

Reconstruction of three-dimensional porous media using generative adversarial neural networksLukas Mosser,^{*} Olivier Dubrule,[†] and Martin J. Blunt[‡]*Department of Earth Science and Engineering, Imperial College London, London SW7 2BP, United Kingdom*

(Received 18 April 2017; published 23 October 2017)

To evaluate the variability of multiphase flow properties of porous media at the pore scale, it is necessary to acquire a number of representative samples of the void-solid structure. While modern x-ray computer tomography has made it possible to extract three-dimensional images of the pore space, assessment of the variability in the inherent material properties is often experimentally not feasible. We present a method to reconstruct the solid-void structure of porous media by applying a generative neural network that allows an implicit description of the probability distribution represented by three-dimensional image data sets. We show, by using an adversarial learning approach for neural networks, that this method of unsupervised learning is able to generate representative samples of porous media that honor their statistics. We successfully compare measures of pore morphology, such as the Euler characteristic, two-point statistics, and directional single-phase permeability of synthetic realizations with the calculated properties of a bead pack, Berea sandstone, and Ketton limestone. Results show that generative adversarial networks can be used to reconstruct high-resolution three-dimensional images of porous media at different scales that are representative of the morphology of the images used to train the neural network. The fully convolutional nature of the trained neural network allows the generation of large samples while maintaining computational efficiency. Compared to classical stochastic methods of image reconstruction, the implicit representation of the learned data distribution can be stored and reused to generate multiple realizations of the pore structure very rapidly.

DOI: [10.1103/PhysRevE.96.043309](https://doi.org/10.1103/PhysRevE.96.043309)**I. INTRODUCTION****Image reconstruction**

The reconstruction and evaluation of the material properties of porous media plays a key role across many engineering disciplines. Many physical processes such as the movement of multiple phases of fluids through sedimentary rocks are controlled by individual pores at the micron and submicron scales [1].

In carbon capture and sequestration (CCS), the long-term storage behavior is controlled by the physical and chemical interaction of supercritical CO₂ with the reservoir brine, as well as the spatial distribution and connectivity of minerals in the pore space [2,3]. The variability of the controlling properties such as the permeability of the host rock is determined by repeated experiments or numerical modeling of these processes.

Using modern computer tomographic methods, it is possible to observe porous materials and evaluate their material properties at the micrometer scale [micro computed-tomography (micro-CT)] under static and transient conditions at high pressures and temperatures in near real time. Performing micro-CT imaging of porous media requires specialized expensive equipment and in the case of CCS, only a single image of the investigated rock type is typically acquired.

To evaluate the variability associated with the geometrical and mineralogical morphology of the pore space, numerous physical experiments using the same rock type would have to be performed to obtain a distribution over larger volumes. Due to time and cost limitations inherent with the experimental acquisition of high-resolution images, this is often deemed unfeasible. Material properties governing the single- and multiphase flow behavior of porous media can be estimated from numerical solution of partial differential equations at a scale larger than that of a representative elementary volume (REV) and verified by experimental results [4].

Many sedimentary rocks consist of granular siliciclastic or carbonate materials. Boolean models use this fundamental characteristic of natural granular materials to emulate the shape of the arising pore space, due to an underlying random process that controls the distribution of the individual grains [5,6]. While for the classical Boolean model the centers of the grains are uniformly distributed in space and grains can arbitrarily overlap, more complicated models with rigid hard-sphere grains and more complex grain interaction functions have been developed [7–10]. The framework of Boolean models also allows extension beyond spherical particles and enables derivation of the properties of material models as a function of the parameters of the underlying random process [11–14].

In clastic rocks, the arrangement of individual grains occurs due to the transport of material from a high-energy source to a low-energy sink. Process models, where depositional mechanisms are simulated, have been shown to reproduce realistic granular reconstructions capturing the pore-space morphology of granular sedimentary rocks [15].

Spatial probabilistic models such as truncated Gaussian processes or sequential indicator simulation have been widely applied in the geosciences to model the spatial distribution of materials [16]. Many of these methods rely on two-point probability functions as a measure of spatial variability,

^{*}lukas.mosser15@imperial.ac.uk[†]o.dubrule@imperial.ac.uk[‡]m.blunt@imperial.ac.uk

whereas recent methods in geostatistics use training images as a basis for sample reconstruction [17–19]. These images are usually assumed to exhibit stationarity of the probability distribution of the properties of interest and rely on higher-order multiple-point statistics (MPS) to reconstruct stochastic random media.

With MPS, the probability distributions are represented by training images and are sampled using a limited multiscale neighborhood that captures the variation on a large scale as well as fine-structural details on smaller scales [20]. Multiple-point-statistics-based methods have been used in two- and three-dimensional conditional simulations of spatial properties in reservoir-scale earth modeling applications [21]. The computational complexity of these methods is highly dependent on individual algorithms as well as the size of the domains used to sample from the training images [22]. Parallelized versions have been developed, reducing the computational time required to perform reconstruction using MPS [23,24].

Three-dimensional porous media have been reconstructed using a modified multiple-point statistics approach based on two-dimensional images of porous media [25–27]. Stochastic methods based on simulated annealing allow the incorporation of arbitrary cost functions of statistical and morphological properties used in unconditional three-dimensional image reconstruction [28,29]. Recent advances have reduced the computational run time of simulated-annealing-based methods for reconstruction of porous media, to the order of tens of hours per realization at the scale of 300^3 voxels [30].

In the following section we introduce a recently developed class of unsupervised machine learning methods called generative adversarial networks (GANs) that allow simulation of probability distributions given a set of training data [31]. Volumetric generative adversarial networks have previously been applied to low-resolution three-dimensional computer-aided design model synthesis and practical applications of three-dimensional (3D) GANs are few compared to their two-dimensional counterparts [32]. Integration of multiresolution data sets incorporating image data across a number of length scales is possible in the GAN framework by using a Laplacian pyramid approach such as LapGAN [33].

We investigate the applicability of GANs to model three-dimensional textures of rocks based on three-dimensional binary representations of porous media acquired at the micrometer scale. We compare statistical, morphological, and transport properties of the simulated images with those of the training images. We evaluate the single-phase directional permeability to show that the synthetic realizations sampled from the learned representation of the input data can capture single-phase flow properties of sedimentary rocks.

Training of these neural networks involves finding a set of hyperparameters that lead to stable training [34]. While this training can take on the order of tens of hours, the sampling of large volumetric domains occurs on the order of seconds on the current generation of graphical processing units. We show that in favorable cases convolutional neural networks incorporated in the GAN framework allow the generation of synthetic reconstructions of porous media that exceed the dimensions of their training images. Contrary to most existing simulation techniques, the set of parameters used to generate

synthetic realizations can be stored once trained, allowing rapid generation of new samples to assess the variability of material properties.

While we apply GANs to a set of micro-CT images of porous media, the method can readily be applied to volumetric images of porous media obtained from other three-dimensional microscopy instruments such as nano- or medical-CT instruments. We discuss the challenges involved in training GANs for stochastic image reconstruction of porous media as compared to other stochastic image reconstruction methods and we evaluate the computational efficiency of GAN-based image reconstruction. Finally, we provide empirical guidelines on the requirements of the input data set to allow successful training of GANs on large three-dimensional voxel representations of natural porous media.

All data used in this study are available in the public domain and we have made the code used for training, as well as example pre-trained models, available online [35]. A public data set of high-resolution micro-CT images made available by the Imperial College Pore-Scale Modelling Group [36], of a spherical bead pack, Berea sandstone, and oolitic Ketton limestone will serve as benchmark cases to study the application of GANs to three-dimensional stochastic image reconstruction.

II. GENERATIVE ADVERSARIAL NETWORKS

In the following section we present GANs for three-dimensional image generation. Generative neural networks have been developed in the context of deep learning by Goodfellow *et al.* as a methodology to learn a representation of a high-dimensional probability distribution from a given data set [31]. In the context of image reconstruction, we refer to this data set as a set of training images that are representative samples of the probability distribution underlying the image space.

Generative adversarial networks learn an implicit representation of the probability density as opposed to explicit density models. The main drawback of explicit density models is their computational cost, which grows with the dimensionality of the samples and often requires sequential simulation. For high-dimensional samples such as volumetric image data, the computational cost is $O(N)$, where N represents the number of voxels in the domain of interest and can easily exceed 10^9 voxels for modern high-resolution micro-CT images. Using any of these methods would make it intractable to generate a large number of very large samples. Generative adversarial networks have been designed to perform fast sampling from the learned density representation and allow full parallel generation, making them an ideal candidate to generate large volumetric images [34].

