

A Framework for the AI-based Visualization and Analysis of Massive Amounts of 4D Tomography Data for End Users of Beamlines

Steffen Kieß¹, Thomas Lang², Tomas Sauer^{2,3}, A. Michael Stock³, Andrei Chernov³, Yipeng Sun⁴, Andreas Maier⁴, Tomáš Faragó⁵, Alexey Ershov⁵, Gabriel Lefloch⁵, Guilherme Silva⁶, Tilo Baumbach^{5,6}, Simon Zabler², Astrid Hölzing², Kilian Dremel², Ali Riza Durmaz⁷, Akhil Thomas⁷, Ingo Manke⁸, Nikolay Kardjilov⁸, Tobias Arlt⁸, Tak Ming Wong⁹, Regine Willumeit-Römer⁹, Julian Moosmann⁹, Berit Zeller-Plumhoff⁹, Dieter Froning¹⁰, Sven Simon¹

¹University of Stuttgart, ITI/CIS, Universitätsstraße 38, 70569 Stuttgart, Germany, e-mail: steffen.kiess@cis.iti.uni-stuttgart.de

²Fraunhofer Institute of Integrated Circuits IIS, Am Wolfsmantel 33, 91058 Erlangen, Germany, email: thomas.lang@iis.fraunhofer.de

³University of Passau, Innstraße 43, 94032 Passau, Germany, email: tomas.sauer@uni-passau.de

⁴FAU Erlangen-Nürnberg, PRL, Martensstr. 3, 91058 Erlangen, Germany, email: yipeng.sun@fau.de

⁵Institute for Photon Science and Synchrotron Radiation, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany, email: tomas.farago@kit.edu

⁶Laboratory for Applications of Synchrotron Radiation, Karlsruhe Institute of Technology, Kaiserstrasse 12, 76131 Karlsruhe, Germany, email: tilo.baumbach@kit.edu

⁷Fraunhofer Institute for Mechanics of Materials IWM, Wöhlerstrasse 11, 79108 Freiburg, Germany, email: ali.riza.durmaz@iwm.fraunhofer.de

⁸Helmholtz-Zentrum Berlin, Hahn-Meitner-Platz 1, 14109 Berlin, Germany, email: manke@helmholtz-berlin.de

⁹Helmholtz-Zentrum Hereon, Max-Planck-Straße 1, 21502 Geesthacht, Germany, email: tak.wong@hereon.de

¹⁰Forschungszentrum Jülich, Wilhelm-Johnen-Straße, 52428 Jülich, Germany, email: d.froning@fz-juelich.de

Abstract

The size of 4D tomography datasets acquired at synchrotron or neutron imaging facilities can reach several terabytes, which presents a significant challenge for their evaluation. This paper presents a framework that allows a compressed dataset to be kept in memory and makes it possible to evaluate and manipulate the dataset without requiring enough memory to decompress the entire dataset. The framework enables the compensation of imaging artifacts, including the compression artifacts of the 4D dataset, through the integration of neural networks. The reduction of imaging artifacts can be performed at the imaging facility or at the user's home institution. This framework reduces the computational burden on the computing infrastructure of large synchrotron and neutron facilities by allowing end users to process datasets on their institution's computers. This is made possible by compressing TBs of data to less than 128 GB, allowing powerful PCs to process TBs of 4D tomography data.

Keywords: 4D computed tomography, CT, synchrotron imaging, neutron imaging, neural network

