

<https://doi.org/10.1038/s42005-025-02249-0>

Implicit neural representation for fast 4D computed tomography of multiphase flow in porous media

Check for updates

Henrik Friis ^{1,2,5}, Håkon Nese ^{2,3,5}, Giacomo Luani ¹, Colin Pryme ⁴, Lars Rennan⁴, Basab Chattopadhyay ¹, Anders Kristoffersen ⁴ , Ole Jakob Mengshoel ³ & Dag Werner Breiby ¹

Time-resolved X-ray computed tomography (4D-CT) enables dynamic processes within objects to be tracked over time. A key application of 4D-CT is the scientifically and societally important study of multiphase flow in porous media. Obtaining high-quality 4D-CT images is challenging because a high data acquisition rate must be combined with bespoke algorithms for faithful reconstruction. Here, we introduce NeCT, a physics-based deep learning model achieving state-of-the-art sparse-view and 4D-CT reconstructions. NeCT enables reconstruction of 4D objects using an implicit neural representation in space and time. With standard micro-CT instruments, NeCT achieves a temporal resolution approaching a few seconds. Reconstructions of benchmark static tomography datasets show that NeCT also outperforms established algorithms in 3D. We demonstrate fast imaging of liquid imbibition in sandstone, with the high spatiotemporal resolution allowing salt dissolution and pore-filling events to be directly observed. Finally, opportunities beyond classic CT with our dynamic continuous approach are discussed.

Multiphase flow dynamics within opaque and porous 3D materials or devices is a complicated and ubiquitous phenomenon in Nature. Such processes are key to understanding multiphase flow in porous biological and geological media^{1,2}, in the process industry, and in man-made devices such as batteries and fuel cells³. Micro-computed tomography (μ CT) based on X-ray radiation has revolutionized 3D imaging of microscopic structures, including porous media, during the last two decades. Dynamic or 4D (i.e., 3D + *time*) CT denotes imaging of internal structures in objects that change over time. In geosciences, 4D-CT is used to study a wide range of processes, including mechanical deformation of rocks^{1,4}, multiphase flow^{5,6}, the interaction between CO₂ and minerals under reservoir conditions^{7,8}, and drainage and imbibition in rocks⁹. CT reconstruction with high spatiotemporal resolution can give fundamental insights into pore-level events^{5,10}. The importance of CT in medicine is well known, and numerous efforts are made to reduce the radiation dose with shorter exposures while maintaining a high spatial resolution¹¹.

The analytical and non-iterative Feldkamp-Davis-Kress (FDK) algorithm¹² offers fast and accurate high-quality CT reconstructions if given sufficiently many projections, making it the standard workhorse for 3D CT image reconstruction¹³. Note that FDK requires the number of projections

to fulfill the Shannon-Nyquist sampling for faithful image reconstruction. Iterative methods, such as the “simultaneous algebraic reconstruction technique” (SART) algorithm¹⁴, are also increasingly used, as they facilitate the incorporation of a priori information, at the cost of a significantly increased computational burden. While these iterative methods can reconstruct sparse-view input with reduced artifacts, they unfortunately also tend to reduce the contrast in the reconstructed images. The naive and simplest approach to 4D-CT is to perform a series of independent subsequent 3D CT scans to create multiple time steps. The natural time limitation of this approach is equal to the time it takes to acquire each full tomogram, typically of the order of 1 h¹.

Several groups have during the last decades developed methods for fast CT with time resolution on the order of seconds to tens of seconds, using laboratory-based instruments^{6,15–21}. A common approach is to reduce the number of projections and compensate for the lacking information through prior information or regularization techniques. In several cases, a temporal resolution of about half a minute is reported, often at the expense of image artifacts and reduced spatial resolution.

Using 3rd and 4th generation synchrotron radiation sources²², the beam brilliance is 10–14 orders of magnitude higher than for laboratory

¹Department of Physics, Norwegian University of Science and Technology (NTNU), Trondheim, Norway. ²Aker BP ASA, Lysaker, Norway. ³Department of Computer Science, Norwegian University of Science and Technology (NTNU), Trondheim, Norway. ⁴Equinor ASA, Ranheim, Norway. ⁵These authors contributed equally: Henrik Friis, Håkon Nese. e-mail: ankris@equinor.com; ole.j.mengshoel@ntnu.no; dag.breiby@ntnu.no

sources, facilitating X-ray imaging, microscopy, and CT experiments with high spatial resolution (\sim nm)^{23,24}, short exposures²⁵, beam coherence opening for quantitative phase contrast²³, diffraction contrast²⁶, energy scanning, and nano-focusing. These developments enable subsecond temporal resolution with thousands of radiographic exposures recorded per second. Such sampling frequencies are not possible with current home laboratory technology, owing to photon flux limitations. To achieve 4D-CT with substantially higher time resolution than the second scale, also with synchrotron radiation, the fast-spinning sample approach becomes impractical, and other methods, such as repeated motion^{10,27}, must be invoked.

Machine learning (ML) is a promising way to reconstruct images from sparse-view input, thus enabling a higher temporal resolution. One approach is to use U-Nets²⁸ trained to perform a mapping from a subsampled artifact-rich FDK reconstruction to a final image with the characteristic artifacts removed²⁹. U-Nets can also be used for sinogram inpainting where the U-Net is used to learn the mapping from a subsampled sparse-view sinogram to a complete sinogram³⁰. Generative Adversarial Networks (GANs)³¹ have also been used to reconstruct CT images^{32,33}, often incorporating U-Nets in the generator network. U-Nets are a type of convolutional neural network (CNN), which require extensive training datasets and hyperparameter fine-tuning, and are also prone to hallucinating³⁴. These black-box procedures are effectively uninformed about the (known) physics and constraints of CT.

In the field of computer graphics, novel-view synthesis, denoting the ability to predict views of a 3D scene based on a (limited) number of other camera views, has made great progress in recent years^{35–39}. A key success has been the development of implicit neural representation (INR), a neural network-based continuous functional representation of the object itself^{35,39}. INR-based reconstruction methods promise to have significant advantages over CNN-based reconstruction methods. Most importantly, they are instance-specific, meaning that no additional external dataset is required for training, strongly reducing their tendency of hallucinating. Moreover, INRs have been found to give realistic reconstructions with few artifacts^{35,40}. However, the reconstruction process is iterative and requires guidance by physics modeling.

Inspired by the progress in novel-view synthesis, INR with Fourier features encoding for correctly capturing sharp features³⁷ has been used to represent 3D CT images⁴¹. The INR object representation is refined iteratively using a penalty (loss) function based on comparing the physically modeled CT intensity predictions with the actual CT measurements. See also the study by Zheng and Hatzell⁴². Zha et al. in their 3D INR implementation “Neural Attenuation Fields” (NAF)⁴³ used “multiresolution hash encoding”³⁶ rather than “Fourier features” to decrease the CT reconstruction time by orders of magnitude while further improving the reconstruction quality. For 4D-CT, a parametric motion field can be learned to warp a 3D INR template⁴⁴, but this approach does not fully utilize the representational power of INR in 4D. We note that time-resolved lensless imaging has been reported⁴⁰, as well as 4D INRs for novel-view synthesis^{45–47}. Arguably, the full power of INR is invoked when the entire collection of harvested experimental data is used as input to model the full duration of spatiotemporal dynamics, rather than a limited subset of the data to reconstruct each time step independently^{41,48}.

To the best of our knowledge, we report the first holistic, fully INR-based dynamic image reconstruction algorithm for time-resolved 4D-CT, coined NeCT for Neural Computed Tomography. Essentially, the measured CT data together with an idealized physics-based model of the CT instrument used for data acquisition are used to train a neural network to represent the object. Specifically, NeCT predicts the key parameter for attenuation-contrast CT, the local attenuation constant $\mu = \mu(\mathbf{r}, t)$ represented as a continuous field in space $\mathbf{r} = (x, y, z)$ and time t by a neural network. As we shall discuss, NeCT enables a temporal resolution of the order of seconds in the home laboratory, combined with excellent spatial resolution, suppressed artifacts, and low storage requirements. We demonstrate both sparse-view 3D and 4D reconstructions of high

resolution, surpassing previously reported results, and opening avenues for practical use of ML for CT.