Generative adversarial networks consist of two differentiable functions: a discriminator D and a generator G . The discriminator receives samples of the “real” data set (label 1) $x \sim p_{\text{data}}$ and “fake” samples $G(\mathbf{z})$ (label 0) created by the generator from the hidden latent space Z (see Fig. 1). The latent space Z is composed of independent real random variables, typically normally or uniformly distributed, that represent the random input to the generator G . The generator G maps random variables from the latent space into the space

of images. The discriminator’s role is to assign a probability that a random sample is from the real data distribution p_{data} . The discriminator tries to label each sample correctly, while the generator tries to “fool” the discriminator into labeling the fake images as part of the true data distribution and therefore achieving $D[G(\mathbf{z})]$ close to one.

More formally we can define the loss, i.e., the cost function for GANs as a minimization-maximization problem

$$\min_G \max_D [\mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} \{\log[D(\mathbf{x})]\} + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} \{\log\{1 - D[G(\mathbf{z})]\}\}]. \quad (1)$$

Solutions to this optimization problem have been shown to be Nash equilibria, where each player achieves a local minimum of their loss function with respect to their parameters [34].

In practice we represent G and D by convolutional neural networks that are trained by a gradient-descent-based optimization method. Training is performed in two steps. First, the discriminator is trained to maximize

$$J^{(D)} = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} \{\log[D(\mathbf{x})]\} + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} \{\log\{1 - D[G(\mathbf{z})]\}\} \quad (2)$$

while the parameters of the generator are fixed. This improves the ability of the discriminator to distinguish between real and fake images. Second, we generate synthetic samples $G(\mathbf{z})$ by

drawing samples \mathbf{z} from an N -dimensional normal distributed latent space and train the generator to minimize

$$J^{(G)} = \mathbb{E}_{\mathbf{z} \sim p_z} \{\log\{1 - D[G(\mathbf{z})]\}\} \quad (3)$$

while keeping the discriminator fixed.

By minimizing Eq. (3) the generator tries to fool the discriminator into believing that the samples $G(\mathbf{z})$ are real data samples. In this way the generator learns to represent a distribution $p_g(\mathbf{x})$ that is as close as possible to the real data distribution $p_{\text{data}}(\mathbf{x})$. When convergence is reached $p_g(\mathbf{x}) = p_{\text{data}}(\mathbf{x})$ and the value of the discriminator becomes $\frac{1}{2}$ as it cannot distinguish between the two anymore.

Initially, the discriminator D outperforms the generator significantly, making the gradient used to train the generator close to zero. Therefore, instead of minimizing $\log\{1 - D[G(\mathbf{z})]\}$ for the generator, it is helpful to maximize $\log\{D[G(\mathbf{z})]\}$ [34].

Generative adversarial networks show highly unstable behavior during training and a large number of trial and error runs are required to find an optimal set of hyperparameters that allow stable training. A number of heuristics have been published that have been shown to stabilize GAN training, such as one-sided label smoothing and adding white noise to the input layer of the discriminator [37,38].

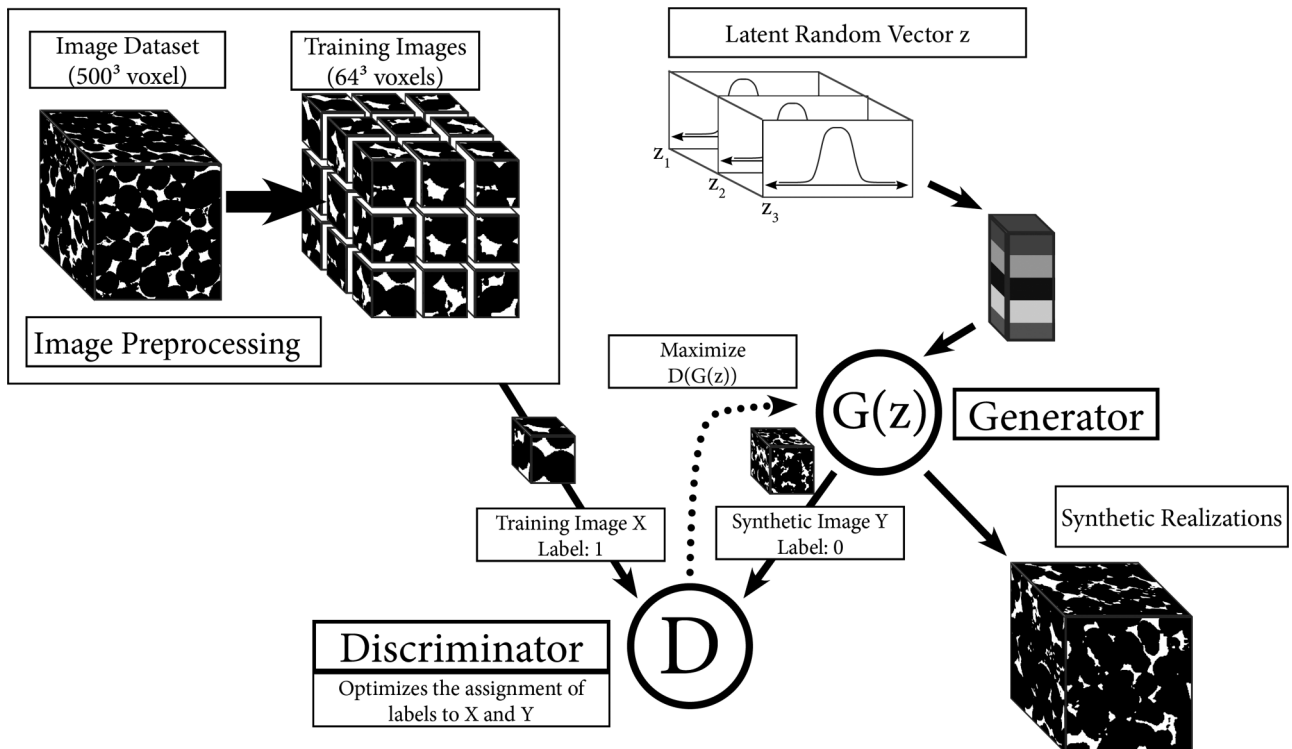


FIG. 1. Overview of the GAN training process. Segmented volumetric images are split into 64^3 - or 128^3 -voxel training images. The generator G is a function that is applied to a sample from a latent random space Z and creates a synthetic realization. We assume that samples drawn from the hidden latent space Z are normally distributed (see Sec. II). The discriminator’s role is to determine whether a sample is part of the training image data set (label 1) or from the generator (label 0). The misclassification error is computed as a binary cross-entropy criterion and the error backpropagated to improve the discriminator’s ability to distinguish real from fake images. Then the generator is updated to improve the quality of the produced samples and fool the discriminator. When sufficient image quality is obtained, training is stopped and the discriminator may be discarded. The generator can now be used to create new samples. By providing larger latent vectors than used initially for training, larger output images can be produced.




FIG. 2. Comparison of $S_2(r)$ for a Boolean model and a packing of hard spheres at a porosity $\phi = 0.5$. Here S_2 exhibits exponential decay for the Boolean model, whereas a dampened oscillation is characteristic for packings of spheres. The mean chord length can be found at the intersection of the slope of S_2 at the origin with the x axis [see Eq. (6)].

We provide a more detailed overview of the neural networks used in this study in Sec. III B (Table I) and later provide suggestions on how to facilitate efficient training (see Sec. VI) for volumetric image data sets of porous media.

III. METHODOLOGY

In the following section we outline the criteria used to evaluate the quality of simulations based on the training image data sets. We treat all images under the assumption of stationarity and the existence of a representative elementary volume.

A. Evaluation criteria

1. Two-point statistics

We characterize the second-order structure of the porous media by calculating the two-point probability function of

the pore phase. By assuming stationarity, this function is equivalent to the noncentered covariance [7]

$$S_2(\mathbf{r}) = \mathbf{P}(\mathbf{x} \in P, \mathbf{x} + \mathbf{r} \in P) \quad \text{for } \mathbf{x}, \mathbf{r} \in \mathbb{R}^d, \quad (4)$$

which is the probability \mathbf{P} that two points \mathbf{x} and $\mathbf{x} + \mathbf{r}$, separated by the lag vector \mathbf{r} , are located in the pore phase P . At the origin, $S_2(0)$ is equal to the porosity ϕ ; S_2 stabilizes around ϕ^2 as $r \rightarrow \infty$ (Fig. 2). Due to the anisotropic nature of many porous media, we compute $S_2(r)$ along the three Cartesian directions, as well as the radial average of $S_2(r)$.