1 Introduction

The 4D tomography at synchrotron sources and neutron sources is extremely data-intensive, as it is based on time series of 3D volume datasets with a storage space of e. g. 50 GB for each time step, so that the entire dataset of a 4D tomography scan reaches the one-digit TB range. These data volumes will continue to increase due to increasing temporal and spatial resolution. Currently, there is no software available, whether open source or commercial, that is capable of handling (visualizing, analyzing, and evaluating) such large 4D datasets on standard computers (not computing clusters). As a result, the potential of 4D tomography for science and industry has not yet been fully exploited. Thus far, individual users have developed individual solutions, which are often of limited quality or scope, for their specific applications. These solutions are not transferable for use by third parties, for example, due to software quality concerns. The results of recent years in the field of artificial intelligence have demonstrated the considerable potential of neural networks, particularly deep convolutional neural networks (CNN), for the improvement of tomography datasets. This is true even for large datasets, provided that suitable learning data is available. Image artifacts occur during computed tomography (CT) scans, e. g. due to beam variations, and for 4D CT in particular, many types of artifacts occur more frequently due to the shorter exposure times, the movement during the scan or residual phase contrast, which occurs due to longer propagation distances caused by the sample environments used. To date, there are isolated solutions that attempt to reduce these artifacts, but these are generally not made available to the end user. In most cases, the end user returns to their institute after a measurement with reconstructed data and is then dependent on the image quality made available. Consequently, a framework for 4D tomography data will be presented here as a contribution to the digital infrastructure for synchrotron and neutron sources. More information about the framework can be found at [1].

2 AI-based imaging methods for integration into the framework for 4D tomography and results

The authors are currently engaged in an ongoing project addressing various aspects of 4D tomography and developing a framework for handling 4D tomography data. This chapter provides an overview of the preliminary results of the project and the AI-based imaging methods and applications that are being addressed.

2.1 Compensation of 4D-CT noise and motion artifacts

Compensating for noise and motion artifacts in 4D-CT is increasingly critical in industrial applications such as automotive safety systems, and aerospace inspections. This can be effectively achieved through the deployment of hybrid machine learning techniques that seamlessly integrate established analytical operators with advanced deep neural network architectures [2]. By embedding proven motion-compensation formulas and mathematically robust reconstruction algorithms within these deep learning frameworks, industries can significantly enhance image quality, reduce ambiguities, and efficiently correct complex temporal variations without the need for extensive training datasets. These integrated approaches not only improve structural accuracy, clarity, and contrast but also minimize distortions caused by dynamic movements and inherent scanning noise, resulting in higher fidelity volumetric images that consistently capture intricate details. Recent industry studies [2] demonstrate that such hybrid models streamline the learning process, offering more reliable and computationally efficient outcomes. This provides a stable and adaptable platform for precise 4D-CT image reconstructions, essential for applications requiring high accuracy and reliability. Additionally, by incorporating well-established imaging operators into meticulously designed deep architectures, these hybrid solutions ensure robust generalization across diverse operational scenarios. The combination of theoretical rigor and adaptive learning mechanisms creates a powerful framework where prior knowledge and data-driven insights collaboratively enhance the representation of structural and temporal patterns. Consequently, industries benefit from refined motion artifact corrections and a training environment resilient to overfitting and random noise, enabling the handling of varied imaging geometries, scanning protocols, and specific physiological characteristics. This translates to greater confidence for engineers in detecting subtle abnormalities or changes over time. Furthermore, recent advancements highlight that strategically selected regularization techniques and the integration of domain-specific knowledge can further bolster these methods [3], expanding their applicability and fostering continuous innovation in merging classical reconstruction theories with deep learning approaches.

2.2 Generation of 4D CT datasets using simulation

Syris [4] simulates X-ray wavefields generated by various synchrotron radiation sources. Objects are described either analytically or by triangular meshes and their interaction with the wavefield is given by the transmission function [5]. The free-space propagation between objects is modeled by the scalar diffraction theory [6] and the detection process including various noise sources follows the EMVA 1288 standard [7]. Figure 1 shows an example of a simulation at different stages of the image formation. All elements of the imaging process are time-dependent, enabling dynamic experiments such as time-resolved tomography. Written in Python, syris provides a flexible, modular interface, while computationally demanding tasks are handled by OpenCL for fast GPU-based calculations. This combination of flexibility and speed helps to optimize imaging methods and measurement conditions, thus improving the synchrotron beam time efficiency. Accurate simulations are also essential for the development of image processing algorithms, especially for AI-based processing of large 4D datasets. Figure 2 shows a comparison of 3D reconstruction quality with different imaging artifacts.




Figure 1: Simulation of a capillary filled with spheres, projected thickness (left), X-ray projection with noise (middle), X-ray projection with phase contrast, noise and scintillator spots (right).




