



3D Grain Shape Generation in Polycrystals Using Generative Adversarial Networks

Devendra K. Jangid¹ · Neal R. Brodnik² · Amil Khan¹ · Michael G. Goebel¹ · McLean P. Echlin³ · Tresa M. Pollock³ · Samantha H. Daly² · B. S. Manjunath¹

Received: 26 July 2021 / Accepted: 20 November 2021 / Published online: 21 January 2022
© The Minerals, Metals & Materials Society 2021

Abstract

This paper presents a generative adversarial network (GAN) capable of producing realistic microstructure morphology features and demonstrates its capabilities on a dataset of crystalline titanium grain shapes. Alongside this, we present an approach to train deep learning networks to understand material-specific descriptor features, such as grain shapes, based on existing conceptual relationships with established learning spaces, such as functional object shapes. A style-based GAN with Wasserstein loss, called M-GAN, was first trained to recognize distributions of morphology features from function objects in the ShapeNet dataset and was then applied to grain morphologies from a 3D crystallographic dataset of Ti–6Al–4V. Evaluation of feature recognition on objects showed comparable or better performance than state-of-the-art voxel-based network approaches. When applied to experimental data, M-GAN generated realistic grain morphologies comparable to those seen in Ti–6Al–4V. A quantitative comparison of moment invariant distributions showed that the generated grains were similar in shape and structure to the ground truth, but scale invariance learned from object recognition led to difficulty in distinguishing between the physical features of small grains and spatial resolution artifacts. The physical implications of M-GAN's learning capabilities are discussed, as well as the extensibility of this approach to other material characteristics related to grain morphology.

Keywords Microstructure generation · 3D grain morphology · Machine learning · Deep learning · Generative adversarial networks · Polycrystalline materials

Introduction

One of the most pervasive challenges of materials discovery and development is that the process demands an enormous amount of time, labor, and capital. Reducing these demands has been a central driver for a number of international initiatives [1–7]. This challenge stems from the inefficient use of a linear approach to explore high-dimensional materials spaces, which is compounded by critical metrics for material classes that account for both processing variations and complexity of application environment. For example, there are a wide range of metal alloy compositions that can be used in turbine engine components, and the type of alloy being used can vary widely depending on the location and operating demands of the specific turbine section. But regardless of composition chosen, any alloy must be able withstand the operating environment, and evaluations of its performance must cover thousands or even millions of manufacturing iterations. The sheer size of these composition and application

Devendra K. Jangid and Neal R. Brodnik have contributed equally to this work.

✉ Neal R. Brodnik
nbrodnik@ucsb.edu

Tresa M. Pollock
tresap@ucsb.edu

Samantha H. Daly
samdaly@ucsb.edu

B. S. Manjunath
manj@ucsb.edu

¹ Electrical and Computer Engineering, University of California, Santa Barbara, USA

² Mechanical Engineering, University of California, Santa Barbara, USA

³ Materials Department, University of California, Santa Barbara, USA

spaces places materials science at a unique tipping point for the integration of machine learning to rapidly accelerate design and discovery.

Deep learning methods have attracted attention as cost reducers in fields such as voice, text, and image processing. Neural networks perform strongly both in accuracy and efficiency when trained on large collections of well-labeled ground-truth data. These networks may reliably outperform traditional algorithmic approaches and can make further important distinctions, like the identification of specific voices, faces, or fingerprints, with only a small amount of additional input. Neural networks can accelerate and refine the research process by reducing the quantity of required experimental iterations and by enhancing the accuracy and efficiency of material analyses. However, their performance depends on the quality of material data available for training. For materials data in particular, providing sufficient amounts of meaningful information is difficult, and finding ways to address this challenge is the goal of this study.