It is a well known result that the specific surface area S_V of a porous medium can be expressed as a function of S_2 [39]. In the case of an isotropic porous medium and in three-dimensions S_V is related to S_2 by

$$S_V = -4S_2'(0), \quad (5)$$

where $S_2'(0)$ is the derivative of $S_2(r)$ at the origin. Furthermore, the average chord lengths within the pore and the grain phases are [10]

$$\bar{l}_c^{\text{pore}} = \frac{\phi}{S_2'(0)}, \quad (6a)$$

$$\bar{l}_c^{\text{grain}} = \frac{1 - \phi}{S_2'(0)}, \quad (6b)$$

which for the pore phase can be readily found from the intersection of the slope of $S_2(r)$ with the x axis (Fig. 2).

In favorable cases, it is possible to find analytical expressions of $S_2(r)$ from the spatial distribution and geometry of the grains. A Boolean model of overlapping spherical grains of uniform spatial distribution exhibits an exponential decay of the covariance until the lag distance is equal to the diameter of the grains where it becomes zero [7]. For porous media that can be well described by a Boolean model, we can estimate the size of the elementary Boolean grain from the decay of S_2 .

Semianalytical expressions for more complex models, such as for a packing of hard spheres, have been developed [40]. Models of $S_2(r)$ for spherical packings exhibit a dampened oscillation (also called hole effect). The shape of the estimated covariance therefore allows us to obtain information on the structure of the porous medium (Fig. 2).

The covariance $S_2(r)$ was estimated for the training images and the stochastic reconstructions generated by the trained GAN model. For each GAN model, we evaluate the noncentered covariance S_2 and the specific surface area S_V [Eq. (5)]

TABLE I. Neural network configurations and hyperparameters used to train on voxelized image subsets.

Parameter	Training image data set		
	Bead pack	Berea	Ketton
training image size	128 ³ voxels	64 ³ voxels	64 ³ voxels
latent space \mathbf{z} dimension	100	512	100
generator filters N_G	64	64	64
discriminator filters N_D	8	16	16
optimizer	generator + discriminator: ADAM		
learning rate; momentum	2×10^{-4} ; 0.5	2×10^{-4} ; 0.5	2×10^{-4} ; 0.5
stabilization	white noise ($\sigma = 0.1$)	label smoothing ($\epsilon = 0.1$)	white noise ($\sigma = 0.1$)

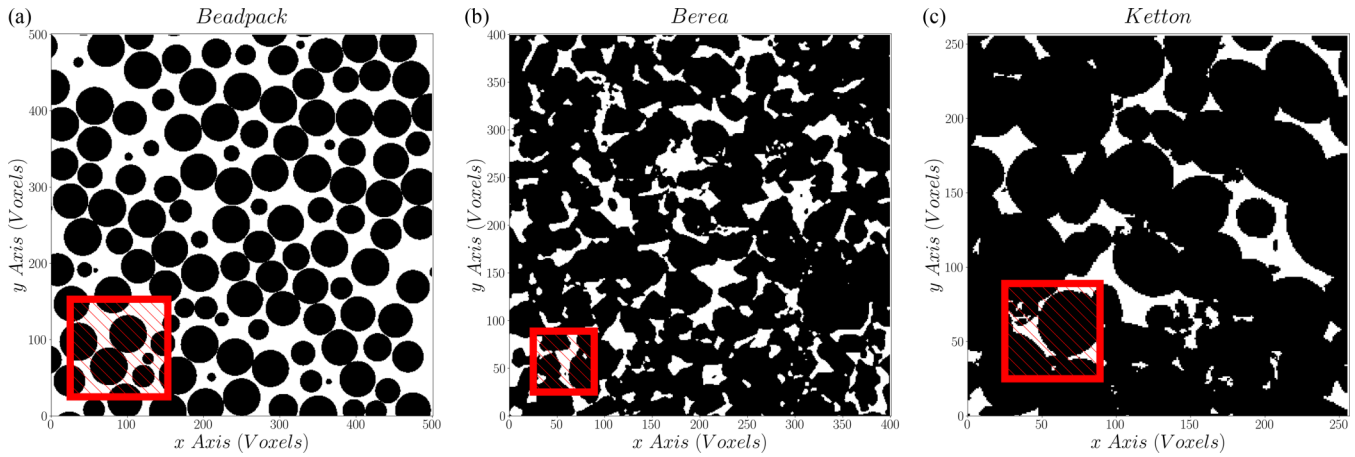


FIG. 3. Cross sections of the three image data sets. The bordered regions indicate the size of the training images extracted from the full data set. (a) The bead pack consists of spheres of equal diameter ($d = 50$ voxels). (b) The Berea sandstone is an angular granular sandstone that shows traces of dispersed clay. (c) The oolitic Ketton limestone consists of ellipsoidal grains showing inter- and intragranular porosity. The voxel sizes are $3 \mu\text{m}$ for the bead pack and Berea sandstone and $15.2 \mu\text{m}$ for the Ketton sample.

and compare these to the values obtained from the original training images.

In our discussion of the required training image sizes (Sec. VI), we will use the average chord length and the specific surface area as possible indicators of the necessary training image size.

2. Morphological measures

It has been shown that flow properties at the pore scale can be related to morphological characteristics of the void-solid

interface of a porous medium [41]. Hadwiger’s theorem states that any continuous rigid motion invariant valuation on compact-convex subsets of \mathbb{R}^d can be described by a linear combination of $d + 1$ independent parameters characterizing the body. In three dimensions we can therefore define four so-called Minkowski functionals that characterize the topology of a three-dimensional object. We compute estimates of three Minkowski functionals: the porosity ϕ , the specific surface area S_V , and the Euler characteristic χ_V corresponding to the zeroth-, first-, and third-order functionals. We compute

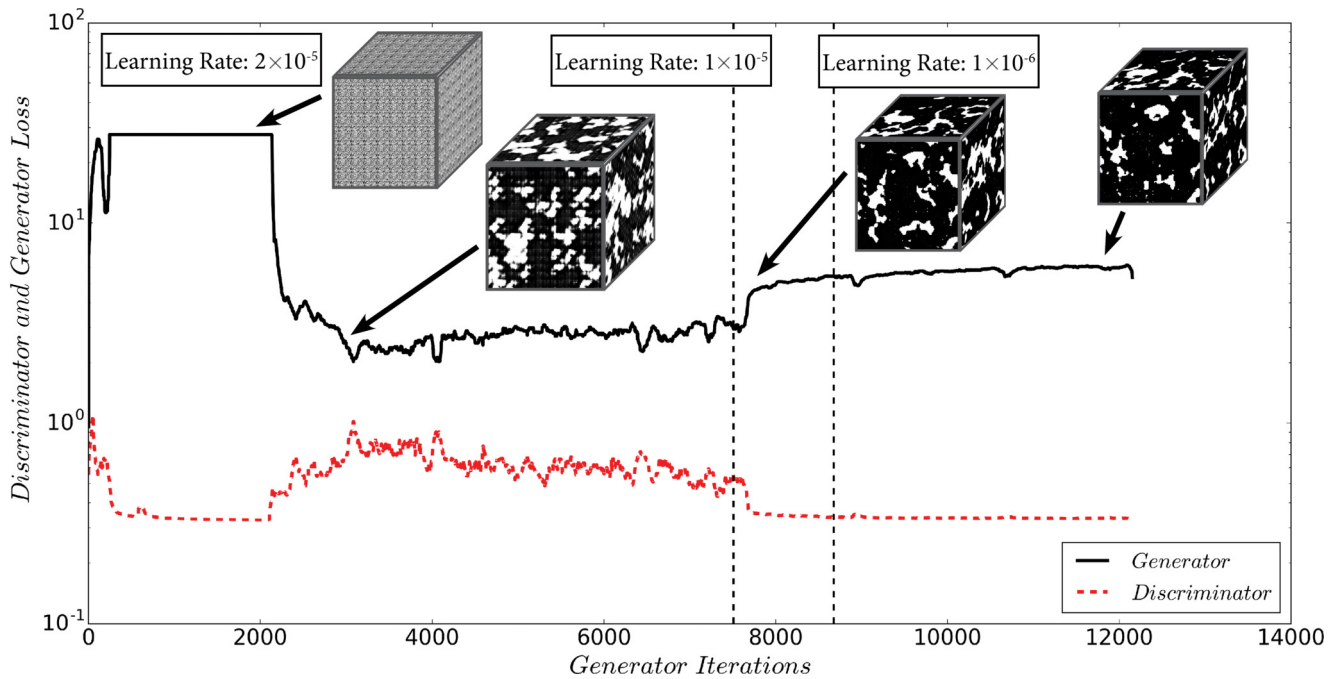


FIG. 4. Value of the cost function, i.e., loss of the discriminator and generator [Eqs. (2) and (3)] for the GAN trained on the Berea sandstone. The samples shown were computed with the same random number seed, showing the evolution of a single realization during training. Initially, image quality is very low and random noise can be observed. After 2000 generator iterations a drop in the generator loss function is observed and coarse structures can be identified in the resulting sample. Loss functions in GAN models do not reflect improvement in image quality, which can be observed from samples. Learning rates [see Eq. (13)] were reduced after sufficient image quality was reached and training stopped based manual inspection of Minkowski functionals and two-point statistics.




