Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials

Tim Hsu, William K. Epting, Hokon Kim, Harry W. Abernathy, Gregory A. Hackett, Anthony D. Rollett, Paul A. Salvador & Elizabeth A. Holm

JOM

The Journal of The Minerals, Metals & Materials Society (TMS)

ISSN 1047-4838 Volume 73 Number 1

JOM (2021) 73:90-102 DOI 10.1007/s11837-020-04484-y



Your article is protected by copyright and all rights are held exclusively by The Minerals, Metals & Materials Society. This e-offprint is for personal use only and shall not be selfarchived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".





AUGMENTING PHYSICS-BASED MODELS IN ICME WITH MACHINE LEARNING AND UNCERTAINTY QUANTIFICATION

Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials

TIM HSU,^{1,2,6} WILLIAM K. EPTING,^{1,3} HOKON KIM,^{1,2} HARRY W. ABERNATHY,^{4,5} GREGORY A. HACKETT,⁴ ANTHONY D. ROLLETT,^{1,2} PAUL A. SALVADOR,^{1,2} and ELIZABETH A. HOLM ^{1,2,7}

1.—US DOE National Energy Technology Laboratory, Pittsburgh, PA 15236, USA. 2.—Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA. 3.—Leidos Research Support Team, Pittsburgh, PA 15236, USA. 4.—US DOE National Energy Technology Laboratory, Morgantown, WV 26505, USA. 5.—Leidos Research Support Team, Morgantown, WV 26505, USA. 6.—Lawrence Livermore National Laboratory, Livermore, CA 94550, USA. 7.—e-mail: eaholm@andrew.cmu.edu

Using a large-scale, experimentally captured 3D microstructure data set, we implement the generative adversarial network (GAN) framework to learn and generate 3D microstructures of solid oxide fuel cell electrodes. The generated microstructures are visually, statistically, and topologically realistic, with distributions of microstructural parameters, including volume fraction, particle size, surface area, tortuosity, and triple-phase boundary density, being highly similar to those of the original microstructure. These results are compared and contrasted with those from an established, grain-based generation algorithm (DREAM.3D). Importantly, simulations of electrochemical performance, using a locally resolved finite element model, demonstrate that the GAN-generated microstructures closely match the performance distribution of the original, while DREAM.3D leads to significant differences. The ability of the generative machine learning model to recreate microstructures with high fidelity suggests that the essence of complex microstructures may be captured and represented in a compact and manipulatable form.

INTRODUCTION

As a component of integrated computational materials engineering (ICME), methods to generate realistic simulation volumes are essential for modeling and simulating materials with complex microstructures.¹ While obtaining microstructures experimentally guarantees realistic simulation volumes, the cost or difficulty of microstructural characterization often limits the size or number of microstructures that can be sampled, particularly in 3D. Thus, to support computational design and performance surveys, a common goal is to synthesize statistically representative sets of microstructural realizations.² A number of successful approaches have been developed, such as those based on n-point correlation functions,^{3,4} ellipsoid packing,^{5,6} physical

descriptors,⁷ Gaussian random field,⁸ and Markov random field.⁹ Bostanaband et al.¹⁰ provide an extensive review of existing microstructure reconstruction techniques. However, despite their successes, most of them require assumptions about the underlying structure and thus are often limited to either binary (two-phase) microstructures, uniform microstructures consisting of primitive/regular geometries, or a combination of both; there is an intrinsic limit of generality that prohibits the use of these statistical methods for a wide range of heterogeneous, complex material systems.

In recent years, deep learning-based generative models, such as variational autoencoders¹¹ or generative adversarial networks (GANs),¹² have attracted the interests of materials scientists for their potential in microstructure image generation.^{13–16} Relative to other reconstruction methods, generative models based on deep convolutional

⁽Received June 24, 2020; accepted October 28, 2020; published online December 2, 2020)

neural networks make no assumption regarding the underlying structure of the image data. Such models can therefore be vastly general. Additionally, these models have the ability to learn the underlying data distribution of their input data set. In other words, a well-trained generative model can (theoretically) synthesize an infinite number of unique microstructure realizations from the learned data distribution at relatively trivial speeds and compute costs.

Most generative models for microstructure focus on 2D representations.^{13–17} For example, Yang et al.,¹³ Cang et al.,¹⁴ and Singh et al.¹⁶ successfully applied them to represent 2D binary (two-phase) microstructure images. Iyer et al.¹⁵ explored the use of a conditional GAN,^{18,19} and Fokina et al.¹⁷ used a stvle-based GAN²⁰ for capturing complex, multiphase, grayscale 2D microstructure images. There is some prior work focused on 3D microstructure representation: Mosser et al.^{21,22} utilized a GAN to generate 3D porous media microstructures, including from binary image information²² and from uniform microstructures made of closely packed spherical grains (i.e., simple or regular geometries).²¹ Most recently, Gayon-Lombardo et al.² applied GANs to generate 3D, three-phase microstructural volumes with periodic boundary conditions that permit tiling to arbitrary system sizes and achieved good agreement with the target structures. These initial works confirm the general utility of generative models in representing microstructures.

In this article, we demonstrate that the GAN framework with deep convolutional layers²⁴ can learn and generate grayscale, heterogeneous, complex multiphase microstructures in 3D at a scale appropriate for simulation of material properties. Our model material is a three-phase commercial solid oxide fuel cell (SOFC) anode, as shown in Fig. 1. The original microstructure image data were experimentally captured using a Xe plasma focused

ion beam combined with scanning electron microscope (Xe PFIB-SEM)^{25,26} and represents the largest 3D volume of SOFC material characterized to date. Although the microstructure consists of three phases, we train our GAN model with the microstructure data in grayscale (SEM signal) format, without prior segmentation, to emphasize and test the generality of the GAN framework. The synthetic microstructures produced by the GAN model are both visually and statistically realistic when compared with the experimental microstructure and with microstructures synthesized by an established, grain-based generation algorithm (DREAM.3D).^{5,6} Importantly, simulations of electrochemical performance, using a locally resolved finite element model,^{27,28} demonstrate that the GAN-generated microstructures closely match the performance distribution of the real SOFC material. The ability of the machine learning model to recreate microstructures with high fidelity makes it a valuable addition to the ICME tool set.

