

1 **Quantifying complex microstructures of earth**
2 **materials: Reconstructing higher-order spatial**
3 **correlations using deep generative adversarial networks**

4 **Hamed Amiri¹, Ivan Vasconcelos¹, Yang Jiao², Pei-En Chen³, Oliver**
5 **Plümer¹**

6 ¹Department of Earth Sciences, Utrecht University, Utrecht, the Netherlands

7 ²Materials Science and Engineering, Arizona State University, Tempe, U.S.A.

8 ³Mechanical and Aerospace Engineering, Arizona State University, Tempe, U.S.A.

9 **Key Points:**

- 10 • We train a GAN model with Wasserstein-loss and gradient penalty (WGAN-GP)
11 to accurately reconstruct 2D images of rock samples with complex structures and
12 morphology.
- 13 • Polytope functions are used as a subset of n -point spatial correlations to assess the
14 reconstruction performance.
- 15 • Compared with a stochastic reconstruction method, we show that our model is able to
16 naturally reconstruct higher-order correlation functions which is important for many
17 geological post-modelling purposes.

Corresponding author: Hamed Amiri, h.amiri@uu.nl

Abstract

Key to most subsurface processes is to determine how structural and topological features at small length scales, i.e., the microstructure, control the effective and macroscopic properties of earth materials. Recent progress in imaging technology has enabled us to visualise and characterise microstructures at different length scales and dimensions. An approach to characterisation is the sampling of n-point correlation functions - known as statistical microstructural descriptors (SMDs) - from images. SMDs can then be used to generate statistically equivalent structures having larger sizes and additional dimensions - this process is known as *reconstruction*. We show that a deep-convolutional generative adversarial network trained with Wasserstein-loss and gradient penalty (WGAN-GP) results in a stable training and high-quality reconstructions of two-dimensional electron microscopy images of complex rock samples. To evaluate reconstruction performance, n-point polytope functions are calculated in both reconstructed and original microstructures and mean square error between them is used as a quality metric. These n-point polytope functions provide statistical information about symmetric, user-oriented higher-order geometrical patterns in microstructures. Our results show that GANs can naturally capture these higher-order statistics at short and long ranges. Furthermore, we compare our model with a benchmark stochastic reconstruction method based solely on two-point correlation. Our findings indicate that although yielding the same two-point statistics, two microstructures can be morphologically and structurally different, emphasising the need for coupling higher-order correlation functions with reconstruction methods. This is a critical step for future schemes that aim to reconstruct complex heterogeneous systems and couple microstructures to macroscopic phenomena.

Plain Language Summary

In our work, we try to train a computer how very small structures, called microstructures, hidden within rocks look like. We show the computer thousands of electron microscope images and use spatial statistics as well as artificial intelligence to have the computer regenerate realistic microstructures. We find that our artificial intelligence is extremely good in generating *fake* microstructures that are indistinguishable from the real ones, giving us hope that in the future we can make artificial three-dimensional rocks although only having information from 2-dimensional microscope images, called 2D to 3D reconstruction. Developing such a reconstruction provides great technical and economical benefits as 2D images are more affordable, widely available, and can cover larger area of a sample in higher resolutions. This approach also paves the way to understand how these microstructures at small scales within rocks control e.g., the failure of rocks to generate earthquakes or the transport of water and gas for renewable subsurface energy.

1 Introduction

Many geological phenomena within the Earth result from physicochemical processes occurring at length scales ranging from nanometers to micrometres. For example, the motion of tectonic plates is associated with the movement of atomic imperfections (i.e., dislocations) within individual mineral grains (e.g., Wallis et al. (2021)). Likewise, the transport of fluids within the crystalline lithosphere (Plümper, John, et al., 2017; Plümper, Botan, et al., 2017) and reservoirs for the storage of CO₂ (Bourg et al., 2015) and hydrogen (Mouli-Castillo et al., 2019) are governed by processes occurring at grain contacts and within microscopic pore spaces. However, it is not the individual dislocation nor pore that controls phenomena at larger length scales but the interaction of many dislocations and pores in tandem. Hence, to understand and model geological processes at larger length scales, we need to (i) quantitatively characterise microstructures that both includes existing samples and generalises beyond them and (ii) establish a link between these microstructures and the physical properties of rocks across length scales.

68 Two steps need to be taken to achieve a comprehensive description of microstructures.
69 On the one hand, samples of complex earth materials need to be imaged with sufficient
70 resolution to capture the smallest features of any given system. On the other hand, reliable
71 mathematical descriptions - likely coupled to machine-learning algorithms - need to be
72 developed to describe complex microstructures quantitatively and generate a larger yet
73 statistically equivalent sample, allowing us to upscale both in lengths and dimensions and
74 go beyond the limitations of available imaging techniques.

75 In the past decades, rapid advances in imaging technologies have made it possible to
76 characterise earth materials at different length scales. For example, electron backscattered
77 diffraction is a 2D imaging technique used to describe polycrystalline rocks and gives valu-
78 able information such as the lattice-preferred orientation of minerals in multiphase systems,
79 grain shape, and distribution at different scales (Britton et al., 2016; Prior et al., 2009).
80 Scanning electron microscopy (SEM) utilising backscattered electron (BSE) imaging is an-
81 other technique that can be used to acquire high-resolution images of large areas of rock
82 surfaces. However, applications of such 2D imaging techniques remain limited as many geo-
83 logical processes such as rock deformation and fluid transport in porous media are inherently
84 volumetric.

85 X-ray tomography is a widely used technique to obtain three-dimensional images of
86 rock microstructures, which provides detailed information about internal structures with a
87 maximum pixel size of $\pm 0.5 \mu\text{m}$. However, a resolution of $\pm 0.5 \mu\text{m}$ would not be sufficient to
88 resolve much smaller features observed in complex heterogeneous media such as carbonates
89 (Dehghan Khalili et al., 2013), shales (Wu et al., 2019) or dense crystalline rocks. In
90 such circumstances, techniques such as focused ion beam nanotomography (e.g., Holzer and
91 Cantoni (2012); Liu et al. (2016)) can be employed, which acquires images with nanometer
92 pixel size but at the expense of the field of view (FoV). This tradeoff highlights the inherent
93 limitation of imaging technologies; resolution and FoV are in direct competition. Moreover,
94 while tiny structures may control the overall behaviour of a given medium, the modelling
95 domain (i.e., FoV) needs to be sufficiently large to be representative of the whole system
96 (Niu et al., 2020).

97 Another challenge in heterogeneous systems is that the microstructural properties can
98 significantly vary from one sample to another, so their variabilities also need to be evaluated
99 to have a more realistic model. The variability assessment is typically done by conducting
100 several imaging experiments or numerical simulations on different samples, allowing us to
101 obtain a distribution over larger samples and capture the heterogeneity of the medium (e.g.,
102 Mosser et al. (2017)). However, acquiring large image datasets comes with high costs and a
103 severe time penalty.

104 A pragmatic approach is to reconstruct synthetic but still realistic images of a given
105 microstructure using reconstruction methods with no experimental limitations. The funda-
106 mental underlying assumption of this approach is that the geometrical patterns of a limited
107 number of samples are representative of a large class of materials sharing those patterns.
108 As such, these patterns, often represented in terms of spatial statistics, must be implicitly
109 or explicitly exploited in the reconstruction process. Thus, reconstruction of heterogeneous
110 media is an inverse problem in which a limited amount of microstructural information is
111 used to (re)construct realistic microstructures and evaluate macroscopic properties (Jiao
112 et al., 2007; Yeong & Torquato, 1998). Consequently, image reconstruction has become
113 an essential aspect of digital rock physics to produce representative samples for upscaling,
114 multi-scale modelling, and uncertainty assessment.

