
GENERATING 3D STRUCTURES FROM A 2D SLICE WITH GAN-BASED DIMENSIONALITY EXPANSION

PREPRINT - FEBRUARY 16, 2021

Steve Kench ^{*1} and Samuel J. Cooper ^{†1}

¹*Dyson School of Design Engineering, Imperial College London, London SW7 2DB*

ABSTRACT

Generative adversarial networks (GANs) can be trained to generate 3D image data, which is useful for design optimisation. However, this conventionally requires 3D training data, which is challenging to obtain. 2D imaging techniques tend to be faster, higher resolution, better at phase identification and more widely available. Here, we introduce a generative adversarial network architecture, SliceGAN, which is able to synthesise high fidelity 3D datasets using a single representative 2D image. This is especially relevant for the task of material microstructure generation, as a cross-sectional micrograph can contain sufficient information to statistically reconstruct 3D samples. Our architecture implements the concept of uniform information density, which both ensures that generated volumes are equally high quality at all points in space, and that arbitrarily large volumes can be generated. SliceGAN has been successfully trained on a diverse set of materials, demonstrating the widespread applicability of this tool. The quality of generated micrographs is shown through a statistical comparison of synthetic and real datasets of a battery electrode in terms of key microstructural metrics. Finally, we find that the generation time for a 10^8 voxel volume is on the order of a few seconds, yielding a path for future studies into high-throughput microstructural optimisation.

1 Introduction

The properties and behaviour of a material depend strongly on its microstructure, which in turn is controlled by conditions of manufacture. Physical simulations play an important role in offering insight into these relationships to help inform the design of next generation materials. Importantly, many material behaviours, such as deformation under stress or fluid flow through a porous medium, are inherently volumetric and cannot be accurately modelled using 2D data alone. The fidelity of simulation techniques used to extract these properties is thus partly determined by the quality of available 3D microstructural datasets. However, as compared to their 3D equivalents, 2D micrographs are often easier to obtain, higher resolution and can contain a more representative distribution of features based on a larger field of view. Accurate methods for the statistical reconstruction of 3D volumes using 2D micrographs are thus highly desirable.

Generative adversarial networks (GANs) are a promising candidate model for this task. A GAN consists of two neural nets: a generator, G , to synthesise fake samples, \mathbf{f} , and a discriminator, D , to distinguish between \mathbf{f} and real samples, \mathbf{r} , from the dataset. During training, G and D are updated iteratively, enabling the generator to capture features of the real dataset and ultimately, if successful, produce realistic images^{5,9,19}. Since their introduction, extensive work has led to improved image quality, feature diversity and training stability². The motivation behind this paper is to apply this powerful technique to the task of dimensionality expansion. Crucially, in order to achieve this, the incompatibility between a 3D generator versus 2D training data must be solved. Furthermore, important considerations regarding the effect of information density on image quality are raised. The key contributions of this work are summarised as follows:

^{*}s.kench19@imperial.ac.uk

[†]samuel.cooper@imperial.ac.uk

1. In section 3, we introduce a novel GAN architecture, *SliceGAN*, which is designed to address the challenge of generating 3D volumetric data from 2D slices of an isotropic material. The ability to statistically reconstruct anisotropic microstructures with a simple extension is also demonstrated.

2. In section 4, the low quality regions found at the outer edges of GAN generated images is discussed. This issue is found to be associated with non-uniform information density in the generator. The authors define a set of requirements for the parameters of transpose convolutional operations to avoid this issue.

3. In section 5, *SliceGAN* is applied to a range of 2D microstructures, and shown to produce 3D volumes with visually indistinguishable cross sections from training data. The seven example training sets include a synthetic crystalline microstructure, an anisotropic polymer membranes and a three phase battery electrode material. This final example is further validated through a comparison of three key metrics to demonstrate the statistical similarity between synthetic volumes and the original dataset.

2 Background

2.1 GANs and their application in material science

Deep convolutional generative adversarial networks use a layered architecture to identify and reproduce the hierarchical features of an image dataset. When trained successfully, they offer high resolution, diverse image generation. Since the introduction of GANs, numerous adaptations have been proposed, falling broadly into two categories. First, there are those dedicated to improved training of adversarial nets, such as Wasserstein GANs (WGANs), which increase training stability and generated image diversity by addressing the challenge of vanishing gradients^{2,12}. Second, there are approaches that aim to bring new capabilities such as conditional GANs (cGANs), which are used to condition a generator based on image labels, thus allowing the synthesis of multiple image categories using a single generator¹⁷. A further relevant example of this latter category developed a method for spatial model generation from a set of 3D training samples²³. The diverse and rapid evolution of these different methodologies indicates the flexibility and wide applicability of GANs.

It is thus unsurprising that GANs have already proven useful in the field of material science¹⁸. A recent study by Gayon-Lombardo *et al.* used GANs to generate the

microstructure of energy storage materials based on 3D training data⁷. The synthetic volumes were not only much larger than the training set, but also had periodic boundaries, which makes them particularly well suited to generating representative domains for multi-physics simulation. GANs have also been applied to the task of material design. Li *et. al* trained a generator to synthesise realistic micrographs of a semi-conductor surface²⁶. A batch of instances can then be produced by the generator, and the property of interest, in this case absorption co-efficient, calculated for each sample. A Gaussian processes algorithm was then used to optimise for this property through an informed exploration of the microstructural space. This offers an efficient alternative compared to a conventional experimental materials design approach. These studies are an indication of the broad potential applications of GANs in the field of material science.

2.2 Current microstructural reconstruction approaches

Due to the widespread potential applications of 2D to 3D reconstruction algorithms, a range of techniques have been developed in previous studies. For example, given a particle size distribution measured from a 2D image, simulated annealing algorithms can be used alongside ellipsoid packing to emulate powder processing²⁴. Equally, tessellation algorithms are incorporated in software tools such as Dream3D, which uses extracted 2D properties such as grain size and orientation to simulate 3D grain growth¹¹. Unfortunately, these techniques are designed for, and thus limited to, powder and polycrystalline microstructures respectively. Furthermore, a loss of geometric information is inevitable, as irregular particles are approximated to ellipsoids, and complex grain shapes are characterised by mean diameters and aspect ratios.

Other approaches have been developed based on correlation statistics. Given an image containing i phases, the two-point correlation function S_2 describes the probability that two randomly selected pixels, separated by distance r , are of the same phase²². Importantly, it has been shown analytically that for a given isotropic material, the same two-point correlation function is recovered whether calculated for a 2D or 3D dataset²¹. This led to the development of stochastic algorithms which take initially random 3D volumes and iteratively attempt to minimise the difference between the correlation functions of the real and generated data^{13,15}. Whilst successful in some cases, the degeneracy of two-point correlations can lead to unrealistic results⁸.

Furthermore, even state-of-the-art methods are computationally expensive with simulation times to synthesise a modest 10^6 voxel sample on the order of hours using a 2.8 GHz, 16GB RAM computer, making this approach unsuitable for materials optimisation approaches²⁷.