Figure 2: slice through the 3D reconstructed volume of the capillary in Figure 1, no artifacts (left), noise present (middle), noise and ring artifact stemming from scintillator spots (right).

Working with 4D data introduces intrinsically higher memory and computational demands. In our case, the time required to sample the thickness of time-resolved meshes along a wavefront increases significantly. To address this, we leverage a real-time, GPU-optimized data acceleration structure called the Linear Bounding Volume Hierarchy (LBVH) [8] alongside an enhanced Ray/Triangle intersection algorithm [9]. The LBVH dramatically reduces computation times, achieving speedups of up to 100 times. Meanwhile, the improved Ray/Triangle algorithm enhances accuracy by virtually eliminating false positives and negatives. It also optimizes memory usage and transfer overhead by adaptively employing higher precision only when necessary, ensuring optimal performance without compromising accuracy. With a fast and accurate sample thickness computation, datasets with individual isolated artifacts can be generated at first in order to disentangle the common 4D imaging problems. Subsequently, datasets with all artifacts combined are created, which pose a significant challenge and can in turn serve to benchmark and improve the 4D image processing and analysis algorithms.

2.3 Micro-CT digital volume correlation & volumetric optical flow

To track deformations and strains in order to elucidate morphology-function relationships of tomography volumetric data, an optical flow neural network is developed for digital volume correlation (DVC) analysis. This network, named VolRAFT [10], estimates the dense 3D displacement field between the reference volume and the deformed volume, using a supervised training approach on synthetic (see Figure 3) and measured datasets (see Figure 4). Experiments show that VolRAFT performs well in estimating different displacement fields compared to the cutting-edge variational DVC approach [11] for high-resolution tomography data of bone-implant interfaces measured at synchrotron radiation-based computed tomography beamline.

Meanwhile, the volumetric optical flow study focuses on the computation and visualization of 4D optical flow between two 3D time steps. In our application, we consider a scan of a zinc-air battery during its discharge cycle. The 3D scans were processed to estimate the motion vectors within the volumetric data over time. The implementation leverages the Total Variation L1 (TV-L1, [12]) optical flow method, extended to three spatial dimensions (X, Y, Z) and one temporal dimension (T), to capture the displacement vector field. For enhanced efficiency, this method is applied at each level of a multiresolution image pyramid. This methodology enables detailed tracking and visualization of material motion within tomographic data, offering insights into the internal dynamics of an object. The approach combines robust preprocessing, accurate alignment, and high-fidelity visualization techniques, demonstrating its utility in materials science and related domains. The visualization step is utilized to render the 3D




Figure 3: Evaluation of the synthetic displacement field based on the measured volume and the synthetically deformed volume. The specimen consists of a magnesium-5%-gadolinium alloy and surrounding bone. The field component in the direction of the screw's symmetry axis is shown. Adapted from [10].




Figure 4: Evaluation of the displacement field based on measured volume pairs, comparing to the variational optical flow method. The specimen consists of a titanium screw and surrounding bone. The field component in the direction of the screw’s symmetry axis is shown. Adapted from [10].




Figure 5: The evolution of a Zink-air battery neutron scan at time t (a) towards the state $t + \Delta t$ after some time (c) is efficiently captured by a dense volumetric optical flow computation (b), confirming the general intention of the movement. The length of the depicted arrows indicate the magnitude of the movement

optical flow field. Figure 5a shows the volumetric data of the considered battery at some point in time t and its evolution during a discharging (Figure 5c). The optical flow field Figure 5b captures the implicit movement in the material efficiently and confirms in a visually appealing manner the intuition about the movement. Moreover, the length of the depicted arrows indicate the magnitude of change, i.e. the speed with which the contained structures are moving.