FIG. 5. Radial averaged covariance $S_2(r)$ for the bead pack sample and 20 synthetic realizations generated by the GAN model. The specific surface area S_V and mean chord lengths \bar{l}_C are derived from the slope of the covariance at the origin [see Eqs. (5) and (6)].

the densities of the Minkowski functionals by dividing by the volume V .

The Minkowski functional of order zero is the porosity, defined as the ratio of volume of the void space to the bulk volume of the sample

$$\phi = \frac{V_{pore}}{V}, \quad (7)$$

and is therefore a measure of the ability of a porous medium to store fluids. The Minkowski functional of rank one is the specific surface area S_V ,

$$S_V = \frac{1}{V} \int dS, \quad (8)$$

where integration occurs over the void-solid interface S . The specific surface area S_V has dimensions of $\frac{1}{\text{length}}$ and its inverse allows us to define a characteristic pore size.

TABLE II. Chord lengths \bar{l}_C for the pore and grain phase [Eq. (6)] determined from the radial averaged covariance $S_2(r)$ of each training image and corresponding realizations generated by the GAN model. The specific surface area S_V and porosity ϕ were evaluated for each of the samples using direct image morphological computation and derived from the covariance. Close agreement between estimates of the porosity and specific surface area can be observed for values determined by direct image morphological estimation and derived values obtained from the radial averaged covariance.

Parameter	Bead pack				Berea				Ketton			
\bar{l}_C^{pore} (voxel)	20				10				9			
\bar{l}_C^{grain} (voxel)	36				41				64			
	Training image		Synthetic		Training image		Synthetic		Training image		Synthetic	
Minkowski functional	$S_2(r)$	direct	$S_2(r)$	direct	$S_2(r)$	direct	$S_2(r)$	direct	$S_2(r)$	direct	$S_2(r)$	direct
porosity ϕ	0.363	0.359	0.368	0.366	0.196	0.198	0.199	0.197	0.127	0.119	0.119	0.119
$S_V \times 10^{-2}$ ($\frac{1}{\text{voxel}}$)	7.0	7.3	6.9	7.5	7.5	8.2	7.9	8.5	5.2	5.2	4.7	5.2

The specific Euler characteristic is closely related to the order-3 Minkowski functional and is defined as

$$\chi_V = \frac{1}{4\pi V} \int \frac{1}{r_1 r_2} dS, \quad (9)$$

where r_1 and r_2 are the principal radii of curvature of the void-solid interface. To compute χ_V we do not directly evaluate the integral in Eq. (9) but instead make use of a relationship for the Euler characteristic of arbitrary polyhedra

$$\chi = V - E + F - O, \quad (10)$$

where V is the number of vertices, E the number of edges, F the number of faces, and O the number of objects [1]. This expression is the basis for efficient algorithms to compute Minkowski functionals of arbitrary geometric bodies represented as volumetric voxelized domains [42]. To compute these three Minkowski functionals we have used the open-source image morphological software library MorphoLibJ [43].

While the porosity expresses the ability to store fluids in a porous medium, adsorption and dissolution processes are controlled by the specific surface area. The Euler characteristic allows the connectivity of the porous medium to be characterized, which is a critical component in the ability of fluids to flow. Reconstructions of porous media should therefore closely match the observed Minkowski functionals to represent the behavior of relevant physical processes at the pore scale. The direct computation of the specific surface area S_V and porosity ϕ from images allows us to perform a comparison with the values obtained from estimates obtained by computing the empirical noncentered covariance $S_2(r)$ [see Eq. (5)].

3. Single-phase permeability

To evaluate the single-phase permeability of the porous media and their generated synthetic reconstructions we solve the Stokes equations for slow incompressible flow assuming small inertial forces

$$\nabla \cdot \mathbf{v} = 0, \quad (11a)$$

$$\mu \nabla^2 \mathbf{v} = \nabla p. \quad (11b)$$

The Stokes equations are solved on the domain that is connected to the fluid inlet and outlet. This allows us to define

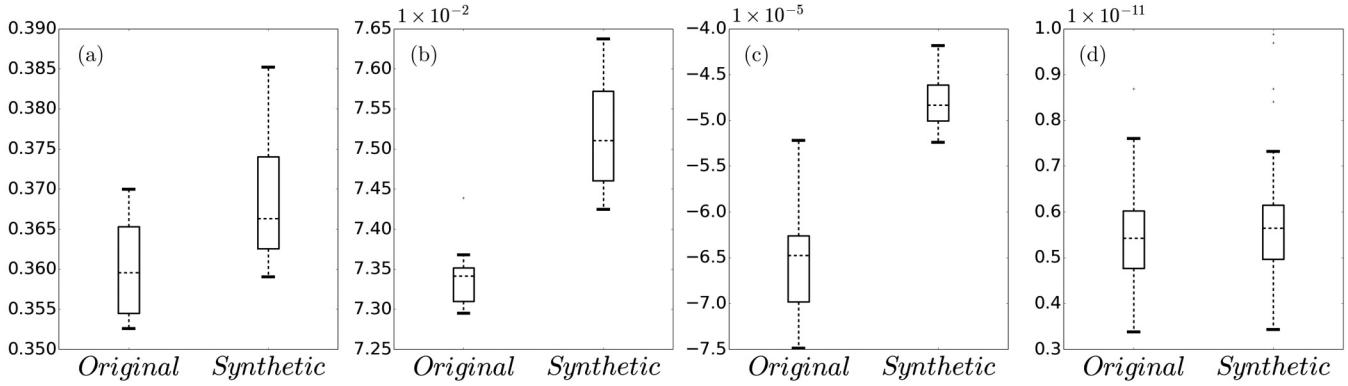


FIG. 6. Comparison of three Minkowski functionals [(a) porosity, (b) specific surface area, and (c) Euler characteristic] for the bead pack evaluated on 200^3 -voxel subdomains of the original training image and realizations of the GAN model. An error of less than 5% can be observed for the porosity and surface area. (d) Comparison of the estimated values of permeability.

an effective porosity where only the fraction of the pore space that also contributes to flow is considered

$$\phi_{\text{eff}} = \frac{V_{\text{flow}}}{V}. \quad (12)$$

A finite-difference method to solve Eqs. (11a) and (11b) on pore-space representations has been implemented as a parallel flow solver, in the free open source numerical framework OpenFOAM [4,44].

B. Neural network architecture

The neural network architecture used for the three-dimensional image reconstruction corresponds to a volumetric version of the DCGAN network [45]. The network consists of two independent fully convolutional neural networks, the generator G and the discriminator D . Upsampling from the input latent vector \mathbf{z} is performed by volumetric transposed convolution, followed by batch normalization and a rectified linear unit (ReLU) activation function in all layers except the last [46,47].

The discriminator D receives images sampled from the latent space by the generator $G(\mathbf{z})$ and images from the set of training images representing $p_{\text{data}}(\mathbf{x})$. Therefore, the size of the input layer of the discriminator corresponds to the dimensions of the input training images. The discriminator consists of volumetric convolution layers combined with LeakyReLU activation functions [48]. The final convolutional layer of the discriminator is followed by a hyperbolic tangent activation function.

This combination of generator and discriminator neural network architectures has previously been applied to subsets of the Imagenet and CIFAR-10 data sets [45]. The hyperparameters for the generator to be used in the optimization of the neural network architecture are the number of trainable convolutional filters in each layer of the neural network $N_{G,F}$ and $N_{D,F}$ and the size of the latent vector \mathbf{z} .

The generator and discriminator are optimized using a gradient-descent-based method where the parameters w are changed by taking k steps in the gradient

$$w^{k+1} = w^k - \alpha \nabla f(w^k), \quad (13)$$




FIG. 7. Comparison of the directional covariance [(a) x direction, (b) y direction, and (c) z direction] of the bead pack and the average covariance of GAN-based synthetic realizations. A clear hole effect can be observed in the original data set, which is captured by the GAN model.




