3D MICROSTRUCTURE RECONSTRUCTION OF AEROGELS VIA CONDITIONAL GANS

Prakul Pandit*
Institute for Frontier Materials on Earth and in Space
German Aerospace Center (DLR), Cologne, Germany
prakul.pandit@dlr.de

Sugan Kanagasenthinathan*
Institute for Frontier Materials on Earth and in Space
German Aerospace Center (DLR), Cologne, Germany
sugan.kanagasenthinathan@dlr.de

Ameya Rege
Department of Mechanics of Solids, Surfaces & Systems, University of Twente, The Netherlands
Institute for Frontier Materials on Earth and in Space, German Aerospace Center (DLR), Cologne, Germany
ameya.rege@utwente.nl

ABSTRACT

Aerogels are low-density and highly porous materials (90–99% porosity) with exceptional thermal and mechanical properties, governed by their intricate nanoporous microstructure. Understanding their structure-property relationships is essential for optimizing their performance across industrial applications. A significant challenge appears in precisely identifying the complete pore space and thus mapping their microstructural morphology of aerogels. This work presents a deep learning-driven digital twin framework for aerogels, leveraging Conditional Generative Adversarial Networks (cGANs) and Convolutional Neural Networks (CNNs) for 3D microstructure reconstruction and predictive modeling. Our approach reconstructs 3D aerogel microstructures from synthetic 2D scanning electron microscopy (SEM) images that mimic real samples by incorporating depth effects. A CNN predicts key microstructural parameters, including pore radius, relative density, and pore size distribution, with minimal error. A 3D cGAN then generates aerogel microstructures by capturing global spatial features and conditioning on the extracted parameters.
We demonstrate that conditioning improves the fidelity of reconstruction by enforcing physically meaningful constraints. This method provides a scalable, data-driven approach for microstructure modeling, enabling efficient structure-property predictions, and guiding aerogel design for targeted applications.

1 INTRODUCTION

Aerogels are low-density, highly open-porous materials synthesized by drying and are characterized by their unique nanostructured morphology and exceptional thermal properties. Rather than referring to a specific material, the term ”aerogel” encompasses a broad class of materials developed via the sol-gel process, where the liquid phase is replaced by a gas through a controlled drying process. This transformation retains the original shape and volume while producing a highly porous structure with upto 99.98% porosity (Kistler, 1931). Due to their extremely low thermal conductivity, low bulk densities, and high surface area, aerogels are widely used in aerospace, automotive, and insulation applications. Their macroscopic properties are predominantly governed by their intricate nanoporous microstructure, making precise microstructure characterization essential for material optimization (Rege, 2023).

Despite their promising applications, establishing accurate structure-property relationships for aerogels remains a significant challenge due to their stochastic gelation process and complex pore morphology. One critical aspect of aerogel characterization is the precise estimation of pore size distribution, which directly influences key material properties such as thermal conductivity, mechanical

*Equal contribution.

strength, and adsorption capacity. However, due to aerogels' hierarchical and highly porous nature, obtaining reliable pore size measurements is nontrivial.

Experimental techniques for 3D characterization such as X-ray computed tomography (CT) face limitations when applied to aerogels due to their poor resolution in the mesosporous region (2-50 nm), making it difficult to capture fine structural details (Tannert et al., 2017). While CT is effective for materials with larger and well-defined features, its inability to resolve pores below 50 nm leads to incomplete or inaccurate microstructural reconstructions of aerogels. Similarly, gas adsorption-based porosimetry methods, such as the Brunauer-Emmett-Teller (BET) and Barrett-Joyner-Halenda (BJH) methods, rely on simplified cylindrical pore models and assume uniform pore geometries. These assumptions often fail for aerogels, where pore shapes remain highly irregular, interconnected, and exhibit multi-scale heterogeneity. Furthermore, these techniques primarily estimate pore sizes based on adsorption-desorption isotherms, which can overlook macropores and inter-particle voids, leading to significant discrepancies when compared to direct imaging techniques such as SEM and transmission electron microscopy (TEM) (Horvat et al., 2022). Although SEM and TEM allow for direct visualization of the aerogel microstructure, these techniques are limited to two-dimensional projections. Consequently, they do not adequately capture the three-dimensional connectivity and interconnectivity of the pore network. It remains to be seen if more advanced techniques such as electron tomography in the so-called advanced mode can be applied to investigate diverse aerogels, as preliminary reports on silica aerogels have been promising (Roiban et al., 2016).