METHODS

Experimental Microstructure Data Set

As shown in Fig. 1, the SOFC anode microstructure is a three-phase, 3D composite of highly interconnected yttria-stabilized zirconia (YSZ), nickel (Ni), and pores. This material comprises sintered grains with characteristic length scales on the order of 0.5 μ m and exhibits substantial variability in volume fractions over the length scale of 15 μ m (or 30 times the characteristic length scale).^{25,26} The extent and connectivity of triplephase boundaries (TPB), where YSZ, Ni, and pore phases meet in a line junction, governs the performance of the material. The TPB network topology is complex and interconnected, necessitating a 3D microstructural representation for accurate performance modeling.



Fig. 1. Complex, multiphase microstructure of a commercial solid oxide fuel cell (SOFC) anode containing yttria-stabilized zirconia (YSZ, light gray), nickel (Ni, dark gray), and pore phases (transparent). The $110 \times 124 \times 8 \ \mu m^3$ 3D volume was obtained by PFIB-SEM; bright spots are imaging artifacts. The inset shows an example 2D cross-section, imaged by SEM, where YSZ is bright gray, Ni is dark gray, and pores are black.

The microstructure training data set originates from the active anode layer of a commercial SOFC sample supplied by Materials and Systems Research, Inc. (Salt Lake City, UT), which was imaged using PFIB-SEM previously \mathbf{as} described.^{25,26} The resulting microstructure datum is a large-scale, 3D volumetric grayscale image whose dimension is $1697 \times 1900 \times 124$ voxels, or $110 \times 124 \times 8 \ \mu m^3$; thus, the voxel size is $65 \times 65 \times$ 65 nm³. Each phase images with a characteristic grayscale intensity, and microscopy artifacts may appear as well. Subvolumes of this microstructure data set are randomly sampled during training sessions.

DREAM.3D Synthetic Microstructures

As described elsewhere, 26,29 synthetic microstructures were constructed using the DREAM.3D code package⁶ (BlueQuartz Software, Springboro, OH), which generates multiphase microstructures via an ellipsoid packing scheme. The parameters specified were phase fraction and the first and second moments of the log-normal particle size distribution, which were chosen to match as closely as possible the measured vales for each phase in the anode microstructure. Forty distinct microstructures of $12.48 \times 12.48 \times 12.48 \ \mu m^3$ were first generated, and they were then split into octants to produce 320 subvolumes of $6.24 \times 6.24 \times 6.24 \ \mu m^3$ ($96 \times 96 \times 96$ voxels).

WGAN Implementation

Convolutional neural networks (CNNs) are a class of artificial neural networks particularly well suited to image data and computer vision applications.³⁰ In its forward mode, a CNN takes in an image and applies a series of signal processing steps (typically filter convolutions, rectification operations, and pooling or downsampling) to generate a compact representation of the visual information contained in the image, termed the feature vector. The feature vector can then be used for a variety of image processing tasks, including visual similarity and classification.³¹ The architecture of the CNN, including the type, number, order, and connectivity of the operations, is user-defined, and the training process attempts to optimize the many internal variables for best performance.

CNNs can also operate in reverse, taking in a feature vector and transforming it into an image; this is the same basic concept as an auto-encoder.³² As shown in Fig. 2, we use a variant of GAN called the Wasserstein generative adversarial network (WGAN).³³ The model consists of two CNN components: The generator, in reverse mode, creates candidate images, which the critic, in forward mode, attempts to output in an approximated measure of dissimilarity (Wasserstein loss) between the syn-

thetic and the real batches of microstructure images. Both networks participate in a feedback loop, so the generator learns to make more convincing synthetic images, and the critic learns to better tell apart synthetics from real. This training dynamic eventually leads to realistic images generated by the generator.^{12,33} Although the distinction of GAN and WGAN is important to some, in the following sections, we will refer to our model as GAN and WGAN interchangeably.

We implemented the WGAN model using the Tensorflow platform³⁴ with the Keras API (https:// keras.io/). The model architecture, given in Table I, largely follows the work by Miyato et al.³⁵ on spectral normalization and is specifically based on their standard architecture. However, for the generator, we used upsampling and convolutional layers. rather than standard deconvolutional (transposed convolution) layers, to avoid potential checkerboard artifacts.³⁶ We used a public code repository at https://github.com/IShengFang/Spectr alNormalizationKeras to implement spectral normalization applied to convolutional and dense layers in the critic network. The Adam optimizer³² was used for training the model. Relevant hyperparameters are shown in Table II.

During training, for each critic update, a batch of subvolumes randomly sampled from the experimental 3D microstructure and a batch of 3D images randomly generated from the generator are fed into the critic network to compute the Wasserstein loss³³ and the subsequent gradient update. The sampling process involves cropping and symmetry operations (rotation and mirror-flipping) of the cropped subvolumes to increase the perceived training data variability for the critic. Note that performing such symmetry operations will result in the GAN model not being able to learn microstructural anisotropy. The microstructure studied in this work does not visually exhibit any strong anisotropy.

