3D Denoising Diffusion Probabilistic Models for 3D microstructure image generation of fuel cell electrodes

Abdelouahid Bentamou, Stéphane Chrétien, Yann Gavet

Computational Materials Science 248 (2025) 113596

Contents lists available at ScienceDirect

Computational Materials Science

journal homepage: www.elsevier.com/locate/commatsci

Full length article

3D Denoising Diffusion Probabilistic Models for 3D microstructure image generation of fuel cell electrodes

Abdelouahid Bentamou [a,*], Stephane Chretien [b], Yann Gavet [a]

[a] Mines Saint-Etienne, Univ Lyon, CNRS, UMR 5307 LGF, Centre SPIN, F - 42023 Saint-Etienne, France
[b] Laboratoire ERIC, Université Lyon 2, France

ARTICLE INFO

Keywords:
3D microstructure image
Solid Oxide Fuel Cells (SOFCs)
O_2 fuel cell electrode
3D Denoising Diffusion Probabilistic Models

ABSTRACT

The generation of realistic 3D microstructure images is crucial for understanding and optimizing materials in various fields, including fuel cell technology. In this article, we present a novel approach for generating synthetic 3D microstructure images using 3D Denoising Diffusion Probabilistic Models (3D DDPM). This approach extends to n-phase materials. Unlike conventional image generation techniques, our method leverages the principles of diffusion in three-dimensional space to simulate the intricate evolution of microstructures. By incorporating stochastic processes and diffusion equations, 3D DDPMs enable a more realistic and controlled representation of the dynamic processes occurring within materials. This approach generates synthetic microstructures that capture the spatial complexities inherent in real-world materials across multiple phases. Through experimental evaluation, we demonstrate that our approach generates realistic 3D microstructure images of O_2 fuel cell electrodes for two or three phases.

1. Introduction

Solid Oxide Fuel Cells (SOFCs) [1] represent a highly promising and innovative technology in the realm of electrochemical energy conversion. Unlike conventional fuel cells, SOFCs operate at elevated temperatures, typically between 500 to 1000 °C, enabling them to efficiently convert a variety of fuels, including hydrogen, natural gas, and even liquid hydrocarbons, into electricity. The distinctive feature of SOFCs lies in their solid electrolyte, which is typically composed of ceramic materials such as Yttria-stabilized zirconia (YSZ) [2]. This solid electrolyte allows the direct movement of oxygen ions (O^{2-}) from the cathode to the anode, facilitating the electrochemical reactions at the fuel cell interfaces. The high operating temperatures also contribute to improved electrochemical kinetics, resulting in higher efficiency and power density compared to other fuel cell types. SOFCs find applications in various sectors, including stationary power generation, distributed energy systems, and auxiliary power units for vehicles [3,4]. Additionally, their ability to directly utilize hydrocarbon fuels without requiring a separate reforming step makes them versatile and attractive for diverse energy scenarios.

SOFCs represent a cutting-edge technology for clean energy conversion; however, their practical implementation faces inherent limitations, primarily linked to the intricacies of electrode microstructure. The durability of SOFCs, a central concern, arises from the challenging operational conditions that induce material degradation [5]. The high temperatures, pronounced electrochemical potential gradients, and thermal and redox cycling inherent in SOFC operation contribute to material breakdown through mechanisms such as reactivity changes, poisoning, crystallographic alterations, and microstructural evolution. An in-depth examination reveals that the 3D microstructure of both hydrogen H_2 and oxygen O_2 electrodes plays a pivotal role in influencing SOFC performance.

Activation and concentration overpotentials, crucial factors in SOFC efficiency, are intricately linked to microstructural features. The availability of active sites and the distribution of pores within the electrode significantly influence electrochemical reactions and mass transport processes. Suboptimal microstructure can exacerbate issues related to electrode polarization, limiting overall cell performance and contributing to reduced durability. Therefore, a comprehensive understanding of the microstructure becomes imperative in addressing the limitations associated with SOFCs [6]. By scrutinizing and optimizing electrode microstructure, researchers and engineers can tailor materials and designs to enhance the resilience of SOFCs, ultimately paving the way for more robust and durable fuel cell systems in practical applications.

Various methods exist for the 3D reconstruction of microstructures, each offering distinct advantages and drawbacks. Traditional imaging techniques, such as scanning electron microscopy (SEM) and transmission electron microscopy (TEM) [7], provide high-resolution 2D images

* Corresponding author.
E-mail address: abdelouahid.bentamou@gmail.com (A. Bentamou).

https://doi.org/10.1016/j.commatsci.2024.113596
Received 10 September 2024; Received in revised form 4 December 2024; Accepted 5 December 2024
Available online 14 December 2024

that can be stacked to form 3D reconstructions. While these methods offer exceptional spatial resolution, they are often time-consuming, labor-intensive, and may suffer from limited sample sizes. X-ray tomography [8], including synchrotron X-ray tomography, enables non-destructive 3D imaging with good resolution, making it suitable for a range of materials. However, these methods can be resource-intensive, requiring access to specialized facilities. Advanced imaging techniques, such as focused ion beam scanning electron microscopy (FIB-SEM) [9] and serial block-face imaging, allow for automated serial sectioning and imaging, facilitating 3D reconstructions. Although these techniques offer improved efficiency, they may compromise spatial resolution compared to traditional methods. Additionally, the associated equipment and processing requirements can be expensive. Each method has its strengths and limitations, and the choice depends on the specific goals of the study, considering factors such as resolution, sample size, and resource availability.

To address the challenges associated with experimental methods for 3D reconstruction of microstructures, an alternative and complementary approach involves the generation of synthetic microstructures through computational modeling. This approach offers several advantages, including cost-effectiveness, scalability, and flexibility. By leveraging computational tools and algorithms, researchers can simulate complex microstructures with controlled parameters, allowing for a systematic exploration of various configurations and conditions. Synthetic microstructure generation provides a platform to investigate a wide range of scenarios that may be impractical or challenging to achieve experimentally. Furthermore, this approach facilitates the exploration of the impact of specific microstructural features on material properties, enabling a more targeted and efficient optimization process. Synthetic microstructures also serve as valuable benchmarks for validating and refining experimental techniques. Overall, the generation of synthetic microstructures complements experimental methods, offering a powerful tool for gaining insights into the intricacies of materials at the microscale while mitigating some of the constraints associated with traditional imaging and reconstruction techniques.

While Generative Adversarial Networks (GANs) [10] have demonstrated remarkable capabilities in generating realistic images, they are not without certain limitations. One primary hurdle stems from the inherent instability of the training process, a consequence of the adversarial nature of the loss function that pits the generator against the discriminator. This instability manifests as mode collapse, restricting GAN models to generating samples within a limited range of probabilities. Consequently, addressing this issue amplifies the complexity of fine-tuning hyperparameters and refining the training process. To address these limitations, an alternative and promising approach for synthetic image generation is the utilization of diffusion probabilistic models [11]. Unlike GANs, diffusion probabilistic models explicitly leverage stochastic processes and diffusion equations to simulate the gradual transformation of a simple noise distribution into a more complex target distribution, closely mimicking the evolution of real-world microstructures. By iteratively transforming a noise distribution, these models offer a nuanced and realistic simulation of microstructure evolution, making them a compelling choice for synthetic image generation in materials science applications.

The initial applications of diffusion models for 2D microstructure generation were pioneered by Lee et al. [12] and Düreth et al. [13]. Recently, Huang et al. [14] extended this paradigm to encompass 3D microstructure generation from 2D images. Their contribution, the Micro3Diff framework, introduces a unique perspective on 2D-to-3D reconstruction of microstructures using diffusion models. However, these methods are limited to binary images, addressing only two phases.

In this article, we advocate for the adoption of 3D diffusion models as a novel and effective approach for the generation of synthetic 3D microstructure images, extending the application to n-phases. Unlike conventional image generation techniques, our proposed method leverages the principles of diffusion in three-dimensional space to simulate the intricate evolution of microstructures. By incorporating stochastic processes and diffusion equations, the 3D diffusion models enable a more realistic and controlled representation of the dynamic processes occurring within materials. This approach allows for the generation of synthetic microstructures, it captures the spatial complexities inherent in real-world materials across multiple phases.