FIG. 8. Comparison of the numerically estimated single-phase permeability of the bead pack for 128^3 -voxel subdomains of the original image and equal-size GAN-based realizations shows a slight overestimation of the effective porosity for the synthetic models. The mean and variance of both permeability distributions are in close agreement (see Fig. 6).

where α is the learning rate. We have used the gradient-descent-based optimizer ADAM for optimization of both the generator and discriminator [49].

Generative adversarial networks have been shown to exhibit unstable behavior during training. The addition of Gaussian noise to the input of the discriminator is an effective way to stabilize the training process [38]. An additional stabilization measure called one-sided label smoothing, wherein the class label of 1 for real images is replaced by a new value of $1 - \epsilon$, has been empirically shown to improve training of GANs [37].

Both label smoothing and white noise addition to the input of the discriminator have been used in this study to stabilize the training based on the volumetric image data sets. Table I gives an overview of the neural network hyperparameters used for each evaluated sample, the hyperparameters, and the stabilization measure used during training.

Images generated by the GAN were postprocessed using a 3^3 median filter to remove single-pixel noise. The resulting images are grayscale images with all voxel values close to zero or one. To compare the resulting images to the binary training images, we segment the generated images using Otsu's method [50].

IV. EXPERIMENTAL DATA

A. Image data and processing

To evaluate the applicability of GANs for reconstruction of natural porous media we use three previously acquired data sets. All images have been segmented into a three-dimensional binary voxel representation of the pore space (white) and

grain (black) (Fig. 3). We create a training database of images by extracting subvolumes from the voxelized binary images. Ideally, these training images should represent independent domains, but due to the limited size of these images, we extract subsets that overlap.

Training image sizes were chosen based on an estimate of the average grain size for each sample. To be able to match the covariance $S_2(r)$ [Eq. (4)] and image morphological characteristics, training images larger than the structuring element were necessary. We discuss this requirement in more detail in the discussion of our results (see Sec. VI). Due to computational limitations, training image sizes exceeding 128^3 voxels were not considered.

1. Bead pack

The bead pack is based on a real packing of equally sized grains in a disordered close packing [51]. The image consists of 500^3 voxels with a size of $3 \mu\text{m}$. The size of an individual sphere is 50 voxels. 1727 training images were extracted of size 128^3 voxels corresponding to a spacing of 32 voxels between them in the original image.

2. Berea

Berea sandstone is a fluvial sandstone of medium to fine grain size (Wentworth classification) [52]. The individual grains are bonded by clays. The sample analyzed in this study was acquired from an outcropping of the Berea sandstone in a quarry near Berea, Ohio. De Witt showed that the Berea sandstone was deposited in the early Carboniferous [354–323 million years ago (mya)] [53].

The image of Berea sandstone consists of angular grains with no clay present in the intergranular pore space. The image has dimensions of 400^3 voxels with a voxel size of $3 \mu\text{m}$.

To capture the local interaction of grains we have extracted training images at 64^3 voxels, which allows a number of grains to be present in one training image (see Sec. VI). Due to the small image size of 400^3 voxels, subvolumes were extracted at a spacing of 16 voxels. In all, 10 647 training images were used for the image reconstruction.

3. Ketton

The Ketton sample is an oolitic limestone of Jurassic age (201.3–145 mya). The sample was acquired from a quarry of Lincolnshire limestone in the northeast of England. The oolites contained in the Lincolnshire formation are mainly nonferroan calcite grains. The oolitic limestones of the Lincolnshire show a wide variety of cementation, ranging from uncemented oolite sands with no intergranular cement to heavily ferroan spar-cemented oolites with infilled microporosity [54]. Microstructures in the pore space can be observed that lead to a reduction in porosity (Fig. 3).

The Ketton sample chosen for this study consists of large grains compared to the overall image size. The image used for the following evaluation has been downsampled from a 500^3 voxel representation to an image size of 256^3 voxels. This allows more grains to be resolved per training image extracted from the full volume. The downsampled voxel size is $15.2 \mu\text{m}$.




FIG. 9. Twenty realizations of the spherical bead pack (top) generated to evaluate the statistical, image morphological, and transport properties considered in this study. Also shown is a cross-sectional view of the bead pack training image data set (bottom).

Training images were extracted at a subvolume size of 64^3 with a spacing of 8 voxels leading to a total of 15 624 training images. The small spacing of the training images results from the small CT image size of 256^3 voxels.

V. RESULTS

Three GANs were trained based on the network architectures highlighted in Sec. III B. The training time for each data set was 24 h. Manual inspection of synthetic realizations was performed during training to ensure convergence and intermediate evaluation of the covariance and Minkowski functionals.

Figure 4 shows the training curve for the Berea sandstone data set. Initially, the generator loss function [see Eq. (3)] is very high and no structural components can be observed in the samples. After a large reduction in the loss function of the generator, initial structures are observed. Image reconstruction quality significantly improves with the number of generator iterations, but cannot be linked to the loss function of the generator. This can be observed from the increase in generator loss at the end of training while image quality improves significantly.

The final GAN models were subsequently evaluated in terms of their directional and radial averaged noncentered

covariance $S_2(r)$, Minkowski functionals, and the single-phase permeability. For all data sets, 20 realizations were generated using the trained GAN model. In the following section we present the results of the evaluation of the properties outlined in Sec. III A and compare these to the properties of the original input training image.

A. Bead pack

The evaluation of the noncentered covariance $S_2(r)$ for the bead pack (Fig. 5) shows a strong hole effect reflecting the spherical nature of the grains. A GAN model was trained for 24 h on the bead pack training image data set. The GAN model achieves a small error in the porosity of the generated images with a tendency towards higher porosities (Fig. 6). A bias can be observed for the specific surface area and the Euler characteristic of the microstructure (Table II).

This bias can be explained by the deviation of the grains from a perfect spherical shape in the synthetic realizations. Due to the smooth nature of the spherical particles in the training image, any deviation from this geometry will lead to an increase in the surface area. This is reflected by a higher specific surface area for the synthetic realizations. In addition, we observe a reduction in connectivity, represented by a less negative Euler characteristic.




FIG. 10. Radial averaged covariance $S_2(r)$ for Berea sandstone training images and 20 synthetic realizations generated by the GAN model.

The directional covariance S_2 measured on the generated samples shows excellent agreement up to the training image size of 128^3 voxels and stabilizes at ϕ^2 (see Fig. 7). As expected, no directional variation of the covariance is observed and the sample is therefore assumed to be isotropic.

Single-phase permeability shows close agreement in both magnitude and variance between the measured training image and the synthetic realizations (Fig. 6). Figure 8 shows a cross-plot of the effective porosity ϕ_{eff} , i.e., the porosity open to flow [Eq. (12)], and the single-phase permeability exhibiting a similar trend in the distribution of values computed on training images and synthetic realizations.

We provide a comparison of all 20 realizations generated by the GAN model in cross sections through the x - y plane of the original model and a synthetic realization in Fig. 9.

Many of the grains show a circular to ellipsoidal shape, for which, considering the fact that *a priori* the GAN model does not have any knowledge of the geometry of the grains, learning a representation of a perfect sphere can be considered challenging (see Sec. VI). The complex grain-grain interface where individual beads are in contact at single points can be observed for numerous grain arrangements in the generated realizations.

B. Berea

The radial averaged covariance $S_2(r)$ in Fig. 10 shows a near exponential decay and stabilization occurs at a lag distance of 30 voxels for both covariance functions obtained from the Berea training image and synthetic realizations generated by the GAN model. Additionally, Fig. 11 shows that the directional two-point statistics characterized by the directional covariances is captured in the generated images. In all three directions, the GAN model shows excellent agreement and closely follows the trend of the empirical estimates of S_2 .

The results of the direct computation of the Minkowski functionals are presented in Fig. 12 and show comparable distributions for the porosity ϕ , the specific surface area S_V , and the Euler characteristic χ_V of the training images and the synthetic realizations. A comparison of the specific surface area S_V obtained from the covariance and the direct computation of the Minkowski functional show nearly equal values (Table II).

The estimates of the single-phase permeability show a similar distribution covering the range of effective permeability measured on the training images. Figure 13 shows the computed values of permeability and the corresponding effective porosity. The permeability of the synthetic realizations capture the values, variability, and trend obtained from the Berea training image data set.