2.4 Interactive data-driven segmentation

Due to the large variability of samples and scanning protocols at synchrotron facilities, there is a general lack of fully annotated training data for obtaining robust deep learning segmentation methods. Therefore, we propose the extension of an interactive segmentation procedure based on active learning [13]. In the latter, the volumetric data is visualized to the user, e.g., by displaying slice images, and the user is prompted to annotate individual voxels. In our scheme, voxels belonging to the material/component that shall be segmented are annotated with a positive label by a left mouse click (cf. blue crosses in Figure 6a) whereas voxels belonging to other categories are marked with a negative label by a right mouse click (cf. red crosses in Figure 6a). From local regions around these voxels, feature vectors are extracted and a sparse probabilistic classification method is trained [14].




Figure 6: Active learning based segmentation of a Zink-air battery scan. The user provides ground truth labels interactively (a) from which a segmentation model is trained, yielding a probabilistic classification into particles (b).

Hyperparameter tuning is performed using a parallel grid search approach. Applying the trained model on the volumetric dataset (conceptually for all voxels in parallel, neglecting padding areas), yields a probability field indicating the probability of a voxel belonging to the desired material/component. The active learning component of this workflow then utilizes the human-in-the-loop approach where the user can decide if the segmentation is good enough. If not, the user can augment the set of annotated voxels interactively and run the algorithm again. This workflow enables the segmentation of large volumetric datasets without any a priori annotated training data and in a relatively fast manner. We demonstrate this workflow in Figure 6 on a neutron tomography scan of a Zink-air battery [15]: As described, a user annotated interactively relatively few voxels by either a positive or negative label. In this case, a positive label (blue cross) identifies voxels belonging to particles while negative labels (red crosses) identify voxels not being particles (Figure 6a), the zoomed region shows this more clearly. After executing the algorithm, probabilistic segmentation is produced (Figure 6b) in which each voxel is colored by the probability of it being a particle. The result data may then be further postprocessed to yield the final segmentation.

2.5 Compression for archiving of 4D tomography data

At high resolution, scanning even a medium-sized sample results in very large datasets, necessitating efficient compression methods. Based on the three-dimensional discrete wavelet transform (dwt3d) [16], data-adaptive and error-bound schemes with better generalizability and control of compression quality can be realized compared to state-of-the-art methods, e.g., compressing each slice of the dataset individually using 2D image compression. Our starting point for 4D compression is temporal interpolation between two successive time steps (“frames”) by applying a 3D (Haar) wavelet transform to each frame and compress the time steps individually. While interpolation is possible in both the standard basis and the wavelet basis, locally varying levels of detail can result in high-resolution noise-like artifacts. Hence, we follow a different approach in the presented 4D compression prototype and select *keyframes* as every n -th frame, e.g., $n=2,4,6$, and predict the frames in between by locally using data from one of the keyframes, inspired by video encodings (lifted to 3D). To that end, denote by j the one-dimensional index identifying the current time step and by i the three-dimensional spatial index of any voxel of one time step. Any voxel of the original 4D data is thus identified by $x_i^{(j)}$, resulting in the corresponding compressed and regularized 3D frame in the Haar wavelet basis $\text{dwt3d}_{\text{reg}}(x_i^{(j)})$, and its according decompressed frame $y_i^{(j)} = \text{idwt3d}(\text{dwt3d}_{\text{reg}}(x_i^{(j)}))$. As regularization method, we use the heuristic sparsification method presented in [17] such that a high peak signal-to-noise ratio (PSNR) is retained while producing a strong compression. For $j=kn$, where k is some natural number, the frames at time j are independently 3D-compressed. In the other cases, e.g., for $n=2$, we now consider the keyframes $x_i^{(j)}$ and $x_i^{(j+2)}$ to compress the frame $x_i^{(j+1)}$. To exclude time-independent reconstruction noise, we use the regularized frame data $y_i^{(j)}$ or $y_i^{(j+2)}$ as prediction $p_i^{(j+1)}$ to predict the $(j+1)$ -th frame at coordinate i , and we thus store the compressed and regularized prediction error $\text{dwt3d}_{\text{reg}}(x_i^{(j+1)} - p_i^{(j+1)})$ in its wavelet representation. A keyframe at time (j) or $(j+2)$ is locally selected as prediction reference for the voxel at coordinate i if the variance of a local environment of size $3 \times 3 \times 3$ voxels centered around i in one keyframe is smaller than in the other keyframe. Intuitively, this prefers more homogeneous, i.e., less noisy regions. As an example application of this compression procedure,