Results

INR architecture

The NeCT pipeline, which consists of three main conceptual parts, is illustrated in Fig. 1. Virtual X-rays are computationally traced from the source through the sample to each pixel on the area detector. Points along each ray are sampled and fed through the neural network, returning the attenuation coefficient field $\hat{\mu}(\mathbf{r}, t)$ associated with each coordinate in space and time. The predicted attenuation coefficients at the points along the ray are summed using Beer-Lambert’s law, giving an estimate of the intensity value \hat{I}_j for the ray r_j that hits the detector pixel j . Finally, the simulated intensity \hat{I}_j is compared to the measured intensity I_j , and their absolute difference (L1 loss) serves as the cost function. The neural network minimizes this cost through backpropagation, effectively learning the attenuation coefficient field $\mu = \mu(\mathbf{r})$ for the static case and $\mu = \mu(\mathbf{r}, t)$ for the dynamic case. The network architecture for 3D CT is an improved version of NAF⁴³ (see also Supplementary Notes 3 and 4), optimized for reconstruction of sparsely sampled static 3D structures.

Our INR architecture optimized for time-resolved studies, here coined “NeCT QuadCubes”, is inspired by reported efforts to visualize dynamic scenes^{45–47}. Formally, our INR is a continuous function $\Phi : \mathbb{R}^4 \rightarrow \mathbb{R}$. Using four 3D hash encoders instead of either a 3D+4D hash encoder⁴⁵ or six 2D encoders^{46,47}, can be considered a compromise between a large number of hash collisions and low coupling between the input dimensions in the encoder. The four 3D multiresolution hash encoders, with respective dimensions xyz , xyt , xzt , and yzt , encode the input coordinate and time-stamp into a higher dimensional space. The multilayer perceptron (MLP) predicts the attenuation coefficient $\hat{\mu}(\mathbf{r}, t)$ at any desired spatial coordinate and time point.

Conventional cone beam CT proceeds by stepwise angular increments until a complete 360° scan has been performed, giving only marginally new information for each subsequent projection. With the philosophy of harvesting as much complementary information as possible for each new recorded projection, the “golden section” sampling scheme has been proposed⁴⁹. With golden section scanning, the angular interval between subsequent projection angles is approximately 137.5°. In practice, however, this large angular step has the disadvantage that a significant amount of time is spent on unproductive instrument motor movements. We have thus devised a “hybrid golden section” procedure where a fixed number of equidistant projections (say 25) are collected for each full revolution of the sample, and the golden section is used to decide the starting angle for each subsequent revolution, see also Methods and Supplementary Note 1.

With NeCT QuadCubes, in favorable cases of sufficiently large dynamically changing objects, a temporal resolution approaching the experimental time between subsequent projection acquisitions is achieved. All projections from a dynamic CT experiment comprising multiple sample rotations are used to optimize a single 4D INR object that describes the whole experiment. Using a 4D continuous representation of the object, the algorithm combines the information from all the acquired projections to give the best possible estimate $\hat{\mu}(\mathbf{r}, t)$. Static structures will benefit from the many unique projections to become sharply defined, whereas short-lived features will tend to be less precisely reconstructed as they may be supported by just a couple of projections. Comparisons of NeCT QuadCubes with conventional FDK, demonstrating the large improvement in spatiotemporal resolution, are provided in Supplementary Note 5.

NeCT simulations

Simulations of idealized static 3D sparse-view CT were carried out to benchmark NeCT against the reconstruction schemes FDK¹², “ordered subsets - simultaneous algebraic reconstruction technique - total variation” (OS-SART-TV)^{50,51}, and NAF⁴³. For this purpose, several CT datasets were obtained from the Open SciVis Datasets project⁵². An example of a reconstruction based on 49 projections is shown in Fig. 2a for the *Stagbeetle*

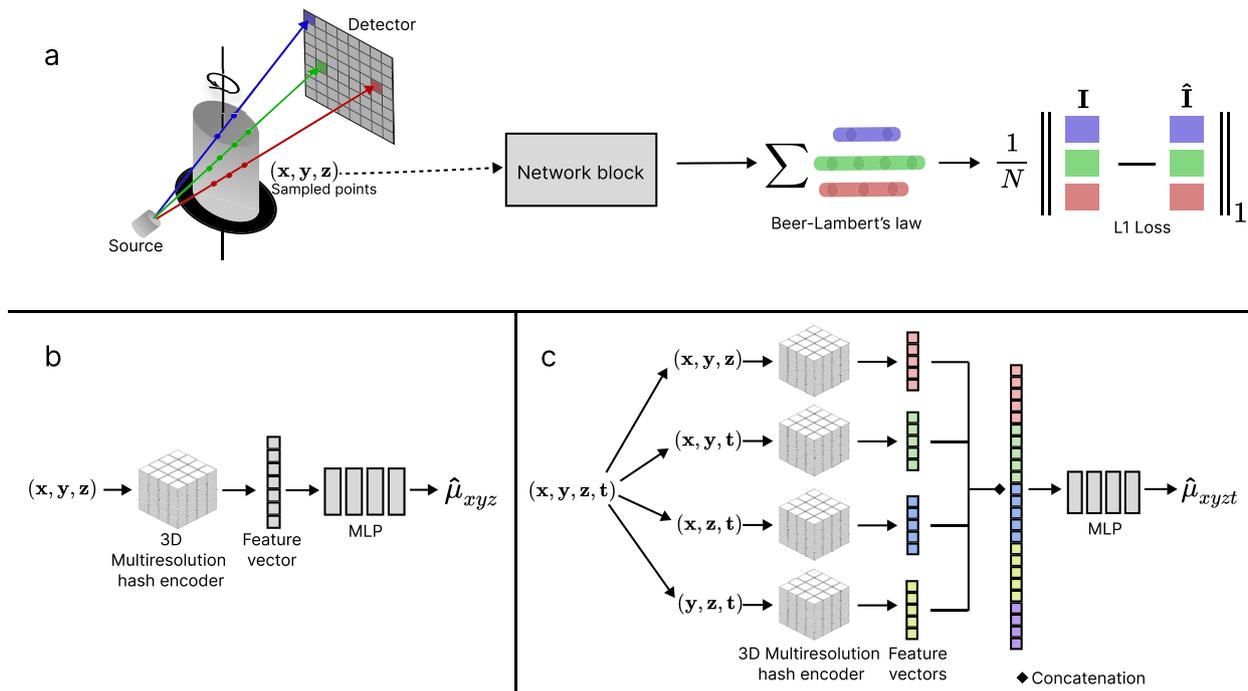


Fig. 1 | Overview of the NeCT architecture. **a** Sampling and intensity comparison. In the illustration, three rays (color coded red, green, blue) propagate from the source, through the object, to the detector. The number of sampled points per ray depends on the length traveled within the object as the distance between each point is set to be constant. **b** Network block for sparse-view 3D CT image reconstruction, similar to NAF⁴³, with the forward call returning $\hat{\mu}(x, y, z)$. **c** Network block for 4D-

CT reconstruction, named NeCT QuadCubes. The input data (\mathbf{r}, t) is split into 4 separate multiresolution hash encoders, and the feature vectors are concatenated together and sent into a multilayer perceptron (MLP). A forward call to NeCT QuadCubes gives the estimated attenuation coefficient $\hat{\mu}$ for the space-time position (\mathbf{r}, t) .

dataset (see also Supplementary Note 4 for several more examples and how the projections are obtained). Using the standard metrics of peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM)⁵³, we demonstrate that NeCT performs better than the competitors for both metrics across all the tested datasets, as illustrated in Fig. 2b.