Figure 14 shows a comparison of 20 realizations of the GAN model trained on the Berea data set. A smaller training image size of 64^3 voxels was used, as compared to the bead pack (128^3 voxels). This is due to the smaller size of the structuring elements observed in the training image. A smaller training




FIG. 11. Directional noncentered covariance comparison for Berea sandstone. The trained GAN model shows good agreement with the noncentered covariance S_2 of the training image.

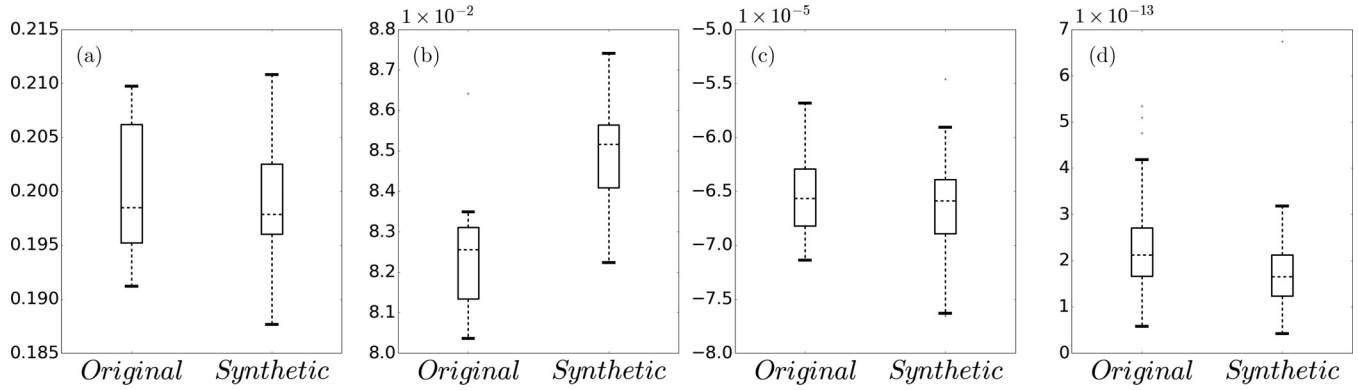


FIG. 12. Comparison of three Minkowski functionals [(a) porosity, (b) specific surface area, and (c) Euler characteristic] for Berea sandstone. The porosity, specific surface area, and specific Euler characteristic show good agreement between the training image and samples from the trained GAN model. (d) Estimates of permeability for the Berea sandstone and generated samples.

image size was therefore sufficient to capture the long- and short-range correlations found in the Berea sample.

C. Ketton

The covariance $S_2(r)$ of the Ketton limestone shown in Fig. 15 shows a pronounced hole effect due to the ellipsoidal oolitic grains. Due to the hole effect observed in the radial averaged covariance (Fig. 15), we relate the Ketton sample to a hard-sphere model (Fig. 2). Figure 16 indicates that the images generated by the GAN model trained on the Ketton image capture the oscillatory and anisotropic behavior of the covariance observed in Ketton. The specific surface area S_V derived from the generated images is in close agreement with

the training data. An error of approximately 1% was achieved in the porosity of the GAN generated images compared to the original Ketton data set (Fig. 17).

The measured specific surface area of the synthetic images shows a higher variance compared to the original training images. Nevertheless, the average values of the porosity ϕ and specific surface area S_V derived from the noncentered covariance $S_2(r)$ [see Eq. (5)] are in good agreement with values obtained from direct image morphological estimation (see Table II).

The distribution of single-phase permeability estimates of the synthetic GAN realizations overlies the permeability values of the Ketton training images (Fig. 18). The Euler characteristic χ_V and the permeability of the Ketton training data set are closely matched by the synthetic images and therefore capture the connectivity observed in the oolitic Ketton limestone. We present an overview of the 20 realizations generated by the GAN model trained on the Ketton data set in Fig. 19.

VI. DISCUSSION

This paper presents a method for three-dimensional stochastic image reconstruction based on generative adversarial neural networks trained on three-dimensional segmented images. To summarize, the objectives of this contribution are threefold. First, generate stochastic reconstructions of porous media such as sedimentary rocks exceeding the size of the acquired image data sets. Second, evaluate the ability of GAN models to capture the image morphological and physical properties of microscale porous media. Third, establish a method of stochastic image reconstruction that allows a probabilistic treatment of pore-scale properties such as permeability without the need to acquire numerous images of a single rock type.

The first objective stems from technical limitations of micro-CT data acquisition. Images are acquired as a trade-off between sample size, i.e., how many representative structures can be captured in one image versus the resolution at which these pore-scale structures are resolved. The generation of large porous domains based on high-resolution images enables this gap in scales to be bridged and microscale features to be incorporated in macroscale models.




FIG. 13. Distribution of numerically obtained permeability values on 128^3 -voxel subdomains and sampled realizations obtained from a GAN model trained on the Berea sandstone data set show close agreement in the effective porosity, as well as the evaluated permeability.




FIG. 14. Realizations generated by the GAN model (top) compared to training images (bottom) for Berea sandstone.

Our findings show that GANs can learn an implicit representation of the image space given a limited number of training images subsampled from larger images. These subdomains were extracted based on characteristic length scales (see Sec. III A 1) and serve as a training set for the GAN model. For the Ketton limestone, a small spacing of the extracted subdomains was required to increase the size of the training image data set. While we did not find any evidence of an introduced bias by using correlated subdomains, we believe that these extracted training images should represent independent regions.

We have evaluated the ability to train GANs for a number of training image sizes less than and up to twice the size of the structuring elements. We have found that models trained on images smaller than the average grain size results in artifacts and distorted shapes occurring in the generated microstructures. For the bead pack, the size of an individual sphere is 50 voxels. A training image of 64^3 voxels would typically only contain parts of an individual grain and only capture the interaction of the particles, but not the geometry of the structuring element. For the bead pack, models trained on 64^3 voxels were successful in learning a representation of the short-scale microstructure but failed to reproduce the long-distance correlation. A larger training image of 128^3 voxels, as was used to model the bead pack, has a much higher chance to represent the full geometry of the particles and

therefore learn not only interactions, but also the shapes of grains.

We therefore suggest that training images extracted from large data sets must be larger than the average grain size. For models that are well described by a Boolean model, the size of the structuring element can be readily estimated from stabilization of the covariance $S_2(r)$. For more complex samples a different measure must be used to estimate the size of the required training image.

The chord length is one additional measure that can be obtained to characterize the grain space of porous media. While we have found that the mean chord length of the grain space \bar{l}_C^{grain} is always less than or equal to that of our chosen training image size, \bar{l}_C^{grain} increases with decreasing porosity. This contradicts the need to have the largest training domain for the bead pack sample, which also has the highest porosity. A better estimate may be derived from the representative elementary volume of the specific surface area, which by definition is the same for the grain and pore space and is therefore more representative of the morphology of the porous medium [55]. Based on the properties we have evaluated we could not find a measure derived from two-point statistical or image morphological properties that is closely related to the required training image size and we see a theoretical discussion of this as possible future work.




FIG. 15. Radial averaged covariance $S_2(r)$ for the oolitic Ketton limestone training image and 20 synthetic realizations generated by the GAN model.

Conceptually, the simplest model considered in this study, the spherical bead pack, has proven to be the most challenging as a training image for the GAN model (Sec. V A). While we observe spherical and ellipsoidal shapes in the resulting realizations (see Fig. 9), the shape is exactly defined by the spherical nature of the grains. Any deviation from this shape, which for GANs is learned implicitly from the data itself, will lead to a misrepresentation of the effective properties. Random hard-sphere models with spherical grains will efficiently capture the nature of this data set. Therefore, we suggest a fit-for-purpose application of GANs, for training images that exhibit variability of grain sizes and shapes, which are not readily captured by a simpler model.

While for many sedimentary granular rocks representative volumetric images can be obtained, this may be more challenging for carbonate samples with complex pore-grain structures. The three training images considered in this study were all treated under the assumption of stationarity, i.e., we do not expect a systematic variation in the mean and variance of the averaged properties as a function of location. In theory, GANs are not limited to learning representations of stationary data sets. This is shown by the many successful applications for two-dimensional image and texture synthesis of nonstationary domains, such as learned image representations of human faces [56] or galaxies [57,58]. Therefore, a model that incorporates nonstationarity for a single rock type would technically be possible in the GAN framework but would require the acquisition of many images of the same porous medium.

A valid representation of the microscale variability and connectivity of the pore space is critical to assess the single- and multiphase flow behaviors of porous media. Therefore, any stochastic reconstruction method used in the process of deriving or evaluating the variability of microscale properties must capture the statistical and image morphological characteristics of the reconstructed porous medium. While we have shown that for the evaluated data sets the GAN-based image reconstructions capture the variation and characteristics of these porous media, a number of challenges remain in this task that are different from those encountered by classical stochastic methods of image reconstruction.