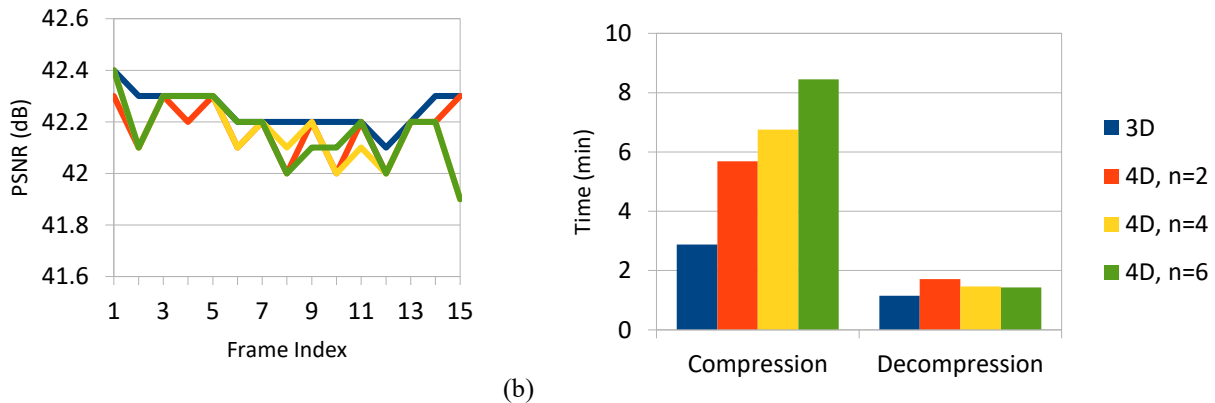


Figure 7: (a) PSNR values after decompression for 3D time steps in the Zinc-air battery scan after independent (“3D”) and 4D compression with varying distance n between keyframes. (b) Total compression and decompression times (in minutes) for the tested algorithms applied to all 15 Zinc-air battery scan frames. The right legend applies to both plots in (a) and (b).

Algorithm	3D	4D, $n=2$	4D, $n=4$	4D, $n=6$
Compressed file size (MiB)	63.2	54.9	59.2	62.2
Compression ratio	63.4	73.0	67.7	64.4

Table 1: Compressed file sizes in MiB and compression ratios relative to uncompressed data when the compressed data is persisted using the specified compression algorithm.

we consider a Zinc-air battery scan [15] (512x512x267 voxels, 32-bit floating-point) with a total number of 15 frames, yielding input data of roughly 4 GiB in size. Figure 7a shows the PSNR values of the independent 3D and keyframe-based 4D approaches after decompression, whereas Figure 7b illustrates the cost of replacing more keyframes by predicted frames. Although the 4D methods require more computational time, there are no significant differences in image quality. Yet, there are measurable improvements when considering the size of the compressed files. Table 1 shows the file sizes of the compressed data and compression ratios for the aforementioned algorithms, indicating that the 4D method with $n=2$ is best in terms of file size and thus compression ratio, while the 3D method is clearly the fastest in this example. While the SCR file format [18] is used for the 3D compression, work on the SCR-based 4D prototype is still ongoing.

2.6 Online and real time decompression of 4D tomography datasets

To visualize a 4D dataset consisting of multiple timesteps with a 3D voxel dataset in each timestep, two possible strategies can be used. One possibility is to load the entire 4D dataset into memory, which makes it easy to visualize changes in volume. However, this approach is often not feasible due to the large amount of memory required. The other option is to load only one or a small number of 3D volumes into memory. However, this means that evaluations that should be performed at each time step have to be repeated several times, once for each time step. This process is often so time consuming that only a limited number of time steps are evaluated and the information in the 4D dataset is not fully exploited. To address this challenge, the authors have implemented a scheme where the data is kept in memory in a compressed form [19]. During visualization, for example, when a 2D slice through the volume is displayed, the data required for that slice is decompressed and displayed. The decompressed data is cached, so that the decompression is fast enough to work with the data in real time, i. e. as the slice is moved through the volume the display will still be updated several times per second, even though the data is decompressed on the fly.