The remainder of this article is organized as follows. Section 2 reviews related works on microstructure image generation. Section 3 introduces our proposed approach, 3D Diffusion Models, for generating 3D microstructure images with multiple phases. Section 4 details the dataset used in this study, presents the experimental results, and provides a comparative analysis. Finally, Section 5 offers conclusions and discusses future perspectives.

2. Related works

Microstructure image generation is a rapidly evolving field with the potential to revolutionize material science and engineering. Machine learning has emerged as a powerful tool for this task, enabling the creation of realistic microstructure images from a variety of data sources.

2.1. Stochastic geometrical modeling techniques

Stochastic geometrical modeling techniques have been developed for generating 3D microstructure images. These methods leverage randomness to create realistic microstructures that capture the inherent variability and intricate patterns observed in real materials.

Random models provide a general framework for generating microstructures by imposing statistical constraints on the distribution of microstructural features. Poisson Point Processes [15], where the models generate microstructures by randomly placing points in a domain according to a Poisson distribution. Boolean Random Sets [16,17] represent microstructures as collections of objects with random shapes and sizes. The interactions between these objects are governed by Boolean logic operations, such as union, intersection, and exclusion. Gaussian Random Fields [18,19] are characterized by a mean and covariance function that determine the overall distribution and spatial correlations of microstructural features. Random graphs [20] provide a network-based representation of microstructures, where nodes represent microstructural features and edges represent connections between them. These graphs can be used to capture the topology and connectivity of microstructures, which play a crucial role in determining their physical properties.

2.2. Generative Adversarial Networks (GANs)

GANs [10] have emerged as a powerful tool for microstructure image generation, surpassing the capabilities of traditional stochastic methods in terms of both realism and flexibility. GANs consist of two competing neural networks: a generator which aims to produce realistic microstructure images, and a discriminator which attempts to distinguish between real and fake images. This adversarial training process forces the generator to learn the underlying patterns in the training data and produce images that are indistinguishable from real ones.

Numerous variants of GANs have been developed for microstructure image generation. While most works initially employed standard GANs, researchers have explored specific variants to enhance performance [21]. For instance, Gayon-Lombardo et al. [22] implemented DC-GAN to generate realistic n-phase microstructures with periodic boundaries. Kishimoto et al. [23] utilized a conditional GAN for synthesizing porous microstructures, introducing a volume fraction loss for control during training. Lambard et al. [24] showcased the use of StyleGAN for inflating scanning electron microscopy image

datasets and reproducing coarse and fine microstructural details. Fokina et al. [25] employed the StyleGAN architecture for microstructure synthesis, proposing image quilting to merge fixed-sized samples.

Some works have gone further by modifying GANs to generate more accurate microstructure images. Zhang et al. [26] introduced an RNN-based generative model incorporating GANs for 2D-to-3D reconstruction of porous media, showcasing an innovative approach. Nguyen et al. [27] combined GANs with actor-critic reinforcement learning for synthesizing realistic microstructures with controlled properties.

A notable trend in microstructure image generation involves generating 3D microstructures from 2D images. Slice-wise Generative Adversarial Networks (SliceGANs) [28] have been designed specifically for this purpose. SliceGANs decompose 3D microstructures into a series of 2D slices, generating each slice individually using a 2D GAN. Kench et al. [29] introduced SliceGAN, capable of generating high-fidelity 3D microstructure images from a single 2D image. Subsequently, they leveraged SliceGAN for generating comprehensive 3D microstructural datasets from 2D micrographs. Sugiura et al. [30] applied SliceGAN to generate 3D microstructure images from 2D images in ferrite-martensite dual-phase steels.

In addition to SliceGAN, other approaches involve DC-GAN, such as Liu et al.'s application for the generation of a 3D microstructure of a fuel cell catalyst layer [31], and Zhao et al.'s proposal for generating high-quality 3D microstructural images from 2D images of hardened cement paste using solid texture synthesis [32]. Kononov et al. [33] presented an algorithm using GANs for reconstructing 3D microstructures from 2D images, incorporating an invariant reconstruction error.

In summary, the versatility and adaptability of GANs and their variants in microstructure image generation, coupled with innovations like SliceGANs, underscore their pivotal role in advancing materials science and computational imaging.

2.3. Denoising Diffusion Probabilistic Models (DDPM)

In recent times, DDPM have emerged as promising tools for microstructure image generation. These models operate by initiating the generation process with a noisy image and systematically denoising it until it transforms into a pristine microstructure image. Guided by a learned probabilistic distribution, this process ensures the creation of realistic images that align with the patterns observed in the training data.

Initially, DDPM found application in 2D microstructure generation. Lee et al. [34] introduced a denoising diffusion model specifically tailored for microstructure reconstruction. The method's robustness was demonstrated across various microstructure types, including polycrystalline alloys, carbonates, ceramics, copolymers, and fiber composites, each possessing distinct morphological characteristics. Düreth et al. [13] extended the use of DDPM to generate 2D microstructure images, applying them to diverse and intricate real micrograph data as well as a smaller fiber dataset.

Expanding the scope to 3D microstructure generation, Lee et al. [14] proposed the Micro3Diff framework for 2D-to-3D reconstruction of microstructures using diffusion-based generative models (DGMs). This innovative approach relies solely on pre-trained DGMs for generating 2D samples, with the dimensionality expansion (2D-to-3D) occurring dynamically during the generation process. Notably, the framework incorporates the concept of multi-plane denoising diffusion, which effectively transforms noisy samples from different planes into a coherent data structure while preserving spatial connectivity in the 3D space. This advancement enhances the fidelity and spatial coherence of the generated 3D microstructure images, showcasing the potential of diffusion models in bridging the gap between 2D and 3D representations. Recently, Lyu et al. [35] used DDPM to reconstruct 2D microstructures of various composite materials. Then, they extended to 3D microstructure reconstruction with defined permeability range.

All aforementioned works focus on generating binary (two-phase) microstructure images. In contrast, our research pioneers the use of 3D DDPM for the generation of 3D microstructure images with multiple phases, thereby generalizing the approach for n-phases. This novel method leverages the principles of 3D diffusion to simulate the intricate evolution of microstructures, enabling a more realistic and controlled representation of the dynamic processes occurring within materials. By incorporating stochastic processes and diffusion equations in three-dimensional space, our approach captures the spatial complexities inherent in real-world materials, offering a significant improvement in the fidelity and applicability of generated microstructure images.

3. Proposed approach

3.1. Preprocessing of the training set

In our study, the initial image data, stored as 8-bit grayscale elements, uses voxel values as material labels to indicate the material they represent. For the O_2 electrode with 2 phases, this translates to black (0) for pores and white (255) for Lanthanum Strontium Cobalt Ferrite (LSCF). For the O_2 electrode with 3 phases, black (0) denotes pores, gray (127) represents Gadolinium-doped Ceria (GDC), and white (255) corresponds to LSCF.

When using diffusion models for material analysis, materials with two phases could rely on a single number expressing confidence in a specific phase. However, for materials with three or more phases, this approach becomes problematic due to potential misinterpretation. For example, a voxel with low confidence in deciding between black or white might produce gray (i.e., a value halfway between black and white), leading to potential misinterpretation.

A solution is "one-hot" encoding, which adds a dimension to create a 4D volume, transforming the initial 3D cubic volume of $w \times h \times d$, where w, h and d denote width, height and depth respectively, into a 4D $c \times w \times h \times d$ volume, where c represents the number of material phases present. Each material phase has a '1' at its location and '0' elsewhere. Thus, black, gray and white become [1,0,0], [0,1,0], and [0,0,1] respectively. Decoding these 4D volumes back to 3D grayscale is straightforward, taking the maximum value as the label, even in cases of labeling uncertainty [22].

This technique offers the advantage of extending DDPM implementation to generate materials with multiple phases. Therefore, the method introduced in this study is not limited to three-phase data; it can be implemented for materials with any number of phases, denoted as n-phase materials.