When considering available material data for use in deep learning approaches, the most significant obstacles are limitations in data scope and dimensionality. Here, data scope is defined as the breadth of readily available data, and what range of information that data describes. Large amounts of materials data have been captured, but much of these data are either difficult to access or are tailored to a specific purpose that limits its broader context. In other deep learning applications, such as text and voice, a great deal of data is made available through literary and audio collections and social media. In materials, not all data are readily available due to the expense of data generation or proprietary concerns, and when data are available, the wide array of collection methods can make establishing broader connections challenging. When discussing data dimensionality, limitations can be considered in terms of both variable and spatial dimensions. Variable dimensions refer to the number of different descriptors or characteristics represented in a given dataset. This challenge is particularly prevalent in engineering materials like the ones considered in this investigation. The vast majority of engineering materials are polycrystalline, so the bulk material is typically composed of billions of crystals that are rarely uniform in their crystallographic, morphological, or chemical properties. The effects of external environment and test conditions (e.g., specimen geometry, surface treatment, loading conditions) must also be considered, which adds further complexity to these material descriptors. In addition to the number of data variables, it is also critical to consider the dimensions of physical space being represented. In materials analysis, an important point to note is the difficulty of generating and subsequent lack of three-dimensional experimental measurements, which are essential for physically accurate neural network training. Recent experimental advances have greatly improved

volumetric materials characterization [8–15], but the quality and quantity of available data are limited by inherent restrictions, such as costly and specialized equipment, reduced spatial/temporal resolution, limitations on sample material/geometry, and sample destruction during characterization.

Deep learning capabilities that have been developed in other domains can be leveraged to address the problem of limited data in the materials domain. Adapting these capabilities for the characteristics of materials requires establishing a relationship between the processing–structure–properties space of materials science and the classification spaces used in other fields of deep learning. One such established space in deep learning is object recognition, where semantically labeled objects are recognized by neural networks based on the collection of features that compose that object. Approaches in this space are readily extensible to materials structure, where relationships can be established between microstructural morphology and other image recognition approaches that focus specifically on object shape. For example, in polycrystalline metals, much of the relationship between material processing, structure, and properties manifests through the shape and arrangement of the crystalline subdomains we call grains. In the case of uniform crystallization, the physics are well understood. Grains take the shape of faceted polyhedrons, where the number of facets is dictated by impinging neighbors, and facet size and shape are dictated by interfacial energies. Morphologically, these polyhedral shapes are relatively simple, but variations in thermal and mechanical processing can drastically alter the distribution of possible shapes, as demonstrated in Fig. 1. Under normal growth conditions, the most probable grain shapes are polyhedrons, but orientationally biased mechanical or thermal processing can create conditions where anisotropic grains that are flattened, twisted, or directionally aligned become more probable. When considered from the standpoint of deep learning, these types of processing–structure relationships can be seen as changes to the distribution of expected grain morphologies for a given material system. While learning just the standard class of polyhedrons might be a relatively simple task, additional processing possibilities present changes to the shape distribution that vary based on the functional constraints of the applied process. This type of constraint-based change in morphology is fundamentally similar to the learnable constraints for other functional object classes, and building off this conceptual relationship is the premise for this investigation.

Here, we develop a novel network architecture that learns the shape of volumetric data with functional constraints and applies it to the problem of grain morphology of polycrystalline metals. This approach integrates recent advances in generative adversarial networks (GANs) [16] into the materials domain to generate new grain morphologies based on the distributions of training data. As grains are stochastic




Fig. 1 Grain distributions across materials and processes: 3D datasets collected from Ti–6Al–4V (top left), strontium titanate (top right), and additively manufactured (AM) Inconel 718 superalloy (bottom) samples with detailed examples of specific grains contained within. Data shown were gathered with 3D electron backscatter diffraction (EBSD). Average grain size is much larger relative to voxel size reso-

lution in the strontium titanate dataset compared the Ti–6Al–4V dataset, resulting in fewer total grains, but better defined grain facets. The Inconel 718 sample and the grain contained within the melt pool on the top surface of the sample are elongated due to the directional heat flux in the AM process

and valid morphologies are not always recognizable, network training and refinement was initially performed on functional objects (e.g., chair, guitar, etc.) from a publicly available dataset built for deep-learning-based object recognition, and the network was subsequently trained to generate Ti–6Al–4V microstructures. The resultant grain morphologies were evaluated for adherence to the shape constraints of actual polycrystalline metals. The results of this analysis are discussed in detail, and the limitations and potential benefits of this type of network-based approach are discussed in the context of the physics of crystalline materials.