The compression scheme used is JPEG, which was chosen because it is a well-known compression scheme and because of its fast compression and decompression. To allow decompression of only part of the volume, the entire volume is divided into blocks of 16x16x16 voxels. Each of these blocks then contains 16 slices of 2x2 JPEG blocks, for a total of 64 JPEG blocks to be compressed. The resulting bitstreams for all blocks are concatenated, but an additional table gives the position of the bitstream for each block, allowing decompression of individual blocks. Figure 8 shows the effect of compression with a compression ratio of 10 on the quality of the data. As can be seen from the profile line Figure 8d, the deviation is very small and does not cause any visible artifacts.

On an AMD Ryzen 9 7900X3D processor, it was possible to extract slices from the compressed data at between 5 MPixel/s and 10 MPixel/s, which is sufficient for interactive display of slices and smooth scrolling through the volume. The code for the voxel viewer can be found at [20].




Figure 8: Slice through a volume of a battery scanned on the synchrotron at the Helmholtz-Zentrum Berlin, (a) original, (b) compressed by a factor of 10, (c) profile line through (a) and (b), (d) magnified detail of (c).

2.7 Processed 4D synchrotron data for NN training

The quality of the segmentation/classification of the 4D data is highly dependent on the accuracy of the available training data from real samples. A multimodal, multidimensional characterization approach using a wider range of additional complementary imaging tools, such as electron microscopy or FIB/SEM tomography [21], is used for detailed quantification/identification of the material compositions in the samples or devices (e. g. batteries) analyzed by 4D synchrotron tomography by means of complex, time-consuming hand-made segmentation and "human-in-the-loop" techniques. This allows us to obtain reliable ground truth and training data for improved quantification accuracy and training of NNs. An important application of 4D tomography is the operando examination of batteries. Suitable datasets for algorithm development are identified and selected. As a rule, this data must first be processed (selection, formatting, new reconstruction if necessary) and converted into a suitable format. Ground truth data from real data is required to train the NNs, for example to identify deviations between the simulation-based training and the real data.




Figure 9: 4D tomography of a zinc-air battery during discharge; dissolution of the zinc particles [15].

A typical example for 3D data quantification is the analysis of active particles and their embedding in the surrounding matrix electrodes (see Figure 9). Various quantitative analyses are performed based on the fully reconstructed and segmented 4D data. The aim here is to demonstrate and verify the advantages of the concepts developed in this project on the basis of this use case and to compare them with conventional procedures with data that is massively affected by artifacts - as has been the case to date. Statistical analyses and comparisons, for example of the temporal structural and morphological changes as well as the particle size and particle shape distributions, can be used to quantitatively verify the concepts developed for the specific application and provide feedback for further optimization. For this reason, quantitative 4D analyses are also carried out here, in which the datasets are analyzed using typical complex 3D evaluation algorithms (such as particle and pore size distributions, tortuosity distributions, crack analysis, etc.).

The entire concept is also applied to 4D neutron tomography data, too. Datasets are also selected from the existing archives of the Helmholtz-Zentrum Berlin. Here too, the data will first be converted into a format suitable for processing within the scope of this project. Ground truth data will be extracted and used for training of the NNs as it was done for the synchrotron data.

2.8 Application of the Framework for material characterization: dynamic water content analysis.

Gas diffusion layers (GDLs) are slim components in fuel cells which are responsible for fluid transport between the gas channel structures on anode and cathode sides of the fuel cell and the electrochemical active layer between them. The GDL structure of paper-type material is based on fibers with a thickness of 7 to 10 μm , arranged almost planar. The typical thickness of such GDLs ranges from 200 to 300 μm , resulting in 26 fiber layers in the case of planar fiber arrangements, 7.5 μm fiber thickness and 195 μm GDL thickness. Figure 10a shows a 2D slice of a geometry model which is stochastic equivalent to the 3D structure. The real structure, however, looks like Figure 10b when taken from a Zeiss Xradia Versa nano-CT device, and like Figure 10c when taken from the BESSY synchrotron. The 2D images represent a layer of the resolutions' thickness of the 3D micro-structure of the GDL.