3.2. 3D denoising diffusion probabilistic models (3D DDPM) for 3D microstructure image generation

The adoption of 3D DDPM as a novel approach involves the integration of stochastic processes and diffusion equations within three-dimensional space. By simulating how microstructural elements diffuse and interact in 3D, these models capture the spatial complexities inherent in real-world materials. The use of diffusion in three dimensions allows for a more faithful representation of the dynamic processes occurring within the microstructure.

The 3D DDPM involves generating a 3D image, denoted as x_0, by progressively moving a starting 3D image $x_T \sim \mathcal{N}(0, I)$ closer to the data distribution through multiple denoising steps $x_{T-1}, \ldots, x_0$. 3D DDPM consist of two diffusion processes: forward diffusion and reverse diffusion.

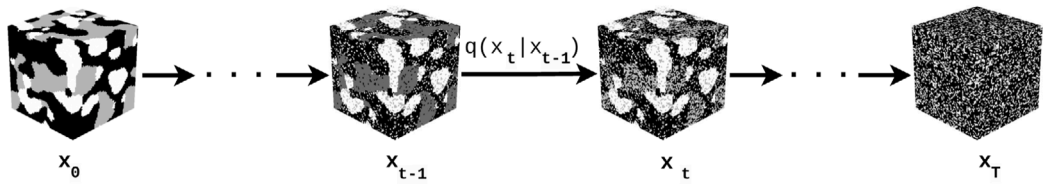

Fig. 1. Forward diffusion process.

3.2.1. Forward diffusion

3D DDPM are constructed as a Markov chain comprising T steps, where the dependency at each step is solely on the preceding one, an assumption of moderate nature (c.f. Fig. 1). Commencing with a 3D data point x_0 drawn from the real data distribution $q(x)$ ($x_0 \sim q(x)$). In our case, x_0 represents the one-hot encoded representation of the original 3D microstructure data.

The forward diffusion process involves progressively corrupting x_0 by adding Gaussian noise at each step of the Markov chain, generating a sequence of 3D latent variables x_t. At each timestep t, the new variable x_t is drawn from a Gaussian distribution with mean $\mu_t = \sqrt{1-\beta_t}x_{t-1}$ and variance $\Sigma_t = \beta_t I$. This process is expressed mathematically as:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \mu_t = \sqrt{1-\beta_t}x_{t-1}, \Sigma_t = \beta_t \mathbf{I}) \tag{1}$$

Here, $\mathcal{N}(\mu_t, \Sigma_t)$ denotes a Gaussian distribution defined by its mean μ_t and covariance matrix Σ_t. In the multidimensional context, $\mathbf{I}$ denotes the identity matrix, implying equal standard deviation β_t for each dimension. The forward diffusion process gradually increases noise through the variance parameter β_t, which is either constant or follows a schedule. We adopt a cosine schedule for β_t [36], gradually increasing from $\beta_0 = 0$ to $\beta_T = 1$ over T timesteps, ensuring controlled noise addition at each step.

It is important to note that at each timestep, x_t is sampled from the Gaussian distribution $q(x_t|x_{t-1})$, meaning a new value for x_t is generated by perturbing x_{t-1} based on the specified mean and variance. This ensures the noise is systematically applied to the entire dataset, transitioning x_t to a fully noisy state at $t = T$.

The complete forward diffusion process can be represented as:

$$q(x_{1:T}|x_0) = \prod_{t=1}^{T} q(x_t|x_{t-1}) \tag{2}$$

The symbol : in $q(x_{1:T}|x_0)$ indicates the repeated application of q from timestep $t = 1$ to $t = T$, commonly known as a trajectory. To streamline the process, we can directly obtain 3D x_t from 3D x_0 in a single step using the closed form:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon \tag{3}$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$, and $\alpha_t = (1-\beta_t)$. As β_t remains a hyperparameter, we can precompute α_t and $\bar{\alpha}_t$ for all timesteps, allowing for the sampling of noise at any timestep t and obtaining 3D x_t in one step. This serves as our target for calculating the tractable objective loss L later in the process.

At $t = 0$, the mean of the distribution corresponds to the original data x_0, while at $t = T$, the mean becomes zero, indicating a fully noisy state where all traces of the original microstructure have been removed. This forward diffusion process is central to training, where the model learns to reverse this noise addition, ultimately enabling the generation of new images from pure noise during inference.

3.2.2. Reverse diffusion

The reverse process in 3D DDPM is a crucial aspect, aiming to unravel the denoising steps performed during the forward process and recover the less noisy 3D image x_{t-1} from the more noisy 3D image x_t as illustraded in Fig. 2. This intricate procedure involves determining the posterior distribution for x_{t-1} given x_t and the initial 3D image x_0.

The posterior distribution is defined as follows:

$$q(x_{t-1}|x_t) \sim \mathcal{N}(x_{t-1}; \tilde{\mu}_t(x_t, t), \tilde{\Sigma}_t(x_t, t)) \tag{4}$$

where $\tilde{\Sigma}_t = \tilde{\beta}_t \mathbf{I}$. $\tilde{\mu}_t$ and $\tilde{\beta}_t$ are calculated as follows:

$$\tilde{\mu}_t = \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}x_0 + \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}x_t \tag{5}$$

$$\tilde{\beta}_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t \tag{6}$$

Furthermore, from Eq. (3), we can represent x_0 as:

$$x_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t - \sqrt{1-\bar{\alpha}_t}\epsilon) \tag{7}$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. By combining the last two equations, each timestep now has a mean $\tilde{\mu}_t$ (our target) dependent only on x_t:

$$\tilde{\mu}_t(x_t) = \frac{1}{\sqrt{\alpha_t}}(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon) \tag{8}$$

Thus, we can utilize a 3D neural network $\epsilon_\theta(x_t, t)$ to approximate ϵ and consequently the mean:

$$\tilde{\mu}_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}}(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(x_t, t)) \tag{9}$$

The training loss (L) is formulated to optimize the 3D neural network's prediction of the noise ϵ. Ho et al. [11] introduced a simple loss function for enhanced stability during training:

$$L = \mathbb{E}_{t,x_0}\|\epsilon - \epsilon_\theta(x_t, t)\|^2 \tag{10}$$

Here, $\mathbb{E}$ represents the expectation operator, averaging over all timesteps t and input data samples x_0. In this equation, ϵ represents the true noise added during the forward diffusion process, and $\epsilon_\theta(x_t, t)$ is the neural network's prediction of the noise at timestep t.

The complete training and sampling procedures are detailed in Algorithm 1 and Algorithm 2, respectively. An essential part of the training process is the optimization step described in Step 5 of Algorithm 1. At this stage, the 3D neural network model parameters θ are updated using gradient descent to minimize the loss function L (Eq. (10)). The gradient of this loss function, $\nabla_\theta L$, is computed, and a step is taken in the negative gradient direction to update the model parameters:

$$\theta \leftarrow \theta - \eta\nabla_\theta L \tag{11}$$

where η is the learning rate. This iterative optimization enables the neural network ϵ_θ to learn to predict the noise accurately, which is crucial for effectively reversing the diffusion process during inference. By minimizing this loss across all timesteps t, the model achieves a refined ability to reconstruct the input data or generate new data samples starting from random noise. In this paper, we take $T = 1000$.

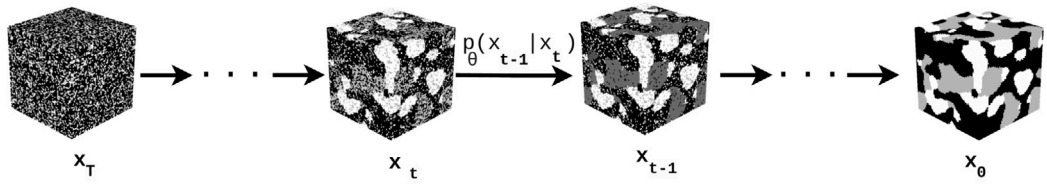

Fig. 2. Reverse diffusion process. During the training phase, the objective is to restore the original image from noise, enabling the model to learn the mapping between noisy and denoised representations. In the inference phase, only the reverse diffusion process is used to generate a new image, starting from pure noise and iteratively reconstructing the synthetic microstructure.