115 Several methods have been developed for image reconstruction in recent years, which
116 can be grouped into two main approaches: stochastic methods and deep-learning-based
117 reconstructions. Despite their differences, both methods should be supplemented with mi-
118 crostructure characterisations of the system. Characterisation refers to statistically quan-
119 tifying and representing the morphology of a system using spatial correlation functions,

also known as statistical microstructure descriptors (Bostanabad et al., 2018). Traditional stochastic methods take such correlation functions as an input and construct a synthetic microstructure with the same characteristics. On the other hand, although deep-learning methods do not require such information a priori, it is necessary to compute the correlation functions from original and reconstructed images to evaluate the reconstruction accuracy.

Most research on stochastic methods formulates the image reconstruction as an optimisation problem in which n -point correlation functions, defined as the probability of n random points to lie in a phase of interest (e.g., solid, liquid, or void), are calculated from original images and used as target functions. Next, these methods seek to reconstruct a medium for which the calculated correlation function(s) matches the target function(s) derived from the original image. This match can be obtained by applying stochastic optimisation techniques such as simulated annealing (SA) (Jiao et al., 2007, 2008; Sheehan & Torquato, 2001). While such a framework is shown to be able to reconstruct single-scale microstructures such as Fontainebleau sandstone using a simple two-point correlation function (Jiao et al., 2008), it fails in the case of multi-scale complex heterogeneous systems since it only captures the largest scale features (Gommes et al., 2012b; Jiao et al., 2010). Recently, Karsanina and Gerke (2018) proposed a novel hierarchical optimisation approach to incorporate two-point correlations of different scales in SA for reconstructing coarse and fine microstructures in a single image.

An alternative to two-point correlation is to employ multi-point statistics (MPS), or high-order n -point correlation functions ($n \geq 3$). These methods have been used for 3D image reconstruction from 2D images, showing to be more effective in long-range connectivity (Okabe & Blunt, 2005; Strebelle, 2002). While reconstructing more realistic images, these methods and improved variants (Hajizadeh et al., 2011; Tahmasebi & Sahimi, 2012, 2013) are computationally costly and limited to isotropic media. Chen et al. (2019) developed a set of hierarchical descriptors, termed n -point polytope functions, which successively capture higher-order correlations of a given phase in an image. In contrast to MPS, these polytope functions can be computed significantly faster since only the probability of n vertices of a randomly- or user-selected regular polytope (e.g., triangle, square, hexagon for $n = 3, 4$ and 6 , respectively) is considered, i.e., they can be seen as a subset of n -point correlation functions with a fixed edge length. It has been shown that incorporating these higher-order correlation functions in SA as target functions will improve the reconstruction accuracy. However, adding more correlation functions increases the computational costs and makes convergence harder to achieve during SA (Chen et al., 2020).

In recent years, the advent of deep learning has opened up unprecedented opportunities and insight into image reconstruction. Several studies have shown the successful application of these techniques for 2D (Guan, 2018) and 3D (Mosser et al., 2017, 2018; S. Liu et al., 2019). Notably, a growing body of literature has recently investigated 2D to 3D image reconstructions intending to infer 3D morphological and structural properties using features extracted from 2D images in specific orientations (Chung & Ye, 2021; Feng et al., 2020; Kench & Cooper, 2021; Volkhonskiy et al., 2019).

Despite showing promising results, training GANs stably and efficiently is a non-trivial task. In this work, we employ a GAN coupled to an Earth-mover or Wasserstein-loss with a gradient penalty (WGAN-GP) (Arjovsky et al., 2017; Gulrajani et al., 2017) to improve the accuracy of reconstructions and overcome the common training challenges associated with original GANs (see section 3.3.1). We apply both WGAN-GP and stochastic methods to reconstruct two-dimensional electron microscopy images taken from two of the most common fluid-rock interactions within the Earth’s lithosphere; (1) the hydrothermal alteration of feldspar in igneous rocks and (2) the hydration of upper mantle rocks to induce serpentinisation. While only two-point correlations are used in SA as a target function, polytope functions are calculated to evaluate the accuracy of both methods in reproducing higher-order structural information in the systems. To our knowledge, previous studies only used two-point correlation and a subset of the so-called Minkowski functionals (e.g., specific

173 surface area and Euler connectivity) to assess image reconstruction performance - none of
 174 these choices, however, have explicitly addressed how machine-learning-based reconstruc-
 175 tions perform in reproducing high-order spatial correlations. When determining how well
 176 we can reconstruct higher-order complexity, our results show that while a reconstructed
 177 microstructure can have the same two-point correlation as the original one, they can be
 178 morphologically different. Therefore, it is necessary to couple higher-order correlation func-
 179 tions with reconstruction methods. We also show that GANs can naturally reproduce the
 180 higher-order correlations at all ranges without being explicitly trained. This is important
 181 for reconstructing the higher-order geometry in complex heterogeneous systems and linking
 182 the microstructures to macroscopic phenomena.

183 2 Theoretical approach

184 2.1 Rock samples and dataset

185 We focus on two commonly occurring rock types affected by fluid-rock interaction.
 186 The first example is an altered igneous rock representative for fluid-rock interactions of
 187 the Earth's crust (?, ?) and the second example is a partially serpentinised peridotite
 188 representative for alteration within the Earth's uppermost mantle (Plümper et al., 2012).
 189 For simplicity we refer to these two rock samples as meta-igneous rock and serpentinite from
 190 here on.

191 For the meta-igneous rock, a small core was drilled with a diameter and height of 2.5
 192 mm and 1 cm, respectively. The core was cut, and the surface was imaged in backscattered
 193 electron (BSE) mode using the Zeiss Atlas software installed on a Zeiss Gemini 450 SEM.
 194 Zeiss Atlas allows large-area BSE imaging of up to several centimetres. Acquisition condi-
 195 tions were 20 kV acceleration voltage and 2 nA beam current. The pixel size was set to 50
 196 nm. Subsequently, a region of interest with the dimensions of 17920 by 54784 pixels (0.9
 197 mm by 2.7 mm) was imaged.

198 For the partially serpentinised peridotite, the Atlas software-based BSE imaging ap-
 199 proach was utilised on various samples from selected Norwegian peridotites previously de-
 200 scribed in Plümper et al. (2012) and Plümper et al. (2014). The rock samples are charac-
 201 terised by a lizardite-mesh texture with remaining olivine and secondary magnetite. The
 202 serpentinisation process produces a fracture network creating pathways for fluid flow. In
 203 this case, a small part of a thin section was scanned with an acceleration voltage of 15 kV,
 204 2 nA current, and pixel size of 500 nm, resulting in a large image of dimensions 15000 by
 205 30000 pixels (7.5 mm by 15 mm).