Froning et al. [22] developed a CNN to predict the permeability of porous material from 3D image series. The artificial data was taken to train a CNN to predict the permeability of a GDL from a 3D image set (like in Figure 10a). The CNN was applied to real data obtained from the BESSY synchrotron (Figure 10b) and from a Zeiss Xradia Versa nano-CT (Figure 10c). Although both kinds of real data images look somehow similar to human eyes with their image artifacts, the predictions of the CNN were more accurate for data from the BESSY synchrotron than for data from a Zeiss Xradia Versa nano-CT device.

The dynamic water content in electrolyzers at Forschungszentrum Jülich was analyzed by Colliard-Granero et al. [23] using deep learning, however presently based on 2D image series.




Figure 10: Image data representing a slice of Toray GDL; (a) geometry model, 512x512 image, resolution; 1.5 $\mu\text{m}/\text{px}$; (b) image of size 1250x1250 from BESSY synchrotron, resolution: 1.5 $\mu\text{m}/\text{px}$; (c) image of size 694x670 from a Zeiss Xradia Versa nano-CT, resolution: 1.797 $\mu\text{m}/\text{px}$. Source of the images: [22], license: CC-BY <https://creativecommons.org/licenses/by/4.0>.

Conclusion

The processing and evaluation of large 4D tomography datasets represents a significant challenge due to their considerable size. The proposed framework addresses several key aspects of working with 4D datasets. It provides support for the compensation of 4D-CT noise and motion artifacts, the generation of 4D CT datasets using simulation, digital volume correlation, compression with high compression ratios for archiving, and NN training based on 4D tomography data. It allows storing a compressed dataset in memory, making it possible to work with large 4D datasets without the need for extensive hardware resources. By allowing users to perform the evaluation on their institution's computers, it reduces the burden on the computing infrastructure of facilities that provide synchrotron and neutron tomography. More information about the framework can be found at <https://ki4d4e.eu/>.

Acknowledgements

This work was supported by the German Federal Ministry of Education and Research (BMBF) under the grant number 05D2022 and the title “Ein KI-basiertes Framework für die Visualisierung und Auswertung der massiven Datenmengen der 4D-Tomographie für Endanwender von Beamlines” (KI4D4E, “An AI-based framework for the visualization and analysis of massive amounts of 4D tomography data for end users of beamlines”).