The GANs described in this work take a similar time span to train (c. 4 hours on an NVIDIA Titan Xp GPU); however, once trained, the generator is able to synthesise 10^8 voxel instances in seconds. This 10^5 acceleration in generation speed compared to²⁷ enables both high-throughput optimisation and large scale modelling. Furthermore, GANs are highly flexible compared to the tessellation algorithms described above. Indeed, a key strength of deep convolutional neural nets is their ability to efficiently capture the diverse and complex features of an image without any user defined statistical metrics as inputs. Another important advantage of using GANs is the breadth of available variants that could be incorporated. Adaptations might aim to decrease training time, for example through transfer learning, or to facilitate new capabilities, for example using conditional GANs to allow the generation of different microstructures with specific features indicated by a label. *SliceGAN* is thus a building block for a powerful tool set of machine learning methods for microstructural characterisation, modelling and optimisation.

3 SliceGAN

The fundamental role of *SliceGAN* is to resolve the dimensionality incompatibility between a 2D training image and 3D generated volumes. This is achieved by incorporating a slicing step before fake instances from the 3D generator are sent to the 2D discriminator. For a generated cubic volume of edge length l voxels, $3l$ 2D images are obtained by taking slices along the x , y and z directions at 1 voxel increments. During the

training of D , for each fake 2D slice, a real 2D image is also sampled and fed to the discriminator, such that it learns equally from real and fake instances. To train G , the same slicing procedure is followed; however, a larger batch size (m_G) compared to the discriminator (m_D) is used. This rebalances the effect of training D on a large number of slices per generated sample. We find that $m_G = 2m_D$ typically results in the best efficiency. Importantly, the particular architecture of G used in this study means that each slice in any given set of 32 adjacent planes is synthesised through a unique combination of kernel elements (see section 4). As such, a minimum of 32 slices in each direction must be shown to D in order to ensure each path through the generator is trained. In practice, we find training to be both more reliable and efficient when D is applied to all 64 slices in each direction, despite the doubling of discriminator operations that this entails. After slicing, a standard Wasserstein loss function is used as proposed in² to encourage stable training. The complete training process for *SliceGAN* is described in Algorithm 1, with the data sampling, generation and slicing stages depicted in Figure 1.

Importantly, the above approach is only suitable for isotropic materials, where a single 2D image can be used to train slices from the x , y and z orientations of \mathbf{f} . An extension of the architecture is required for anisotropic materials. For example, given a microstructure containing a matrix with rods oriented parallel to the z axis, two training images are needed: one perpendicular to the z axis containing circular features and another parallel to the z axis containing long rectangular features (whilst these are the simplest orientations for this example, any set of perpendicular planes could be chosen in principal). During training, images \mathbf{r}_z are sampled from the former and images $\mathbf{r}_{x,y}$ from the latter, whilst slices from the fake samples are also divided into those taken along the z axis, \mathbf{f}_z , and those

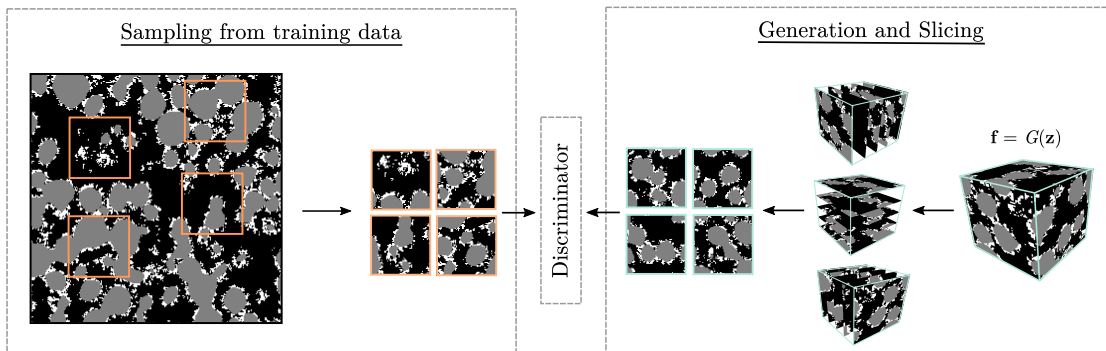


Figure 1: *SliceGAN* training procedure. First, real images are sampled from a 2D training sample. Second, a fake volume \mathbf{f} is generated and sliced along x , y and z . This yields a compatible pair of datasets, which can both be fed to a 2D discriminator.

Algorithm 1 SliceGAN algorithm for isotropic materials

Require: w and θ , the trainable parameters for D and G respectively; the number of discriminator iterations per generator iteration n_D ; the batch sizes m_D and m_G for D and G respectively; the gradient penalty coefficient, λ ; volume edge length, l ; and Adam optimizer hyperparameters¹⁶, α, β_1, β_2 .

- 1: **while** θ has not converged **do**
- 2: Discriminator training:
- 3: **for** $t = 0, \dots, n_D$ **do**
- 4: **for** $i = 1, \dots, m_D$ **do**
- 5: Sample a latent vector from a normal distribution $\mathbf{z} \sim p(\mathbf{z})$
- 6: $\mathbf{f} \leftarrow G_\theta(\mathbf{z})$ generate a 3D volume
- 7: **for** $a = 1, 2, 3$ **do**
- 8: **for** $d = 1, \dots, l$ **do**
- 9: $\mathbf{f}_s \leftarrow$ 2D slice of \mathbf{f} at depth d along axis a
- 10: Sample an $l \times l$ image from the real dataset $\mathbf{r} \sim \mathbf{P}_r$
- 11: Sample a random number $\epsilon \sim \mathbb{U}[0, 1]$
- 12: $\mathbf{k} \leftarrow \epsilon \mathbf{f}_s + (1 - \epsilon) \mathbf{r}$
- 13: $L_D^{(a,d)} \leftarrow D_w(\mathbf{f}_s) - D_w(\mathbf{r}) + \lambda(\|\nabla_{\mathbf{k}} D_w(\mathbf{k})\|_2 - 1)^2$
- 14: $w \leftarrow \text{Adam}(\nabla_w \frac{1}{m_D} \sum_{a=1}^3 \sum_{d=1}^l L_D^{(a,d)})$
- 15: Generator Training:
- 16: **for** $j = 1, \dots, m_G$ **do**
- 17: Sample a latent vector from a normal distribution $\mathbf{z} \sim p(\mathbf{z})$
- 18: $\mathbf{f} \leftarrow G_\theta(\mathbf{z})$ generate a 3D volume
- 19: **for** $a = 1, 2, 3$ **do**
- 20: **for** $d = 1, \dots, l$ **do**
- 21: $\mathbf{f}_s \leftarrow$ 2D slice of \mathbf{f} at depth d along axis a
- 22: $L_G^{(a,d)} \leftarrow -D_w(\mathbf{f}_s)$
- 23: $\theta \leftarrow \text{Adam}(\nabla_\theta \frac{1}{m_G} \sum_{a=1}^3 \sum_{d=1}^l L_G^{(a,d)})$

taken from x or y , $\mathbf{f}_{x,y}$. Separate discriminators are then taught to capture the distribution of features along the different orientations, and thus teach G to synthesise volumes that look like \mathbf{r}_z along the z axis, and $\mathbf{r}_{x,y}$ along the x and y axis. The adjusted *SliceGAN* algorithm for anisotropic materials is given in Supplementary Information S1. Successful examples of this approach are presented in section 5.