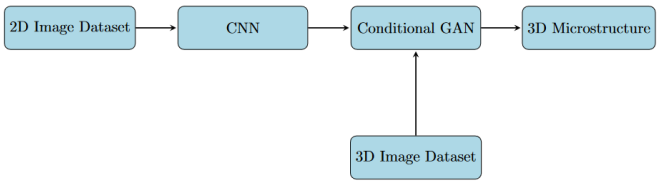

Figure 1: Hierarchy diagram that depicts the flow from the 2D and 3D datasets to the generated 3D microstructure

One potential approach to addressing these challenges is leveraging generative adversarial networks (GANs) to develop a mesoscopic digital twin for aerogels. Specifically, we explore the application of a conditional GAN (cGAN) for reconstructing 3D aerogel microstructures from a batch of 2D microstructure images. By incorporating convolutional neural networks (CNNs) to predict key structural parameters—such as pore radius, relative density, and pore size distribution—our framework aims to enable accurate structure-property predictions in a data-driven manner. Prior works have demonstrated the use of GANs for 3D microstructure reconstruction (Hsu et al., 2020; Kench & Cooper, 2021; Li et al., 2023); however, a key limitation of these approaches is the stochastic nature of their outputs, which can lead to variability in generated microstructures and hinder consistency in reconstruction fidelity. By conditioning the generative model on morphological features, our approach ensures that the generated 3D microstructures preserve essential physical characteristics. Other studies have applied similar microstructure conditioning techniques (Iyer et al., 2019); however, these efforts were confined to aligning the material with the processing parameters.

While physics-based simulations such as diffusion-limited cluster cluster aggregation (DLCA) or coarse-grained molecular dynamics can model microstructural formation, they can become computationally intensive—especially when the underlying physics includes fine-grained interaction potentials or time-resolved aggregation behavior. These methods are typically not well-suited for inverse design tasks or rapid sampling of diverse yet morphology-consistent structures across a parameter space. The proposed framework, in contrast, provides a scalable and energy-efficient alternative to traditional experimental characterization techniques and simulation-heavy approaches, offering a practical route toward generative microstructure modeling with structural fidelity. Figure 1 presents a schematic representation of the workflow.

2 METHODS

2.1 MICROSTRUCTURE DATA GENERATION

To provide sufficient training data, the CNN and the cGAN were both trained on microstructures produced using the diffusion-limited cluster-cluster aggregation (DLCA). The DLCA algorithm is well-established in the literature for accurately simulating the structural features and the kinetic characteristics involved in the formation of certain aerogels, such as silica aerogels (Abdusalamov et al., 2021). Furthermore, the structures derived from this algorithm can be easily processed to ascertain properties such as porosity, pore size distribution, and the pore network associated with the specified aerogel. The microstructure dataset was generated using the DLCA algorithm with particle radii between 2–7 nm, relative densities of 0.04–0.1, and a domain size of 300 nm. The DLCA model outputs particle coordinates, which are voxelized based on their radii to construct the final image array. This voxelization process generates a three-dimensional array with a predefined resolution (128^3) based on the domain size, particle coordinates, and their respective radii. Within this 3D array, each voxel is assigned a value of 1 if occupied by a particle and 0 if left unoccupied, resulting in the formation of a 3D binary image representation of the aerogel structure at the specified resolution. To generate a training dataset for the CNN while mitigating the resource-intensive nature of acquiring large-scale SEM images, we synthesize 2D SEM-like images from 3D microstructures. We utilize PORESPY visualization module (Gostick et al., 2019) to generate the image dataset. For a 3D binary structure of size 128^3, we apply a stacking factor of 8, producing 16 SEM-like grayscale images per sample. These synthesized images retain depth information, closely resembling experimentally acquired SEM images.