For porous media, many flow-related properties can be related to the porosity. Classical stochastic methods are able to capture the porosity efficiently by defining a specific proportion of the grain and pore domain. The GAN-based model presented in this study initially has no knowledge of the porosity. The porosity therefore arises as a feature of the training image data. Matching the porosity distribution of the training image data set was found to be the main challenge in training a GAN model. An error of 3% in porosity, for instance, could lead to a significant mismatch in the permeability of the synthetic images. It is therefore necessary to continuously monitor the derived properties such as the




FIG. 16. The directional covariance of the Ketton sample shows (a) oscillating behavior in the x direction, whereas a nearly exponential decrease can be observed for the (b) y and (c) z directions. This anisotropy in $S_2(r)$ is also reflected in the covariance of the samples obtained from the GAN model.

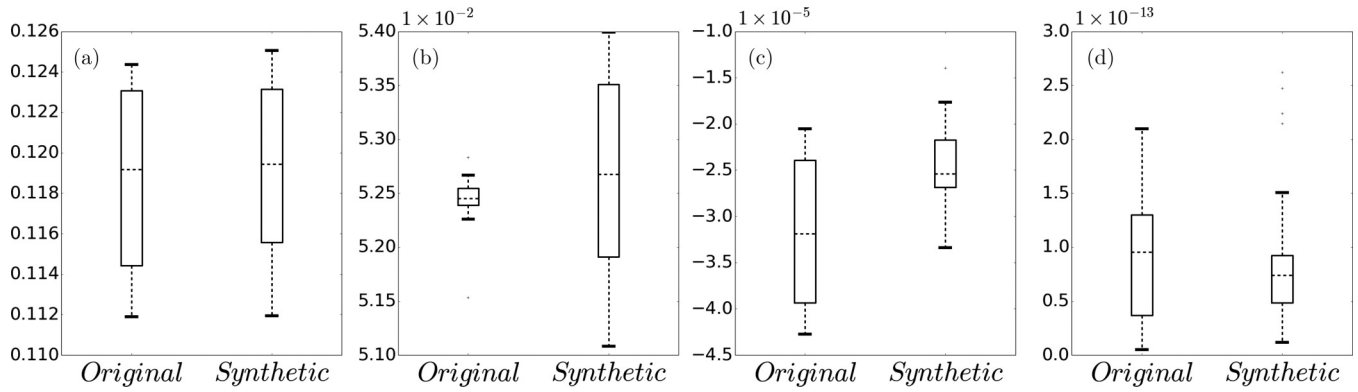


FIG. 17. Comparison of the Minkowski functionals [(a) porosity, (b) specific surface area, and (c) Euler characteristic] for the Ketton training image. The three evaluated Minkowski functionals show good agreement. The evaluated Euler characteristic indicates that the sampled synthetic realizations show a degree of connectivity similar to that of the training image. (d) Estimated permeability values for the Ketton sample.

Minkowski functionals or estimates of the permeability in the course of training the neural networks to ensure that synthetic realizations created by the GAN model are able to capture the effective properties of the microscale domains.

While this can be considered one of the main challenges in the application of GANs for synthetic image reconstruction, learning an implicit representation of the training data itself can be seen as a strength. Many classic stochastic methods rely on the formulation of an objective function that ensures that statistical properties are captured in the generated realizations, e.g., matching $S_2(r)$ and the specific surface area S_V of the stochastic reconstructions to a desired precision. The GAN approach does not require an

explicit objective function *a priori*. The objective function is encoded in the discriminator and adapted in the course of training.

During adversarial training both the generator and discriminator are continuously improved. The discriminator's sole purpose is to be able to distinguish real training data from generated synthetic data. On the other hand, the generator tries to generate synthetic data that the discriminator is not able to distinguish from the training data. Due to the multiscale representation of the convolutional neural networks, these features must be learned across the full range of length scales present in the training data, leading to a high-resolution image that captures small- and large-scale features of the image data set. A number of stacked GAN models can be trained on, e.g., low-resolution medical-CT data and high-resolution micro-CT allowing incorporation of spatial information across multiple length scales [59].

Once the GAN model has successfully learned to create physically representative samples of the porous medium, one possible application is to evaluate the variability in the flow properties by evaluating the properties of a large number of samples. This not only requires a physically valid representation of the porous medium but also requires a method that allows fast image reconstruction. In Sec. V we have shown that training was performed for approximately 24 h and may vary due to the need for manual inspection of the generated samples in the training process. Figure 20 shows the CPU time required for generation of images at increasing image size. The fully convolutional nature of the GAN architecture allows very large images, exceeding the size of the original sample to be generated very efficiently and at low computational cost and run-time.

While training requires considerable time and computational resources in the form of modern graphics processors as well as optimized neural network frameworks, image reconstruction requires little computational effort and scales linearly in the total number of voxels of the generated images. This therefore enables the generation of ensembles of large domains based on volumetric images acquired from 3D microscopy that capture the physical behavior of the porous medium. The learned representation of the generator consists

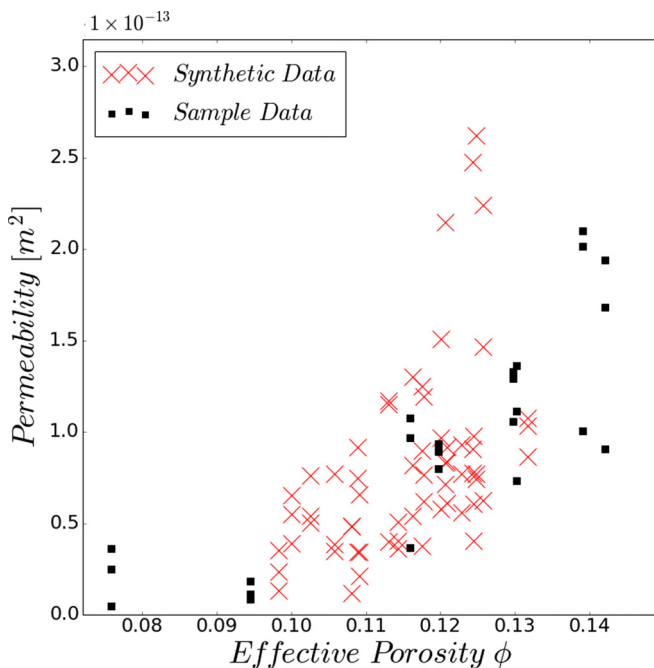


FIG. 18. Evaluated single-phase permeability for the Ketton training image. The synthetic realizations show effective porosity and permeability similar to those of the Ketton sample.




FIG. 19. Realizations generated by the GAN model (top) compared to training images (bottom) for Ketton limestone.

of the weights of the convolutional filters learned in the training process and can therefore be stored for future use once training has finished.

VII. CONCLUSION

We have evaluated the application of generative adversarial neural networks for stochastic image reconstruction of porous

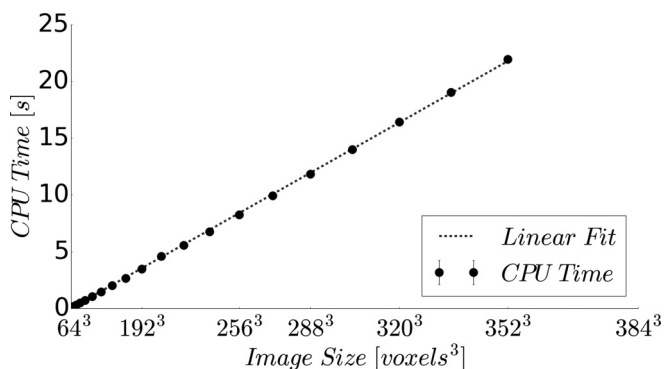


FIG. 20. Measured CPU time for generating synthetic realizations of Berea sandstone at increasing image size. One hundred realizations were computed at each image dimension and CPU time averaged. Computational cost increases linearly with the number of voxels in the generated image.

media based on previously acquired images of sedimentary rocks. Three image data sets were used as training images: a bead pack, a Berea sandstone, and an oolitic Ketton limestone.

By evaluating two-point statistical measures and image morphological features and computing the single-phase effective permeability we have shown that the synthetic images generated by the GAN model are able to match key characteristic statistical and physical parameters of these porous media. While a large computational effort is required to train the GAN model, the generation of samples from the learned representation is highly efficient and learned models are easily stored for future use.