4 Generator information density

During the development stages of *SliceGAN*, significantly lower visual quality was observed near sample edges compared to the bulk. This may not be problematic for GANs applications such as face generation, where key objects often lie in the centre of an image and the edges are of less importance. However, for microstructural synthesis this is not the case, as although the relative position of features is important, their absolute location is often arbitrary. Edges are thus of equal importance to the center, and can strongly influence the measured physical properties of a microstructure. To understand the cause of edge artifacts, a simple test case is proposed to explore the nature of transpose convolutions, which are the key operator within the convolutional generator neural nets used in this work. We

first define a 3×3 input with a single channel, and pixel values set to 1. The transpose convolution parameters are selected with kernel size of $k = 4$; stride, $s = 2$; padding, $p = 0$ (NB. In a generator, p refers to the removal of near edge layers, rather than their addition) and output channels $C_{\text{out}} = 1$. The kernel weights are set to 1. Figure 2a depicts this simple scenario, and shows how the transpose convolution operates on the input pixel by pixel. We first observe the output from the top left pixel, which is distributed over a region equal to the kernel size. As the contribution from subsequent pixels is added, an overlap of information is seen. The final output shows a central region with flat information density of 4, and lower information content towards the edges and corners. Figure 2b shows the additive effect of this phenomenon in a typical generator with multiple convolutional operations. Indeed, the number of parameters contributing to central pixels is 3 orders of magnitude greater than for edge pixels, which explains the poor image quality in this region.

In order to avoid artifacts of this nature, three rules can be used to define acceptable combinations of k, s and p such that a uniform information density is ensured (Although an important consideration for image quality, C_{out} does not effect information density gradients):

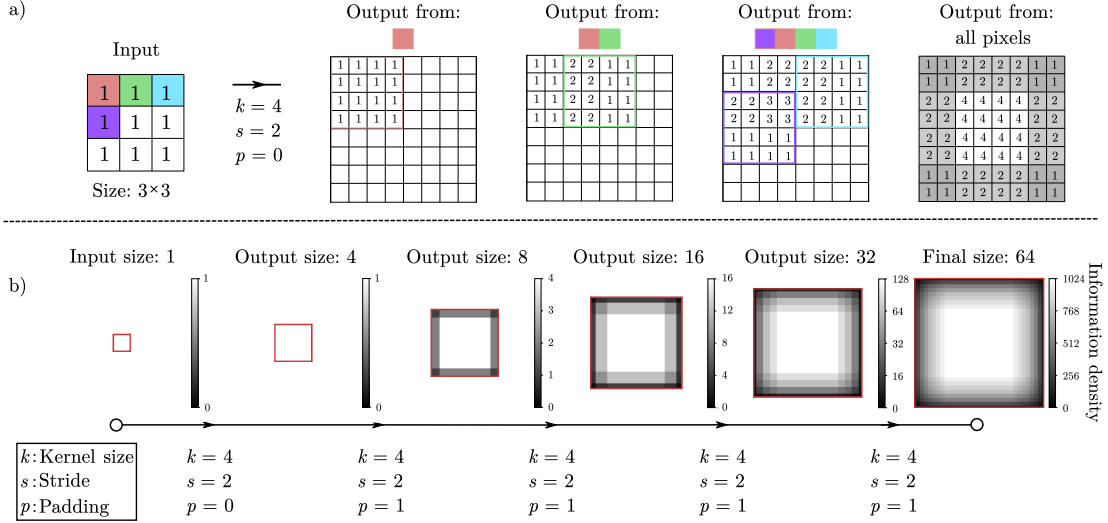


Figure 2: a) An example of how information density is calculated for a single transpose convolution with input size = 3×3 , $k = 4$, $s = 2$ and $p = 0$. The overlap from the outputs of individual pixels is depicted, showing a information density gradient at the edges of the final image. b) Illustration of how information density gradients propagate and are exacerbated in a full neural net as an input is passed through multiple transpose convolutional layers.

1. $s < k$ (Although $s \leq k$ is a sufficient requirement for uniformity, $s < k$ ensures that there is kernel overlap which is crucial for avoiding feature discontinuities in the output.)
2. $k \bmod s = 0$ (If k is not divisible by s , the kernel output regions from nearby input pixels never lie flush to each other, resulting in a checkerboard information density²⁰. An example of this behaviour with $k = 4$, $s = 3$ is shown in Supplementary Information S2.)
3. $p \geq k - s$ (The region of edge voxels with less than the central uniform information density is of width $k - s$, and must thus be removed by applying a padding of this value or greater.)

These constraints leave a small set of practical transpose convolutions. We first discard any prime value of k , as in this scenario rule 1 and 2 together only allow for $s = 1$. Such configurations result in no expansion of the spatial dimensions and are thus not useful. In general, $k > 6$ is also impractical, especially in 3D convolutional nets, as the number of parameters scales with k^3 . Assuming p is chosen to be as small as possible to minimise the amount of cropping at each layer (which equates to wasted computational expense), this leaves three practical sets for $\{k, s, p\}$: $\{4, 2, 2\}$, $\{6, 3, 3\}$, $\{6, 2, 4\}$. In the work presented here, the $\{4, 2, 2\}$ set of parameters are used for most transpose convolutions to avoid the computational expense of $k = 6$. An alternative approach to avoid edge artifacts could be the use of resize-convolution generator architectures with a different set of parameter

rules (most importantly, the use of any zero-padding results in undesirable information gradients). Unfortunately, under these constraints resize-convolution has significantly larger memory requirements for 3D volume generation than transpose convolutions. A sacrifice in either training time (by reducing batch size) or image quality (by reducing the number of filters) thus becomes necessary, as shown in Supplementary Information S3.

Another crucial attribute when designing a GAN architecture for material science applications is the ability to generate different size volumes during inference. This is important as often the standard training volume (normally 64 voxels cubed to allow reasonably large batch sizes under memory constraints) is not large enough for use in multiphysics simulations. A previous study suggests that after training, bespoke volumes can be synthesised by varying the spatial size of the latent vector \mathbf{z} ⁷. However, this approach results in lower quality microstructures with distortions. This observation can be explained by the fact that during training, \mathbf{z} has a spatial dimension of 1, making the first transpose convolution of G unique in that it operates on a single voxel per channel. In this scenario, no kernel overlap occurs in the first layer output. Increasing the spatial dimension of \mathbf{z} during inference changes this behaviour and introduces overlap. To avoid this problem, we choose to train into the first generator layer an understanding of overlap; thus an input vector with spatial size 4 is used. The resulting information spread is de-

Table 1: Table showing the architecture of the 3D generator (memory size \approx 50 MB) and 2D discriminator (memory size \approx 11 MB) used in this work.