This optimization step is a cornerstone of the training process and ensures that the diffusion model learns the underlying data distribution effectively.

Algorithm 1 Training

repeat
  $x_0 \sim q(x_0)$
  $t \sim \text{Uniform}(\{1, \dots, T\})$
  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$
  Take gradient descent step on $\nabla_\theta \|\epsilon - \epsilon_\theta(x_t, t)\|^2$
until converged

Algorithm 2 Sampling

$x_T \sim \mathcal{N}(0, \mathbf{I})$
for $t = T$ to 1 **do**
  $z \sim \mathcal{N}(0, \mathbf{I})$ if $t > 1$, else $z = 0$
  $x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) + \sigma_t z$
end for
return x_0

3.2.3. 3D U-Net architecture

While DDPM offer considerable flexibility and compatibility with various deep learning architectures, the U-Net architecture [37] is particularly favored for the iterative denoising process within DDPM due to its matching input and output dimensions. U-Net takes a 3D noised image at a particular time step and return the predicted noise. Among the array of architectural enhancements explored for DDPM, integrating residual blocks and attention modules, including self-attention and cross-attention, has shown to yield superior results in data synthesis quality. Accordingly, this study adopts a 3D U-Net architecture featuring residual blocks and multi-head attention, as depicted in Fig. 3.

The 3D U-Net architecture, depicted in Fig. 3, is a deep learning model specifically designed for volumetric image segmentation. By extending the traditional U-Net architecture to include three-dimensional convolutional operations, the 3D U-Net is particularly effective for tasks requiring spatial context in three dimensions.

The network architecture is structured as follows:

- **Initial Convolution**: The network starts with a convolutional layer [38], which processes the input 3D noisy data to extract initial features. In parallel, positional embeddings are computed to encode information about the noise levels, enabling the network to differentiate between timesteps during training.
- **Downsampling Stages**: The input features are passed through several downsampling stages to capture information at different resolutions. Each stage includes:
  - Convolutional blocks to learn feature representations.
  - Normalization layers to stabilize training by ensuring consistent feature scaling [39].
  - Attention mechanisms to focus on important spatial and channel-wise features in the data [40].
  - Residual connections to retain key information across layers and avoid the loss of detail [41].
  - A downsampling operation to reduce spatial dimensions while increasing the depth of feature maps.
- **Middle Blocks**: At the network's center, additional convolutional blocks and attention mechanisms are applied to refine the learned features and integrate context from the entire input volume.
- **Upsampling Stages**: After processing in the middle blocks, the network reconstructs the original resolution through upsampling stages. Each stage mirrors the downsampling stages but includes an upsampling operation to increase spatial dimensions. This process progressively refines feature maps, preparing them for the final output.
- **Final Convolution**: The network concludes with a final convolutional block, followed by a convolutional layer, to produce the desired output.

This architectural design enables the 3D U-Net to efficiently capture and utilize spatial information across multiple scales, making it highly effective for segmenting complex 3D microstructures.

4. Experimental results

4.1. Dataset

Our research has greatly benefited from privileged access to a confidential database graciously provided by CEA. This exclusive repository encompasses intricately detailed 3D images meticulously focused on the oxygen (O_2) electrode, featuring both two and three distinct phases. These high-fidelity images were acquired through the utilization of the FIB-SEM technique.

For the O_2 electrode with two phases, the original training image comprises 750^3 voxels, subsequently segmented into a binary voxel representation. Herein, black voxels delineate the pore space, while white voxels depict the LSCF material structure. Similarly, for the O_2 electrode with three phases, the original training image, also consisting of 750^3 voxels, underwent segmentation into a trinary voxel representation. In this case, black voxels denote the pore space, gray voxels signify the GDC material, and white voxels represent the LSCF material structure.

To construct the training dataset, subvolumes of dimensions 64^3 were extracted from these voxelized images. Ideally, the training images would be independent domains, but the constrained size of the original image necessitated the use of overlapping subvolumes. With a spacing of 25 voxels, this process yielded a total of 23,234 training images for each case, whether it involved 2 phases or 3 phases.

4.2. Evaluation criteria

4.2.1. Morphological measures

Minkowski Functionals [42], a set of mathematical measures derived from differential geometry and topology, offer a powerful framework for quantifying and characterizing the geometric properties of

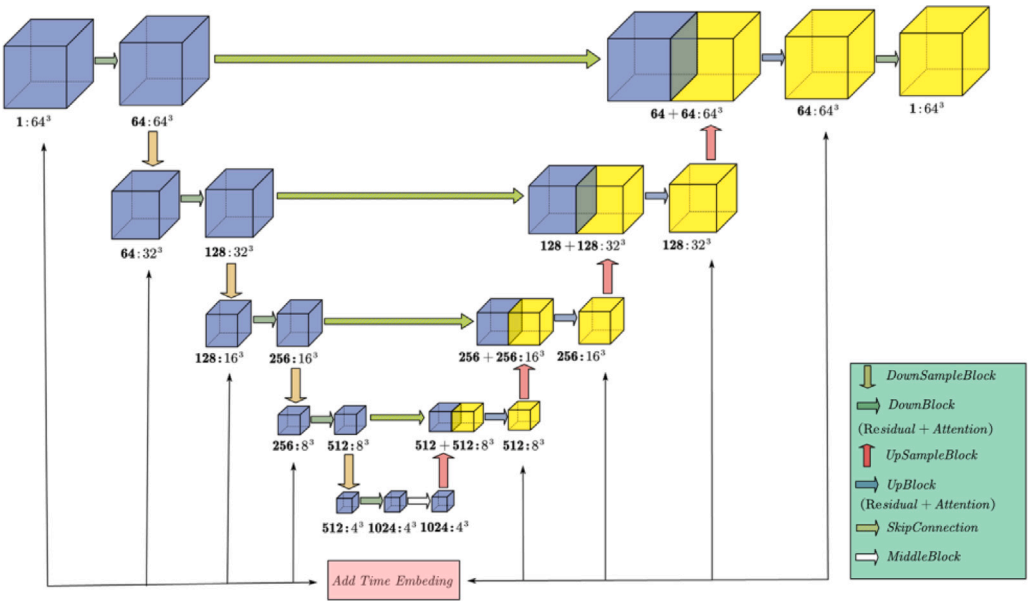

Fig. 3. 3D U-Net network architecture in 3D DDPM [35].

multi-phase structures. Defined by Hermann Minkowski, these functions provide insightful metrics to analyze the shape, structure, and complexity of objects in three-dimensional space. The four primary Minkowski Functionals include the volume fraction ϕ, the specific surface area S_V, the mean integral curvature C, and the Euler characteristic χ_V. In this study, we use ϕ, S_V, and χ_V.

- **Volume Fraction** ϕ_i indicates the proportion of the total volume occupied by a particular phase, reflecting its contribution to the overall structure.
- **Specific Surface Area** S_{Vi} measures the surface area per unit volume of a given phase, impacting interfacial processes such as adsorption and reaction rates.
- **Euler Characteristic** χ_{Vi} provides insight into the connectivity and topology of the phase, influencing transport and flow properties.

The densities of the Minkowski functionals were computed by dividing by the volume of the microstructure V_m [43]. These can be mathematically expressed as:

$$\phi_i = \frac{V_i}{V_m} \tag{12}$$

$$S_V^{(i)} = \frac{1}{V_m} \int_{\text{surface of phase } i} dS \tag{13}$$

$$\chi_V^{(i)} = \frac{1}{V_m}(Ve_i - E_i + F_i - O_i) \tag{14}$$

where Ve_i is the number of vertices, E_i is the number of edges, F_i is the number of faces, and O_i is the number of objects in phase i.