We trained the model in a distributed fashion with 64 high-performing GPUs (Nvidia Tesla P100 PCIe) on a supercomputer (NETL Joule) using the Horovod framework.³⁷ The training took about 39 h, or 32,565 iterations of generator update. Due to the nature of the synchronous distributed training model implemented by Horovod, the effective minibatch size is the original minibatch size multiplied by the number of GPU workers, or $8 \times 64 = 512$. Although Goyal et al.³⁸ have shown that under synchronous distributed training, the learning rate can be scaled linearly with respect to the minibatch size without accuracy loss for classification tasks using ResNet-50,³⁹ we have observed that the WGAN training dynamic is highly sensitive to adjustments made to the learning rate. Therefore, we did not scale the learning rate in order to preserve training stability.

Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials



Batch of synthetic images

Fig. 2. Schematic of the WGAN model. The generator transforms a random noise vector *z* into a 3D image, and the critic outputs the dissimilarity metric (Wasserstein loss) between real and synthetic images. The real microstructural volume images are sampled from the experimental data, and the synthetic images are generated by the generator network.

Layer	Operation(s)	Output dimensions	Kernel size	Stride	Batch normalization	Activation function
Genera	tor: input = $z \in \mathbb{R}$	$1^{100} \sim \mathcal{N}(0,1)$				
G_{in}	Dense reshape	(6, 6, 6, 512)	-	-	_	Linear
G_1	2× Upsample Conv3D	(12, 12, 12, 256)	4	1	Yes	ReLU
G_2	2× Upsample Conv3D	(24, 24, 24, 128)	4	1	Yes	ReLU
G_3	2× Upsample Conv3D	(48, 48, 48, 64)	4	1	Yes	ReLU
G_4	$2 \times$ Upsample Conv3D	(96, 96, 96, 32)	4	1	Yes	ReLU
G_5	Conv3D	(96, 96, 96, 1)	3	1	No	Tanh
Critic: i	nput = gravscale	image $x \in \mathbb{R}^{96 \times 96 \times 96 \times 1}$				
D_1	ConvSN3D	(48, 48, 48, 32)	3	1	No	lReLU
-	ConvSN3D		4	2	No	lReLU
D_2	ConvSN3D	(24, 24, 24, 64)	3	1	No	lReLU
	ConvSN3D		4	2	No	lReLU
D_3	ConvSN3D	(12, 12, 12, 128)	3	1	No	lReLU
	ConvSN3D		4	2	No	lReLU
D_4	ConvSN3D	(6, 6, 6, 256)	3	1	No	lReLU
	ConvSN3D		4	2	No	lReLU
D_5	ConvSN3D	(6, 6, 6, 512)	3	1	No	lReLU
D_{out}	DenseSN	(1)	-	-	_	Linear

Table I. A	rchitecture of	f the	WGAN	model f	for 3D	microstructure	generation
------------	----------------	-------	------	---------	--------	----------------	------------

All convolutional layers use the "same" padding. The slope for all leaky ReLU (lReLU) activations is 0.1. Upsample and Reshape operations are not associated with kernel size, stride, batch normalization, and activation function. Batch normalization operation is applied after activation function, not before.

Segmentation

Experimental and GAN synthetic grayscale microstructures were segmented into three phases using a custom watershed-based segmentation code, which could be implemented in a high-throughput fashion; the major steps are as follows:

- 1. Sobel filters were applied to an input grayscale image to approximate image gradients.
- 2. A 2D density map was generated, plotting the grayscale intensity in voxels (x-axis) versus the gradient between voxels (y-axis).
- 3. Three high-density regions in the density map were then manually bounded to indicate what

Table II. Model hyperpa	rameters
-------------------------	----------

Symbol	Description	Value
α	Learning rate for both generator and critic (initial)	0.00005
α	Learning rate for both generator and critic (after 27 epochs)	0.00001
α	Learning rate for both generator and critic (after 78 epochs)	0.000005
α	Learning rate for both generator and critic (after 147 epochs)	0.000001
β_1	First decay rate for Adam optimizer	0.1
β_2	Second decay rate for Adam optimizer	0.9
n _{dis}	Number of critic update(s) per generator update	1
γ_{BN}	Batch normalization momentum term	0.8

voxels would be selected for seeds in the watershed algorithm. For example, voxels with gradients < 1 and intensity values < 40 would be selected as seeds for phase 1. This manual step was done once, incorporating the entire experimental data set in the selection process. The same bounds were then used for all 646 subvolumes during batch segmentation.

4. A watershed dilation algorithm was then applied using the three groups of voxel markers determined above. The algorithm accounts for image gradients while dilating/growing the markers (seeds). The output of this step is a segmented image with three phases. When using the same groups of markers in the commercial AVIZO package, statistically identical results were produced, which we take as validation that the code operates as intended.

As described in our earlier works,^{40–42} custom inhouse Python codes were used to calculate relevant microstructural properties for each subvolume, including: volume fractions; the average and standard deviation of diameter from the volumeweighted size distribution for each phase (based on an inscribed sphere method); interfacial surface area between each pair of phases; geometric tortuosity factors for each phase; and total and active (connected) TPB densities. For each phase in each subvolume, we also calculated the formation factor K_i , which represents the ratio between the effective and bulk diffusivity or conductivity in phase *i*. Here we use the fairly common estimation for the formation factor as the ratio between the phase fraction and geometric tortuosity factor: $K_i = \theta_i / \tau_i$.⁴¹

Electrochemistry Simulations

Simulations of the electrochemical performance of 90 3D microstructures were performed, as described previously.^{27,28} Briefly, the application ERMINE—Electrochemical **R**eactions in Microstructural Networks-was used, which is instantiated within MOOSE, the open-source finite element framework developed by Idaho National Laboratory. Thirty subvolumes for each of the original, GAN, and DREAM.3D microstructures were subjected to the workflow described elsewhere,²⁸ including appending a 4.16 μ m YSZ layer electrolyte. Morphology-preserving as an microstructural meshes were generated using a commercial meshing software (Simpleware ScanIP + FE 7.0, Synopsys, Inc., Mountain View, CA) and used as the simulation domains. All simulations were carried out on the Joule supercomputer (National Energy Technology Laboratory, Morgantown, WV). One hundred twenty cores were used in parallel for simulation of any subvolume, and simulations of many individual subvolumes also could be run in parallel on JOULE. The typical time to complete a simulation for a single subvolume, which involved simulations from 0 V to 0.4 V overpotential, was approximately 12 min. The total clock time for simulation of all 90 subvolumes was less than an hour using Joule.