Simulations of dynamic CT, based on 4D datasets of liquid invasion in porous media created with the open source Python package *PoreSpy*⁵⁴, were used to estimate the spatiotemporal resolution of NeCT. A slice of the porous rock taken at time $T = 0.8$, on a normalized time scale ranging from 0 to 1, is shown in Fig. 3. When comparing the ground truth with the reconstruction, there are only small visible differences. The difference image shows that NeCT typically misses by at most 1 pixel when reconstructing the pore boundary edges. Evidently, the dynamic process is almost fully resolved, with NeCT lagging slightly behind the ground truth in a single pore, which happened to be in the middle of a pore-filling event (cf. Supplementary Note 5).

In an effort to better understand and quantify the performance of NeCT QuadCubes, we performed two sets of simulations in the same artificial porous framework, with both single *on-off* and periodically switching domain structures (cf. Supplementary Note 5.3). These simulations demonstrate that NeCT QuadCubes is able to learn periodicity, giving a better resolution for the periodic events. For a cubic domain of size 12^3 voxels, the dynamics could be discerned if lasting for the time of two projections, and essentially fully resolved if lasting for ten projections.

Dynamic CT of spontaneous imbibition in sandstone

To support the development of NeCT, an in situ dynamic μ CT experiment was performed. Rather than a traditional core flooding sequence, a simpler and potentially more dynamic process was selected for this purpose: spontaneous imbibition of brine into a dry porous rock under ambient conditions. Bentheimer sandstone⁵⁵ is a homogeneous outcrop rock with relatively uniform properties, making it well suited for laboratory testing of

multiphase flow. In this study, a cylindrical core sample of Bentheimer, drilled to a diameter of 8.8 mm and a length of 19.6 mm, was used for the dynamic experiments. The permeability of the measured sample was (2.3 ± 0.3) D and the porosity $(23.0 \pm 2.3)\%$. The small size of the core plug enabled a μ CT setup with the sample placed close to the X-ray source, allowing sufficient geometric magnification for inspecting pore scale effects. The Bentheimer core plug was put in a rubber (VitonTM) sleeve inside an X-ray transparent core holder with an applied confining pressure of about 2 MPa. The sample holder was mounted on the rotary CT stage, with inlet lines for the doped brine supplied by a pump and the confining pressure connected at the bottom of the core holder. To keep the setup simple, no lines were connected at the top of the core holder.

The brine was a synthetic saltwater solution with a concentration of 237 g/L, resembling reservoir conditions. In addition, the brine was doped with CsI (182 g/L) to increase its X-ray attenuation. The pore space of the rock was partially filled with precipitated grains of the salt NaI, to a filling ratio of approximately $(20.3 \pm 1.0)\%$, in order to create a richer dynamics with dissolution events to be observed during the brine invasion.

The scanning protocol involved a total of 1400 projections distributed across 14 revolutions using our hybrid golden section sampling scheme, with alternating rotation direction for subsequent revolutions to avoid tangling of the inlet liquid supply lines. After the first full sample revolution of X-ray scanning, the pumping of brine at 5 μ L/min was started. Capillary contact between the invading brine and the rock was established during the second rotation. The whole 14-revolution sequence was completed within 40 min, when breakthrough of brine was observed at the top of the core sample. All projections were pre-processed with flat-field and dead pixel corrections. See Methods and Supplementary Notes 1 and 2 for full experimental details.

Because NeCT is a continuous model of the dynamic object represented by its attenuation coefficient in space and time, $\mu = \mu(\mathbf{r}, t)$, explicit 2D or 3D discretized (pixel or voxel-based) instances of the reconstructed object can be realized for any chosen time. Figure 4 gives a sample-scale view with a

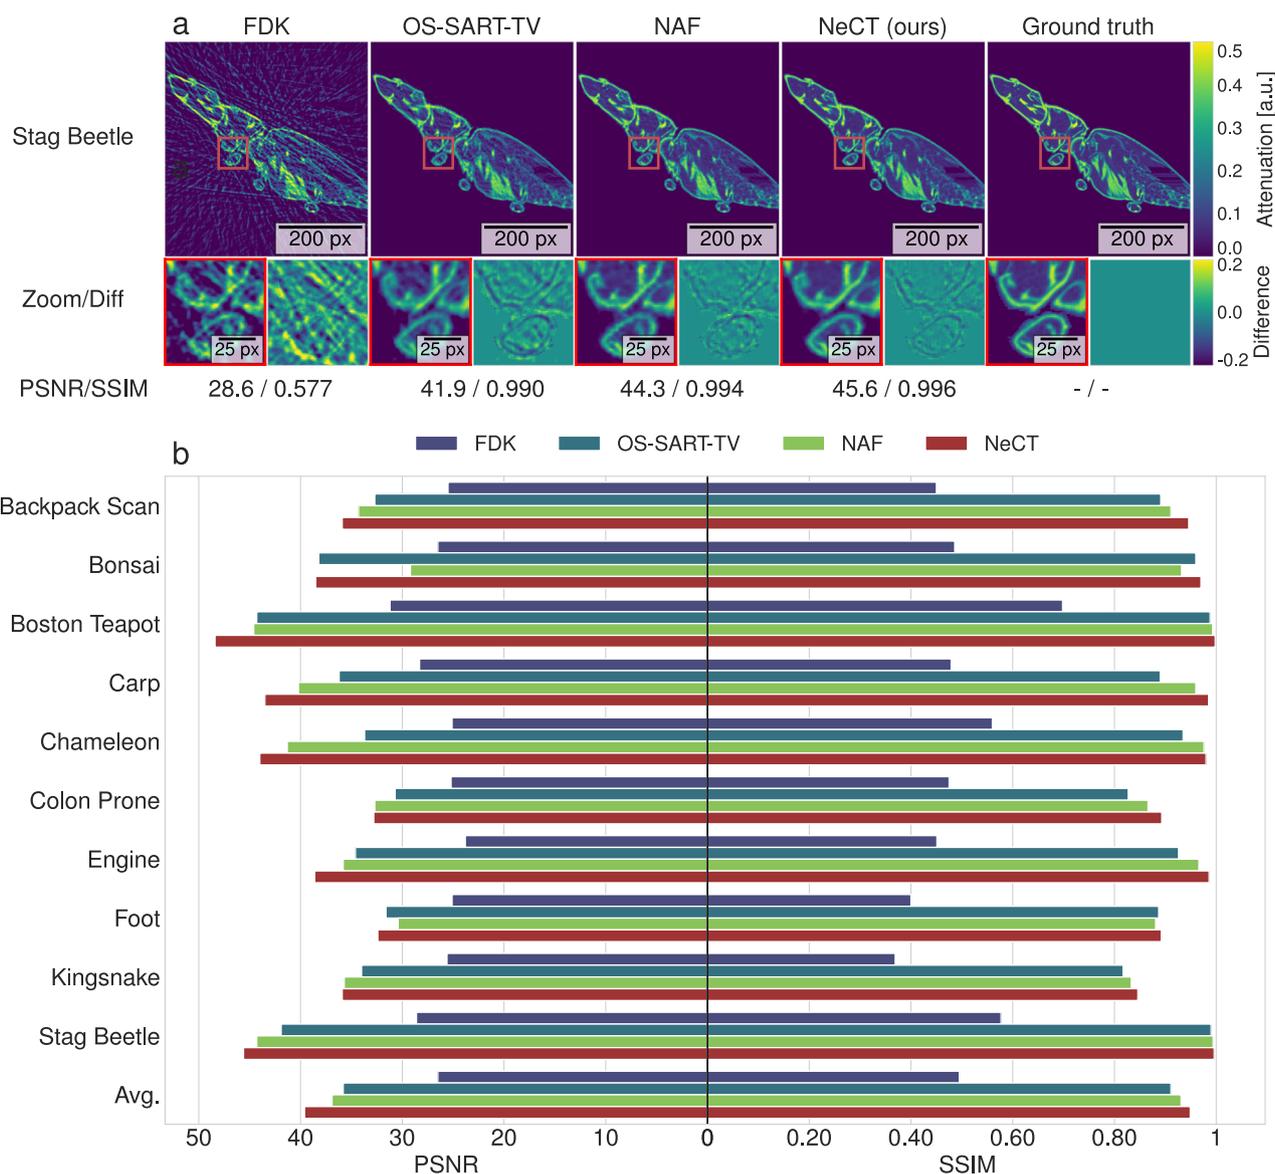


Fig. 2 | Static sparse-view 3D reconstructions. **a** Panoramic and zoomed sections, and accompanying difference images with respect to the ground truth, are shown for the Feldmann-Davis-Kress (FDK), ordered subsets - simultaneous algebraic reconstruction technique - total variation (OS-SART-TV), neural attenuation field (NAF), and NeCT (our) algorithms. The reconstructions are based on 49 projections

of the benchmark *Stagbeetle* CT dataset⁵². **b** Peak signal-to-noise ratio (PSNR) and structural-similarity index measure (SSIM) values for all static datasets are shown. NeCT is seen to outperform the other methods both visually and in terms of PSNR and SSIM for all datasets.

segmented reconstructed slice of the Bentheimer sandstone 700 seconds into the spontaneous imbibition process (see also Supplementary Movie 1). The areas of highest attenuation are segmented as salt (red), while the difference with respect to t_0 was used to segment the incoming brine (blue).