4.2.2. Two-Point Statistics

The quality of the generated image was assessed using the two-point probability function, which is equivalent to the non-centered covariance function $S_2(r)$. This metric was compared to values derived from the original training images to evaluate the accuracy of the generated image across multiple phases.

$$S_2^{(i)}(r) = \mathbf{P}(x \in P_i, x + r \in P_i) \quad \text{for } x, r \in \mathbb{R}^d \tag{15}$$

This function represents the probability **P** that two points, x and $x + r$, separated by the lag vector r, are both within phase P_i. At the origin, $S_2^{(i)}(0)$ corresponds to the volume fraction (or porosity in the case of porous media) of phase P_i. As r increases to infinity, $S_2^{(i)}(r)$ stabilizes around the ϕ_i^2.

4.3. Results

In this section, we present the evaluation of our approach using 3D diffusion models to generate 3D microstructure images. The dataset comprises images with either two or three phases. We evaluate the quality of the generated images by comparing the two-point statistics $S_2(r)$ and Minkowski functionals for each phase, followed by a comparison of our method with state-of-the-art methods based on Minkowski functionals.

4.3.1. Two-phase images

The Fig. 4 provides a comparative overview of the generated 3D microstructure images by the diffusion models (a) alongside the training 3D microstructure images (b). On a visual level, the images on both sides exhibit remarkable similarity.

The visual comparison between (a) and (b) underscores the efficacy of the diffusion models in replicating the essential characteristics of the training images. The generated microstructures maintain the overall structural integrity and detailed features observed in the real images. The distribution of phases, the geometric arrangements, and the overall textural features are captured, demonstrating the potential of diffusion models for generating realistic 3D microstructures.

This visual analysis confirms that, on a superficial level, the diffusion models are proficient in creating synthetic images that are remarkable similarity with their real counterparts, paving the way for further quantitative and qualitative evaluations.

Two-Point Statistics $S_2(r)$

Following the visual assessment of the microstructures, we now analyze the Two-Point Statistics $S_2(r)$ function for both the porous and LSCF phases (c.f. Fig. 5). For the porous phase, a faster decay of $S_2(r)$ in the original data compared to the synthetic data indicates a higher probability of finding interconnected pores at larger separation distances in the original microstructure. This suggests a network of smaller and potentially more isolated pores in the original phase. Similarly, for the LSCF phase, a faster decay of $S_2(r)$ in the original data suggests that the original microstructure has a more spread-out or interconnected arrangement of LSCF particles compared to the synthetic data. This might be due to limitations in the training data; the training data might not have included a wide enough variety of phase distributions. Evaluating the $S_2(r)$ function for both phases highlights the strengths and limitations of the DDPM in capturing the detailed microstructural characteristics of these phases.

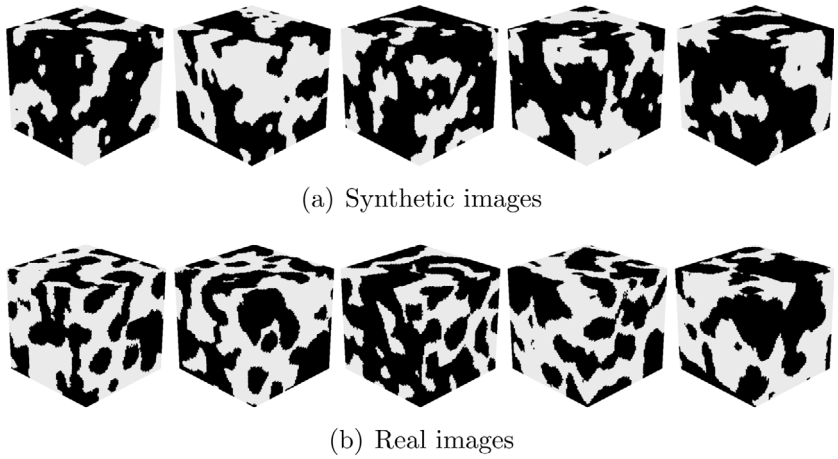

(a) Synthetic images

(b) Real images

Fig. 4. Examples of 3D generated microstructure images by the diffusion models (top) compared to 3D training microstructure images (bottom). The images depict two phases: black representing the porous phase and white representing the LSCF phase.

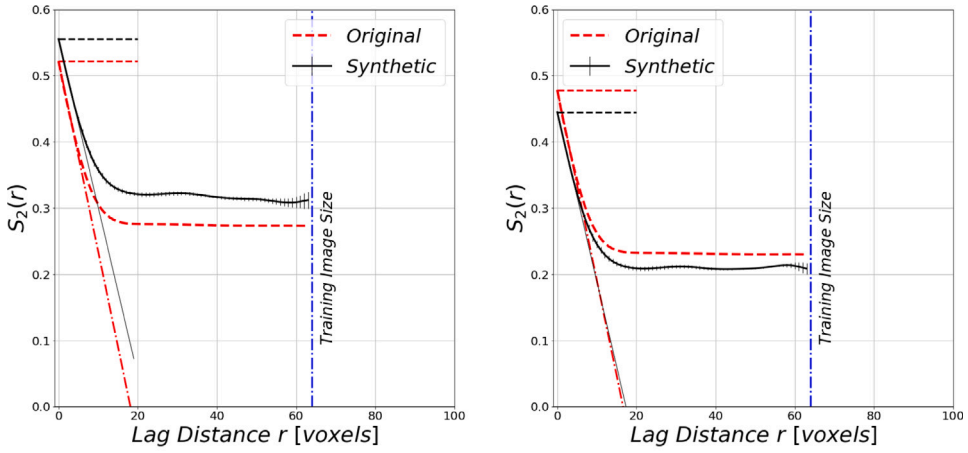

(a) Porous phase (b) LSCF phase

Fig. 5. Two-Point Statistics $S_2(r)$ functions for both the porous (top) and LSCF phases (bottom) comparing original and synthetic microstructure images.

Morphological Measures

The analysis of the Minkowski functionals (c.g. Fig. 6) for both the porous and LSCF phases within the fuel cell electrode microstructure reveals key insights into the geometric characteristics captured by the 3D DDPM. For the porous phase, the distributions of volume fraction, specific surface area, and specific Euler characteristic indicate that the 3D DDPM effectively captures the overall porosity but shows a narrower distribution in specific surface area and specific Euler characteristic. This suggests that while the model replicates the general porosity, it tends to generate pores of more uniform size and less varied network complexity compared to the original microstructures. Similarly, for the LSCF phase, the 3D DDPM captures the overall volume fraction but produces a narrower distribution in specific surface area and specific Euler characteristic, indicating a more uniform size and shape of LSCF particles with less varied connectivity in the synthetic data. Overall, the 3D DDPM demonstrates partial success in replicating the geometric properties of both phases, effectively capturing the overall volume fractions but not fully replicating the detailed variations in size, shape, and network complexity of the original microstructures.

4.3.2. Three-phase images

In Fig. 7, we provide an overview of the generated 3D microstructure images by 3D DDPM juxtaposed with the training 3D microstructure images. On a visual level, the images on both sides (a: generated images) and (b: real images) are generally similar, showcasing the model's capability to replicate the complex structures of the three phases: black for the porous phase, gray for the GDC phase, and white for the LSCF phase.

Two-Point Statistics $S_2(r)$

The Two-Point Statistics Analysis $S_2(r)$ for the porous, GDC, and LSCF phases (c.f. Fig. 8) highlights the differences in pore-size distribution and particle connectivity between the original and synthetic data generated by the 3D DDPM. For the porous phase, the faster decay of $S_2(r)$ in the original data indicates a higher probability of interconnected pores over larger distances, implying a more connected microstructure compared to the synthetic data. Similarly, for the GDC phase, the original data shows a faster decay in the $S_2(r)$ function, suggesting that GDC particles in the original microstructure form a more interconnected network across larger distances than in the synthetic microstructure. The LSCF phase also exhibits this trend, with the original data's $S_2(r)$ function decaying more rapidly, indicating a higher probability of interconnected LSCF particles at larger separation distances. Overall, the original microstructures consistently demonstrate greater connectivity and more extensive networks than their synthetic counterparts.