RESULTS AND DISCUSSION

Visual Similarity

The 3D geometry and topology of the three-phase SOFC microstructure is highly complex and interconnected, which is by design to have good electrochemical performance.⁴³⁻⁴⁶ This complexity is visually apparent in Fig. 3 (top). A rendering of the whole $124 \times 110 \times 8 \ \mu m^3$ volume is given in Fig. 3 (top left), as are 21 grayscale images of $6.24 \times$ $6.24 \times 6.24 \ \mu m^3$ volumes cropped from the original (top right). The grayscale intensities in the experimental data reflect phase contrast, with some microscopy artifacts (e.g., charging). Three hundred twenty-three volumes can be cropped without overlap from the experimental data, though the training sets used here include overlaps and data augmentation.

Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials



Fig. 3. Volume-renderings of representative microstructures cropped from the original experimental data (top) or generated by the trained generator using random noise vector inputs (bottom). Visual similarity between the two sets of images is notable. The cropped/generated volumes have the dimensions of $96 \times 96 \times 96$ voxels, with voxel sizes of $65 \times 65 \times 65$ nm³. Therefore, the cropped volumes have physical dimensions of $6.24 \times 6.24 \times 6.24 \mu m^3$.

Upon sufficient training, the generator component of the GAN model produces microstructures that visually capture the complexity of the original with striking fidelity. Twenty-one rendered grayscale images produced by the generator are given in Fig. 3 (bottom right). The generator can, in theory, create an infinite number of unique microstructure volumes, and indeed all 21 generated volumes appear to be different. While the GAN-generated structures and original data look very similar, close inspection reveals that the original volumes display more variability than do the generated volumes. For example, some original volumes show strong SEM surfacecharging artifacts (which appear as bright white regions), yet the generated volumes contain few or none of these artifacts. This is consistent with observations that GANs tend to reproduce less variation than the data used to train them;^{47,48} they learn the average structure more readily than the outliers. Nevertheless, the GAN microstructures appear similar to those in the original distribution.

Figure 4 offers a deeper visual comparison between the original (top group) and GAN-generated (middle group) microstructures. Four volumes of each are shown in the left column, and six slices (2D images) from them are shown as grayscale images in the middle column. The 2D slices clarify that the anode microstructures consist of three grayscale intensities, corresponding to the three component phases: pores (black), nickel (gray), and yttria-stabilized zirconia (bright gray or white). Similarly, the GAN-generated microstructures display three grayscale intensities, and the morphology visually resembles that of the original microstructure. Close visual inspection again uncovers some differences. The end slices (e.g., slice 0) of the generated volumes are blurrier. Also, the finer features of the original volumes contain sharper angles, which are less common in the GAN-synthetic volumes. These visual discrepancies indicate that there is still room to improve the GAN implementation, should more precise similarity be required.

In prior work,^{26,28,29} we used an ellipsoid packing method built into the DREAM.3D software to create microstructures approximating the original PFIB data sets. The DREAM.3D microstructures are generated as segmented data, having three voxel values (colors); examples are shown in Fig. 4 (bottom group). Visually, the GAN-synthetic morphology resembles the original more closely than the DREAM.3D-synthetic morphology does. For example, compared to the DREAM.3D synthetics, the original structures and the GAN synthetics have greater phase connectivity and less agglomeration of the solid (gray) phases. This is not surprising, considering that DREAM.3D input was limited to phase fraction and the first and second moments of the log-normal particle size distribution. While it is possible that the DREAM.3D structures could be improved by iterative optimization or specifying additional microstructural parameters, important aspects of microstructural topology, including clustering and connectivity, are constrained by the ellipsoid packing scheme and are not directly tunable.

Statistical Similarity

The microstructural parameters that are known to impact electrochemical properties in three-phase

Author's personal copy



2D slice view



Fig. 4. Comparison of original (top), GAN-synthetic (middle), and DREAM.3D-synthetic (bottom) microstructures. Four random volumes (left column) for each category were sampled, and six 2D slices are shown for each, either in grayscale (center column) or segmented/labeled format (right column). The original microstructures consist of three phases: pores (black), nickel (gray), and yttria-stabilized zirconia (bright gray or white). For the original and GAN structures, segmentation of grayscale images was performed as a post-processing step. DREAM.3D generates microstructures with labeled fields, which are segmented by definition.

electrodes include volume fractions, particle sizes, tortuosity factors, and formation factors (volume fraction/tortuosity factor) for each phase as well as surface area for each pair of phases, and TPB density. These distributions are shown in Fig. 5 for the original (blue), GAN-synthetic (orange), and DREAM.3D-synthetic (green) microstructure volumes, and the mean and standard deviation of each distribution are tabulated in Table III. It was shown previously that these parameters can vary significantly across the electrode and between phases,^{26,29} as supported by the distributions of the original data in Fig. 5. It is impressive that the mean values for the GAN-synthetic volumes match the original volumes within a standard deviation and that there is a very high degree of overlap for each of the 16 distributions shown in Fig. 5. The GAN has clearly learned the underlying distributions of features inherent in the original microstructural volumes.