A detailed visualization of the dynamics in selected regions of the Bentheimer sample is provided in Fig. 5. In low porosity regions, the local attenuation stayed essentially constant with time, as it should in the absence of physicochemical reactions. As the brine front reached and filled open pores, the local attenuation increased. The 4D reconstruction allowed us to investigate the pore-filling events in great detail. In some larger pores, we observed that trapped air bubbles formed when brine entered the pore and wetted the rock surface. Interestingly, NeCT enabled us to directly observe the dissolution of the salt deposits in the sample as the brine front arrived, with a typical dissolution time of the order of 10 s (see also Supplementary Movie 2). CT studies of reactive transport and dissolution are a topic of high current interest; see, e.g., the recent work on carbonate dissolution by Agrawal et al.⁵⁶.

Note that in both the 3D and 4D NeCT reconstructions, the common CT image artifacts of streaks (arising for undersampled data) and rings (for detector pixel errors)¹³ are nonexistent (cf. Supplementary Note 5.4). A partial explanation for the smooth appearance of the reconstructions may be that the INR effectively acts as a regularizer by being spectrally biased towards low frequencies³⁷, and presumably also because the information from all the projections is exploited to inform the entire time-span of the model $\hat{\mu}(\mathbf{r}, t)$ ^{40,48}. For non-smooth events unfolding over a too low number of projections, NeCT may still yield unphysical solutions to the inverse problem due to the insufficient information provided by the small set of 2D images to fully reconstruct the complete 4D data.

Discussion

Using NeCT, we have successfully repurposed a standard commercially available CT instrument designed for 3D CT scans of typical duration ~1 hour to enable dynamic “4D-CT” scans with a temporal resolution approaching the time between subsequent projections, i.e., an improvement

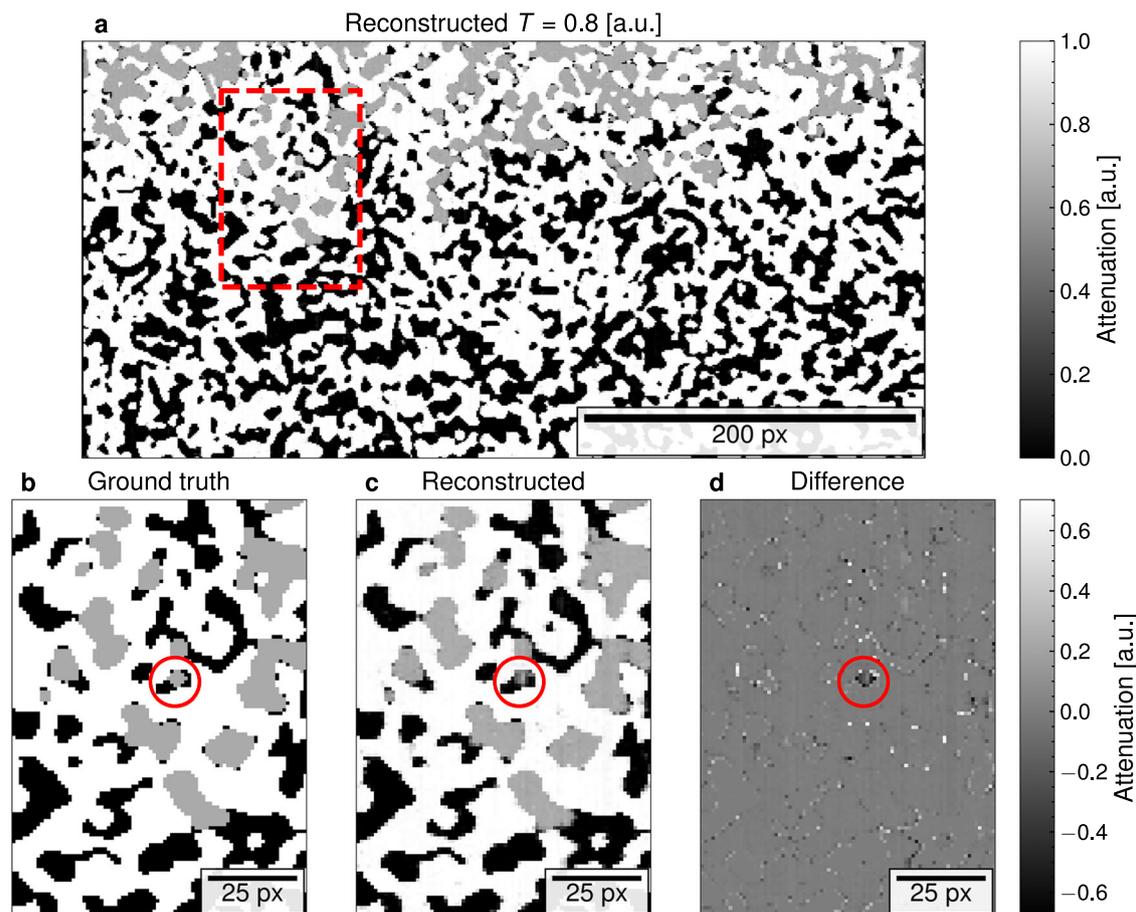


Fig. 3 | Simulated porous medium with fluid invasion reconstructed using NeCT. Results at normalized time $T = 0.8$ are shown. Black = pore w/air, gray = fluid, white=matrix. **a** Large overview. **b–d** A region centered on a pore-filling event is shown (the region highlighted in (a)), comparing the ground truth to the

reconstruction. There are no errors in the reconstruction except for a few places at the pore boundaries and the highlighted pore. The pore was temporarily in the middle of a pore-filling event and lagged slightly behind the ground truth.

by 2–3 orders of magnitude as compared to FDK-based reconstruction. This achievement was obtained using a combination of our “hybrid golden section” sampling scheme and the continuous INR representation NeCT QuadCubes of the sample, which exploits all the temporally preceding and succeeding information to concertedly inform the reconstruction at any instance of time. The NeCT architecture is original and has many possibilities for further extensions and improvements. Arguably, NeCT also benefits from being conceptually closer to conventional reconstruction algorithms like SART than to e.g., trained GAN networks, which likely lowers the barrier to user adoption, further convincing us that our approach represents an important step towards new physics-informed measurement methods in the natural sciences.

The capability of resolving 3D movies of events inside an opaque material on a time scale of seconds using home laboratory equipment is a substantial advance of the state-of-the-art in dynamic imaging. Indeed, conventional μ CT sampling and reconstruction methods, which give a temporal resolution on the order of an hour, make it practically infeasible to carry out many experiments of high scientific or industrial importance. Many research groups are currently engaged in resolving this challenge, with NeCT providing a continuous representation of the dynamically evolving object on a length scale of micrometers and a time scale of seconds. The spontaneous imbibition experiment with Bentheimer presented in this article revealed dynamics that can currently only be inferred by synchrotron imaging methods.