206 To remove the noise and artefacts that are often present in raw images, the acquired
 207 grayscale BSE images were first denoised by applying an edge-preserving denoising algo-
 208 rithm known as bilateral filtering. This filter smooths an image by averaging pixels based
 209 on their spatial distance and pixel value similarities. For more details, please see the original
 210 work by Tomasi and Manduchi (1998) and the scikit-image Python package's documenta-
 211 tion (Van der Walt et al., 2014). The filtered images were subsequently segmented into
 212 binary images where a pixel value of 1 corresponds to the phase of interest, which is the
 213 reaction-induced pore network in feldspar and the fracture network in serpentinite samples,
 214 respectively. Image segmentation employed a convolutional neural network (CNN) with
 215 U-Net architecture (Ronneberger et al., 2015). CNN is a supervised deep-learning method
 216 and thus, requires previously annotated images for training. These labelled images were
 217 created using the ilastik software; an interactive software developed for image classification
 218 and segmentation (Berg et al., 2019). It is worth mentioning that although one can do the
 219 image segmentation using ilastik alone, the benefit of training a CNN is that, once trained,
 220 it can be used to quickly segment future images either directly (if the sample and imaging
 221 conditions are the same) or using transfer-learning.

222

2.2 n -point correlation functions

223

224

225

226

227

Correlation functions have been proposed as an effective means to describe complex heterogeneous microstructures mathematically. The most widely used microstructural correlation function is the two-point correlation, $S_2(r)$, which is the probability P of two random points of distance r to occur in the same region of phase i , V_i , within a d -dimensional space R^d (Torquato & Haslach Jr, 2002):

$$S_2^{(i)}(r) = P(x \in V_i, x + r \in V_i) \text{ for } x \text{ and } V_i \in R^d \quad (1)$$

228

229

230

231

232

where x is an index showing the location of a pixel in the microstructure image. This is a radial form of two-point correlation calculated by averaging the functions in horizontal and vertical orientations for a statistically isotropic and homogeneous system. According to this definition, the probability of one random point ($r=0$) to occur in phase i is a one-point correlation corresponding to the volume fraction of that phase, i.e., $S_1^{(V_i)} = S_2^{(V_i)} = \phi_i$.

233

234

235

236

237

238

239

240

241

242

As mentioned earlier, two-point statistics is not sufficient to uniquely characterise complex systems containing higher-order spatial correlations in different locations of an image. In such cases, an n -point correlation function S_n can be defined as a probability of n random points to lie in a given phase. While a complete description of the medium can be achieved by a set of S_n with $n = 1, 2, 3, \dots, \infty$, calculating and storing probabilities of all possible n -points statistics is computationally intractable. Thus, to sample high-correlations from digitally-sampled microstructures, a compromise must be made between geometric completeness and algorithmic practicality. One such approach to capture complex microstructures uses n -point polytope functions, defined as a probability of n vertices of a random regular n -point polytope having a given edge length that occurs within the same phase (Fig.1).




Figure 1. Illustration of polytope functions in meta-igneous rock (a) and serpentinite (b) samples. P_{3H} , and P_{3V} indicate the probability of vertices of horizontal and vertical triangles, and P_4 and P_6 those of a square and hexagon, respectively.

243

244

245

246

247

The polytope functions capture partial higher-order n -point correlations of a specific phase, and can be directly computed from 2D or 3D images and incorporated in a stochastic optimisation method for image reconstruction (Chen et al., 2019). As shown in Fig. 1, for $n = 2$, the two-point polytope function is the same as S_2 . However, for $n \geq 3$, the higher-order polytope functions can be seen as a subset of the n -point correlation functions, which

248 can be efficiently computed as the edge length (r) is the only variable. A scaled version of
 249 these functions, known as scaled autocovariance function, has been introduced by Torquato
 250 and Haslach Jr (2002) and is related to the P_n functions via:

$$F_n(r) = \frac{P_n(r) - \phi^n}{\phi - \phi^n} \quad (2)$$

251 where ϕ is the phase fraction. According to this equation, $F_n(r = 0) = 1$ and $F_n(r \rightarrow$
 252 $\infty) = 0$. The latter is obtained as for $r \rightarrow \infty$, we have $P_n \approx \phi^n$. Although both F_n and
 253 P_n functions show the same behaviour for a given microstructure, it is more convenient to
 254 use the scaled autocovariance as it is normalised by the phase fraction - meaning that they
 255 describe geometric patterns independently from a given phase's volume fraction.

256 2.3 Stochastic reconstruction

In this study, we use the stochastic Yeong-Torquato-based reconstruction algorithm developed by Jiao and co-workers as a benchmark (Jiao et al., 2007, 2008). This approach formulates the image reconstruction as an optimisation problem in which reconstruction is performed by minimising a cost function using simulated annealing (SA) optimisation (Kirkpatrick et al., 1983). The algorithm starts with a two-point correlation S_2 of an original image as a target function and a random initial system configuration with the same volume fraction. An energy function is defined as the sum of the square errors between the target function and correlation function of the proposed configuration \hat{S}_2 :

$$E = \sum_r \left[\hat{S}_2 - S_2(r) \right]^2 \quad (3)$$

To evolve the random reconstruction towards the original image, the values of two random pixels associated with different phases (i.e., black and white pixels) are exchanged, ensuring that the volume fraction of both phases are preserved. Hence, a new energy E_{new} corresponding to the new configuration, and the energy difference between two successive configurations $\Delta E = E_{\text{new}} - E_{\text{old}}$ are calculated. Finally, the pixel exchange is accepted according to the Metropolis acceptance rule:

$$p(\Delta E) = \begin{cases} 1, & \text{if } E_{\text{new}} < E_{\text{old}} \\ \exp(-\Delta E/T), & \text{if } E_{\text{new}} \geq E_{\text{old}} \end{cases} \quad (4)$$

257 where $p(\Delta E)$ is the acceptance probability of the pixel exchange, T is an imaginary tem-
 258 perature that is initially set to a high value and decreases by a factor of α (selected to
 259 be less than but close to 1) after each annealing stage of the algorithm i.e., $T = \alpha \times T_0$.
 260 Accordingly, when the temperature is high at the initial steps, the acceptance probability
 261 of the pixel exchange can be higher even if $E_{\text{new}} \geq E_{\text{old}}$, i.e., the error between target and
 262 sampled two-point correlation in the new configuration is higher than the old one. Thus,
 263 the probability of accepting a bad configuration is higher. This helps to explore the whole
 264 solution space and prevents the algorithm from trapping in local minima. These steps are
 265 repeated until the energy (error) of the reconstructed image is less than a predefined thresh-
 266 old value or a maximum number of iterations is reached. In this study, 0.0000001, 0.97,
 267 and 250 were selected for the initial temperature (T_0), the decreasing factor (α), and the
 268 number of iterations, respectively.

269 Although the Yeong-Torquato SA is a well-known, flexible optimisation method that
 270 allows the incorporation of correlation functions to improve the reconstruction process, the
 271 computational costs significantly increase by including additional functions or dimensions -
 272 because they must be sampled at any iteration. And, similar to most stochastic optimisation
 273 methods, inference (reconstructions) accuracy is achieved after relatively large numbers of
 274 iterations. Hence, we only use the two-point correlation function as a benchmark against
 275 our GAN models.