Future work in the application of GANs to stochastic image reconstruction of porous media will include improving the quality of the image reconstruction by evaluating various generator-discriminator architectures, the use of grayscale and multichannel training images, and the application of large-multiscale domains of porous media to evaluate the ensemble behavior of single- and multiphase flow properties in porous media. Recent advances in the understanding of GANs should lead to a more stable and consistent training process [60,61].

ACKNOWLEDGMENT

O.D. thanks Total S.A. for support as a Visiting Professor at Imperial College.

- [1] M. J. Blunt, *Multiphase Flow in Permeable Media: A Pore-Scale Perspective* (Cambridge University Press, Cambridge, 2017).
- [2] R. Juanes, E. J. Spiteri, F. M. Orr, Jr., and M. J. Blunt, *Water Resour. Res.* **42**, W12418 (2006).
- [3] Q. Kang, P. C. Lichtner, H. S. Viswanathan, and A. I. Abdel-Fattah, *Transp. Porous Media* **82**, 197 (2010).
- [4] P. Mostaghimi, M. J. Blunt, and B. Bijeljic, *Math. Geosci.* **45**, 103 (2013).
- [5] G. Matheron, *Random Sets and Integral Geometry* (Wiley, New York, 1975).
- [6] J. Serra, *Comput. Graphics Image Processing* **12**, 99 (1980).
- [7] G. Matheron, *Elements pour une Theorie des Milieux Poreux* (Masson, Paris, 1971).
- [8] C. H. Arns, M. A. Knackstedt, and K. R. Mecke, *Phys. Rev. E* **80**, 051303 (2009).
- [9] P. A. Rikvold and G. Stell, *J. Chem. Phys.* **82**, 1014 (1985).
- [10] S. Torquato, *Random Heterogeneous Materials: Microstructure and Macroscopic Properties* (Springer Science + Business Media, New York, 2013), Vol. 16.
- [11] T. Bretheau and D. Jeulin, *Rev. Phys. Appl.* **24**, 861 (1989).
- [12] D. Jeulin, *Stat. Comput.* **10**, 121 (2000).
- [13] C. Lin and M. H. Cohen, *J. Appl. Phys.* **53**, 4152 (1982).
- [14] C. L. Y. Yeong and S. Torquato, *Phys. Rev. E* **57**, 495 (1998).
- [15] P.-E. Øren and S. Bakke, *J. Petrol. Sci. Eng.* **39**, 177 (2003).
- [16] M. J. Pyrcz and C. V. Deutsch, *Geostatistical Reservoir Modeling* (Oxford University Press, Oxford, 2014).
- [17] J. Caers and T. Zhang, *AAPG Mem.* **80**, 383 (2004).
- [18] G. Mariethoz, P. Renard, and J. Straubhaar, *Water Resour. Res.* **46**, 11 (2010).
- [19] E. Meerschman, G. Pirot, G. Mariethoz, J. Straubhaar, M. Van Meirvenne, and P. Renard, *Comput. Geosci.* **52**, 307 (2013).
- [20] P. Tahmasebi, M. Sahimi, and J. Caers, *Comput. Geosci.* **67**, 75 (2014).
- [21] A. Comunian, P. Renard, J. Straubhaar, and P. Bayer, *J. Hydrol.* **405**, 10 (2011).
- [22] G. Mariethoz and J. Caers, *Multiple-Point Geostatistics: Stochastic Modeling with Training Images* (Wiley, New York, 2014).
- [23] J. Straubhaar, P. Renard, G. Mariethoz, R. Froidevaux, and O. Besson, *Math. Geosci.* **43**, 305 (2011).
- [24] T. Huang, X. Li, T. Zhang, and D.-T. Lu, *Comput. Geosci.* **57**, 13 (2013).
- [25] H. Okabe and M. J. Blunt, *Phys. Rev. E* **70**, 066135 (2004).
- [26] H. Okabe and M. J. Blunt, *J. Petrol. Sci. Eng.* **46**, 121 (2005).
- [27] H. Okabe and M. J. Blunt, *Water Resour. Res.* **43**, 3 (2007).
- [28] W. E. Smith, H. H. Barrett, and R. G. Paxman, *Opt. Lett.* **8**, 199 (1983).
- [29] D. I. Svergun, *Biophys. J.* **76**, 2879 (1999).
- [30] L. M. Pant, Stochastic characterization and reconstruction of porous media, Ph.D. thesis, University of Alberta, 2016.
- [31] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran, Red Hook, 2014), pp. 2672–2680.
- [32] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling, in *30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), pp. 82–90.
- [33] E. L. Denton, S. Chintala, A. Szlam, and R. Fergus, Deep generative image models using a Laplacian pyramid of adversarial Networks, in *29th Annual Conference on Neural Information Processing Systems (NIPS 2015), Montreal, Canada*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., 2015), pp. 1486–1494.
- [34] I. Goodfellow, [arXiv:1701.00160](https://arxiv.org/abs/1701.00160).
- [35] <https://github.com/LukasMosser/PorousMediaGan>.
- [36] <http://www.imperial.ac.uk/earth-science/research/research-groups/perm/research/pore-scale-modelling/micro-ct-images-and-networks/>.
- [37] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran, Red Hook, 2016), pp. 2226–2234.
- [38] C. Kaae Sønnderby, J. Caballero, L. Theis, W. Shi, and F. Huszár, [arXiv:1610.04490](https://arxiv.org/abs/1610.04490).
- [39] P. Debye, H. R. Anderson, and H. Brumberger, *J. Appl. Phys.* **28**, 679 (1957).
- [40] S. Torquato and F. Lado, *J. Phys. A* **18**, 141 (1985).
- [41] C. Scholz, F. Wirner, M. A. Klatt, D. Hirneise, G. E. Schröder-Turk, K. Mecke, and C. Bechinger, *Phys. Rev. E* **92**, 043023 (2015).
- [42] C. Lang, J. Ohser, and R. Hilfer, *J. Microsc.* **203**, 303 (2001).
- [43] D. Legland, I. Arganda-Carreras, and P. Andrey, *Bioinformatics* **32**, 3532 (2016).
- [44] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, *Comput. Phys.* **12**, 620 (1998).
- [45] A. Radford, L. Metz, and S. Chintala, [arXiv:1511.06434](https://arxiv.org/abs/1511.06434).
- [46] S. Ioffe and C. Szegedy, in *Proceedings of the 32nd International Conference on Machine Learning (JMLR, 2015)*, pp. 448–456.
- [47] V. Nair and G. E. Hinton, in *Proceedings of the 27th International Conference on Machine Learning, Haifa, 2010* (to be published), pp. 807–814.
- [48] A. L. Maas, A. Y. Hannun, and A. Y. Ng, *Proc. ICML*. **30**, 1 (2013).
- [49] D. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [50] N. Otsu, *Automatica* **11**, 23 (1975).
- [51] J. Finney, *Proc. R. Soc. London Ser. A* **319**, 479 (1970).
- [52] J. F. Pepper, *Geology of the Bedford Shale and Berea Sandstone in the Appalachian Basin: A Study of the Stratigraphy, Sedimentation and Paleogeography of Rocks of Bedford and Berea Age in Ohio and Adjacent States* (U.S. Government Printing Office, Washington, DC, 1954), Vol. 259.
- [53] W. de Witt, Jr., *U.S. Geol. Surv. Bull.* **1294**, A23 (1970).
- [54] D. Emery, *J. Geol. Soc.* **145**, 621 (1988).
- [55] J. Bear, *Dynamics of Fluids in Porous Media* (Courier, North Chelmsford, 2013).
- [56] J. Gauthier, Class project for Stanford CS231n: Convolutional neural networks for visual recognition, 2014 (to be published).
- [57] K. Schawinski, C. Zhang, H. Zhang, L. Fowler, and G. Krishnan Santhanam, *Mon. Not. R. Astron. Soc.: Letters*, **467**, L110 (2017).

- [58] S. Ravanbakhsh, F. Lanusse, R. Mandelbaum, J. G. Schneider, and B. Póczos, *Enabling Dark Energy Science with Deep Generative Models of Galaxy Images* (AAAI, Menlo Park, CA, 2017), pp. 1488–1494.
- [59] H. Zhang, T. Xu, H. Li, S. Zhang, X. Huang, X. Wang, and D. Metaxas, [arXiv:1612.03242](https://arxiv.org/abs/1612.03242).
- [60] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, and Z. Wang, [arXiv:1611.04076](https://arxiv.org/abs/1611.04076).
- [61] M. Arjovsky, S. Chintala, and L. Bottou, [arXiv:1701.07875](https://arxiv.org/abs/1701.07875).