach addresses the challenges of scanning isotropic CT data with circular apertures in the sampling mask, it requires separate training for each anisotropic volume. Although the training time is relatively short compared to the time saved in scanning, future work could focus on improving computational efficiency to further reduce overhead.

Our self-supervised framework enables the acquisition of high-resolution, isotropic volumes with EI micro-CT using slit-shaped apertures in the sample mask, which would normally only be possible with circular apertures and thus incurring long scan times. We believe that the framework will benefit applications of EI micro-CT for which both a isotropic resolution and fast scans are required.

CRediT authorship contribution statement

Charlotte K. Hagen: Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Daniel M. Pelt:** Writing – review & editing, Writing – original draft, Supervision, Software, Resources, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Amir R. Zekavat:** Data curation, Conceptualization. **Louisa Brown:** Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jiayang Shi:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by an EPSRC New Investigator Award (EP/T029080/1). Jiayang Shi and Daniël M. Pelt are supported by the European Union H2020-MSCA-ITN-2020 under Grant Agreement No. 956172 (xCTing). Charlotte K. Hagen is supported by the Royal Academy of Engineering under the Research Fellowship scheme. The authors would like to thank Dr. Gabriel Galea (UCL Institute of Child Health) and co-workers for the preparation of the mouse embryo sample, and UCL's Advanced X-Ray Imaging (AXIm) group for their support in acquiring the experimental datasets.

References

[1] P.C. Hansen, J. Jørgensen, W.R. Lionheart, Computed Tomography: Algorithms, Insight, and Just Enough Theory, SIAM, 2021.
[2] P.J. Withers, C. Bouman, S. Carmignato, V. Cnudde, D. Grimaldi, C.K. Hagen, E. Maire, M. Manley, A. DuPlessis, S.R. Stock, X-ray computed tomography, Nat. Rev. Methods Prim. 1 (2021) 18.
[3] S. Wilkins, Y.I. Nesterets, T. Gureyev, S. Mayo, A. Pogany, A. Stevenson, On the evolution and relative merits of hard x-ray phase-contrast imaging methods, Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci. 372 (2014) 20130021.
[4] A. Olivo, Edge-illumination x-ray phase-contrast imaging, J. Phys.: Condens. Matter 33 (2021) 363002.
[5] A.R. Zekavat, G. Lioliou, O.R. i Morgó, C.M. Jones, G. Galea, E. Maniou, A. Doherty, M. Endrizzi, A. Astolfo, A. Olivo, et al., Phase contrast micro-ct with adjustable in-slice spatial resolution at constant magnification, Phys. Med. Biol. 69 (2024) 105017.
[6] T.M. Buzug, Computed tomography, in: Springer Handbook of Medical Technology, Springer, 2011, pp. 311–342.
[7] P.C. Diemoz, F.A. Vittoria, A. Olivo, Spatial resolution of edge illumination x-ray phase-contrast imaging, Opt. Express 22 (2014) 15514–15529.
[8] C.K. Hagen, F.A. Vittoria, M. Endrizzi, A. Olivo, Theoretical framework for spatial resolution in edge-illumination x-ray tomography, Phys. Rev. Appl. 10 (2018) 054050.
[9] F. Orieux, J.-F. Giovannelli, T. Rodet, Bayesian estimation of regularization and point spread function parameters for Wiener–Hunt deconvolution, JOSA A 27 (2010) 1593–1607.
[10] W.H. Richardson, Bayesian-based iterative method of image restoration, JoSA 62 (1972) 55–59.
[11] C. Peng, W.-A. Lin, H. Liao, R. Chellappa, S.K. Zhou, Saint: spatially aware interpolation network for medical slice synthesis, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 7750–9.
[12] Q. Liu, Z. Zhou, F. Liu, X. Fang, Y. Yu, Y. Wang, Multi-stream progressive up-sampling network for dense CT image reconstruction, in: Proceedings of the Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23, Springer, 2020, pp. 518–28.
[13] P. Yu, H. Zhang, H. Kang, W. Tang, C.W. Arnold, R. Zhang, RPLHR-CT dataset and transformer baseline for volumetric super-resolution from CT scans, in: Proceedings of the Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VI, Springer, 2022, 344-353.
[14] C. Fang, L. Wang, D. Zhang, J. Xu, Y. Yuan, J. Han, Incremental cross-view mutual distillation for self-supervised medical ct synthesis, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 20677–86.
[15] J. Shi, D.M. Pelt, K.J. Batenburg, Multi-stage deep learning artifact reduction for computed tomography, arXiv preprint arXiv:2309.00494, 2023.
[16] C. Zhao, B.E. Dewey, D.L. Pham, P.A. Calabresi, D.S. Reich, J.L. Prince, Smore: a self-supervised anti-aliasing and super-resolution algorithm for mri using deep learning, IEEE Trans. Med. Imaging 40 (2020) 805–817.
[17] H. Xie, Y. Lei, T. Wang, Z. Tian, J. Roper, J.D. Bradley, W.J. Curran, X. Tang, T. Liu, X. Yang, High through-plane resolution ct imaging with self-supervised deep learning, Phys. Med. Biol. 66 (2021) 145013.
[18] J. Shi, D.M. Pelt, K.J. Batenburg, SR4ZCT: self-supervised through-plane resolution enhancement for CT images with arbitrary resolution and overlap, in: Proceedings of the International Workshop on Machine Learning in Medical Imaging, Springer, 2023, pp. 52–61.
[19] D.M. Pelt, J.A. Sethian, A mixed-scale dense convolutional neural network for image analysis, Proc. Natl. Acad. Sci. 115 (2018) 254–259.
[20] D.P. Kingma, A method for stochastic optimization, in: Y. Bengio, Y. LeCun, (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings, 2015.⟨http://arxiv.org/abs/1412.6980⟩.
[21] D.M. Pelt, A.A. Hendriksen, K.J. Batenburg, Foam-like phantoms for comparing tomography algorithms, J. Synchrotron Radiat. 29 (2022).
[22] J. Shi, O. Elkilany, A. Fischer, A. Suppes, D. Pelt, K. Batenburg, Lodoind: introducing a benchmark low-dose industrial ct dataset and enhancing denoising with 2.5D deep learning techniques, in: Proceedings of the 13th Conference on Industrial Computed Tomography (iCT) 2023, School of Engineering, Wels Campus, Austria, 2024. ⟨https://doi.org/10.58286/29228⟩.
[23] A.L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B.A. Landman, G. Litjens, B. Menze, et al., A large annotated medical image dataset for the development and evaluation of segmentation algorithms, arXiv preprint arXiv:1902.09063, 2019.
[24] P.C. Diémoz, F.A. Vittoria, C.K. Hagen, M. Endrizzi, P. Coan, A. Bravin, U. H. Wagner, C. Rau, I.K. Robinson, A. Olivo, A single-image retrieval method for edge illumination x-ray phase-contrast imaging: application and noise analysis, Phys. Med. 32 (2016) 1759–1764.
[25] E. Maniou, F. Farah, A.R. Marshall, Z. Crane-Smith, A. Krstevski, A. Stathopoulou, N.D. Greene, A.J. Copp, G.L. Galea, Caudal FGFR1 disruption produces localised spinal mis-patterning and a terminal myelocystocele-like phenotype in mice, Development 150 (2023).
[26] N. Wiener, Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications, The MIT Press, 1949.
[27] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Image Process. 13 (2004) 600–612.