Generator			Discriminator		
Layer	k, s, p	Output shape	Layer	k, s, p	Output shape
Input z	-	$64 \times 4 \times 4 \times 4$	Input	-	$3 \times 64 \times 64$
1	4,2,2	$512 \times 6 \times 6 \times 6$	1	4,2,1	$64 \times 32 \times 32$
2	4,2,2	$256 \times 10 \times 10 \times 10$	2	4,2,1	$128 \times 16 \times 16$
3	4,2,2	$128 \times 18 \times 18 \times 18$	3	4,2,1	$256 \times 8 \times 8$
4	4,2,2	$64 \times 34 \times 34 \times 34$	4	4,2,1	$512 \times 4 \times 4$
5	4,2,3	$3 \times 64 \times 64 \times 64$	5	4,2,0	$1 \times 1 \times 1$
softmax	-	$3 \times 64 \times 64 \times 64$			

picted in Supplementary Information S4, with the final network architectures given in Table 1. Interestingly, this alteration leads to periodicity within the generator, whereby sets of every 32nd plane are generated from the same combination of kernel elements, as shown in Supplementary Information S5. This does not make the synthetic volume itself periodic, as each plane is associated with a different region of the input seed. However, it does mean that it is not possible to train G to capture anisotropy within a given plane (e.g a graded density).

5 Empirical results

5.1 Data pre-processing

Whilst it is possible to train *SliceGAN* using RGB or greyscale micrographs (for which no pre-processing is necessary), the approach has been most successful with one-hot encoded representations of segmented n -phase microstructural data. For a training dataset, the encoding process entails the creation of n separate channels, each containing 1s where the material phase they correspond to is present and 0s otherwise, as shown in Supplementary Information S6. In generated instances, n channels are again required, though they now represent the probability of finding a given phase at a location in an image. This is implemented using a *softmax* function as the final layer of the generator. In either case, an n -phase microstructure can always be recovered from a one-hot array through a *max* function over the channel axis.

5.2 Application to microstructures

To demonstrate the capabilities of *SliceGAN*, a range of micrographs were chosen to be statistically reconstructed, as shown in Figure 3. The first column shows a region of the 2D training data used. The second column shows an example of a generated 3D instance, with the volume edge lengths selected to best display the microstructural features. Next, slices of a synthetic volume taken i) parallel to a discriminated plane to show the dataset quality and ii) at a 45° angle to the discriminators to demonstrate the consistency of the microstructure in directions not involved in training. For each case, the material, imaging technique and training details are given in Table 2. The first example presented in row A of Figure 3 is a synthetic grain structure, chosen for its well defined features: perfectly straight grain boundaries and exclusively convex grains. *SliceGAN* performs reasonably well on this data, generating realistic looking volumes, though some disconnected and curved grain boundaries are observed. These shortcomings are not specific to the *SliceGAN* approach, as similar features are also observed in synthetic 2D images from a conventional GAN. The next training image (row B) is a real micrograph from surface analysis of a ceramic, where the black phase correspond to areas with high Caesium content¹⁰. Again, a visually realistic sample is generated through *SliceGAN*.

Two images of fibre reinforced composite taken from perpendicular views demonstrate the ability to syn-

Table 2: Synthesis details for the examples presented in Figure 3.

Label	Material	Imaging technique	Image type	Isotropy	Ref
A	Polycrystalline grains	Synthetic	Two-phase	Isotropic	-
B	Ceramic (perovskite)	Kelvin probe force topography	Two-phase	Isotropic	10
C	Carbon fibre rods	Secondary electron microscopy	Two-phase	Anisotropic	3
D	Battery separator	X-ray tomography reconstruction	Three-phase	Anisotropic	6
E	Steel	Electron back-scatter microscopy	Colour	Isotropic	1
F	Grain boundary	Synthetic	Three-phase	Anisotropic	-
G	NMC battery cathode	X-ray tomography	Three-phase	Isotropic	14