While significant advantages of the NeCT approach include that the data model is continuous in space and time, does not require external training data, and is memory efficient for storage, these features also come

with several limitations that will require significant future efforts to fully understand and mitigate. Most importantly, the spatiotemporal resolution and how the algorithm chooses solutions from solution space are complex questions that require further research. Similarly, the choice of data collection settings, the NeCT hyperparameters, and the training schemes appear to be well-behaved, but their exact influence on the reconstruction performance should be explored further. The long training, i.e., reconstruction times, in particular for 4D datasets, being of the order of tens of hours on advanced graphical processing units (GPUs), can be problematic for some applications. To this end, we envision that perhaps a two-step reconstruction process where a computationally cheaper method, maybe along the lines of van Eyndhoven et al.¹⁸ or by Goethals et al.²¹, can be used to initialize NeCT.

In conventional 4D-CT, the temporal resolution is given by the time it takes to complete each full scan (360°), since each time step is given by an independent reconstruction. This restriction is lifted with the NeCT approach, and because NeCT opens for capturing dynamics with less than a full rotation for each time step, also the fast spinning requirement mentioned in the Introduction is significantly relaxed. Estimating the ability of NeCT to resolve 4D dynamics in space and time is not straightforward. The resolution depends on the sampling scheme, and attention must be paid to the optimization of angular velocity, exposure time, and number of projections per scan. The optimal configuration gathers the most complementary information per time. When reconstructing a 4D model, it is necessary to discuss the spatiotemporal resolution rather than only the temporal resolution. Presumably, the size, location, and type of event will influence the dynamic results, see also Goethals et al.¹⁹. Quantifying the

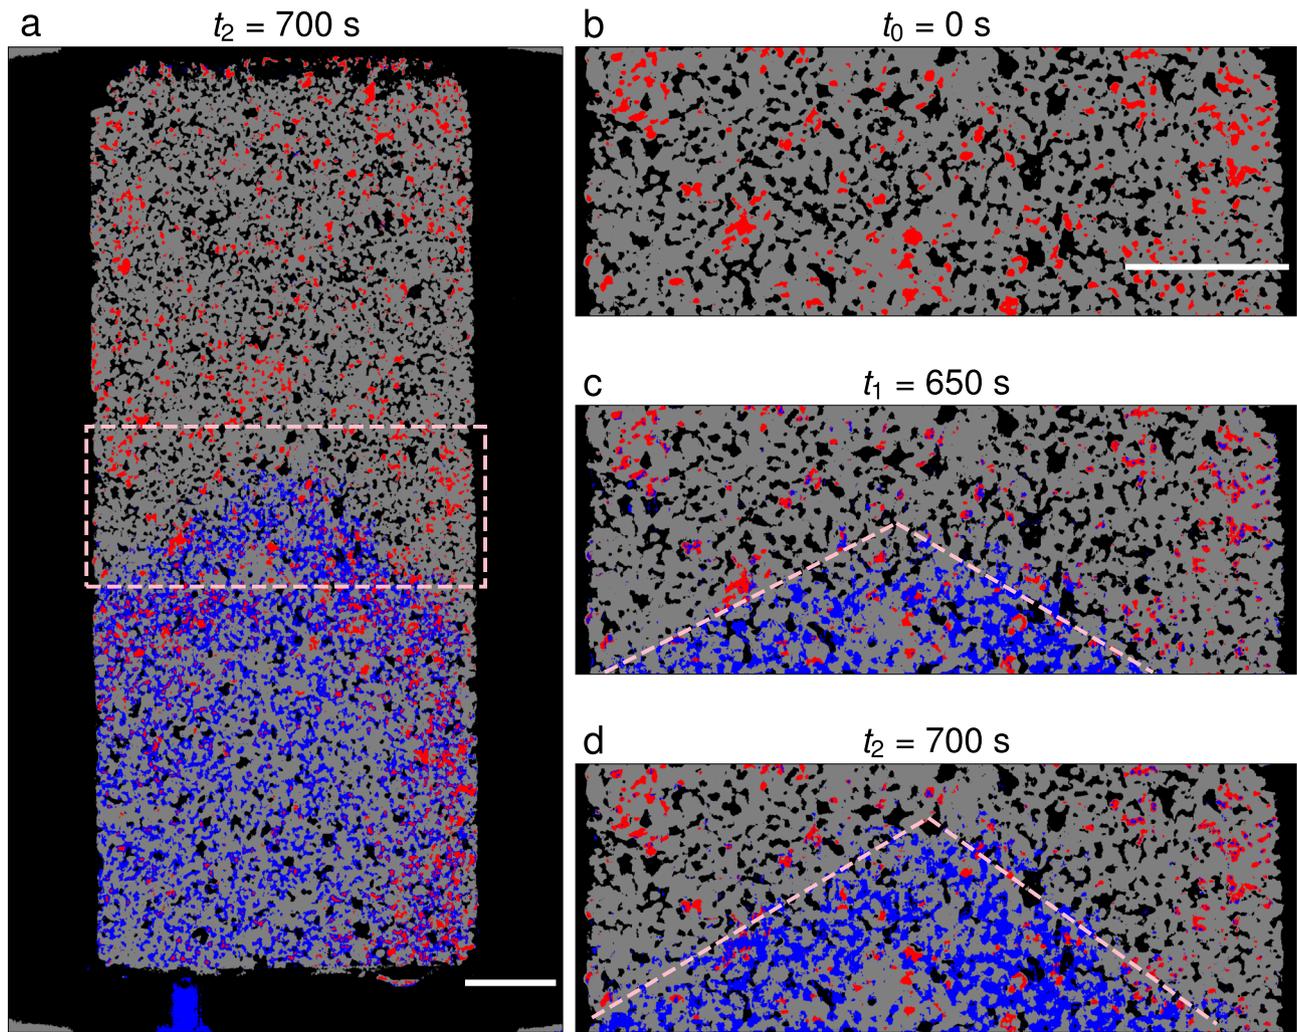


Fig. 4 | Segmented dynamic NeCT reconstruction of Bentheimer sandstone, panoramic view. The shown sections are explicitly decoded representations of the NeCT model for the specified times. Black = pore space, gray = matrix, red = salt,

blue = brine. **a** Overview of the full sample at time $t_2 = 700$ s. **b–d** Magnified view corresponding to the rectangle in **(a)**, at $t_0 = 0$, $t_1 = 650$ and $t_2 = 700$ seconds, respectively. The fluid front is highlighted with broken lines. Scale bars are 2 mm.

exact spatiotemporal resolution and further optimization, performance, and limitations of NeCT are topics to explore in future studies.

In NeCT, the underlying physical model assumes that the X-ray radiation can be described as effectively monochromatic and that scattering can be ignored. These simplifications are commonly made in CT analysis and are known to work well in practice¹³, as also seen in the current work. Still, it is reasonable to assume that enhancing the physics-based forward model to better capture additional phenomena that can give systematic modifications of the measured signal, such as scattering, polychromatic radiation, and detector nonlinearity, will further improve the reconstructions. Although we have exclusively focused on cone-beam CT in this article, a parallel-beam forward model has been implemented, making NeCT equally applicable to synchrotron and neutron tomography.

An interesting extension of NeCT would be to account for the fact that spectrally (energy) resolved CT is already widely used for medical CT⁵⁷, and developments are underway for laboratory μ CT⁵⁸. As in the original INR work by Mildenhall et al.³⁵, where three color channels (RGB) are returned for each coordinate, we envision that rather than a single-valued attenuation coefficient $\hat{\mu}$ from each space-time coordinate (\mathbf{r}, t) , an energy-dependent $\hat{\mu} = \hat{\mu}(\mathbf{r}, t, E)$ can be foreseen, opening for chemical contrast, see also Wu et al.⁵⁹. Similar lines of thought apply to the recent developments within tensorial CT, where orientation information about sub-resolution anisotropic texture is retrieved from scattering signals^{26,60}. INR-models could be

key to substantially simplifying the heavy data analysis associated with such datasets. Following these ideas, NeCT can be extended to return continuous multi-valued objects for each space-time coordinate (\mathbf{r}, t) .