The DREAM.3D-synthetics, on the other hand, do not match the original data as well as the GANsynthetics do; for metrics shown in Fig. 5a, b, c, d,

Author's personal copy

Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials



Fig. 5. Histograms of original (blue), GAN-synthetic (orange), and DREAM.3D-synthetic (green) microstructure volumes with respect to 3D microstructural metrics: (a) volume fraction of each phase, (b) mean particle size for each phase, (c) tortuosity factor for each phase, (d) interfacial surface area between each pair of phases, (e) formation factor for each phase, and (f) total triple phase boundary (TPB) density. Phase 1 indicates pores; phase 2 is Ni; phase 3 is yttria-stabilized zirconia (YSZ). These statistics are based on 323 original (cropped without overlap), 323 GAN-synthetic, and 320 DREAM.3D-synthetic microstructure volumes of the same size. The curves show kernel density estimations (Color figure online).

Original GAN DREAM.3D **Microstructural metric** $\mu \pm \sigma$ $\mu \pm \sigma$ $\mu \pm \sigma$ Volume fraction 0.22 ± 0.01 Pore 0.21 ± 0.02 0.22 ± 0.02 0.36 ± 0.01 0.32 ± 0.03 0.37 ± 0.01 Ni YSZ 0.42 ± 0.02 0.42 ± 0.01 0.46 ± 0.02 Particle size (μm) 0.50 ± 0.02 0.45 ± 0.01 Pore 0.51 ± 0.06 0.55 ± 0.01 Ni 0.55 ± 0.01 0.74 ± 0.12 YSZ 0.50 ± 0.02 0.50 ± 0.02 0.55 ± 0.01 Tortuosity factor Pore 2.30 ± 0.61 2.29 ± 0.55 2.05 ± 0.40 Ni 1.22 ± 0.03 1.23 ± 0.04 1.49 ± 0.20 1.11 ± 0.02 1.10 ± 0.02 YSZ 1.12 ± 0.06 Interfacial area $(um^2/\mu m^3)$ Pore-Ni 0.55 ± 0.03 0.54 ± 0.03 0.56 ± 0.04 Pore-YSZ 0.87 ± 0.05 0.89 ± 0.06 1.07 ± 0.09 1.82 ± 0.05 Ni-YSZ 1.79 ± 0.07 1.36 ± 0.07 Formation factor 0.11 ± 0.02 Pore 0.10 ± 0.03 0.10 ± 0.02 Ni 0.30 ± 0.02 0.29 ± 0.02 0.22 ± 0.04 YSZ 0.37 ± 0.02 0.38 ± 0.02 0.42 ± 0.03 3.16 ± 0.20 3.06 ± 0.23 TPB density $(\mu m/\mu m^3)$ 2.92 ± 0.20

Table III. Mean μ and standard deviation σ for the distributions of microstructural metrics for the original, GAN-generated, and DREAM.3D-constructed volumes

and e, one or more DREAM.3D distributions have significant mismatch or very little overlap with the original distribution. DREAM.3D's methodology targets count-weighted particle statistics rather than the volume-weighted statistics presented here, which could contribute to some of the discrepancy seen here. Additionally, the ellipsoid packing construction algorithm constrains the topology of the resulting structures. The measured size distributions from the original data could also be missing Ni-Ni and YSZ-YSZ grain boundaries, which are not detected by the data collection technique; this could be another cause of discrepancies between the measured particle size distribution and that obtained by ellipsoid packing with similar input parameters. As noted previously, iterative tuning of the DREAM.3D parameters could yield a better matching microstructure, but the scope of DREAM.3D in this work was to compare GAN's imitation to a user's honest attempt at using DREAM.3D to create a similar microstructure. Although they do not match the original data as well as the GAN-synthetics do, the distributions from the DREAM.3D-synthetics are within a reasonable range relative to those of the original and have the same relative order of magnitude as those of the original, consistent with past observations.^{26,29}

To further evaluate the statistical similarity between the original and GAN microstructures, the distributions in Fig. 5 are presented as boxplots in Fig. 6. All pair-wise boxplots have significant overlaps of the interquartile range (IQR), or the box region, as well as the "minimum"-"maximum" range (from $Q_1 - 1.5IQR$ to $Q_3 + 1.5IQR$). This reinforces the assertion that the GAN is able to generate microstructures that are statistically representative of the original microstructures, especentral around the cially portion of the distributions. We define outlier values as those outside of the statistical "minimum" and "maximum" values, shown as diamonds in Fig. 6. There are more outliers for the original microstructures than for the GAN-synthetics, and these outliers cover a wider range of values. Thus, while the current GAN implementation captures the central portions of the distributions, further work is necessary to capture the full variability, particularly the outliers, of the original microstructures.

Electrochemical Performance

The microstructure metrics quantified in the previous section have been used in effective medium theory models to describe electrochemical properties or performance of electrodes.^{49–51} The statistical similarity between the original and GAN structures suggest that their effective medium properties will match as well. For a more direct comparison of local performance, microstructurally resolved, finite element simulations of electrochemistry were carried out on 30 volumes from each microstructure type: original, GAN-synthetic, and DREAM.3D-synthetic.

These 90 sub-volumes, which are randomly sampled subsets of the microstructure data in Fig. 5, were subjected to a finite element computational workflow described in detail elsewhere.^{27,28,52} In essence, the voxelated microstructure images were meshed to produce morphology-conforming finite element computational domains, and a standard reaction-andtransport electrochemistry model was implemented across the meshed microstructures. The simulation model only works on microstructural features fully interconnected to the external source/sink locations of the transport species and therefore only includes active triple-phase boundaries. We note that, though not shown in Fig. 5, the original and GAN-synthetics have similar values of inactive TPB densities, but the DREAM.3D has significantly higher inactive TPB densities, highlighting again the difference between GAN-synthetics and DREAM.3D synthetics for capturing locally resolved topological features.