Morphological Measures

Fig. 9 illustrates the Minkowski functionals of the generated alongside the original training images for each phase. The analysis of Minkowski functionals across the porous, GDC, and LSCF phases reveals that the 3D DDPM effectively replicates the overall volume fractions,

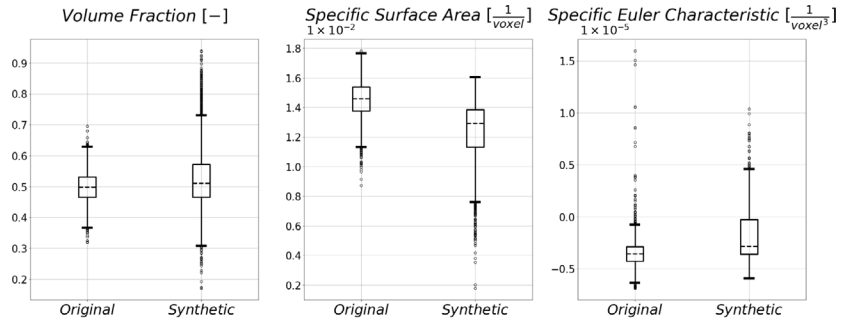

(a) Porous phase

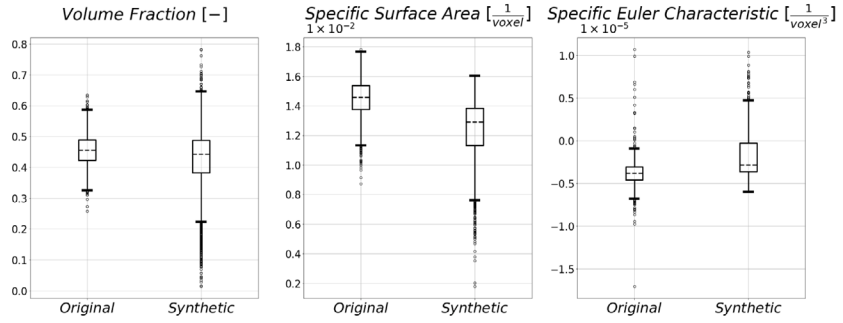

(b) LSCF phase

Fig. 6. Comparison of three Minkowski functionals for the original images and synthetic images. The top subfigure represents the porous phase, while the bottom subfigure represents the LSCF phase.

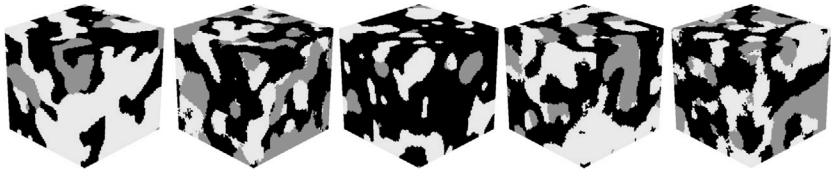

(a) Synthetic images

(b) Real images

Fig. 7. Examples of 3D generated microstructure images by the diffusion models (top) compared to 3D training microstructure images (bottom). The images depict three phases: black representing the porous phase, gray representing the GDC phase and white representing the LSCF phase.

indicating accurate phase proportioning. However, the Specific Surface Area and Specific Euler Characteristic metrics show notable discrepancies. For all phases, the synthetic data exhibit less variability and complexity compared to the original data, suggesting a more uniform and simplified structure. This is particularly evident in the GDC and LSCF phases, where the synthetic specific surface areas are more consistent, and the Euler characteristics indicate simpler connectivity networks. While the 3D DDPM captures the general structure, it falls short in replicating the intricate details and variability of the original microstructures, highlighting areas for improvement in modeling complex material properties.

4.4. Comparison with state-of-the-art methods

The comparative analysis based on the Table 1 of the geometric and topological properties of two phases, Porous and LSCF, across real and simulated images (GAN, SliceGAN, and 3D DPPM), reveals key differences and similarities. For the Porous phase, the volume fraction is most accurately captured by 3D DPPM (0.53) compared to the real images (0.497), with GAN and SliceGAN also showing close approximations. The specific surface area for the Porous phase is slightly lower in 3D DPPM (0.013) than in real images (0.014), while GAN and SliceGAN slightly overestimate it (0.015). The Euler characteristic indicates that

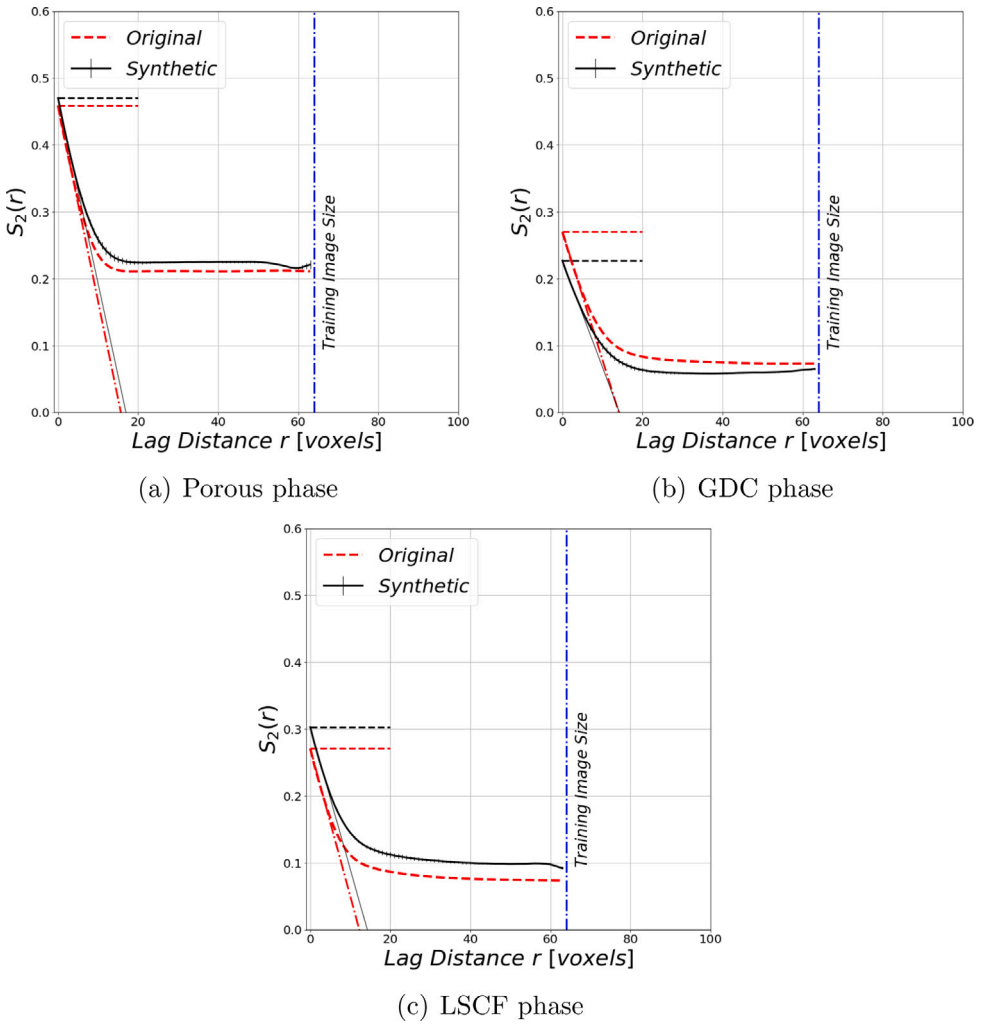

Fig. 8. Two-Point Statistics $S_2(r)$ functions for the porous (top), GDC (medium) and LSCF phases (bottom) comparing original and synthetic microstructure images.

3D DPPM produces less complex structures (-2×10^{-5}) than the real images (-3.48×10^{-5}), with GAN and SliceGAN also generating structures of higher complexity but still differing from the real images. For the LSCF phase, the volume fraction is best matched by SliceGAN (0.477) compared to real images (0.456), while 3D DPPM and GAN show minor deviations. The specific surface area in 3D DPPM (0.013) is again slightly lower than in real images (0.014), with GAN and SliceGAN both overestimating it (0.015). The Euler characteristic suggests that 3D DPPM generates less complex LSCF structures (-2.01×10^{-5}) than real images (-3.77×10^{-5}), while GAN and SliceGAN produce more complex structures. Overall, 3D DPPM appears to offer a closer approximation of the real microstructures, particularly in terms of volume fraction and specific surface area, although it tends to produce less topologically complex structures compared to the real images.