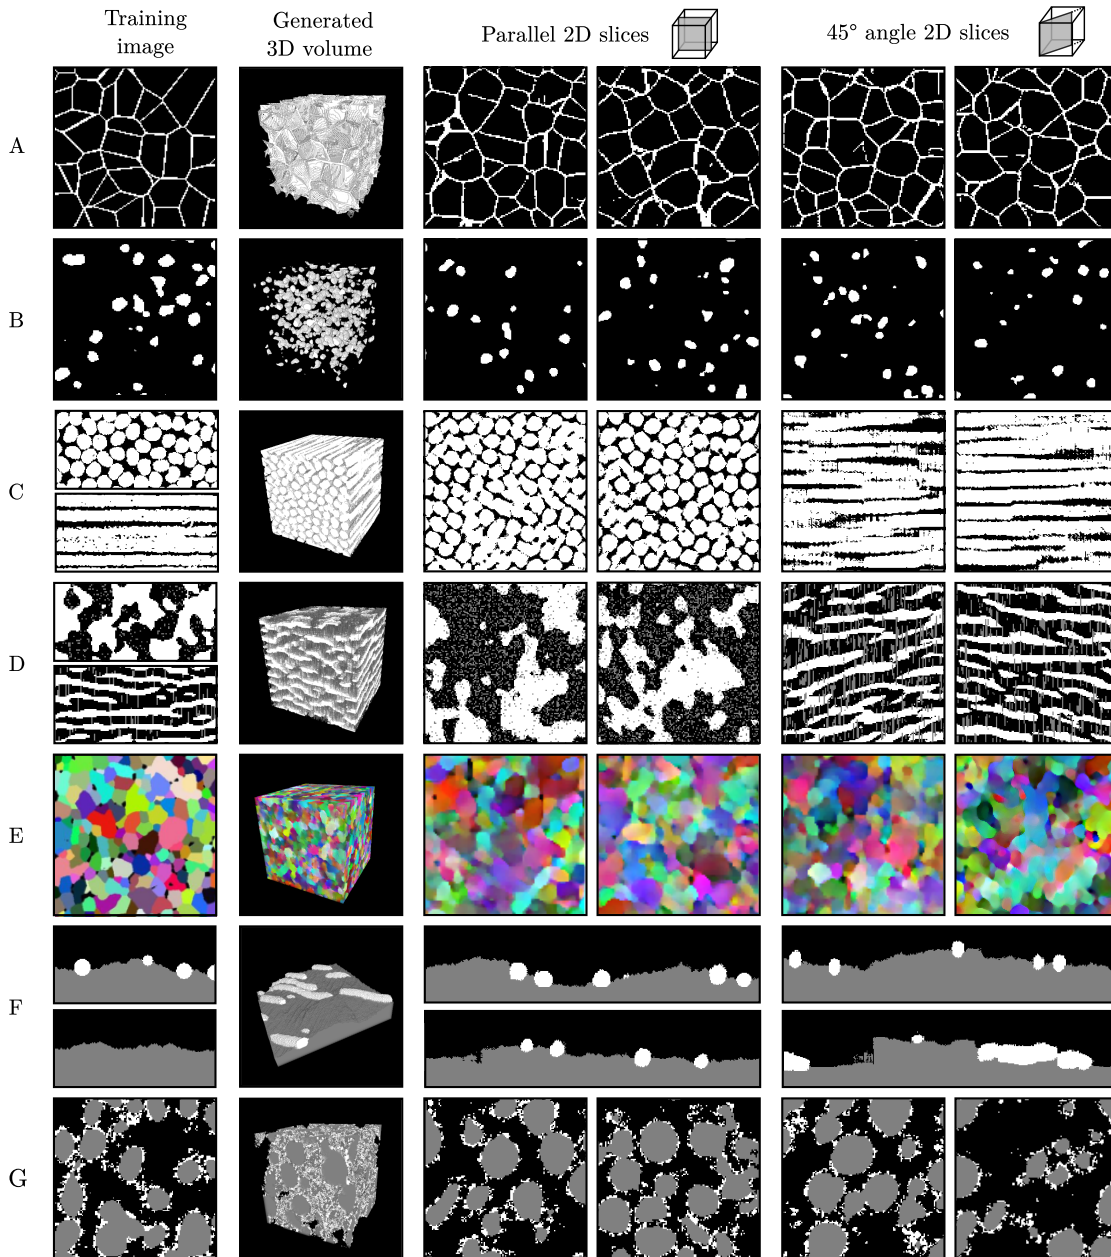


Figure 3: Application of *SliceGAN* to a variety of microstructures, with details specified in Table 2. From left to right, the original training dataset, a 3D statistical reconstruction and 2 examples of slices through an example volume taken at different angles.

these anisotropic images³. A similar approach is taken for a polymer battery separator material^{6,25}. This demonstrates the flexibility of *SliceGAN*, and will enable its application to many common microstructures (rows C and D). A micrograph of polycrystalline single phase steel is considered next (row E)¹. Previous work was focused on segmented 2 or 3 phase materials; here the extension to colour images (where color encodes grain orientation) is presented. Image quality is some-

what reduced compared to previous examples. This is likely because our 12 GB GPU memory constraint means that the networks are too small to capture continuous label spaces (instead of the segmented materials in other examples). These volumes can be made to look more like the training data by applying filters (e.g. mode), but this kind of post-processing is beyond the scope of this study and would need to be selected on a case by case based. In row F, a test on a synthetic

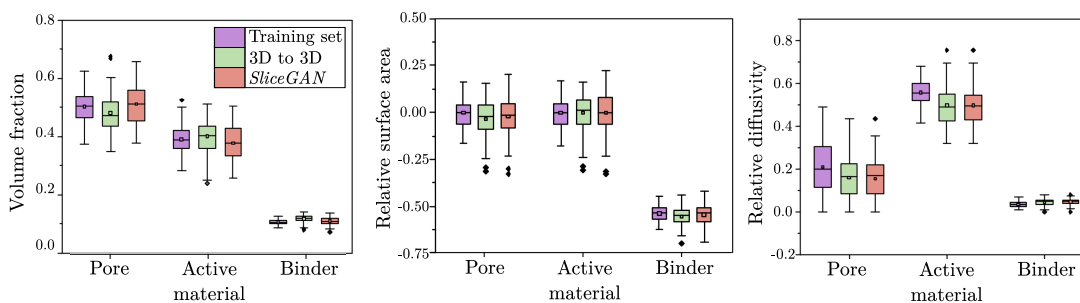


Figure 4: Statistical analysis of three electrochemical and general materials properties, showing that the real properties of the dataset are captured well by both a 3D to 3D GAN and *SliceGAN*. Boxes depict the interquartile range (IQR) from the lower quartile, $Q1$ to the upper quartile, $Q3$. Each box contains a horizontal line and a small square to indicate the median and mean respectively. The upper and lower whiskers show the last datum greater than $Q3 + (1.5 \times IQR)$ and less than $Q1 - (1.5 \times IQR)$ respectively. Outlier data points beyond this range are represented by solid black diamonds.

material interface is presented where a large 128×10^4 pixel image was generated using a bounded random walk algorithm to generate an interface, followed by seeding of randomly size boundary particles. In this case, the 2D example images shown are small crops of the full training dataset. This was an opportunity to test whether 2 discriminators, both perpendicular to the material interface, were sufficient for training. Interestingly, whilst the general surface topology of the synthesised material reflects the real dataset well, the interface particles form long ridges on the surface at a 45° angle. This is due to insufficient constraint without the top view of the material, which allows the generator to effectively trick the discriminator. This could potentially be resolved by implementing discriminated slices at 45° angles.

Finally, open access tomographic data collected from a Li-ion NMC cathode sample is used as shown in row G¹⁴. Although the dataset is 3D, only a random subsample of 2D sections are used for *SliceGAN* training. This experiment is particularly useful as 3D properties of the original dataset can be compared to those of the synthesised microstructure. Furthermore, instances from a 3D to 3D GAN can also be evaluated. Thus, to quantify the performance of *SliceGAN*, 100 sample batches were taken from each of i) the real dataset ii) the 3D to 3D dataset generated by Gayon-Lombardo *et al*⁷ and iii) *SliceGAN*. Three relevant metrics were chosen and calculated using *TauFactor*⁴ (an open-source microstructural analysis package) for each set of microstructures, as shown in Figure 4. Two-point correlation functions and triple phase boundary densities are also compared in Supplementary Information S7.

For all of these conventional microstructural metrics, both GAN architectures reproduce the training set well.

Although the medians correspond very closely in each case, the distributions, in particular in the case of the relative diffusivity, are not exactly the same as the training set. Crucially, unlike the other metrics, the effective diffusivity is an emergent property of the generated 3D volume as a whole, and is not measurable from any individual 2D slice. As such, it’s particularly impressive that *SliceGAN* can reproduce this metric.

6 Conclusions

The results presented in this study show the ability of *SliceGAN* to faithfully synthesise 3D n -phase media from 2D micrographs. This tool has three key uses: First, in the field of material characterisation as a means to visualise a statistically realistic 3D instance of a 2D micrograph. This is especially important considering the limitations of direct 3D imaging, namely in terms of resolution and field of view. Second, for electrochemical simulation; 2D micrographs are insufficient for a range of desirable electrochemical and mechanical simulation techniques. Finally, for materials optimisation purposes, as the speed of the generator allows rapid exploration of the space of possible 3D microstructures. Future work will aim to develop this tool set in combination with other common GANs approaches such as conditional GANs and transfer learning. The former will enable interpolation between microstructures with differing properties, whilst the latter aims to speed the training process. Thus, there is significant potential for expansion of this approach as a tool for material characterisation and optimisation.

Data Availability

The study used open-access training data available from the following sources: ceramic¹⁰, carbon fibre rods³,

battery separator^{6,25}, steel¹, and NMC battery cathode¹⁴.