Background

Previous Studies in 3D Object Generation

As this investigation builds on the foundations of machine learning for 3D object recognition and synthesis, we will briefly describe relevant established approaches. Until recently, the majority of studies on 3D object generation have focused on the retrieval or combination of components of the object of interest, such as in [17–19]. In these approaches, for a given database of shapes, a probabilistic graphical model learns the geometric and semantic relationships that will yield a stylistically compatible object. Taking this a step further, Wu et al. [20] represented geometric 3D shapes as probability distributions of binary variables on a 3D voxel grid and was able to successfully demonstrate shape completion from “2.5-dimensional” depth maps. A related investigation from [21] proposed a network that used the ShapeNet dataset to learn a mapping from 2D images

to their underlying 3D shapes, enabling the generation of a 3D representation of an object from an input 2D image. This led to efforts from [22–24] and [25] to generate 3D representations from 2D images. Although these methods showed encouraging results, most relied on some form of human supervision and did not focus on directly generating 3D objects from latent space, which can be thought of as a compressed data space used by the network to describe different possible output of interest. In this case, a lack of focus on latent space focus means these methods were not as well-suited for the recognition and synthesis of broadly described stochastic objects like grain morphologies. In the context of unsupervised 3D object synthesis, [26] proposed a promising autoencoder-based network to learn a deep embedding of object shapes, which yielded then state-of-the-art shape completion results, but overall, supervised approaches have historically had greater success with libraries of readily distinguishable objects.

When considering GAN-based approaches for shape generation, the most relevant approaches are from [27, 28], which focus on 3D shape generation from a probabilistic latent space with feature-based learning and adversarial loss components. The 3D-GAN network from [27] generates 3D objects from a low-dimensional latent space, thus allowing for both the sampling of objects without a reference image or CAD model as well as the exploration of the 3D object manifold. Similarly, [28] proposed a novel 3D GAN network, but supplemented it with a 2D image enhancer network. This enhancer network was able to effectively learn and feed image features into the 3D model generator to synthesize high-quality 3D data. While both of these networks deliver on their promise to provide a solution to the 3D model

generation problem, they offer somewhat limited resolution and detail in the shapes that they generate. Furthermore, enhancement of three-dimensional object results using two-dimensional supplemental images is not readily achievable with images of materials microstructure because most, if not all, 2D microstructure imaging techniques offer only a planar image of surface microstructure with minimal depth of field, no photographic perspective, and no real means by which underlying 3D microstructure can be inferred.

Materials-Specific Approaches

Within the materials domain, the earliest computational approaches used to generate microstructural morphologies derive from physics-based models, which have been recognized for both their high level of detail and realistic output. A study by [29] used a Voronoi tessellation model that simulated ceramic grain boundary evolution based on well-established equations developed by [30–32]. Additionally, research approaches by [33] combined Monte Carlo simulations and grain growth kinetics to model metal crystallization. These physics-based models are promising, but often require significant computational power and detailed knowledge of the energetics of the system. These are not always readily available, especially in more complex processing scenarios. To avoid this knowledge and computational burden, models have also been developed to generate microstructures based primarily on their statistical properties. For example, [34, 35] used statistical descriptors to generate microstructures using tessellation and ellipsoid coarsening, respectively. A method for the robust comparison of synthetic microstructures (e.g., generated by the model of [35]) with experimental results is described in [36], with the finding that certain morphological descriptors resulted in improved discrimination, and that the starting shapes for grain generation impacted different descriptors. Approaches such as these are computationally efficient and versatile, compared to physics-based models, but tend to result in less realistic grain morphologies and limited accuracy in local grain environment descriptions.

Beyond physics-based models, machine learning has been employed to explore the material microstructure space of non-crystalline materials. [37] used graph neural networks of 3D atomic arrangements to describe general amorphous structures, such as those found in glassy materials. These graphical descriptions were then used as predictive tools to explore how the observed structure affects mobility and resultant glass properties. This approach showed promise for glasses, but does not extend to the crystalline domain. In the context of more general two-phase microstructures, [38] used 3D convolutional neural networks to characterize possible stochastic microstructures made from filtered noise. These types of stochastic microstructures have also been

investigated in 2D by [39] using classification trees and by [40] using GANs.