The Table 2 presents the comparative analysis of geometric and topological properties of three distinct phases (Porous, GDC, and LSCF). The comparative analysis reveals varied performances among the GAN, SliceGAN, and 3D DDPM models. Both GAN and SliceGAN models show strong alignment with real images, particularly in replicating volume fraction and specific surface area across all phases, with minor deviations indicating their effectiveness in capturing the overall structure and surface complexity. However, the 3D DDPM model exhibits more significant deviations, especially in the GDC phase's volume fraction and Euler characteristic, suggesting potential areas for improvement in accurately modeling complex connectivity and distribution. Notably, the 3D DDPM's higher volume fraction for the Porous phase and lower specific surface area for GDC indicate a trend towards generating smoother, less intricate structures. Overall, while all models demonstrate capability in generating realistic microstructures, GAN and SliceGAN show closer fidelity to the real images, highlighting the need for further refinement in the 3D DDPM approach to enhance its accuracy in representing complex microstructural properties.

Table 1
Comparative properties of real and simulated 3D microstructure images for two phases: Porous and LSCF.

	Real images	Simulated images		
		GAN [44]	SliceGAN [29]	3D DPPM
Porous phase				
Volume fraction	0.497	0.503	0.476	0.53
Specific surface area	0.014	0.015	0.015	0.013
Euler characteristic	−3.48e−05	−4.23e−05	−4e−05	−2e−05
LSCF phase				
Volume fraction	0.456	0.451	0.477	0.423
Specific surface area	0.014	0.015	0.015	0.013
Euler characteristic	−3.77e−05	−4.24e−05	−4.18e−05	−2.01e−05

5. Conclusion

The study presents a novel approach for generating 3D microstructure images using 3D Denoising Diffusion Probabilistic Models (3D DDPM). By integrating stochastic processes and diffusion equations

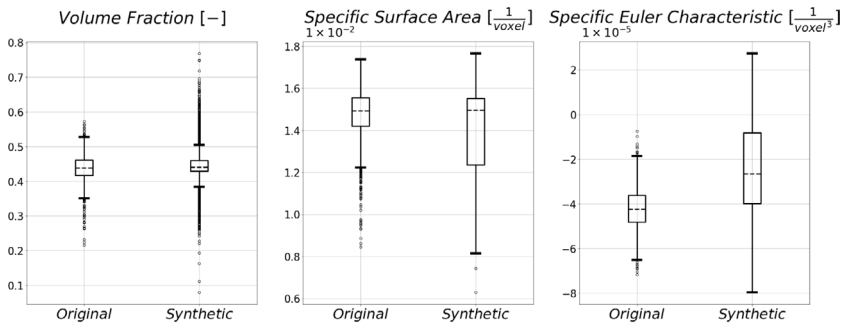

(a) Porous phase

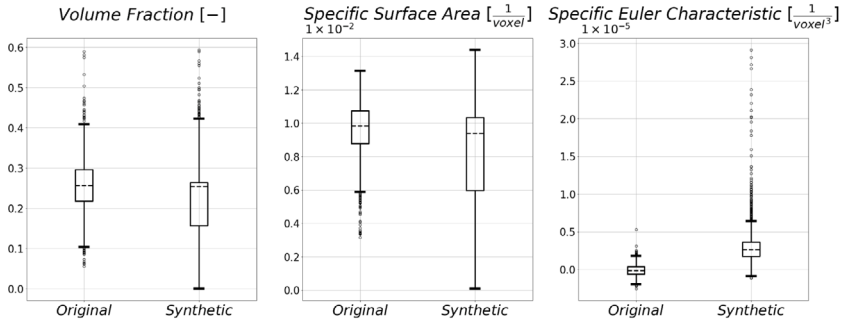

(b) GDC phase

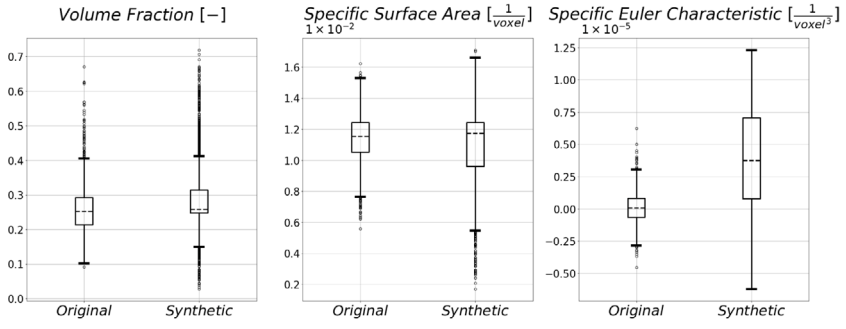

(c) LSCF phase

Fig. 9. Comparison of three Minkowski functionals for the original images and synthetic images. The top represents the porous phase, the medium represents the GDC phase, and the bottom represents the LSCF phase.

Table 2
Comparative properties of real and simulated 3D microstructure images of various methods for 3 phases: Porous, GDC, and LSCF.

Method	Property	Phases		
		Porous	GDC	LSCF
Real images	Volume fraction ϕ_i	0.437	0.259	0.258
	Specific surface area S_V	0.015	0.01	0.011
	Euler characteristic χ_V	−4.21e−05	−1.1e−06	8.91e−07
GAN [44]	Volume Fraction ϕ_i	0.434	0.262	0.258
	Specific Surface Area S_V	0.015	0.01	0.011
	Euler Characteristic χ_V	−3.52e−05	1.16e−05	3.57e−05
SliceGAN [29]	Volume fraction ϕ_i	0.445	0.252	0.257
	Specific surface area S_V	0.014	0.01	0.011
	Euler characteristic χ_V	−3.89e−05	−2.55e−06	3.7e−06
3D DDPM	Volume fraction ϕ_i	0.458	0.17	0.324
	Specific surface area S_V	0.013	0.006	0.01
	Euler characteristic χ_V	−8.86e−06	3.92e−05	8.92e−06

in a three-dimensional space, this method successfully captures the spatial complexities and dynamic processes inherent in real-world materials. The implementation of 3D DDPM extends beyond three-phase data, demonstrating its applicability to n-phase materials. This advancement offers significant potential for material science, allowing for the controlled and realistic simulation of intricate microstructures. Future research will focus on refining these models and exploring their applications across various types of materials.

CRediT authorship contribution statement

Abdelouahid Bentamou: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Stephane Chretien:** Writing – review & editing, Visualization, Validation, Methodology, Formal analysis. **Yann Gavet:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Abdelouahid BENTAMOU reports financial support was provided by National Graduate School of Mines Saint-Etienne. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We express our sincere gratitude to Jérôme Laurencin for generously providing the invaluable 3D microstructure data that served as the foundation for this study.

The authors would also like to express their sincere gratitude to Bpifrance for their valuable support and funding. This work was conducted as part of the UDD project (Usines De Demain) in collaboration with Bpifrance, contributing to advancements in object detection under uncontrolled acquisition environment.

Data availability

The data that has been used is confidential.