2.4 Generative adversarial networks(GANs)

GANs are unsupervised generative algorithms based on game theory (Goodfellow et al., 2014), which can directly learn complex high-dimensional probability distributions from the input data. The term *adversarial* originates from the fact that GANs are composed of two neural networks competing against each other: a generator (G) and a discriminator (D). The generator’s task is to generate realistic images from the data distribution p_{data} . This is done by transforming noise (i.e, random) vectors z into $x = G_{\theta}(z)$, with θ being a set of learnable parameters. These noise vectors, also known as latent space, are random variables usually sampled from a normal distribution. During training, the discriminator D_{θ} , which is a binary classifier, receives both real images (from $p_{data}(x)$) and reconstructed images (from $p_{model}(x)$), and then learns to maximise the probability of correctly labelling reconstructed (with label= 0) and real images (with label= 1). At the same time, the generator is trained in such a way that a chosen discriminator metric e.g., $\log(1 - D(G(z)))$ is minimised (i.e., to ‘fool’ the discriminator into classifying the reconstructed image as real with $D(G(z)) = 1$). Mathematically, the cost function for a GAN is a minimax game with value function $V(G, D)$ (Goodfellow et al., 2014):

$$\underbrace{\min}_G \underbrace{\max}_D V(D, G) = \mathbb{E}_{x \sim p(data)}[\log D(x)] + \mathbb{E}_{Z \sim p(Z)}[\log(1 - D(G(Z)))] \quad (5)$$

2.4.1 Challenges in training GANs

Although GANs have been successfully applied to reconstruct a wide range of images, training GANs stably and efficiently is non-trivial. The training involves achieving a Nash equilibrium to a non-cooperative game between the generator and the discriminator, each of them having its cost function: $J^{(D)}(\theta^{(D)}, \theta^{(G)})$ and $J^{(G)}(\theta^{(D)}, \theta^{(G)})$ for the discriminator D and the generator G, respectively. A Nash equilibrium is reached when a combination of parameters $(\theta^{(D)}, \theta^{(G)})$ is found so that $J^{(D)}$ is minimum with respect to $\theta^{(D)}$, and $J^{(G)}$ is minimum with respect to $\theta^{(G)}$. Finding these parameters to reach Nash equilibrium is difficult because a change in $\theta^{(D)}$ to reduce $J^{(D)}$ may increase $J^{(G)}$, and similarly, a modification to $\theta^{(G)}$ for minimising $J^{(G)}$ can increase $J^{(D)}$. Therefore, although the two players might reach an equilibrium in some cases, updating the parameters of both models does not necessarily lead to stable and convergent training. However, this is not a specific issue of GANs, but it is a general problem with game-theory-based approaches.

Mode collapse is another common issue in GANs, which occurs when the generator collapses to a set of parameters θ that leads to reconstructing the same images, i.e., mapping different noise vectors into the same output (Goodfellow, 2016). The reason is that the discriminator receives and analyses each image independently. Therefore, when it cannot differentiate between real and reconstruction for a specific example, the generator updates its parameters to create more of that example and ‘wins’ the game.

Several heuristic methods are proposed to overcome these challenges and improve training stability. Some effective techniques are: feature matching, minibatch discrimination, historical averaging, one-sided label smoothing, and visual batch normalisations (Salimans et al., 2016). Adding Gaussian noise to the discriminator’s input and label switching can also stabilise the training process (Mosser et al., 2018), though the greater statistical conditions and implications of this approach are still to be researched. Furthermore, some studies have investigated the use of other distance metrics and value functions than binary cross-entropy (BCE) (Eq. 5). Arjovsky et al. (2017) used Earth-mover or Wasserstein-1 distance to measure the distance between the probability functions of real and reconstructed images. This method, known as WGAN, can prevent mode collapse and improve the stability and convergence behaviour by forcing the gradient of the discriminator in a constrained space. This can be done by applying a weight clipping or a gradient penalty, with the latter method known as WGAN-GP (Gulrajani et al., 2017). In this paper, we use WGAN-GP to enforce

the output of discriminator in $[-1,1]$:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p(\text{data})}[D(x)] + \mathbb{E}_{Z \sim p(Z)}[D(G(Z))] + \lambda \mathbb{E}_{\hat{x}}[(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2] \quad (6)$$

311 where λ is a coefficient and here is set to 10, and \hat{x} is a mixture of real and the reconstructed
 312 image from the generator network calculated via $\hat{x} = \epsilon(x) + (1 - \epsilon(x))G(z)$, in which ϵ is a
 313 random number from a uniform distribution $U[0, 1]$.

314 2.4.2 Training workflow

315 While different architectures can be employed in GANs, several studies suggest that
 316 using convolutional networks in generator and discriminator can improve the fidelity of
 317 synthetic images and training performance. Such a network, known as deep-convolutional
 318 GAN (DCGAN), was first introduced by Radford et al. (2015). Our proposed workflow
 319 for image reconstruction is shown in Fig. 2a. We extracted smaller images from the two
 320 samples to prepare sufficient training images. In the meta-igneous sample, a total number
 321 of 14,697 images were created by sliding a window of the size 512 pixel^2 with a stride of
 322 256 pixels over the original large BSE image. For the serpentinite sample, sliding window
 323 size and stride were 1024 pixels and 200, respectively, providing 10150 training images.
 324 These images were then segmented and resized to 128^2 . The smaller window size is selected
 325 for the meta-igneous sample because we observed a significant loss of information while
 326 downsampling larger images to 128^2 which was our target size. More details about the
 327 architecture and hyperparameters used for training out WGAN-GP can be found in table
 328 S1 and table S2, respectively.




Figure 2. Training WGAN-GP with BSE images of the meta-igneous rock. (a) Schematic description of workflow. G and D are the generator and the discriminator with architectures described in table A1.(b) Training dynamics showing how two losses are converging around zero with generator iterations. In practice, the parameters of the discriminator were updated five times for each generator’s update. The same workflow was used for the reconstruction of the serpentinite.

329 Fig. 2b. displays the evolution of loss functions versus generator iterations. It can be
 330 seen that after 250 iterations, generator and discriminator losses start converging stably.
 331 However, since a decrease in generator loss cannot be related to reconstruction quality, an
 332 average S_2 curve was calculated for real and simulated batch images (batch size = 128) at
 333 the end of each iteration. Subsequently, the MSE between two average curves was used as
 334 the reconstruction quality criterion. The best results were obtained at iteration 4300 with
 335 $MSE = 1.66e^{-8}$. The training was performed on two NVIDIA Quadro P6000 GPUs, and the
 336 curves converged after 3 hours.

3 Results

To compare the capability of stochastic and GAN reconstruction, two WGAN-GPs are trained on BSE images of both samples using the workflow shown in Fig. 2a. The reconstruction quality is then evaluated in terms of the error between S_2 and polytope functions of the original and reconstructed images. In all cases, the average functions of 128 images are calculated and compared.

3.1 Characterisation of the original microstructures

In this section, we present the P_n polytope functions corresponding to different geometrical correlations. These functions provide new statistical information about the microstructures. For example, the P_4 function shows the square correlations (patterns) in the sample and adds unique higher-order information to triangular correlations (P_{3H} and P_{3V}), which in turn, are complementary to P_2 . Fig. 3a-b show the average P_n functions of the original microstructures in the two samples. In each case, all functions start with the same probability at $r = 0$, corresponding to area fractions of phases of interest, which are 0.052 and 0.352 for the meta-igneous rock and serpentinite samples, respectively. It can be seen that the correlation functions initially decrease as r increases. However, the reduction rate of each correlation function is faster than its lower-order function. This is because, for example, it is less probable that all vertices of a hexagon lie in the same phase compared to those of a triangle with the same edge length. Furthermore, the r value in which the curves stabilise shows the average size of the features of interest in the samples. Thus, one can infer that the average pore size is <10 pixels ($\sim 0.5 \mu m$) in meta-igneous microstructures (Fig. 3a), and similarly, the average serpentinite fracture width is 20 pixels ($\sim 10 \mu m$). The same trends can be observed in the scaled correlation functions F_n for each sample (Fig. 3c-d).