Continuous scanning can increase the temporal resolution of CT scanning because the sample is imaged without pause while it rotates at a constant speed, avoiding the angular acceleration and homing at each measurement point^{1,24}. Continuous scanning can provide more information about fast events by recording dynamics that could otherwise take place between exposures. A complication is that continuous scanning inevitably introduces motion smearing because of the constant rotation speed. NeCT can be extended to account for spatial averaging, most directly by approximating each projection angular interval as a series of discrete steps, or by using more elaborate integration schemes. Thereby, the contribution from each point within the object during rotation would be accounted for, potentially further increasing the information harvesting rate and hence the effective temporal resolution of the experiment.

The continuous INR representation of the experimentally measured data is highly storage efficient, allowing complete time-resolved datasets to be represented with just a few GB of data, as compared to several TB of data for the 4D data in conventional voxel representation. Consequently, we envision the sparse sampling and data representation to be a partial answer to the growing demands for data storage in CT experiments. We note that NeCT should be easily adaptable to future improved neural network

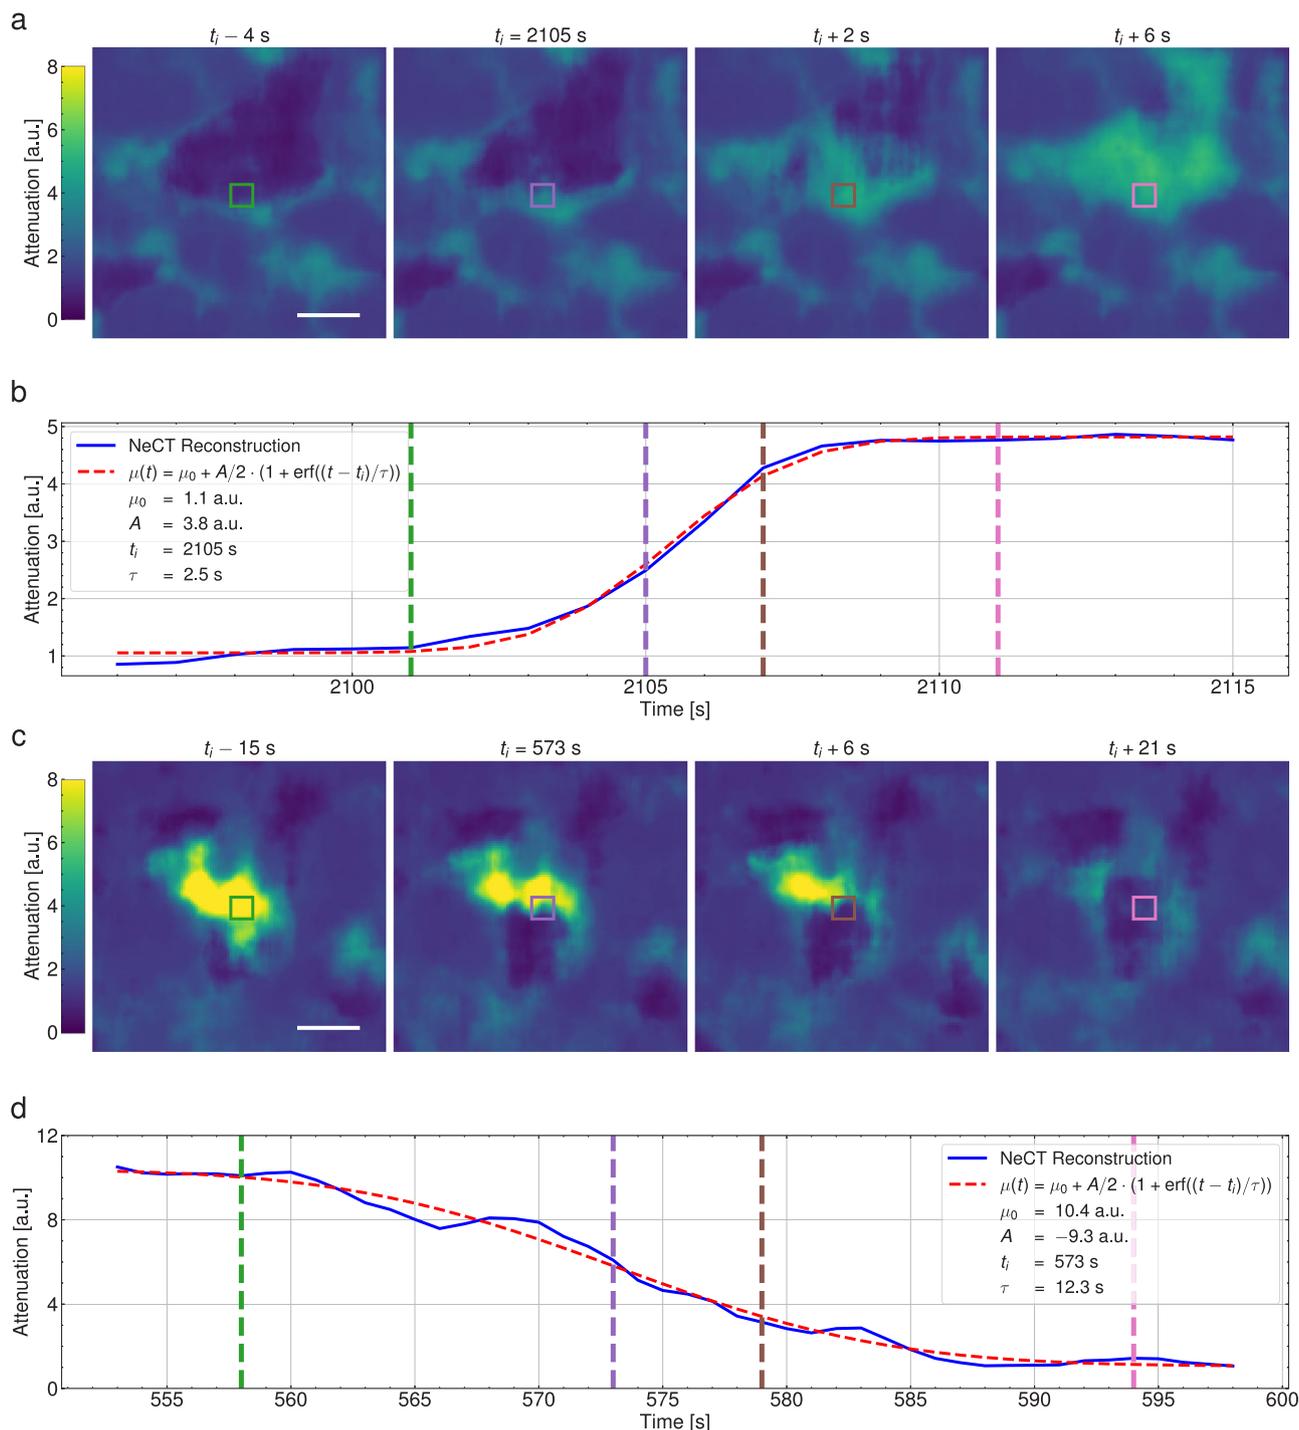


Fig. 5 | Dynamic NeCT reconstruction of Bentheimer sandstone, pore-level view. **a, b** The filling of an initially empty pore with brine. Note the progressive water invasion leading to a homogeneous filling of the whole pore volume during a time span of about 10 s. **c, d** A salt grain dissolves, leaving behind an empty pore as

capillary forces fail to hold the brine. The dissolution of the entire grain lasted ~30 s. The blue lines in **(b)** and **(d)** are NeCT reconstructions, and the red broken lines are error-function fits to the data as described in the legends. Scale bars are 0.1 mm.

architectures, which are expected to be developed as a result of the ongoing massive research into artificial intelligence. Similarly, while NeCT has been developed specifically for operating without prior information, incorporating additional data into the reconstruction data is often desirable^{6,15,40,48}. In the context of dynamic CT, such priors could include static high-resolution scans^{6,15}, complementary simultaneous measurements of pressure, temperature, and acoustic emissions, to mention a few.

In summary, NeCT enables dynamic or 4D μ CT measurements to be carried out with high spatiotemporal resolution, allowing us to resolve liquid

flow in porous media obtained with standard cone-beam CT equipment. Our method is likely to help facilitate laboratory-based measurements of dynamic processes inside 3D structures, which will be of high interest to academia and industry alike.