2.2 DEEP LEARNING FRAMEWORK

2.2.1 CNN MODEL AND TRAINING

To efficiently predict 3D structural features from 2D SEM-like images, we adopt a lightweight CNN architecture instead of computationally expensive models like U-Net, which would consume significant memory from the 3D cGAN. As a baseline, we extend the LeNet (LeCun et al., 1998) architecture by incorporating three convolutional layers interspersed with max pooling layers for feature extraction and dimensionality reduction. The fully connected layers are modified for regression, predicting key structural parameters such as pore radius, relative density, and pore size distribution. For training, synthetic 2D SEM-like images are preprocessed to enhance feature extraction. A dataset class loads images and parses filenames to extract target values. The dataset is initially split into 80% training and 20% testing. The training set is further partitioned into 80% training and 20% validation. The CNN is trained using mean squared error (MSE) loss with backpropagation, and the performance is monitored on the validation set to ensure generalization. The trained CNN is then used to condition the 3D cGAN, improving the fidelity of reconstructed aerogel microstructures.

2.3 CGAN MODEL

The baseline architecture is a Conditional Wasserstein Generative Adversarial Network (cWGAN), comprising a generator and a discriminator conditioned on structural properties, specifically relative density. The generator consists of six ConvTranspose3D layers with batch normalization and ReLU activations, except for a sigmoid activation in the final layer. The one-hot encoded relative density label is transformed via a linear layer to match the latent space before concatenation with the noise vector. The discriminator mirrors the generator, employing Conv3D layers with instance normalization and leaky ReLU activations. The relative density label is transformed into a 3D tensor (128^3) and is element-wise multiplied with the input before convolution. Wasserstein distance and gradient norm statistics are monitored during training to ensure stability. We apply an adaptive binarization to the generator's output, which employs a sigmoid activation in its final layer. The binarization threshold is set as the mean value μ of the generated volume $G(z)$, where z is the input noise. Each voxel v_{ijk} is then transformed as:

$$B(v_{ijk}) = \begin{cases} 1, & G(z)_{ijk} \geq \mu \\ 0, & \text{otherwise} \end{cases}$$

This ensures alignment with the binary training data while improving interpretability and structural consistency.

3 RESULTS

Three CNNs were trained to predict key 3D microstructural properties—radius, relative density, and pore size distribution mean (PSD mean)—from 2D slices. All models shared the same architecture as previously described and were trained on a dataset comprising 2D stacked slices annotated with corresponding global 3D properties. The CNNs mapped the 2D slice features during evaluation to predict the global 3D properties, achieving prediction accuracy within an acceptable range for the given dataset. The error distributions on the test set are visualized in histograms for each property. For radius predictions (Figure 2a), the CNN achieved low errors, with a mean error of approximately 0.1, and the results were within 10% of the true values. Similarly, predictions for the mean pore size distribution (PSD) (Figure 2b) exhibited slightly higher variance but remained within acceptable bounds. The predictions for relative density (Figure 2c) also showed a small spread of errors.