References

- [1] “KI4D4E: An AI-based framework for 4D tomography”. <https://ki4d4e.eu/>
- [2] A. Maier, H. Köstler, M. Heisig, P. Krauss, S. H. Yang, “Known operator learning and hybrid machine learning in medical imaging—a review of the past, the present, and the future”, *Progress in Biomedical Engineering*, vol. 4, 2022, p. 022002, doi:10.1088/2516-1091/ac5b13
- [3] Y. Sun, Y. Huang, L. S. Schneider, M. Thies, M. Gu, S. Mei, S. Bayer, A. Maier (2024). “EAGLE: An Edge-Aware Gradient Localization Enhanced Loss for CT Image Reconstruction”. arXiv preprint. doi:10.48550/arXiv.2403.10695
- [4] T. Faragó, P. Mikulík, A. Ershov, M. Vogelgesang, D. Hänschke, T. Baumbach, “*syris*: a flexible and efficient framework for X-ray imaging experiments simulation”, *Journal of Synchrotron Radiation*, vol. 24, 2017, pp. 1283-1295, doi:10.1107/s1600577517012255
- [5] M. Born, E. Wolf. “Principles of optics: electromagnetic theory of propagation, interference and diffraction of light”, Elsevier, 2013. isbn:9781483103204
- [6] J. W. Goodman. “Introduction to Fourier optics”, Roberts and Company publishers, 2005. isbn:9780974707723
- [7] B. Jähne, “EMVA 1288 Standard for Machine Vision: Objective specification of vital camera data”, *Optik & Photonik*, vol. 5, 2010, pp. 53-54, doi:10.1002/opph.201190082
- [8] A. Prokopenko, D. Lebrun-Grandié. “Revising Apetrei’s Bounding Volume Hierarchy Construction Algorithm to Allow Stackless Traversal”, arXiv, 2024. doi:10.48550/arXiv.2402.00665
- [9] S. Woop, C. Benthin, I. Wald. “Watertight Ray/Triangle Intersection” *Journal of Computer Graphics Techniques (JCGT)*, vol. 2, no. 1, 65-82, 2013.
- [10] T. M. Wong, J. Moosmann, B. Zeller-Plumhoff, “VolRAFT: Volumetric Optical Flow Network for Digital Volume Correlation of Synchrotron Radiation-based Micro-CT Images of Bone-Implant Interfaces”, 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2024, pp. 53-62, doi:10.1109/cvprw63382.2024.00010
- [11] S. Bruns, D. Krüger, S. Galli, D. F. Wieland, J. U. Hammel, F. Beckmann, A. Wennerberg, R. Willumeit-Römer, B. Zeller-Plumhoff, J. Moosmann, “On the material dependency of peri-implant morphology and stability in healing bone”, *Bioactive Materials*, vol. 28, 2023, pp. 155-166, doi:10.1016/j.bioactmat.2023.05.006
- [12] C. Zach, T. Pock, H. Bischof, “A Duality Based Approach for Realtime TV-L 1 Optical Flow”, *Lecture Notes in Computer Science, Pattern Recognition*, doi:10.1007/978-3-540-74936-3_22
- [13] T. Lang, A. Heim, C. Heinzl. “Big Data Analytics for the Inspection of Battery Materials”, 13th Conference on Industrial Computed Tomography (iCT), 2024. doi:10.58286/29226
- [14] T. Lang, T. Sauer. “Geometric Active Learning for Segmentation of Large 3D Volumes”, ArXiv, 2022. doi:10.48550/arXiv.2210.06885
- [15] T. Arlt, D. Schröder, U. Krewer, I. Manke, “*In operando* monitoring of the state of charge and species distribution in zinc air batteries using X-ray tomography and model-based simulations”, *Phys. Chem. Chem. Phys.*, vol. 16, 2014, pp. 22273-22280, doi:10.1039/c4cp02878c
- [16] A. M. Stock, G. Herl, T. Sauer, J. Hiller, “Edge-preserving compression of CT scans using wavelets”, *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 62, 2020, pp. 345-351, doi:10.1784/insi.2020.62.6.345
- [17] T. Sauer, A. M. Stock. “Haar Wavelets, Gradients and Approximate TV Regularization”, arXiv preprint, 2022. doi:10.48550/arXiv.2208.05389
- [18] A. M. Stock, T. Lang, T. Sauer. “The SCR File Format”, 2023. doi:10.24406/publica-1214

- [19] S. Kieß, P. Gänz, G. Yang, X. Liu, C. Cardona Marín, S. Simon. "Visualization of 4D Tomography TB Datasets on Standard Computers with 128GB RAM", International Conference on Tomography of Materials and Structures, 2024.
- [20] "Voxie". <https://github.com/voxie-viewer/voxie>
- [21] M. Osenberg, A. Hilger, M. Neumann, A. Wagner, N. Bohn, J. R. Binder, V. Schmidt, J. Banhart, I. Manke, "Classification of FIB/SEM-tomography images for highly porous multiphase materials using random forest classifiers", Journal of Power Sources, vol. 570, 2023, p. 233030, doi:10.1016/j.jpowsour.2023.233030
- [22] D. Froning, E. Hoppe, R. Peters, "The Applicability of Machine Learning Methods to the Characterization of Fibrous Gas Diffusion Layers", Applied Sciences, vol. 13, 2023, p. 6981, doi:10.3390/app13126981
- [23] A. Colliard-Granero, K. A. Gompou, C. Rodenbücher, K. Malek, M. H. Eikerling, M. J. Eslamibidgoli, "Deep learning-enhanced characterization of bubble dynamics in proton exchange membrane water electrolyzers", Physical Chemistry Chemical Physics, vol. 26, 2024, pp. 14529-14537, doi:10.1039/d3cp05869g