Methods

Experimental CT acquisition

The dynamic Bentheimer sandstone dataset was acquired at Equinor’s Rotvoll laboratory facility using a North Star Imaging X5000 μ CT. The

X-ray source was an XRayWorX 225 kV operated at 124 kV and 100 μ A with a focal spot of 12.4 μ m. The detector was a Varex Imaging PaxScan 2520DX with 1920×1536 pixels, with a pixel pitch of 127 μ m. The sample was imaged at a source-to-detector distance of 518 mm and a source-to-origin distance of 49.5 mm, corresponding to a conventional voxel size of 12.1 μ m. Each radiograph was recorded as an average of 5 exposures, with a total exposure time of 0.42 s. Each full (360°) rotation lasted 165 s.

Raymarcher sampling scheme

The ray marcher was implemented with an equidistant sampling scheme, implying that equidistant points were sampled along each ray within the object. Instead of using a constant number of points per ray, NeCT gradually increases the number of sampled points over time. The number of points per ray n is defined as $n = \min(n_{init} + \lfloor i/j \rfloor, n_{max})$, where i is the number of projections processed, j is the update interval, and n_{init} and n_{max} are the initial and maximum number of points per ray, respectively. This adaptive strategy improves reconstruction efficiency by refining the sampling resolution over time. NeCT employs an adaptive number of points per ray, increasing linearly from 100 to $1.5 \times \max(n_{Detector})$, where $\max(n_{Detector})$ denotes the number of pixels along the longest axis of the detector. For the experimental Bentheimer sandstone reconstruction, the maximum was set to be reached after 30% of the reconstruction time.

Hyperparameters

All NeCT reconstructions were made using a single Nvidia A100 or H100 GPU. We have verified that NeCT also works on a multi-GPU setup when NVLink is available. The encoder configuration, used for both NeCT (static and dynamic) and our reproduced NAF results⁴³, consists of a hashmap with 2^{23} entries, 23 levels, a base resolution of 16, and 4 features per level. The MLP has 4 hidden layers, each with 128 neurons. The activation function used is “Leaky ReLU”.

Learning rate scaling

To maintain consistent reconstructions across different batch sizes (B) and number of GPUs used ($\#GPU$), we apply the square root scaling rule⁶¹ to scale the base learning rate by the batch size. With the base batch size set to 1 million, the scaling factor becomes $\sqrt{\#GPU \times B \times 10^{-6}}$, where B represents the batch size.

Optimization

We used a learning rate scheduler that included a linear warm-up phase for the first 5000 batches, followed by a cosine decay schedule with an ending learning rate set at 1% of the initial value after the warm-up phase. The number of training epochs was set at four times the standard value (see Supplementary Note 2.3) in all experiments, except for the Bentheimer sandstone reconstruction, where we used 40 times the standard number of epochs. Although the reconstruction had already converged well before reaching the end, we set the number of epochs this high to ensure that we achieved the best possible reconstruction. We used the Adam optimizer with a base learning rate of 1×10^{-3} for static reconstructions and 2×10^{-4} for dynamic reconstructions. The batch size was set to 5×10^6 points per batch. Finally, L1 loss was used as the loss function in all experiments.

Data availability

All static datasets are available through the Open SciVis Datasets project⁵². The dynamic dataset that supports the findings of this study is available at <https://doi.org/10.5281/zenodo.16448474>.

Code availability

Code is available at <https://doi.org/10.5281/zenodo.16461048>.

Received: 16 April 2025; Accepted: 29 July 2025;

Published online: 18 August 2025

References

- Withers, P. et al. X-ray computed tomography. *Nat. Rev. Methods Prim.* **1**, 1–17 (2021).
- Cnudde, V. & Boone, M. High-resolution X-ray computed tomography in geosciences: a review of the current technology and applications. *Earth Sci. Rev.* **123**, 1–17 (2013).
- Ziesche, R. F. et al. 4D imaging of lithium-batteries using correlative neutron and X-ray tomography with a virtual unrolling technique. *Nat. Commun.* **11**, 777 (2020).
- Charalampidou, E.-M., Hall, S. A., Stanchits, S., Lewis, H. & Viggiani, G. Characterization of shear and compaction bands in a porous sandstone deformed under triaxial compression. *Tectonophysics* **503**, 8–17 (2011).
- Blunt, M. J. et al. Pore-scale imaging and modelling. *Adv. Water Resour.* **51**, 197–216 (2013).
- Tekseth, K. R. & Breiby, D. W. 4D imaging of two-phase flow in porous media using laboratory-based micro-computed tomography. *Water Resour. Res.* **60**, e2023WR036514 (2024).
- Panduro, E. A. C. et al. Real time 3D observations of Portland cement carbonation at CO₂ storage conditions. *Environ. Sci. Technol.* **54**, 8323–8332 (2020).
- Voltoolini, M. et al. The emerging role of 4D synchrotron X-ray micro-tomography for climate and fossil energy studies: five experiments showing the present capabilities at beamline 8.3.2 at the advanced light source. *J. Synchrotron. Radiat.* **24**, 1237–1249 (2017).
- Singh, K. et al. Time-resolved synchrotron X-ray micro-tomography datasets of drainage and imbibition in carbonate rocks. *Sci. Data* **5**, 180265 (2018).
- Tekseth, K. R., Mirzaei, F., Lukic, B., Chattopadhyay, B. & Breiby, D. W. Multiscale drainage dynamics with Haines jumps monitored by stroboscopic 4D X-ray microscopy. *Proc. Natl. Acad. Sci. USA* **121**, e2305890120 (2024).
- Seeram, E. *Computed Tomography: Physical Principles, Patient Care, Clinical Applications, and Quality Control* (Saunders, Philadelphia, 2022), 5 edn.
- Feldkamp, L. A., Davis, L. C. & Kress, J. W. Practical cone-beam algorithm. *J. Opt. Soc. Am. A* **1**, 612–619 (1984).
- Kak, A. C. & Slaney, M. *Principles of Computerized Tomographic Imaging* (Society for Industrial and Applied Mathematics, Philadelphia, 2001).
- Andersen, A. H. & Kak, A. C. Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrason. Imaging* **6**, 81–94 (1984).
- Chen, G.-H., Tang, J. & Leng, S. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. *Med. Phys.* **35**, 660–663 (2008).
- Myers, G. R., Kingston, A. M., Varslot, T. K., Turner, M. L. & Sheppard, A. P. Dynamic tomography with a priori information. *Appl. Opt.* **50**, 3685–3690 (2011).
- Bultreys, T. et al. Fast laboratory-based micro-computed tomography for pore-scale research: Illustrative experiments and perspectives on the future. *Adv. Water Resour.* **95**, 341–351 (2016).
- Van Eyndhoven, G. et al. An iterative CT reconstruction algorithm for fast fluid flow imaging. *IEEE Trans. Image Process.* **24**, 4446–4458 (2015).
- Goethals, W. et al. Dynamic CT reconstruction with improved temporal resolution for scanning of fluid flow in porous media. *Water Resour. Res.* **58**, e2021WR031365 (2022).
- Makiharju, S. A., Dewanckele, J., Boone, M., Wagner, C. & Griesser, A. Tomographic x-ray particle tracking velocimetry. *Exp. Fluids* **63**, 16 (2022).
- Goethals, W., Bultreys, T., Berg, S., Boone, M. N. & Aelterman, J. Dyrect computed tomography: dynamic reconstruction of events on a