Code Availability

The codes used in this manuscript, are available at: <https://github.com/stke9/SliceGAN>. All other generated data used is available from the authors on request.

Correspondence and requests for materials should be directed to Steven Kench at sk3619@ic.ac.uk.

Acknowledgements

This work was supported by funding from the EPSRC Faraday Institution Multi-Scale Modelling project (<https://faraday.ac.uk/>; EP/S003053/1, grant number FIRG003 received by SK). The Titan Xp GPU used for this research was kindly donated by the NVIDIA Corporation through their GPU Grant program to SC. The authors would also like to thank Aron Walsh, Andrea Gayon Lombardo and Daniel Figueroa for their valuable comments, as well as the rest of the tldr group.

Authorship

SK designed and developed the code for *SliceGAN*, trained the models, performed the statistical analysis and drafted the manuscript. SJC contributed to the development of the concepts presented in all sections of this work, helped with data interpretation in section 5 and made substantial revisions and edits to all sections of the draft manuscript.

Competing interests

The authors declare no competing interests.

References

- [1] Oxford instruments grain size and grain boundary characterisation in sem, 2020.
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [3] D. Armentrout, M. Kumosa, and T. Mcquarrie. Boron-free fibers for prevention of acid induced brittle fracture of composite insulator grp rods. *Power Delivery, IEEE Transactions on*, 18:684–693, 08 2003.
- [4] S. Cooper, A. Bertei, P. Shearing, J. Kilner, and N. Brandon. Taufactor: An open-source application for calculating tortuosity factors from tomographic data. *SoftwareX*, 5:203–210, 2016.
- [5] A. Creswell, T. White, V. Dumoulin, K. Arulkumar, B. Sengupta, and A. A. Bharath. Generative adversarial networks: An overview. *IEEE Signal Processing Magazine*, 35(1):53–65, 2018.
- [6] D. Finegan, S. Cooper, B. Tjaden, O. Taiwo, J. Gelb, G. Hinds, D. Brett, and P. Shearing. Microstructure reconstruction of battery polymer separators by fusing 2d and 3d image data for transport property analysis. *Journal of Power Sources*, 333:184–192, 2016.
- [7] A. Gayon-Lombardo, L. Mosser, N. P. Brandon, and S. J. Cooper. Pores for thought: The use of generative adversarial networks for the stochastic reconstruction of 3d multi-phase electrode microstructures with periodic boundaries, 2020.
- [8] C. J. Gommès, Y. Jiao, and S. Torquato. Microstructural degeneracy associated with a two-point correlation function and its information content. *Physical Review E*, 85(5):051140, 2012.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [10] P. Gratia, I. Zimmermann, P. Schouwink, J.-H. Yum, J.-N. Audinot, K. Sivula, T. Wirtz, and M. K. Nazeeruddin. The many faces of mixed ion perovskites: unraveling and understanding the crystallization process. *ACS Energy Letters*, 2(12):2686–2693, 2017.
- [11] M. A. Groeber and M. A. Jackson. Dream. 3d: a digital representation environment for the analysis of microstructure in 3d. *Integrating materials and manufacturing innovation*, 3(1):5, 2014.
- [12] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. *CoRR*, abs/1704.00028, 2017.
- [13] A. Hasanabadi, M. Baniassadi, K. Abrinia, M. Safdari, and H. Garmestani. 3d microstructural reconstruction of heterogeneous materials from 2d cross sections: a modified phase-recovery algorithm. *Computational Materials Science*, 111:107–115, 2016.
- [14] T. Hsu, W. K. Epting, R. Mahbub, N. T. Nuffer, S. Bhattacharya, Y. Lei, H. M. Miller, P. R. Ohodnicki, K. R. Gerdes, H. W. Abernathy, et al. Mesoscale characterization of local property distributions in heterogeneous electrodes. *Journal of Power Sources*, 386:1–9, 2018.
- [15] H. Izadi, M. Baniassadi, A. Hasanabadi, B. Mehrgini, H. Memarian, H. Soltanian-Zadeh,

- and K. Abrinia. Application of full set of two point correlation functions from a pair of 2d cut sections for 3d porous media reconstruction. *Journal of Petroleum Science and Engineering*, 149:789–800, 2017.
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.
- [17] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [18] L. Mosser, O. Dubrule, and M. J. Blunt. Reconstruction of three-dimensional porous media using generative adversarial neural networks. *Phys. Rev. E*, 96:043309, Oct 2017.
- [19] A. Odena. Semi-supervised learning with generative adversarial networks. *arXiv preprint arXiv:1606.01583*, 2016.
- [20] A. Odena, V. Dumoulin, and C. Olah. Deconvolution and checkerboard artifacts. *Distill*, 2016.
- [21] S. Torquato and H. Haslach Jr. Random heterogeneous materials: microstructure and macroscopic properties. *Appl. Mech. Rev.*, 55(4):B62–B63, 2002.
- [22] S. Torquato and G. Stell. Microstructure of two-phase random media. i. the n-point probability functions. *The Journal of Chemical Physics*, 77(4):2071–2077, 1982.
- [23] J. Wu, C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *CoRR*, abs/1610.07584, 2016.
- [24] H. Xu, D. A. Dikin, C. Burkhart, and W. Chen. Descriptor-based methodology for statistical characterization and 3d reconstruction of microstructural materials. *Computational Materials Science*, 85:206–216, 2014.
- [25] H. Xu, F. Usseglio-Viretta, S. Kench, S. J. Cooper, and D. P. Finegan. Microstructure reconstruction of battery polymer separators by fusing 2d and 3d image data for transport property analysis. *Journal of Power Sources*, Accepted 2020.
- [26] Z. Yang, X. Li, L. Catherine Brinson, A. N. Choudhary, W. Chen, and A. Agrawal. Microstructural materials design via deep adversarial learning methodology. *Journal of Mechanical Design*, 140(11), 2018.
- [27] Y. Zhang, M. Yan, Y. Wan, Z. Jiao, Y. Chen, F. Chen, C. Xia, and M. Ni. High-throughput 3d reconstruction of stochastic heterogeneous microstructures in energy storage materials. *npj Computational Materials*, 5(1):1–8, 2019.

Generating 3D structures from a 2D slice with GAN-based
dimensionality expansion

Supplementary Information

Steve Kench ^{*1} and Samuel J. Cooper ^{†1}

¹*Dyson School of Design Engineering, Imperial College London, London SW7 2DB*

February 16, 2021

*s.kench19@imperial.ac.uk

†samuel.cooper@imperial.ac.uk

A Algorithm for anisotropic microstructures

Algorithm 1 SliceGAN algorithm for anisotropic materials. Note the key change compared to isotropic is the use of different real data and a different discriminator for each axis.