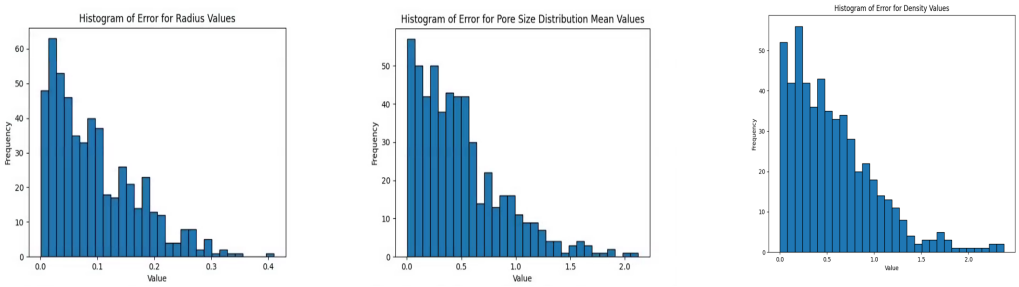

(a) Radius errors (b) PSD mean errors (c) Relative density errors

Figure 2: Histograms of the MAE between predicted and ground truth values for (a) particle radius, (b) mean pore size from the pore size distribution (PSD), and (c) relative density.

As a proof of concept, the cGAN was tested under a single conditioning parameter—the relative density of the 3D microstructure, which was predicted by the CNN from the 2D SEM image. The results demonstrated that the cGAN, when conditioned on a specific relative density, was able to generate aerogel microstructures corresponding to the given condition.

The generated microstructures effectively captured the baseline connectivity and overall structure of the aerogel. However, certain limitations were observed. The generator struggled to reproduce the spherical morphology of particles seen in the dataset, resulting in irregularly shaped blobs instead of well-defined spheres, which is also the case when reconstructing surfaces from CT scans. This discrepancy suggests that while the learned representations preserve global structural features, they may lack the granularity required to replicate fine morphological details seen in experimental samples. A more rigorous comparison with experimental datasets, such as reconstructed tomography images, would help quantify these deviations and refine the fidelity of the generated microstructures. Figure 3 illustrates three such outputs from the trained cGAN when conditioned over the relative density.

4 CONCLUSION

In this study, we introduced a conditional Generative Adversarial Network (cGAN) framework for the generation of 3D aerogel microstructures conditioned on relative density. The effectiveness of conditioning on relative density was demonstrated through the generation of 3D structures that exhibit realistic connectivity and density-dependent variability. While the generated microstructures capture the fundamental structural features, challenges exist in accurately reproducing particle shapes and sizes. Additionally, the cGAN framework has the potential to violate physical constraints inherent to aerogel morphologies, such as geometric packing limitations and inter-particle interactions, due to the model's reliance on data-driven generation rather than explicit physical principles.

This highlights the importance of carefully considering the training data and perhaps the need for incorporating physical constraints into the cGAN model to ensure the generation of physically plausible microstructures. This CGAn-based reconstruction approach, however, is not restricted to just the relative density of the microstructures; it can be extended to other microstructural parameters, such as pore size distribution and particle radius, as predicted by Convolutional Neural Networks (CNNs). A natural progression of this work would involve utilising the generated microstructures to predict material properties, such as thermal conductivity and mechanical characteristics, through the integration of the cGAN framework with physics-based simulations or data-driven predictive models. This work highlights the potential of deep learning-based generative models in advancing the design and characterization of novel materials, offering a data-driven alternative to traditional experimental methods.