At a given global overpotential V, integrating the current I flowing through the electrolyte layer allows one to plot current density versus overpotential for all 90 subvolumes. Those I - V plots can then be fit using a standard activation/ohmic loss model, which yields an effective exchange current density j_0 and an effective ohmic resistance R_{ohmic} for each subvolume (see²⁸). As such, the electrochemical performance of each sub-volume can be described using two distinct parameters. For all 90 volumes, the exchange current density is plotted against the active TPB density in Fig. 7a, and the ohmic resistance is plotted against the inverse of the YSZ formation factor (τ_{YSZ}/θ_{YSZ}) in Fig. 7b.

Based on commonly used approximations, one might expect the exchange current density to be proportional to TPB density and the effective ohmic resistance to be inversely proportional to formation factor.²⁸ Indeed, the exchange current density in Fig. 7a is a nearly perfect linear function of active TPB density (it even extrapolates to the origin). As such, the main difference in performance of these electrodes comes from the variability in the active TPB density. Overall, the original and GAN-synthetic microstructures have highly overlapped scatter distributions. The GAN-synthetics have a slightly lower mean value of active TPB density, and therefore a slightly lower mean value of j_0 , and the range of data is slightly wider for the original sub-volumes than for the GAN-synthetics. The DREAM.3D-synthetic microstructures, on the other hand, have properties that have little overlap with that of the original microstructures. This arises because the distribution of active TPB density of the DREAM.3D-synthetics differs significantly from the original microstructures.

The ohmic resistance in Fig. 7b does not appear to be a strong function of the inverse formation factor of YSZ (the phase through which ions flow and which determines the ohmic resistance). This was shown previously,²⁸ and is thought to arise because internal microstructural heterogeneity impacts

Author's personal copy

Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials



Fig. 6. Boxplots of the distributions of microstructural parameters from the original (blue) and GAN-synthetic (orange) microstructure volumes: (a) volume fraction of each phase, (b) mean particle size for each phase, (c) tortuosity factor for each phase, (d) interfacial surface area between each pair of phases, (e) formation factor for each phase, and (f) total triple phase boundary (TPB) density. Phase 1 indicates pores; phase 2 is Ni; phase 3 is yttria-stabilized zirconia (YSZ). The interquartile range is indicated by the box, the mean by the horizontal line within the box, the minimum (Q1 - 1.5/QR) and maximum ($Q_3 + 1.5/QR$) range by the horizontal lines below and above the box, and outliers are represented by diamonds. Q_3 is the upper quartile, Q_1 is the lower quartile, and IQR is the interquartile range (Color figure online).



Fig. 7. Simulated electrochemical performance metrics from 30 original, 30 GAN-synthetic, and 30 DREAM.3D-synthetic microstructure volumes. (a) Exchange current density plotted against the triple-phase boundary (TPB) density. (b) Ohmic resistance plotted against the inverse formation factor for YSZ, τ_{YSZ}/θ_{YSZ} . For each microstructure category, the mean exchange current density or the mean ohmic resistance is shown in the legend. The microstructure volumes considered here are randomly sampled subsets of the microstructure data for Fig. 5. Note that the simulations were operated directly on meshed microstructures, thus retaining topological information.

local ionic transport and, therefore, impacts ohmic resistance values. Nevertheless, the original and the GAN-synthetic microstructures have highly overlapped scatter distributions: the GAN-synthetics being slightly wider in formation factors and the original being wider in the ohmic resistance values. On the other hand, DREAM.3D-synthetic microstructures have little overlap with the original microstructures: the inverse formation factors are lower and wider, while the ohmic resistance values are slightly higher and wider.

Overall, Fig. 7 demonstrates that the GAN has learned the underlying distribution of features to an extent that the ensemble electrochemical properties are very well matched with the experimental microstructures. While the DREAM.3D-synthetics are somewhat similar to the originals, they do not do well in matching local connectivity at the subvolumes sizes investigated here and therefore result in poorer matches to the performance of the original microstructures. Furthermore, as noted above, the performance simulations reported here presume electrochemical activity is almost entirely restricted to the TPBs. For models that include activity through two-phase boundaries (e.g., through a cathode phase bulk), the performance of the GAN structures would improve further relative to DREAM.3D, since it also captures these features with greater fidelity (Fig. 5d).

CONCLUSION

In this work, we implemented a GAN model for learning and generating microstructural images. This is among the first applications of GAN to 3D microstructure generation at a scale, complexity, and fidelity suitable for ICME applications. We show that the GAN model can generate realistic 3D, topologically complex, multiphase, grayscale microstructures that closely resemble the original, experimental structures in terms of visual appearance, statistical representation of geometric and topological properties, and simulated electrochemical performance. Compared to the commonly used microstructure generation algorithm DREAM.3D, the GAN results are structurally and electrochemically more realistic. Besides its superior fidelity, we consider that the GAN model has the following merits:

- Generality: The GAN framework can be applied to arbitrary materials systems, microstructural morphologies, and imaging modes, making it a general and flexible tool for microstructure generation. This is in contrast to many conventional generation methods that explicitly impose underlying assumptions for the microstructure. The specific example in this article, DREAM.3D, generates microstructures with ellipsoid packing, which can be a greatly simplified assumption.
- Autonomy: Conventional microstructural construction methods require the user to define a set of metrics that characterizes microstructural geometry and topology. The GAN generates realistic structures independent of user assumptions; that is, it reproduces microstructural metrics without needing to know them.
- Throughput: A trained generator can synthesize an arbitrarily large number of unique microstructural volumes.
- Computational efficiency: Although training the GAN model requires substantial HPC resources, using it to generate synthetic microstructures is quite efficient.