Require: w_1, w_2, w_3 and θ , the trainable parameters for three perpendicular discriminators and G respectively; the number of discriminator iterations per generator iteration n_D ; the batch sizes m_D and m_G for D and G backpropagation respectively; the gradient penalty coefficient, λ ; volume edge length, l ; and Adam hyperparameters, α, β_1, β_2 .

```

1: while  $\theta$  has not converged do
2:   Discriminator training:
3:   for  $t = 0, \dots, n_D$  do
4:     for  $i = 1, \dots, m_D$  do
5:       Sample a latent vector from a normal distribution  $\mathbf{z} \sim p(\mathbf{z})$ 
6:        $\mathbf{f} \leftarrow G_\theta(\mathbf{z})$  generate a 3D volume
7:       for  $a = 1, 2, 3$  do
8:         for  $d = 1, \dots, l$  do
9:            $\mathbf{f}_s \leftarrow$  2D slice of  $\mathbf{f}$  at depth  $d$  along axis  $a$ 
10:          Sample an  $l \times l$  image from the real dataset associated with axis  $a$   $\mathbf{r}^a \sim \mathbf{P}_r^a$ 
11:          Sample a random number  $\epsilon \sim \mathbb{U}[0, 1]$ 
12:           $\mathbf{k} \leftarrow \epsilon \mathbf{f}_s + (1 - \epsilon) \mathbf{r}^a$ 
13:           $L_D^{(a,d)} \leftarrow D_{w_a}(\mathbf{f}_s) - D_{w_a}(\mathbf{r}) + \lambda(\|\nabla_{\mathbf{k}} D_{w_a}(\mathbf{k})\|_2 - 1)^2$ 
14:           $w_a \leftarrow \text{Adam}(\nabla_w \frac{1}{m_D} \sum_{d=1}^l L_D^{(a,d)})$ 
15:   Generator Training:
16:   for  $j = 1, \dots, m_G$  do
17:     Sample a latent vector from a normal distribution  $\mathbf{z} \sim p(\mathbf{z})$ 
18:      $\mathbf{f} \leftarrow G_\theta(\mathbf{z})$  generate a 3D volume
19:     for  $a = 1, 2, 3$  do
20:       for  $d = 1, \dots, l$  do
21:          $\mathbf{f}_s \leftarrow$  2D slice of  $\mathbf{f}$  at depth  $d$  along axis  $a$ 
22:          $L_G^{(a,d)} \leftarrow -D_{w_a}(\mathbf{f}_s)$ 
23:          $\theta \leftarrow \text{Adam}(\nabla_\theta \frac{1}{m_G} \sum_{d=1}^l L_G^{(a,d)})$ 

```

B Non-uniform density where s is not divisible by k

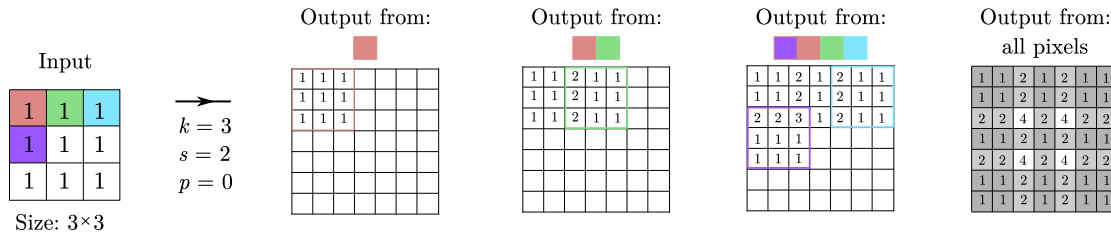


Figure S1: A transpose convolution example where kernel size (k) is not divisible by stride (s), resulting in a checkerboard, non-uniform information density. The fundamental cause of this behaviour is clear in the output from all four colored pixels. The cyan box shows how the convolution skips the central column, which is left with a low information density of 1 compared to its neighbours. This often results in unrealistic periodic features in generated images.

C Information density for a resize-convolution network

Figure S2 shows the edge information gradients associated with resize convolutions when using padding of 1 for each conv. Whilst it’s possible to use no padding, which does yield uniform information gradient, this requires a big input vector (min $8 \times 8 \times 8$) to ensure the spatial dimensions don’t collapse in the first layer, where normally padding would be used to enlarge the output. Large \mathbf{z} significantly limits the number of filters that can be used in the early layers due to computational expense (for a 12GB GPU, the first layer can have ≈ 64 filters, whilst a transpose convolution can have 512). This leads to poorer quality microstructures.

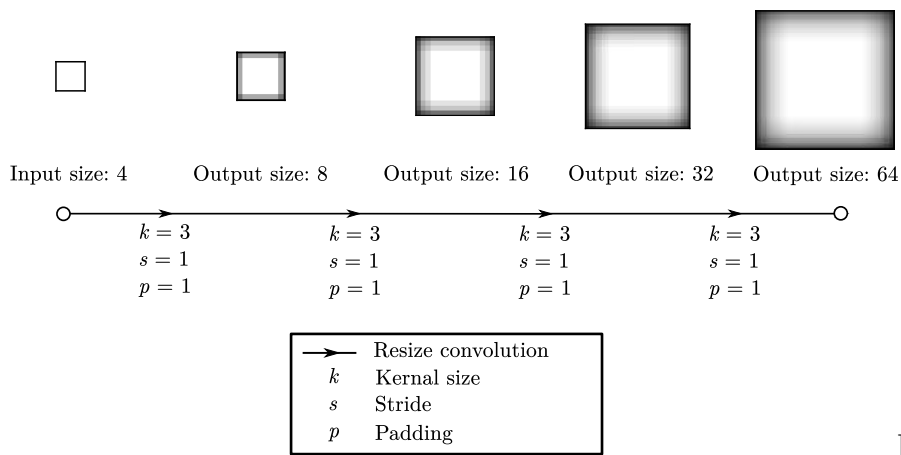


Figure S2: A typical resize-convolution generator using zero-padding and upscaling factor of 2 between each convolution, showing the low information density near the image edges.

D Information overlap during post-training generation

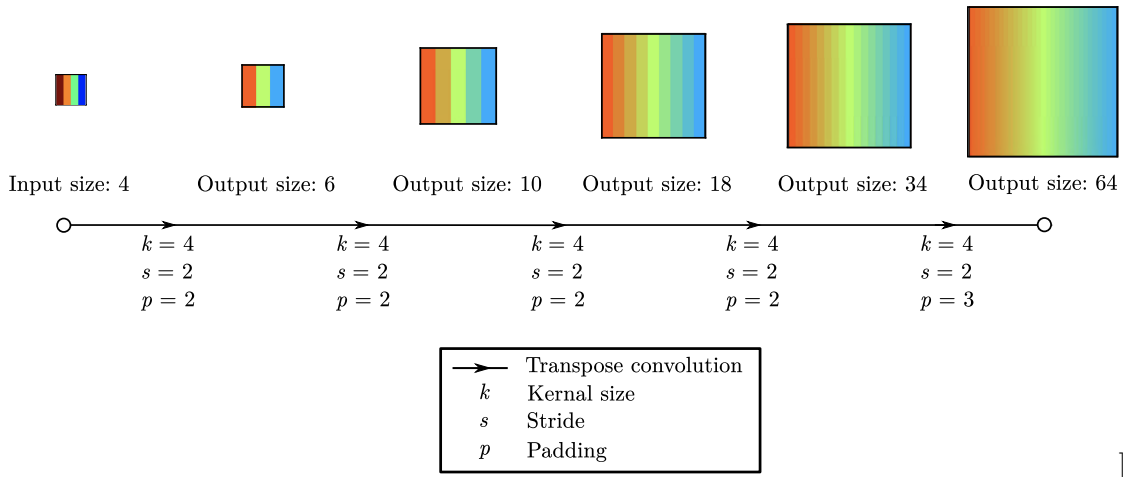


Figure S3: A transpose convolution example using a 4×4 coloured input, showing the overlap of information during generation. The dark red input stripe spreads over to exactly half way across the image, with a linearly decreasing contribution to the output. Note that each of the RGB channels is calculated individually then concatenated and normalised.