Machine learning techniques have also been previously applied specifically to experimental crystalline structures. Investigations by [41] compiled a database of 2D ultra-high-carbon steel micrographs and classified these images into microstructures based on the distributions of their microstructural features [42]. In [43], a Wasserstein GAN with gradient penalty was used to generate 2D microstructures using the database of [41] as training data. Related to this investigation, [44] used StyleGAN to generate various 2D microstructures, and [45] used a convolutional deep belief network to generate 2D microstructures of the same titanium alloy explored in this study. Approaches such as these have strong experimental underpinnings; however, validation of microstructures is difficult in 2D. The challenge in 2D microstructure assessment stems from the fact that material micrographs are planar images of solid material that offer very little or no perspective-based information from which 3D appearance can be inferred. For quantitative evaluation, the challenge becomes virtually impossible, as establishing 3D structure from 2D for anisotropic shape classes is mathematically intractable [46]. In order for the output from machine learning-based models to be fully comparable to material microstructures, generation and evaluation must be made against robust 3D microstructure information.

Noteworthy advances have also been made in 3D crystalline material generation and evaluation in property space. [47] used structural optimization to explore the microstructure space of the iron-gallium alloy Galfenol. This approach focused on techniques for optimizing grain arrangement in the Galfenol microstructure to achieve an orientation distribution that would improve desired properties. While this approach allows for extensive exploration of property space, the basis of the model is theoretical, and there are no means of verifying that the proposed microstructures can be physically realized with available processing techniques. More recently, [48] used GAN-based approaches to generate three-phase microstructures of solid oxide fuel cells, which were then evaluated against experimental results and similar structures generated by the statistical methods of [35]. This approach evaluated microstructures both at the feature and property level and demonstrated strong representation and promising results. As the focus of the network was to generate distributions of the three phases present in the fuel cell anode, namely yttria-stabilized zirconia, nickel, and pore/void, it was tailored around producing phase distributions specifically and did not address features within those phases. Often, deep learning approaches rely on some form of experimental ground-truth data, and for the investigation presented here, the experimental ground truth was gathered using a technique known as electron backscatter diffraction (EBSD). EBSD is a scanning electron microscopy technique

where electrons from the microscope beam are diffracted by a crystalline structure according to Bragg's law [49], forming a pattern that can be indexed to determine information about the underlying structure. This approach is a well established crystallography technique whose capabilities are discussed in [50] and [51]. Typically, the first step in processing EBSD data is indexing, where the diffraction pattern gathered from the microscope is mapped to a particular crystal structure and orientation. The most conventional method of indexing is mapping using a Hough transform, but other more sophisticated mathematical and machine-learning-based approaches, such as those presented by [52–55], allow for more efficient and higher quality indexing with reduced error. Because indexing is a mapping of a diffraction patterns to specific locations, all EBSD data are inherently in pixel or voxel form. After the data are indexed, additional material information can be extracted based on the relative orientations and deformation states of the mapped crystalline grain structures. Because this study is interested in grain morphology, all ground-truth data used in this study are fully indexed and segmented into grains based on continuity of crystalline domains. However, in the investigation presented here, much of the novelty arises not from the demonstrated experimental application, but rather from the extensibility of the approach. By learning morphologies from general distributions of relevant features, a network can be produced that is applicable to a wide variety of materials systems and applications with no needed architectural refinement.

Approach

In this investigation, we demonstrate the GAN-based recognition and synthesis of crystalline grain morphologies using learned feature distributions built on the same fundamental principles used for the generation of other functional objects. We start with readily available databases of easily recognized functional objects to refine the GAN's synthesis

capabilities and evaluate network performance on recognizable objects. We then apply the established architecture to materials systems, more specifically grain morphologies, where the scope of accessible data is limited and recognition is not straightforward. While this is not directly equivalent to transfer learning, it demonstrates how available information databases might be used to teach relevant broader concepts to networks, which, in the right context, can be applied to scientific problems where large-scale datasets are not easily obtained.