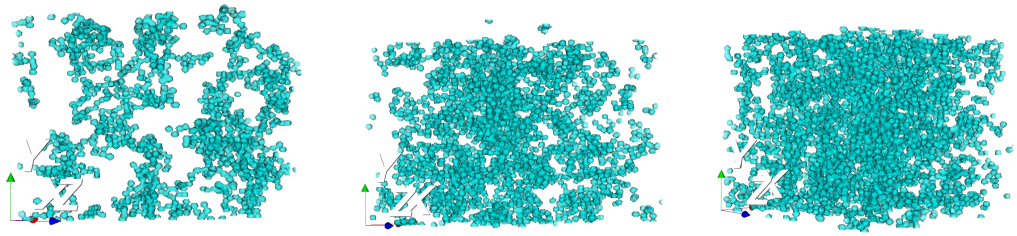

Dataset, $\rho = 0.03$ | Dataset, $\rho = 0.06$ | Dataset, $\rho = 0.09$

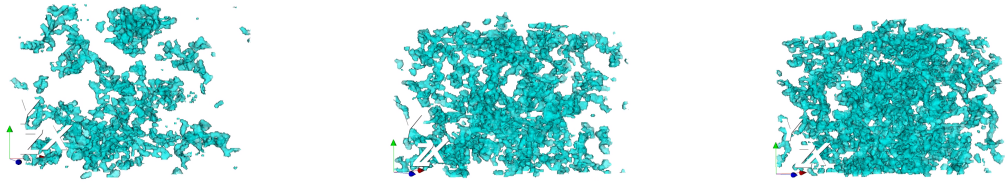

Generated, $\rho = 0.03$ | Generated, $\rho = 0.06$ | Generated, $\rho = 0.09$

Figure 3: Comparison of 3D voxel structures (128^3) at varying relative densities ρ. The top row shows dataset samples, and the bottom row shows corresponding generated structures.

5 ACKNOWLEDGMENTS

The authors would like to acknowledge the projects Fokusanwendungen Fahrzeugstruktur Antriebsstrang Energiemanagement (FFAE) and the Reducing Barriers for AI in Research (ReBAR) project of the German Aerospace Center (DLR) for their financial support and provision of the necessary computational resources. Special thanks are extended to the Federal Ministry for Economic Affairs and Climate Action (BMWK) for their financial backing of the research efforts of the group.

REFERENCES

Rasul Abdusalamov, Christian Scherdel, Mikhail Itskov, Barbara Milow, Gudrun Reichenauer, and Ameya Rege. Modeling and simulation of the aggregation and the structural and mechanical properties of silica aerogels. *The Journal of Physical Chemistry B*, 125(7):1944–1950, 2021.

Jeff T Gostick, Zohaib A Khan, Thomas G Tranter, Matthew DR Kok, Mehrez Agnaou, Mohammadamin Sadeghi, and Rhodri Jervis. Porespy: A python toolkit for quantitative analysis of porous media images. *Journal of Open Source Software*, 4(37):1296, 2019.

G Horvat, M Pantić, Ž Knez, and Z Novak. A brief evaluation of pore structure determination for bioaerogels. gels 8: 438, 2022.

Tim Hsu, William K. Epting, Hokon Kim, Harry W. Abernathy, Gregory A. Hackett, Anthony D. Rollett, Paul A. Salvador, and Elizabeth A. Holm. Microstructure generation via generative adversarial network for heterogeneous, topologically complex 3d materials, 2020. URL `https://arxiv.org/abs/2006.13886`.

Akshay Iyer, Biswadip Dey, Arindam Dasgupta, Wei Chen, and Amit Chakraborty. A conditional generative model for predicting material microstructures from processing methods. *arXiv preprint arXiv:1910.02133*, 2019.

Steve Kench and Samuel J Cooper. Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion. *Nature Machine Intelligence*, 3(4):299–305, 2021.

Samuel Stephens Kistler. Coherent expanded aerogels and jellies. *Nature*, 127(3211):741–741, 1931.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Kai-Qi Li, Yong Liu, and Zhen-Yu Yin. An improved 3d microstructure reconstruction approach for porous media. *Acta Materialia*, 242:118472, 2023.

Ameya Rege. Modeling the structural, fractal and mechanical properties of aerogels. In *Springer Handbook of Aerogels*, pp. 289–305. Springer, 2023.

L Roiban, G Foray, Qian Rong, Anouk Perret, Dris Ihiawakrim, Karine Masenelli-Varlot, E Maire, and Bernard Yrieix. Advanced three dimensional characterization of silica-based ultraporous materials. *RSC advances*, 6(13):10625–10632, 2016.

René Tannert, Marina Schwan, Ameya Rege, Mario Eggeler, Julio Cesar da Silva, Marion Bartsch, Barbara Milow, Mikhail Itskov, and Lorenz Ratke. The three-dimensional structure of flexible resorcinol-formaldehyde aerogels investigated by means of holotomography. *Journal of Sol-Gel Science and Technology*, 84:391–399, 2017.