However, we also note that the GAN implementation has the following relative shortcomings:

• Training data: Obtaining a sufficient volume of

3D microstructural data to train a GAN via experiments or physical simulation can be costly or challenging. In contrast, conventional construction methods such as DREAM.3D can generate microstructures from statistical descriptors, without requiring 3D image data.

- High performance computing: 3D image generation requires high model complexity (number of trainable parameters). As such, high performance multi-GPU training is a requirement.
- Synthetic volume: Typical CNN models, or models with CNN encoding/decoding structures, require fixed input/output sizes. As such, typical GAN models with deep convolutional layers can only generate images at a fixed size corresponding to the image size used during training (although multiple images can be tiled as montages). Training GAN to generate 3D images beyond $64 \times 64 \times 64$ voxels is costly. Due to the computational cost of training, this work was limited to generation of 3D images at the scale of $96 \times 96 \times 96$ voxels. Larger volumes require more computational resources in proportion to system size.
- Tunability: While conventional construction methods allow users to alter input variables to explore different parameter spaces (e.g., phase fraction, particle size), the present GAN implementation is limited to generating microstructures from the parameter space on which it was trained. However, developing tunable GANs is an active area of research.¹⁵

Finally, it is worth commenting on the role of this work in the context of applied machine learning models for tasks in the physical sciences. We have shown that based on the metrics discussed in this article, our GAN model outputs valid and physical microstructures. Thus, a major contribution of our work is the extensive validation of the GAN model output beyond nebulous or subjective measures of visual similarity. This is a necessary step for advancing GAN generative methods as tools supporting ICME and other scientific and engineering endeavors.

ACKNOWLEDGMENTS

This work was performed in support of the US Department of Energy National Energy Technology Laboratory's ongoing research under the RSS contract 89243318CFE000003. EAH acknowledges support by the National Science Foundation under grant CMMI-1826218 and the Air Force D3OM2S Center of Excellence under agreement FA8650-19-2-5209.

FUNDING

This work was funded by the Department of Energy, National Energy Technology Laboratory, an agency of the United States Government, through a

Microstructure Generation via Generative Adversarial Network for Heterogeneous, Topologically Complex 3D Materials

support contract with Leidos Research Support Team (LRST). Neither the United States Government nor any agency thereof, nor any of their employees, nor LRST, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

CONFLICT OF INTEREST

On behalf of all authors, Elizabeth A. Holm states that there is no conflict of interest.

ELECTRONIC SUPPLEMENTARY MATERIAL

The online version of this article (https://doi.org/ 10.1007/s11837-020-04484-y) contains supplementary material, which is available to authorized users.

REFERENCES

- 1. D. Dimiduk, E. Holm, and S. Niezgoda, Integr. Mater. Manuf. Innov. 7, 1 (2018).
- B.L. DeCost and E.A. Holm, in Statistical Methods for Materials Science: The Data Science of Microstructure Characterization, eds. J. Simmons, L. Drummy, C. Bouman, and M. de Graef (CRC Press, Boca Raton, FL, 2019), pp. 241–258. (ISBN: 978-1-498-73820-0).
- Y. Jiao, F. Stillinger, and S. Torquato, *Phys. Rev. E* 77, 031135 (2008).
- S. Torquato and H. Haslach Jr., Appl. Mech. Rev. 55, B62 (2002).
- S. Mandal, J. Lao, S. Donegan, and A.D. Rollett, Scr. Mater. 146, 128 (2018).
- M.A. Groeber and M.A. Jackson, Integr. Mater. Manuf. Innov. 3, 5 (2014).
- H. Xu, D.A. Dikin, C. Burkhart, and W. Chen, Comput. Mater. Sci. 85, 206 (2014).
- Z. Jiang, W. Chen, and C. Burkhart, J. Microsc. 252, 135 (2013).
- R. Bostanabad, A.T. Bui, W. Xie, D.W. Apley, and W. Chen, Acta Mater. 103, 89 (2016).
 R. Bostanabad, Y. Zhang, X. Li, T. Kearney, L.C. Brinson,
- R. Bostanabad, Y. Zhang, X. Li, T. Kearney, L.C. Brinson, D.W. Apley, W.K. Liu, and W. Chen, *Prog. Mater. Sci.* 95, 1 (2018).
- D.P. Kingma and M. Welling, Auto-Encoding Variational Bayes (2013). arXiv:1312.6114 [stat.ML].
- I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, *Gen-*

erative Adversarial Networks (2014). arXiv:1406.2661 [stat.ML].