3.2 Characterisation of the reconstructed microstructures

To compare and evaluate the accuracy of microstructure reconstruction, polytope functions are computed on 128 reconstructed images by the SA and our WGAN-GP model. Fig. 4 shows the quantification of microstructures reconstructed via the SA method (green curves). As can be seen, while there is a good agreement between S_2 curves of the original and reconstructed images, apparent discrepancies are observed between higher-order correlations functions (i.e., for $n > 2$). A small error between S_2 functions was expected as this function is used as the target function in the SA algorithm. Fig. 5 compares the polytope functions calculated from the GAN-reconstructed images with the original images. The close agreements between all polytope functions indicate that reconstructed images via GAN contain higher-order structural information present in the original microstructures.

To quantify and compare the reconstruction accuracy of the two methods, the MSE between F_n functions derived from original and reconstructed images are calculated and presented in table S3. The table highlights that the two reconstruction methods are comparable in capturing S_2 . However, with respect to higher-order correlations, the errors computed from WGAN-GP are between two (for meta-igneous rock) to three (for serpentinite) orders of magnitude less than the stochastic method.

4 Discussion

One of the main aims of this paper was to show that GAN can reconstruct images that are visually realistic and honour different orders of statistical correlations (symmetry) in the original images. Our results show that capturing these geometrical correlations is an inherent capability of GANs. The term *inherent* here refers to the fact that the GAN has not been trained to fit higher-order polytope functions. The information used by GAN during the training were (1) a great number (> 10000 images) of training images, and (2) the




Figure 3. Quantification of microstructures observed in 2D images of the meta-igneous rock (a and c) and the serpentinite (b and d). (c) and (d) are the scaled autocovariance functions calculated by Eq.2.




Figure 4. Comparison of polytope functions calculated from original microstructures (blue) and reconstructed images by the SA method (green).




Figure 5. Comparison of polytope functions calculated from original microstructures (blue) and reconstructed images by our WGAN-GP model (red). The good match between polytope functions shows that our model can accurately reproduce the complex features in the images.

385 MSE error between S_2 of the original and reconstructed images as a stopping criterion. The
 386 above-mentioned capability of GANs may be explained by the fact that there are multiple
 387 layers of convolutions in the discriminator, each of them encodes spatial information from
 388 the images in the form of feature maps. In order for this nonlinear set of parameters to
 389 represent our samples geometry, the training data and chosen loss functions are the user-
 390 supplied priors - which the GAN relies upon during training to achieve network weights
 391 that lead to accurately-inferred reconstructions. In that process, the generator learns to
 392 simulate images containing those spatial correlations to 'fool' the discriminator - informed
 393 by the large prior number of training images extracted from our imaged samples.

394 4.1 Microstructure reconstruction: GAN vs. SA

395 Fig. 6 depicts the image realisations obtained from SA and our GAN and highlights
 396 the importance of capturing higher-order correlation functions for generating realistic im-
 397 ages. Visual inspection of the reconstructions shows that the GAN model can generate more
 398 realistic images with similar geometrical and structural features such as shape, size, and ori-
 399 entation. This is consistent with quantitative analysis of polytope functions (Fig. 5), show-
 400 ing that different levels of morphological symmetries are reproduced in GAN-reconstructed
 401 images. However, to further quantify and compare morphological information of reconstruc-
 402 tions, we use a lineal-path function L (Lu & Torquato, 1992). This function is a statistical
 403 morphological descriptor defined as the probability that an entire line of the length r
 404 occurs in the same phase. Thus, it can provide additional insight into microstructure connectivity
 405 and linear clustering.

406 In the case of the meta-igneous rock, we observe that the pores reconstructed by SA
 407 (Fig. 6c) are circular and smaller than the original pores, lacking the preferred orientation
 408 apparent in the original microstructure. This observation is also confirmed by determining
 409 the average L computed from 128 images, as presented in Fig. 7. Although there is a close
 410 agreement between the L of original and reconstructed images (Fig. 7a), zooming into the
 411 flat part of the curves (Fig. 7c) reveals that SA has underestimated the L (i.e., the linear
 412 connectivity of pores) in the system. In particular, instead of the elongated pores in the
 413 original images, round-shaped isolated pores of smaller diameters are reconstructed. Fig.
 414 8 compares the probability distributions of the area and orientation of the major axis of
 415 pores in the real and reconstructed images of the meta-igneous rock sample, computed by
 416 the label analysis tool within the Thermo Fisher Scientific AVIZO software. From Fig. 7,
 417 we can conclude that GAN reproduces pores of similar size and orientation to the natural
 418 system, whereas there is a significant difference between SA-reconstructions and the natural
 419 microstructure.

420 Moreover, it is also apparent from Fig. 6 that the SA-reconstructed images are entirely
 421 different from original ones in serpentinite fracture networks. Specifically, circular pores of
 422 large diameters are reconstructed instead of reproducing a network of fractures with varying
 423 widths and orientations (Fig. 6d). This observation is in agreement with our earlier results,
 424 which show that SA overestimates all higher-order polytope functions (Fig. 4b) and lineal-
 425 paths (Fig. 7d) with errors up to three-order of magnitudes greater than GAN, leading to
 426 reconstructed images that have no resemblance to the original microstructure.

427 Unrealistic reconstructions obtained by the SA are due to the microstructural degeneracy
 428 of the S_2 function, i.e., a number of different microstructures can be compatible with
 429 a given target function S_2 (Eq. 3), resulting in a near-zero energy (Gommes et al., 2012a).
 430 This means that the S_2 does not contain sufficient information to characterise the system
 431 uniquely - particularly with regards to higher-order structures. However, it can be seen
 432 that SA performs much better in the meta-igneous rock sample than in the serpentinite
 433 sample. This may be explained because the serpentinite microstructure is of higher geo-
 434 metrical complexity, i.e., it contains higher-order correlations at different length scales (Fig.
 435 3b). In contrast, there are not many higher-order correlations in the porous meta-igneous
 436 rock system at longer ranges (Fig. 3a), in which polytope correlations become zero at short
 437 ranges.

438 In addition to higher accuracy, the computational reconstruction time of the GAN is
 439 much less than for the SA method. This time difference is expected given the nature of
 440 stochastic optimisation versus the machine-learning approach. While stochastic optimisa-
 441 tion yields a single reconstruction by numerically sampling a model posterior distribution
 442 iteratively (conducted independently for each reconstruction), GAN directly yields samples
 443 from the model posterior once trained. In the case of GAN, the training data, the choice
 444 of loss functions and architecture parameters, and the upfront time spent in training yield
 445 an implicitly-built prior estimator in the form of trained network weights - allowing for
 446 near-instantaneous inference/reconstruction after training. This process involves finding
 447 the optimal hyperparameters for both generator and discriminator, which requires high-
 448 performance computational resources such as modern GPUs. Although training a GAN can
 449 be challenging, one can save the trained generator and reuse it to generate an unlimited
 450 number of synthetic microstructures swiftly. This is an essential advantage of GANs over
 451 the classical stochastic methods. For example, the time required for training our GAN on
 452 1282 feldspar images was about 3 hours using two 24GB NVIDIA Quadro P6000 GPUs.
 453 However, after training, it only needed 4 seconds to reconstruct 128 images, whereas recon-
 454 struction time of the same number of images via SA was about 2 hours on a system with
 455 24 CPUs (Intel Xeon Gold 6136, 3GHz), showing an acceleration of 1800 times.