Architecture and Training

First proposed by [16], a simple GAN consists of a generator G and a discriminator D . The generator tries to synthesize samples that look like the training data, while the discriminator tries to determine whether a given sample is a real sample that originated from the ground-truth data distribution or is synthesized from the generator. The discriminator D then outputs a confidence value $D(x)$ of whether input x is real or synthetic. A basic layout of the GAN architecture is shown in Fig. 2. StyleGAN, developed by [56], garnered widespread attention for its life-like image quality and unsupervised high-level attribute separation in the generated output. Because of these characteristics, we used StyleGAN as the base architecture for our network. Instead of passing a random noise vector z to the generator, z is first mapped to an intermediate latent space W , which is transformed into spatially invariant styles $\mathbf{y} = (\mathbf{y}_s, \mathbf{y}_b)$. This is then used to control the generator through adaptive instance normalization (AdaIN) at each convolutional layer. The AdaIN is described by [56] using the form:

$$\text{AdaIN}(\mathbf{x}_i, \mathbf{y}) = \mathbf{y}_{s,i} \frac{\mathbf{x}_i - \mu(\mathbf{x}_i)}{\sigma(\mathbf{x}_i)} + \mathbf{y}_{b,i} \quad (1)$$

where $\mathbf{x}_i$ is a feature map normalized separately for each instance i and then scaled and biased using the




Fig. 2 Generative adversarial network (GAN) architecture overview: a layout of the architecture of a traditional GAN, as described by [16]

corresponding scalar components y_s and y_b from style y . Mapping to intermediate latent space allows for disentanglement, which provides clearer tracking of the influence of latent space variations on the generated output. For generative architectures, latent space is the random source noise used as input, and disentanglement, as described by [56], allows for the fluctuations in this space to be separated and applied independently for different generated features. This helps produce better results in scenarios where features are not represented proportionally to the initial sampling, such as for the anisotropic feature distributions expected for both the functional objects and experimental material data investigated here. Additionally, the use of style-based representation allows for better evaluation of network understanding at different feature levels, which is particularly useful in stochastic structures of variable size, such as grains.

In our 3D generative adversarial network (M-GAN), an initial latent space vector z of size 512 is chosen through Gaussian sampling of latent space, similar to the work of [56]. In their work, [56] determined that a size of 512 for z was sufficiently large to keep all components of latent space disentangled. The generator G then maps this latent vector into an intermediate latent space $\mathcal{W}$ using a mapping network of 8 fully connected layers with a LeakyReLU activation function after each layer. The output $\mathcal{W}$ is then converted into the styles previously described as y using a learned affine transformation, denoted as A in Fig. 3, and these styles are used to control adaptive instance normalization (AdaIN) in the synthesis network, as described in Eq. 1. This localization of styles better preserves the small scale physical features of each object or grain, allowing for more detailed morphology output compared to other networks [27, 28].

Our synthesis architecture diverges from that of StyleGAN by using blocks consisting of a 3D deconvolution passed through an AdaIN operation and ReLU activation function. The synthesis network has a total of five blocks, with the first having a constant input vector and normalization, and the fifth having a sigmoid activation function instead of a ReLU. In comparison, StyleGAN blocks use upsampling followed by two alternating convolution layers and AdaIN operations, with noise introduced after each convolution to add stochastic variation. However, in 3D, this noise was found to add instability during training and was therefore not included in the M-GAN network. Architecturally, the discriminator is similar to that of [27] in that it consists of 5 progressive 3D convolution layers with no downsampling. Each of these convolution layers has a LeakyReLU activation function except for the last layer, which has no activation function. The first four layers have a kernel size of 4, stride size of 2, and padding size of 1, and the last block has kernel size of 4 and stride size 2 with no padding. The output of this last block goes directly into the loss function. In this investigation, the discriminator makes its decisions (real vs. generated) based on collections of multiple objects from the same class, rather than just singular instances, similar to the method presented by [57].