- Z. Yang, X. Li, L. Catherine Brinson, A.N. Choudhary, W. Chen, and A. Agrawal, J. Mech. Des. 140(11), art. no. 111416 (2018).
- R. Cang, H. Li, H. Yao, Y. Jiao, and Y. Ren, *Comput. Mater.* Sci. 150, 212 (2018).
- A. Iyer, B. Dey, A. Dasgupta, W. Chen, and A. Chakraborty, A Conditional Generative Model for Predicting Material Microstructures from Processing Methods (2019). arXiv:191 0.02133 [eess.IV].
- R. Singh, V. Shah, B. Pokuri, S. Sarkar, B. Ganapathysubramanian, and C. Hegde, *Physics-Aware Deep Generative Models for Creating Synthetic Microstructures* (2018). arXiv preprint arXiv:1811.09669.
- D. Fokina, E. Muravleva, G. Ovchinnikov, and I. Oseledets, Microstructure Synthesis Using Style-Based Generative Adversarial Network (2019). arXiv:1909.07042 [eess.IV].
- M. Mirza and S. Osindero, Conditional Generative Adversarial Nets (2014). arXiv:1411.1784 [cs.LG].
- A. Odena, C. Olah, and J. Shlens, Conditional Image Synthesis with Auxiliary Classifier GANs (2016). arXiv:1610.0 9585 [stat.ML].
- T. Karras, S. Laine, and T. Aila, A Style-Based Generator Architecture for Generative Adversarial Networks (2018). arXiv:1812.04948 [cs.NE].
- L. Mosser, O. Dubrule, and M.J. Blunt, *Transp. Porous Med.* 125, 81 (2018).
- 22. L. Mosser, O. Dubrule, and M.J. Blunt, *Phys. Rev. E* 96, 043309 (2017).
- A. Gayon-Lombardo, L. Mosser, N.P. Brandon, and S.J. Cooper, Pores for Thought: The Use of Generative Adversarial Networks for the Stochastic Reconstruction of 3D Multi-phase Electrode Microstructures with Periodic Boundaries (2020). arXiv:2003.11632 [cs.NE].
- A. Radford, L. Metz, and S. Chintala, Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks (2015). arXiv:1511.06434 [cs.LG].
- R. Mahbub, T. Hsu, W.K. Epting, N.T. Nuhfer, G.A. Hackett, H. Abernathy, A.D. Rollett, M. De Graef, S. Litster, and P.A. Salvador, *ECS Trans.* 78, 2159 (2017).
 T. Hsu, W.K. Epting, R. Mahbub, N.T. Nuhfer, S. Bhat-
- T. Hsu, W.K. Epting, R. Mahbub, N.T. Nuhfer, S. Bhattacharya, Y. Lei, H.M. Miller, P.R. Ohodnicki, K.R. Gerdes, H.W. Abernathy, et al., *J. Power Sources* 386, 1 (2018).
- T. Hsu, R. Mahbub, J.H. Mason, W.K. Epting, H.W. Abernathy, G.A. Hackett, A.D. Rollett, S. Litster, and P.A. Salvador, *MethodsX* 100822, (2020).
- T. Hsu, J.H. Mason, R. Mahbub, W.K. Epting, H.W. Abernathy, G.A. Hackett, A.D. Rollett, S. Litster, and P.A. Salvador, *Electrochim. Acta* 345, art. no. 136191 (2020).
- R. Mahbub, W.K. Epting, T. Hsu, J.H. Mason, M. Feng, N.T. Nuhfer, H.W. Abernathy, G.A. Hackett, S. Litster, A.D. Rollett, et al., J. Electrochem. Soc. 167, 054506 (2020).
- R. Szeliski, Computer Vision: Algorithms and Applications (Springer, New York, 2010), pp. 1–812.
- B.L. DeCost, T. Francis, and E.A. Holm, Acta Mater. 133, 30 (2017).
- D.P. Kingma and J. Ba, Adam: A Method for Stochastic Optimization (2014). arXiv:1412.6980 [cs.LG].
- M. Arjovsky, S. Chintala, and L. Bottou, Wasserstein GAN (2017). arXiv:1701.07875 [stat.ML].
- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G.S. Corrado, A. Davis, J. Dean, M. Devin, et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems* (2016). arXiv:1603.04467 [cs.DC].
- T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, Spectral Normalization for Generative Adversarial Networks (2018). arXiv:1802.05957 [cs.LG].
- A. Odena, V. Dumoulin, and C. Olah, *Distill*. (2016). http s://doi.org/10.23915/distill.00003.
- A. Sergeev and M.D. Balso, Horovod: Fast and Easy Distributed Deep Learning in TensorFlow (2018). arXiv:1802. 05799 [cs.LG].

- P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour (2017). arXiv:1706.02677 [cs.CV].
- K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition* (2015). arXiv:1512.03385 [cs.CV].
- W.K. Epting, J. Gelb, and S. Litster, *Adv. Funct. Mater.* 22, 555 (2012).
- S. Litster, W. Epting, E. Wargo, S. Kalidindi, and E. Kumbur, *Fuel Cells* 13, 935 (2013).
- W.K. Epting, Z. Mansley, D.B. Menasche, P. Kenesei, R.M. Suter, K. Gerdes, S. Litster, and P.A. Salvador, J. Am. Ceram. Soc. 100, 2232 (2017).
- N.Q. Minh and T. Takahashi, Science and Technology of Ceramic Fuel Cells (Elsevier, Amsterdam, 1995).
- J.R. Wilson, W. Kobsiriphat, R. Mendoza, H.Y. Chen, J.M. Hiller, D.J. Miller, K. Thornton, P.W. Voorhees, S.B. Adler, and S.A. Barnett, *Nat. Mater.* 5, 541 (2006).
- D.R. Rolison, J.W. Long, J.C. Lytle, A.E. Fischer, C.P. Rhodes, T.M. McEvoy, M.E. Bourg, and A.M. Lubers, *Chem.* Soc. Rev. 38, 226 (2009).

- H. Zhang, X. Yu, and P.V. Braun, Nat. Nanotechnol. 6, 277 (2011).
- T. Karras, T. Aila, S. Laine, and J. Lehtinen, Progressive Growing of GANs for Improved Quality, Stability, and Variation (2017). arXiv:1710.10196 [cs.NE].
- T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, *Improved Techniques for Training GANs* (2016). arXiv:1606.03498 [cs.LG].
- 49. S.B. Adler, J. Electrochem. Soc. 143, 3554 (1996).
- J. Mason, I. Celik, S. Lee, H. Abernathy, and G. Hackett, J. Electrochem. Soc. 165, F64 (2018).
- T. Yang, J. Liu, H. Finklea, S. Lee, W.K. Epting, R. Mahbub, T. Hsu, P.A. Salvador, H.W. Abernathy, and G.A. Hackett, J. Hydrog. Energy 43, 15445 (2018).
- T. Hsu, R. Mahbub, W.K. Epting, H. Abernathy, G.A. Hackett, A.D. Rollett, S. Litster, and P.A. Salvador, *ECS Trans.* 78, 2711 (2017).

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