References

[1] R.M. Ormerod, Solid oxide fuel cells, Chem. Soc. Rev. 32 (1) (2003) 17–28.
[2] M.A. Yatoo, F. Habib, A.H. Malik, M.J. Qazi, S. Ahmad, M.A. Ganayee, Z. Ahmad, Solid-oxide fuel cells: A critical review of materials for cell components, MRS Commun. 13 (3) (2023) 378–384.
[3] A. Choudhury, H. Chandra, A. Arora, Application of solid oxide fuel cell technology for power generation—A review, Renew. Sustain. Energy Rev. 20 (2013) 430–442.
[4] S.C. Singhal, Solid oxide fuel cells for stationary, mobile, and military applications, Solid State Ion. 152 (2002) 405–410.
[5] J. Mouginn, J. Laurencin, J. Vulliet, M. Petitjean, E. Grindler, S. Di Iorio, K. Couturier, T. Dejob, B. Gonzalez, G. Cubizolles, et al., Recent highlights on solid oxide cells, stacks and modules developments at CEA, ECS Trans. 111 (6) (2023) 1101.
[6] M. Kleitz, F. Petitbon, Optimized SOFC electrode microstructure, Solid State Ion. 92 (1–2) (1996) 65–74.
[7] B.J. Inkson, Scanning electron microscopy (SEM) and transmission electron microscopy (TEM) for materials characterization, in: Materials Characterization using Nondestructive Evaluation (NDE) Methods, Elsevier, 2016, pp. 17–43.
[8] E. Lay-Grindler, J. Laurencin, J. Villanova, P. Cloetens, P. Bleuet, A. Mansuy, J. Mougin, G. Delette, Degradation study by 3D reconstruction of a nickel–yttria stabilized zirconia cathode after high temperature steam electrolysis operation, J. Power Sources 269 (2014) 927–936.
[9] T.A. Prokop, K. Berent, H. Iwai, J.S. Szmyd, G. Brus, A three-dimensional heterogeneity analysis of electrochemical energy conversion in SOFC anodes using electron nanotomography and mathematical modeling, Int. J. Hydrog. Energy 43 (21) (2018) 10016–10030.
[10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, Adv. Neural Inf. Process. Syst. 27 (2014).
[11] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, Adv. Neural Inf. Process. Syst. 33 (2020) 6840–6851.
[12] H.J. Lim, K.-H. Lee, G.J. Yun, Microstructure design of multifunctional particulate composite materials using conditional diffusion models, Adv. Mater. 7 (2023) 9.
[13] C. Düreth, P. Seibert, D. Rücker, S. Handford, M. Kästner, M. Gude, Conditional diffusion-based microstructure reconstruction, Mater. Today Commun. 35 (2023) 105608.
[14] K.-H. Lee, G.J. Yun, Multi-plane denoising diffusion-based dimensionality expansion for 2D-to-3D reconstruction of microstructures with harmonized sampling, NPJ Comput. Mater. 10 (1) (2024) 99.
[15] J. Escoda, D. Jeulin, F. Willot, C. Toulemonde, Three-dimensional morphological modelling of concrete using multiscale Poisson polyhedra, J. Microsc. 258 (1) (2015) 31–48.
[16] D. Jeulin, Morphology and effective properties of multi-scale random sets: A review, C. R. Méc. 340 (4–5) (2012) 219–229.
[17] F.S. Bourgeois, G.J. Lyman, F. Courteille, Modelling multi-scale microstructures with combined Boolean random sets: A practical contribution, Chem. Eng. Res. Des. 87 (10) (2009) 1390–1399.
[18] H. Moussaoui, J. Debayle, Y. Gavet, P. Cloetens, J. Laurencin, Particle-based model for functional and diffusion layers of solid oxide cells electrodes, Powder Technol. 367 (2020) 67–81.
[19] H. Moussaoui, J. Laurencin, Y. Gavet, G. Delette, M. Hubert, P. Cloetens, T. Le Bihan, J. Debayle, Stochastic geometrical modeling of solid oxide cells electrodes validated on 3D reconstructions, Comput. Mater. Sci. 143 (2018) 262–276.
[20] L. Théodon, J. Laurencin, M. Hubert, P. Cloetens, J. Debayle, A stochastic geometrical 3D model for time evolution simulation of microstructures in SOC-electrodes, Comput. Mater. Sci. 212 (2022) 111568.
[21] A. Henkes, H. Wessels, Generative adversarial networks for three-dimensional microstructure generation, PAMM 22 (1) (2023) e202200064.
[22] A. Gayon-Lombardo, L. Mosser, N.P. Brandon, S.J. Cooper, Pores for thought: generative adversarial networks for stochastic reconstruction of 3D multi-phase electrode microstructures with periodic boundaries, NPJ Comput. Mater. 6 (1) (2020) 82.
[23] M. Kishimoto, Y. Matsui, H. Iwai, Conditional GAN for generation of 3D SOFC electrode microstructure dataset, ECS Trans. 111 (6) (2023) 503.
[24] G. Lambard, K. Yamazaki, M. Demura, Generation of highly realistic microstructural images of alloys from limited data with a style-based generative adversarial network, Sci. Rep. 13 (1) (2023) 566.
[25] D. Fokina, E. Muravleva, G. Ovchinnikov, I. Oseledets, Microstructure synthesis using style-based generative adversarial networks, Phys. Rev. E 101 (4) (2020) 043308.
[26] F. Zhang, X. He, Q. Teng, X. Wu, J. Cui, X. Dong, PM-ARNN: 2D-TO-3D reconstruction paradigm for microstructure of porous media via adversarial recurrent neural network, Knowl.-Based Syst. 264 (2023) 110333.
[27] P.C. Nguyen, N.N. Vlassis, B. Bahmani, W. Sun, H. Udaykumar, S.S. Baek, Synthesizing controlled microstructures of porous media using generative adversarial networks and reinforcement learning, Sci. Rep. 12 (1) (2022) 9034.
[28] S. Kench, I. Squires, A. Dahari, S.J. Cooper, MicroLib: A library of 3D microstructures generated from 2D micrographs using SliceGAN, Sci. Data 9 (1) (2022) 645.
[29] S. Kench, S.J. Cooper, Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion, Nat. Mach. Intell. 3 (4) (2021) 299–305.
[30] K. Sugiura, T. Ogawa, Y. Adachi, Hourly work of 3D microstructural visualization of dual phase steels by SliceGAN, Adv. Theory Simul. 5 (7) (2022) 2200132.
[31] X. Liu, K. Park, M. So, S. Ishikawa, T. Terao, K. Shinohara, C. Komori, N. Kimura, G. Inoue, Y. Tsuge, 3D generation and reconstruction of the fuel cell catalyst layer using 2D images based on deep learning, J. Power Sources Adv. 14 (2022) 100084.
[32] X. Zhao, X. Wu, L. Wang, P. Hou, Q. Li, Y. Zhang, B. Yang, Three-dimensional microstructural image synthesis from 2d backscattered electron image of cement paste, 2022, arXiv preprint arXiv:2204.01645.
[33] E. Kononov, M. Tashkinov, V.V. Silberschmidt, Reconstruction of 3d random media from 2d images: generative adversarial learning approach, Comput. Aided Des. 158 (2023) 103498.
[34] K.-H. Lee, G.J. Yun, Microstructure reconstruction using diffusion-based generative models, Mech. Adv. Mater. Struct. (2023) 1–19.
[35] X. Lyu, X. Ren, Microstructure reconstruction of 2D/3D random materials via diffusion-based deep generative models, Sci. Rep. 14 (1) (2024) 5041.
[36] A.Q. Nichol, P. Dhariwal, Improved denoising diffusion probabilistic models, in: International Conference on Machine Learning, PMLR, 2021, pp. 8162–8171.
[37] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, Springer, 2015, pp. 234–241.
[38] K. Uchida, M. Tanaka, M. Okutomi, Coupled convolution layer for convolutional neural network, Neural Netw. 105 (2018) 197–205.
[39] J.L. Ba, Layer normalization, 2016, arXiv preprint arXiv:1607.06450.
[40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Adv. Neural Inf. Process. Syst. 30 (2017).
[41] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
[42] K.R. Mecke, Additivity, convexity, and beyond: applications of Minkowski functionals in statistical physics, in: Statistical Physics and Spatial Statistics: The Art of Analyzing and Modeling Spatial Structures and Pattern Formation, Springer, 2000, pp. 111–184.
[43] A.M. Boelens, H.A. Tchelepi, QuantImPy: Minkowski functionals and functions with python, SoftwareX 16 (2021) 100823.
[44] L. Mosser, O. Dubrule, M.J. Blunt, Reconstruction of three-dimensional porous media using generative adversarial neural networks, Phys. Rev. E 96 (4) (2017) 043309.