Figure 6. Visual comparison of real (a-b) with reconstructed images using SA (c-d) and GAN (e-f). Red solid lines illustrate how that the lineal-path function L can describe the linear clusters and connectivity in the microstructures.




Figure 7. The average lineal-path functions L calculated from real (blue), SA-reconstructed (green), and GAN-reconstructed (red) images. Grey areas in (c-d) are %95 confidence bounds around real curves. A better match between curves is derived via GAN reconstructions in both samples. CL = confidence level.

4.2 Reconstruction of representative microstructures

The main question that needs to be addressed in microstructure reconstruction is determining representative image size capturing structural elements of the system under consideration. It has been shown that the models trained on small images will create pores with artefacts and unrealistic shapes. Although, the larger the size of the training images, the more computationally demanding and less stable the training.

Mosser et al. (2017) proposed to use average grain size and chord length as the minimum training image size. However, a representative elementary size (RES) analysis should be carried out for heterogeneous and complex samples to find an adequate training image size (Volkhonskiy et al., 2019). RES analysis is a methodology to determine the smallest size of a system that is large enough to capture the system’s heterogeneity as a whole (Bargmann et al., 2018). RES analysis is conventionally performed for a particular rock property such as porosity or permeability and is used in upscaling to evaluate the effective macro-scale properties of rocks from a smaller yet representative sample size. Thus, the RES determined by this method can significantly vary depending on the property of interest. Furthermore, this approach involves plotting sample size versus its corresponding calculated property. A common observation is that the property fluctuates widely at small sizes, but it becomes insensitive to size at some point which can be considered the representative size, i.e., the transition between micro- and macro-scale (Al-Raoush & Papadopoulos, 2010).

Here, we rely instead on the two-point correlation function itself to determine the representative image size - this allows for a material-dependent representative image size.




Figure 8. Comparison of real and reconstructed images in terms of pore properties in the meta-igneous rock. (a) Probability density of pore area calculated from 128 randomly sampled images. Dashed lines show the estimated probability density functions fitting the data. (b) Distribution of pore orientations between -90 (clockwise) and $+90$.

477 Such a representative size is characteristic for porosity (included as volume fraction in S_2),
 478 but it is also structurally and topologically representative, which is important for many post-
 479 reconstruction analyses such as fluid flow simulations. Our approach consists of computing
 480 the average scaled two-point correlation (F_2) for ten images of different sizes randomly
 481 selected from our original large BSE images of both samples (Fig. 9a-b). Then, MSEs
 482 between the largest image (i.e., of size 5000^2 pixels) and smaller images at the overlapping
 483 range is calculated (Fig. 9e-f). It can be seen that the F_2 curves for the images smaller than
 484 2000 pixels show entirely different patterns (Fig. 9c-d), leading to more errors while the MSE
 485 does not decrease significantly beyond 2000 pixels. Thus, an image of size 2000^2 pixels can
 486 be considered representative for both samples - based on the stability of their corresponding
 487 two-point correlations. However, due to the GPUs memory limitations, training images of
 488 2048^2 were first downsampled to 512^2 . We cannot downsample images to any arbitrary sizes
 489 due to the loss of information and aliasing artefacts associated with downsampling. Here,
 490 images were resized using scikit-image package in Python (Van der Walt et al., 2014), and
 491 the image quality was visually inspected.

492 Fig. 10 shows reconstructions of microstructures using GAN trained by representative
 493 images. The results indicate that GAN can reproduce synthetic microstructural images
 494 matching polytope functions of real microstructures. However, SA algorithms did not con-
 495 verge with such a large system because the microstructural degeneracy of the system expo-
 496 nentially increases with the number of pixels. This leads to a rough and complex energy
 497 landscape (i.e., model posterior or solution space) associated with SA optimisation, thereby
 498 exploring the solution space to find the minimum global energy becomes numerically chal-
 499 lenging (Gommes et al., 2012a, 2012b).

500 5 Conclusions and outlook

501 We investigated the use of a Wasserstein-loss GAN with gradient penalty (WGAN-
 502 GP) to reconstruct two-dimensional microstructures of natural rocks. We evaluate its per-
 503 formance in retrieving highly-complex geometries accurately by quantifying higher-order
 504 statistical correlation functions. As inputs, we used electron microscopic images of two het-
 505 erogeneous systems. The first sample is from an altered igneous rock with a mostly isolated
 506 but oriented pore network generated during fluid-driven mineral replacement. The second
 507 sample is from a serpentinised peridotite characterised by a complex and connected fracture
 508 network containing different geometrical patterns. Our results show that GANs are capable
 509 of capturing and reconstructing these topologically complex microstructures without prior




Figure 9. RES analysis by scaled autocovariance function F_n calculated for images of different sizes. (a) and (b) are F_n functions computed for the meta-igneous rock and serpentinite samples, respectively, where (c) and (d) show the magnified views of (a) and (b). Grey shadow indicates the %95 confidence levels around the average values of the largest image (5000^2 pixels). (e-f) MSEs calculated between the average F_n of the largest image and smaller ones.




Figure 10. Characterisation and reconstruction of representative microstructures using our WGAN-GP. (a) and (b) show microstructural characterisation of samples. (c-d) are real microstructures, and (e-f) are GAN-reconstructed images of feldspar pores and serpentinite fracture network, respectively.

510 statistical information - provided one takes steps to ensure training stability and to perform
 511 quality-control after inference with the use of high-order descriptors

512 An additional aim was to compare the performance of GAN with a conventional stochas-
 513 tic reconstruction method that uses two-point correlations as input and SA as an optimi-
 514 sation algorithm. To evaluate and compare the accuracy, n -point polytope correlations are
 515 employed for the original and the reconstructed microstructures, and the MSE used as a
 516 quality metric. Our findings demonstrate that errors associated with WGAN-GP are two
 517 to three orders of magnitude smaller than the stochastic method. Especially in the case
 518 of the serpentinite sample, which displays a high degree of geometrical complexity in its
 519 microstructure, SA fails to reconstruct realistic images. Thus, as a practical implication,
 520 we suggest that before a reconstruction method is selected, the level of complexity (i.e.,
 521 statistical degeneracy) of microstructures under consideration should be evaluated by e.g.,
 522 quantifying higher-order geometrical correlations.

523 We also propose a new methodology for determining the representative image size based
 524 on the S_2 function. This method is more comprehensive than the previous approaches (see
 525 section 4.2) - in terms of how it adapts to the properties of samples in question - because it
 526 considers the structural and morphological information captured by S_2 . In our case, having
 527 determined the representative image size to be larger than 2000^2 pixels, we downsampled the
 528 images of 2048^2 pixels to 512^2 pixels and successfully trained a GAN to reconstruct realistic
 529 microstructures honouring the original n -point polytope statistics. However, we could not
 530 achieve convergence in the SA with images $> 512^2$ pixels by trying different combinations
 531 of parameters.