- continuous timescale. *IEEE Trans. Comput. Imaging* **11**, 638–649 (2025).
22. Raimondi, P. et al. The extremely brilliant source storage ring of the European Synchrotron Radiation Facility. *Commun. Phys.* **6**, 82 (2023).
 23. Aidukas, T. et al. High-performance 4-nm-resolution x-ray tomography using burst ptychography. *Nature* **632**, 81–88 (2024).
 24. Zhang, J., Lee, W.-K. & Ge, M. Sub-10 second fly-scan nanotomography using machine learning. *Commun. Mater.* **3**, 91 (2022).
 25. Wang, Z. et al. Ultrafast radiographic imaging and tracking: an overview of instruments, methods, data, and applications. *Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrometers Detect. Assoc. Equip.* **1057**, 168690 (2023).
 26. Liebi, M. et al. Nanostructure surveys of macroscopic specimens by small-angle scattering tensor tomography. *Nature* **527**, 349–352 (2015).
 27. Walker, S. M. et al. In vivo time-resolved microtomography reveals the mechanics of the blowfly flight motor. *PLoS Biol.* **12**, e1001823 (2014).
 28. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Navab, N., Hornegger, J., Wells, W. M. & Frangi, A. F. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 234–241 (Springer International Publishing, Cham, 2015).
 29. Jin, K. H., McCann, M. T., Froustey, E. & Unser, M. Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.* **26**, 4509–4522 (2017).
 30. Hu, D. et al. Hybrid-domain neural network processing for sparse-view CT reconstruction. *IEEE Trans. Radiat. Plasma Med. Sci.* **5**, 88–98 (2021).
 31. Goodfellow, I. et al. Generative Adversarial Nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. & Weinberger, K. (eds.) *Advances in Neural Information Processing Systems*, vol. 27, 2672–2680 (Curran Associates, Inc., Red Hook, NY, 2014). https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.
 32. Guo, Z. et al. Physics-assisted generative adversarial network for X-ray tomography. *Opt. Express* **30**, 23238–23259 (2022).
 33. Xie, H., Shan, H. & Wang, G. Deep encoder-decoder adversarial reconstruction (DEAR) network for 3D CT from few-view data. *Bioengineering* **6**, 111 (2019).
 34. Patwari, M. et al. Reducing the risk of hallucinations with interpretable deep learning models for low-dose CT denoising: comparative performance analysis. *Phys. Med. Biol.* **68**, 19LT01 (2023).
 35. Mildenhall, B. et al. NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* **65**, 99–106 (2021).
 36. Müller, T., Evans, A., Schied, C. & Keller, A. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* **41**, 102:1–102:15 (2022).
 37. Tancik, M. et al. Fourier features let networks learn high frequency functions in low dimensional domains. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. & Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, 7537–7547 (Curran Associates, Inc., Red Hook, NY, 2020). https://proceedings.neurips.cc/paper_files/paper/2020/file/55053683268957697aa39fba6f231c68-Paper.pdf.
 38. Barron, J. T. et al. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5855–5864 (ICCV, 2021).
 39. Sitzmann, V., Martel, J., Bergman, A., Lindell, D. & Wetzstein, G. Implicit neural representations with periodic activation functions. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. & Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, 7462–7473 (Curran Associates, Inc., 2020). https://proceedings.neurips.cc/paper_files/paper/2020/file/53c04118df112c13a8c34b38343b9c10-Paper.pdf.
 40. Chien, T., Cao, R., Liu, F. L., Kabuli, L. A. & Waller, L. Space-time reconstruction for lensless imaging using implicit neural representations. *Opt. Express* **32**, 35725–35732 (2024).
 41. Shen, L., Pauly, J. & Xing, L. NeRP: Implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems* 1–13 (IEEE, 2022).
 42. Zheng, Y. & Hatzell, K. B. B. Ultrasparse view X-ray computed tomography for 4D imaging. *ACS Appl. Mater. Interfaces* **15**, 35024–35033 (2023).
 43. Zha, R., Zhang, Y. & Li, H. NAF: Neural attenuation fields for sparse-view CBCT reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 442–452 (Springer, 2022).
 44. Reed, A. W. et al. Dynamic CT reconstruction from limited views with implicit neural representations and parametric motion fields. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2238–2248 (IEEE Computer Society, 2021).
 45. Park, S. et al. Temporal interpolation is all you need for dynamic neural radiance fields. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4212–4221 (IEEE Computer Society, 2023).
 46. Fridovich-Keil, S., Meanti, G., Warburg, F., Recht, B. & Kanazawa, A. K-Planes: explicit radiance fields in space, time, and appearance. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12479–12488 (IEEE Computer Society, 2023).
 47. Cao, A. & Johnson, J. HexPlane: A fast representation for dynamic scenes. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 130–141 (IEEE Computer Society, 2023).
 48. Yang, J. & Xie, S. PE-INEr: prior-embedded implicit neural representation for sparse-view CBCT reconstruction. *Appl. Opt.* **63**, 8907–8916 (2024).
 49. Kohler, T. A projection access scheme for iterative reconstruction based on the golden section. In *IEEE Symposium Conference Record Nuclear Science 2004*, vol. 6, 3961–3965 (IEEE, 2004).
 50. Du, Y., Yu, G., Xiang, X., Wang, X. & Deene, Y. D. Convergence of SART + OS + TV iterative reconstruction algorithm for optical CT imaging of gel dosimeters. *J. Phys.: Conf. Ser.* **847**, 12025 (2017).
 51. Biguri, A., Dosanjh, M., Hancock, S. & Soleimani, M. Tigre: A MATLAB-GPU toolbox for CBCT image reconstruction. *Biomed. Phys. Eng. Express* **2**, 55010 (2016).
 52. Klacansky, P. Open SciVis datasets. <https://klacansky.com/open-sci-vis-datasets/> (2017).
 53. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612 (2004).
 54. Gostick, J. et al. PoreSpy: A Python toolkit for quantitative analysis of porous media images. *J. Open Source Softw.* **4**, 1296 (2019).
 55. Peksa, A. E., Wolf, K.-H. A. A. & Zitha, P. L. J. Bentheimer sandstone revisited for experimental purposes. *Mar. Pet. Geol.* **67**, 701–719 (2015).
 56. Agrawal, P. et al. Control of brine composition over reactive transport processes in calcium carbonate rock dissolution: time-lapse imaging of evolving dissolution patterns. *Appl. Geochem.* **161**, 105835 (2024).
 57. Willemink, M. J., Persson, M., Pourmorteza, A., Pelc, N. J. & Fleischmann, D. Photon-counting CT: Technical principles and clinical prospects. *Radiology* **289**, 293–312 (2018).
 58. Egan, C. et al. 3d chemical imaging in the laboratory by hyperspectral X-ray computed tomography. *Sci. Rep.* **5**, 1–17 (2015).
 59. Wu, Q. et al. Unsupervised polychromatic neural representation for CT metal artifact reduction. In Oh, A. et al. (eds.) *Advances in Neural Information Processing Systems*, vol. 36, 69605–69624 (Curran Associates, Inc., 2023). https://proceedings.neurips.cc/paper_files/paper/2023/file/dbf02b21d77409a2db30e56866a8ab3a-Paper-Conference.pdf.

60. Mürer, F. et al. Quantifying the hydroxyapatite orientation near the ossification front in a piglet femoral condyle using X-ray diffraction tensor tomography. *Sci. Rep.* **11**, 2144 (2021).
61. Krizhevsky, A. One weird trick for parallelizing convolutional neural networks. *CoRR*abs/1404.5997 (2014).

Acknowledgements

The authors thank Kim Robert Tekseth and Ruben S. Dragland for valuable discussions regarding the use of hybrid golden section sampling and the U-Net architecture. The laboratory experiments were carried out at the research facilities of Equinor ASA at Rotvoll, Norway. We gratefully acknowledge Equinor ASA and the Norwegian Research Council (project #275182 4D-CT) for financing this study.

Author contributions

H.F. and H.N. implemented the NeCT model and actively contributed to all aspects of the research. H.F., H.N., G.L., C.P., L.R., and A.K. carried out the experiments and analysis. H.F., H.N., G.L., and D.W.B. wrote the first draft of the manuscript. B.C., A.K., O.J.M., and D.W.B. supervised the project. All authors contributed to the finalization of the article.

Funding

Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital).

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02249-0>.

Correspondence and requests for materials should be addressed to Anders Kristoffersen, Ole Jakob Mengshoel or Dag Werner Breiby.

Peer review information *Communications Physics* thanks Adrian Sheppard, Tom Bultreys, and the other anonymous reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

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

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

© The Author(s) 2025