E Periodicity of kernel element combinations

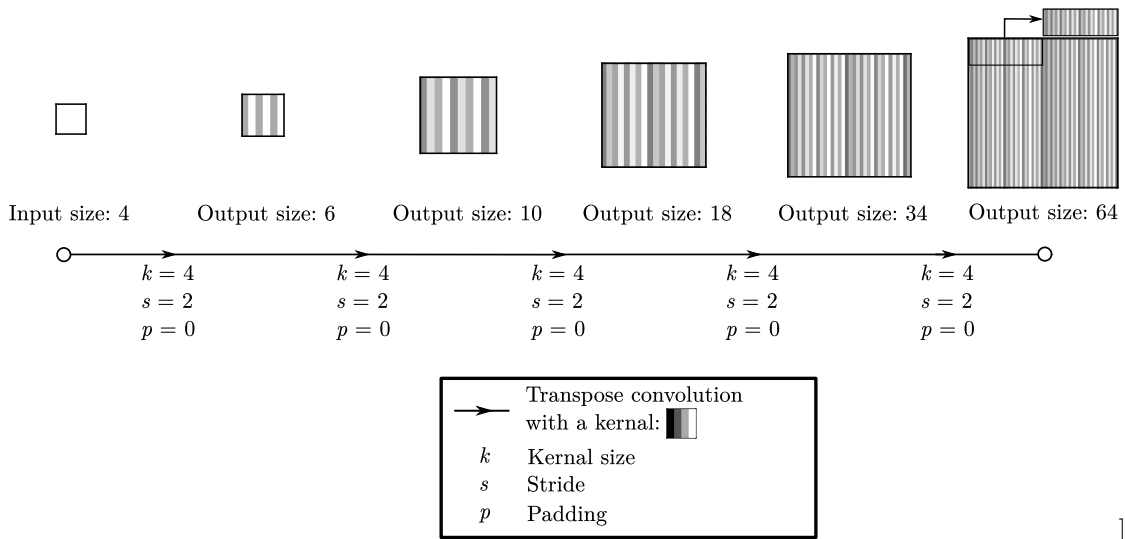


Figure S4: A transpose convolution example using a striped kernel with columns of value 1, 2, 3, and 4, showing the periodic behaviour of the generator. In the final output, only slices that are generated from the exact same sequence of kernel elements are the same darkness - this occurs for every 32nd slice, whilst all other pairs are unique.

F One-hot coded transformation

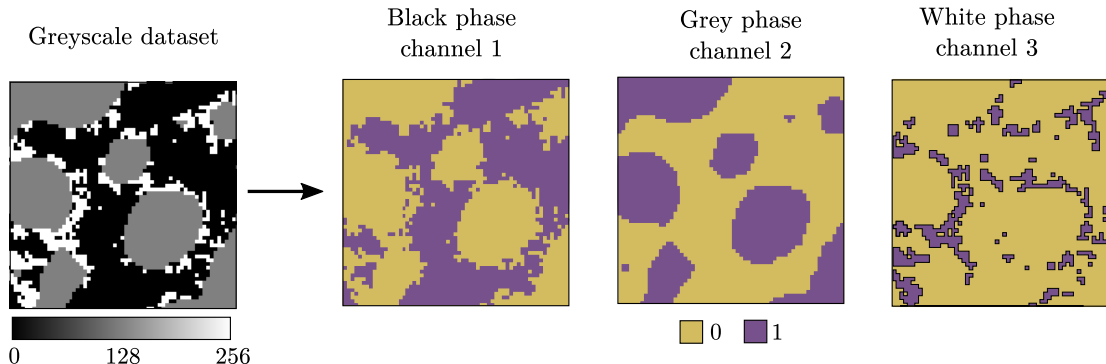


Figure S5: One-hot coding transformation starting from a segmented 8-bit grayscale image where each value corresponds to a particular phase, though their order is typically random. Whilst training image pixels only take binary values, generated samples instead have a probability distribution between 0 and 1 due to the *softmax* function in the final layer.

G Metric comparison of real and synthetic datasets

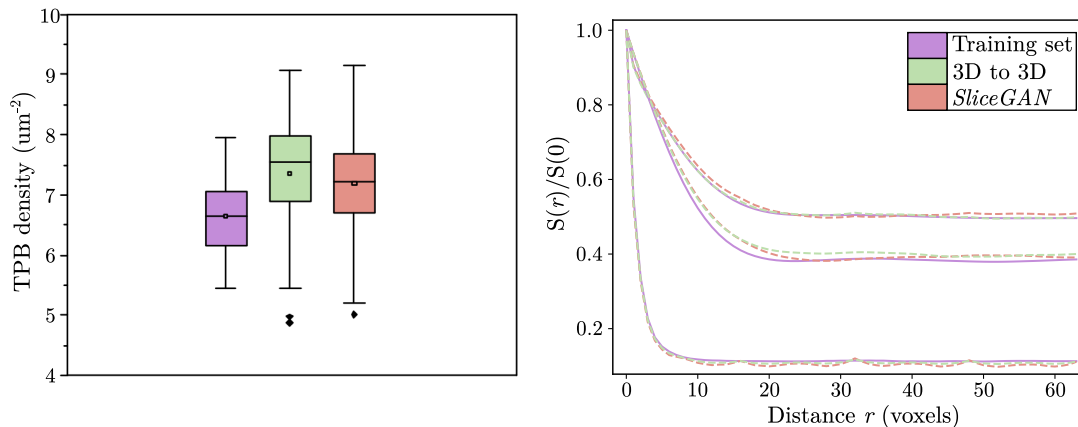


Figure S6: Triple phase boundary (TPB) density and two-point correlation ($S(r)$) functions show the similarity between real and synthetic microstructures. TPB density is a measure of how many voxel vertices are shared by all three phases for a given volume. The median and distributions match well for the 100 samples used in this test. The two-point correlation function describes the probability of finding the same phase at voxels separated by distance r . In this case, we normalise this function against volume fraction, $S(0)$, for ease of plotting. The values shown were measured using a single 250 edge length cube (corresponding to the full training data size). Again, the phase distributions match well.