532 The success of applying GANs for image reconstruction in recent years has resulted in
 533 increased interest and the emergence of new variants of GAN with new capabilities that can
 534 be easily incorporated due to the flexibility of GANs. Examples are Slice-GAN (Kench &
 535 Cooper, 2021; Chung & Ye, 2021), BicycleGAN (Feng et al., 2020), and slice-to-pore GAN
 536 (Volkhonskiy et al., 2019), which all have been developed for 2D to 3D image reconstruc-
 537 tion where 2D images of orthogonal planes are used to generate synthetic but statistically
 538 equivalent 3D microstructures. This type of GANs receives increased attention because 2D
 539 images are easier and more affordable to acquire and usually have higher resolution and
 540 a larger field-of-view (FoV). Another example are progressively growing GANs (PG-GAN)
 541 introduced by Karras et al. (2017), which has been used to reconstruct large images up to
 542 1024^2 pixels from carbonate rocks (You et al., 2021).

543 Despite high-quality image generation, much less is known about the properties of the
 544 latent space (z-space), for example, how image attributes are formed and organised in the
 545 latent space of a well-trained GAN, and the correlation between these attributes. Also,
 546 there is still uncertainty about how GANs can link the latent space to image semantic space
 547 and how the latent space can be interpreted and used for image manipulation (Shen et al.,
 548 2020). The reason is that the generator in GANs is not trained to be invertible, i.e., a
 549 two-way mapping between the dataset (image space) and the latent space is not established
 550 during adversarial training. Instead, GANs learn to produce high-quality synthetic images
 551 indirectly by optimising generator’s weights to imitate the original dataset according to
 552 feedback from discriminator. Learning such mapping between latent and image space allows
 553 us to explore the latent space and to manipulate realistically, edit, and combine the features
 554 in the generated image. Several studies have attempted to address this challenge by using
 555 inverse mapping, i.e., from image space to latent space. More information about such
 556 methods, known as GAN inversion, can be found in the comprehensive survey by Xia et al.
 557 (2021). Other approaches focus on coupling GANs with other generative models, such as
 558 variational autoencoder (Larsen et al., 2016).

559 Employing such methods can be a fruitful area for future works to investigate how
 560 higher-order information (such as polytope functions) are encoded by GAN in latent space
 561 and intermediate semantic space, and how we establish a mapping between these informa-

tion and macroscopic properties of rock. Another potential application of such methods can be image manipulation which can be particularly useful in reaction-induced fluid flow simulations in which a connected pore network exists at the time of reaction. However, most pores are isolated after the reaction ceases (Putnis, 2015), as observed in the meta-igneous rock system. In such a situation, GAN inversion methods may realistically reconnect the isolated pores and simulate the fluid flow at the time of reaction.

6 Data availability statement

Original and segmented BSE images and the data to reproduce the figures are available at Utrecht University Yoda data repository accessible via: <https://public.yoda.uu.nl/geo/UU01/ACSDR4.html>. Python codes are also accessible via <https://github.com/hamediu/GeoWGAN-GP>

Acknowledgments

This research as well as Hamed Amiri and Oliver Plümper were supported by a European Research Council (ERC) starting grant "nanoEARTH" (852069). We also acknowledge the Utrecht University Electron Microscopy Facility and the EPOS-NL MINT facility.

References

- Al-Raoush, R., & Papadopoulos, A. (2010). Representative elementary volume analysis of porous media using x-ray computed tomography. *Powder technology*, *200*(1-2), 69–77.
- Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning* (pp. 214–223).
- Bargmann, S., Klusemann, B., Markmann, J., Schnabel, J. E., Schneider, K., Soyarslan, C., & Wilmers, J. (2018). Generation of 3d representative volume elements for heterogeneous materials: A review. *Progress in Materials Science*, *96*, 322–384.
- Berg, S., Kutra, D., Kroeger, T., Straehle, C. N., Kausler, B. X., Haubold, C., ... others (2019). Ilastik: interactive machine learning for (bio) image analysis. *Nature Methods*, *16*(12), 1226–1232.
- Bostanabad, R., Zhang, Y., Li, X., Kearney, T., Brinson, L. C., Apley, D. W., ... Chen, W. (2018). Computational microstructure characterization and reconstruction: Review of the state-of-the-art techniques. *Progress in Materials Science*, *95*, 1–41.
- Bourg, I. C., Beckingham, L. E., & DePaolo, D. J. (2015). The nanoscale basis of co2 trapping for geologic storage. *Environmental science & technology*, *49*(17), 10265–10284.
- Britton, T., Jiang, J., Guo, Y., Vilalta-Clemente, A., Wallis, D., Hansen, L., ... Wilkinson, A. (2016). Tutorial: Crystal orientations and ebsd—or which way is up? *Materials Characterization*, *117*, 113–126.
- Chen, P.-E., Xu, W., Chawla, N., Ren, Y., & Jiao, Y. (2019). Hierarchical n-point polytope functions for quantitative representation of complex heterogeneous materials and microstructural evolution. *Acta Materialia*, *179*, 317–327.
- Chen, P.-E., Xu, W., Ren, Y., & Jiao, Y. (2020). Probing information content of hierarchical n-point polytope functions for quantifying and reconstructing disordered systems. *Physical Review E*, *102*(1), 013305.
- Chung, H., & Ye, J. C. (2021). Reusability report: Feature disentanglement in generating a three-dimensional structure from a two-dimensional slice with sliceGAN. *Nature Machine Intelligence*, *3*(10), 861–863.
- Dehghan Khalili, A., Arns, J.-Y., Hussain, F., Cinar, Y., Pinczewski, W., Arns, C. H., et al. (2013). Permeability upscaling for carbonates from the pore scale by use of multiscale x-ray-ct images. *SPE Reservoir Evaluation & Engineering*, *16*(04), 353–368.
- Feng, J., Teng, Q., Li, B., He, X., Chen, H., & Li, Y. (2020). An end-to-end three-dimensional reconstruction framework of porous media from a single two-dimensional

- 611 image based on deep learning. *Computer Methods in Applied Mechanics and Engi-*
612 *neering*, 368, 113043.
- 613 Gommes, C. J., Jiao, Y., & Torquato, S. (2012a). Density of states for a specified correlation
614 function and the energy landscape. *Physical review letters*, 108(8), 080601.
- 615 Gommes, C. J., Jiao, Y., & Torquato, S. (2012b). Microstructural degeneracy associated
616 with a two-point correlation function and its information content. *Physical Review E*,
617 85(5), 051140.
- 618 Goodfellow, I. (2016). Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint*
619 *arXiv:1701.00160*.
- 620 Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ...
621 Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information*
622 *processing systems*, 27.
- 623 Guan, K. (2018). Reconstructing pore networks using generative adversarial networks.
- 624 Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. (2017). Improved
625 training of wasserstein gans. *Advances in neural information processing systems*, 30.
- 626 Hajizadeh, A., Safekordi, A., & Farhadpour, F. A. (2011). A multiple-point statistics algo-
627 rithm for 3d pore space reconstruction from 2d images. *Advances in water Resources*,
628 34(10), 1256–1267.
- 629 Holzer, L., & Cantoni, M. (2012). Review of fib-tomography. *Nanofabrication using focused*
630 *ion and electron beams: Principles and applications*, 559201222, 410–435.
- 631 Jiao, Y., Stillinger, F., & Torquato, S. (2007). Modeling heterogeneous materials via two-
632 point correlation functions: Basic principles. *Physical review E*, 76(3), 031110.
- 633 Jiao, Y., Stillinger, F., & Torquato, S. (2008). Modeling heterogeneous materials via two-
634 point correlation functions. ii. algorithmic details and applications. *Physical Review*
635 *E*, 77(3), 031135.
- 636 Jiao, Y., Stillinger, F., & Torquato, S. (2010). Geometrical ambiguity of pair statistics:
637 Point configurations. *Physical Review E*, 81(1), 011105.
- 638 Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of gans for
639 improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- 640 Karsanina, M. V., & Gerke, K. M. (2018). Hierarchical optimization: Fast and robust mul-
641 tiscala stochastic reconstructions with rescaled correlation functions. *Physical review*
642 *letters*, 121(26), 265501.
- 643 Kench, S., & Cooper, S. J. (2021). Generating three-dimensional structures from a two-
644 dimensional slice with generative adversarial network-based dimensionality expansion.
645 *Nature Machine Intelligence*, 3(4), 299–305.
- 646 Kirkpatrick, S., Gelatt Jr, C. D., & Vecchi, M. P. (1983). Optimization by simulated
647 annealing. *science*, 220(4598), 671–680.
- 648 Larsen, A. B. L., Sønderby, S. K., Larochelle, H., & Winther, O. (2016). Autoencoding
649 beyond pixels using a learned similarity metric. In *International conference on machine*
650 *learning* (pp. 1558–1566).
- 651 Liu, King, H. E., Van Huis, M. A., Drury, M. R., & Plümper, O. (2016). Nano-tomography
652 of porous geological materials using focused ion beam-scanning electron microscopy.
653 *Minerals*, 6(4), 104.
- 654 Liu, S., Zhong, Z., Takhiri-Borujeni, A., Kazemi, M., Fu, Q., & Yang, Y. (2019). A case
655 study on homogeneous and heterogeneous reservoir porous media reconstruction by
656 using generative adversarial networks. *Energy Procedia*, 158, 6164–6169.
- 657 Lu, B., & Torquato, S. (1992). Lineal-path function for random heterogeneous materials.
658 *Physical Review A*, 45(2), 922.
- 659 Mosser, L., Dubrule, O., & Blunt, M. J. (2017). Reconstruction of three-dimensional porous
660 media using generative adversarial neural networks. *Physical Review E*, 96(4), 043309.
- 661 Mosser, L., Dubrule, O., & Blunt, M. J. (2018). Stochastic reconstruction of an oolitic
662 limestone by generative adversarial networks. *Transport in Porous Media*, 125(1),
663 81–103.
- 664 Niu, Y., Wang, Y. D., Mostaghimi, P., Swietojanski, P., & Armstrong, R. T. (2020). An
665 innovative application of generative adversarial networks for physically accurate rock

- 666 images with an unprecedented field of view. *Geophysical Research Letters*, *47*(23),
667 e2020GL089029.
- 668 Okabe, H., & Blunt, M. J. (2005). Pore space reconstruction using multiple-point statistics.
669 *Journal of petroleum science and engineering*, *46*(1-2), 121–137.
- 670 Plümper, O., Botan, A., Los, C., Liu, Y., Malthe-Sørenssen, A., & Jamtveit, B. (2017).
671 Fluid-driven metamorphism of the continental crust governed by nanoscale fluid flow.
672 *Nature Geoscience*, *10*(9), 685–690.
- 673 Plümper, O., John, T., Podladchikov, Y. Y., Vrijmoed, J. C., & Scambelluri, M. (2017).
674 Fluid escape from subduction zones controlled by channel-forming reactive porosity.
675 *Nature Geoscience*, *10*(2), 150–156.
- 676 Plümper, O., Røyne, A., Magrasó, A., & Jamtveit, B. (2012). The interface-scale mechanism
677 of reaction-induced fracturing during serpentinization. *Geology*, *40*(12), 1103–1106.
- 678 Plümper, O., Beinlich, A., Bach, W., Janots, E., & Austrheim, H. (2014). Garnets within
679 geode-like serpentinite veins: Implications for element transport, hydrogen production
680 and life-supporting environment formation. *Geochimica et Cosmochimica Acta*, *141*,
681 454–471.
- 682 Prior, D. J., Mariani, E., & Wheeler, J. (2009). Ebsd in the earth sciences: applications,
683 common practice, and challenges. In *Electron backscatter diffraction in materials*
684 *science* (pp. 345–360). Springer.
- 685 Putnis, A. (2015). Transient porosity resulting from fluid–mineral interaction and its con-
686 sequences. *Reviews in Mineralogy and Geochemistry*, *80*(1), 1–23.
- 687 Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with
688 deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- 689 Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomed-
690 ical image segmentation. In *International conference on medical image computing and*
691 *computer-assisted intervention* (pp. 234–241).
- 692 Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016).
693 Improved techniques for training gans. *Advances in neural information processing*
694 *systems*, *29*.
- 695 Sheehan, N., & Torquato, S. (2001). Generating microstructures with specified correlation
696 functions. *Journal of Applied Physics*, *89*(1), 53–60.
- 697 Shen, Y., Gu, J., Tang, X., & Zhou, B. (2020). Interpreting the latent space of gans for
698 semantic face editing. In *Proceedings of the ieee/cvf conference on computer vision*
699 *and pattern recognition* (pp. 9243–9252).
- 700 Strebelle, S. (2002). Conditional simulation of complex geological structures using multiple-
701 point statistics. *Mathematical geology*, *34*(1), 1–21.
- 702 Tahmasebi, P., & Sahimi, M. (2012). Reconstruction of three-dimensional porous media
703 using a single thin section. *Physical Review E*, *85*(6), 066709.
- 704 Tahmasebi, P., & Sahimi, M. (2013). Cross-correlation function for accurate reconstruction
705 of heterogeneous media. *Physical review letters*, *110*(7), 078002.
- 706 Tomasi, C., & Manduchi, R. (1998). Bilateral filtering for gray and color images. In *Sixth*
707 *international conference on computer vision (ieee cat. no. 98ch36271)* (pp. 839–846).
- 708 Torquato, S., & Haslach Jr, H. (2002). Random heterogeneous materials: microstructure
709 and macroscopic properties. *Appl. Mech. Rev.*, *55*(4), B62–B63.
- 710 Van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager,
711 N., . . . Yu, T. (2014). scikit-image: image processing in python. *PeerJ*, *2*, e453.
- 712 Volkhonskiy, D., Muravleva, E., Sudakov, O., Orlov, D., Belozerov, B., Burnaev, E., &
713 Koroteev, D. (2019). Reconstruction of 3d porous media from 2d slices. *arXiv preprint*
714 *arXiv:1901.10233*.
- 715 Wallis, D., Hansen, L. N., Wilkinson, A. J., & Lebensohn, R. A. (2021). Dislocation interac-
716 tions in olivine control postseismic creep of the upper mantle. *Nature communications*,
717 *12*(1), 1–12.
- 718 Wu, Y., Tahmasebi, P., Lin, C., Ren, L., & Dong, C. (2019). Multiscale modeling of shale
719 samples based on low-and high-resolution images. *Marine and Petroleum Geology*,
720 *109*, 9–21.

- 721 Xia, W., Zhang, Y., Yang, Y., Xue, J.-H., Zhou, B., & Yang, M.-H. (2021). Gan inversion:
722 A survey. *arXiv preprint arXiv:2101.05278*.
- 723 Yeong, C., & Torquato, S. (1998). Reconstructing random media. *Physical review E*, 57(1),
724 495.
- 725 You, N., Li, Y. E., & Cheng, A. (2021). 3d carbonate digital rock reconstruction using
726 progressive growing gan. *Journal of Geophysical Research: Solid Earth*, 126(5),
727 e2021JB021687.