The data used to train this network consisted of a combination of object-based data for network refinement and evaluation as well as fully indexed EBSD data for application to material microstructures. For object-based data, six major categories from the ShapeNet Dataset by [58] were used in the network training: car, chair, plane, guitar, sofa, and rifle. Each category has a 128-sample training set. All objects are presented in a voxel-based format, and each individual




Fig. 3 Generator architecture of M-GAN: a mapping network comprised of eight fully connected (FC) layers with Leaky ReLU activation function after each FC layer takes as input a 512-dimensional latent vector z . The output is then mapped to an intermediate latent space $\mathcal{W}$, converted into styles using a learned affine trans-

formation (A), and passed through an AdaIN operation for each of the five blocks in the synthesis network. Block 1 passes constant input through AdaIN and ReLU activation functions, while Blocks 2–5 are deconvolution blocks progressively grown from $8 \times 8 \times 8 \rightarrow 64 \times 64 \times 64$

object instance is contained within of $64 \times 64 \times 64$ voxel volume. The material microstructure data used in this study was 3D EBSD gathered by Wendorf et al. [59] using a method developed by [60] known as the Tribeam system, which performs rapid serial sectioning using ablation with ultrashort pulse femtosecond lasers. The material investigated here is a polycrystalline wrought titanium alloy containing 6.75 wt% aluminum and 4.5 wt% vanadium. This alloy, commonly referred to as Ti–6Al–4V, has applications in turbine engines, aerospace, and medical prostheses. The grain information is gathered at the voxel level using electron backscatter diffraction, with each voxel having a size of $1.5 \mu\text{m} \times 1.5 \mu\text{m} \times 1.5 \mu\text{m}$. Each grain is passed to the network inside a cubic volume of size $64 \times 64 \times 64$ voxels. An image of the full dataset is shown in Fig. 4. This Ti–6Al–4V dataset was originally gathered to study slip behavior and long-range plastic deformation across multiple length scales [61] as well as explore how microtextured regions in titanium alloys affect overall mechanical response [59, 62]. To better capture long-range mechanical response, the sample size is relatively large compared to the grain size, so many grains are relatively small compared to the voxel resolution and detailed grain facets are not easily resolved. Despite these resolution limitations, this set has the major advantage of offering a large number of grains as a good foundation for network training. In total, 84,215 grains from this dataset were used for network training.

Stable training of the GAN was achieved by setting the learning rate for both the generator and discriminator at 0.0002, with a batch size of 16. Similar to [27], the discriminator in this approach is updated five times for each generator update and employs the Adam optimizer described by [63] with $\beta = 0.5$. The output of the network is in a 64×64

$\times 64$ -dimensional space. The Wasserstein loss as described in [64] with a gradient penalty is used for the discriminator and the generator, which are defined as follows:

$$L_{D-\text{Loss}} = -\frac{1}{m} \sum_{i=1}^m D_{\phi}(x^{(i)}) + \frac{1}{m} \sum_{i=1}^m D_{\phi}(G_{\theta}(z^{(i)})) + \frac{1}{m} \sum_{i=1}^m \lambda \left(\left\| \nabla_{\hat{x}^{(i)}} D_{\phi}(\hat{x}^{(i)}) \right\|_2 - 1 \right)^2 \quad (2)$$

$$L_{G-\text{Loss}} = -\frac{1}{m} \sum_{i=1}^m D_{\phi}(G_{\theta}(z^{(i)})) \quad (3)$$

Here, G_{θ} is the generator network, D_{ϕ} is the discriminator network, $z^{(i)}_{i=1}^m \sim p(z)$ is a batch of random noise from latent space ($p(z)$ = population distribution for latent space), $x^{(i)}$ is instance i of real data from a batch of size m such that $x^{(i)}_{i=1}^m \sim \mathbb{P}_r$ ($\mathbb{P}_r$ = population distribution of real data), ∇ is the gradient operator, λ is a gradient penalty coefficient as described by [65], and $\hat{x}^{(i)}$ is defined as $\epsilon x^{(i)} + (1 - \epsilon)G_{\theta}(z^{(i)})$ where ϵ is a uniform random variable in $[0, 1]$. This Wasserstein loss metric, also known as the “earth-mover” metric, quantitatively compares the difference between two distributions by considering the area under each distribution curve and measuring how much area would have to be moved from one distribution for it to match the other. When applied to feature distributions, it is a measure of how well the features in the generated data match those found within the feature distribution of the ground-truth dataset. This type of comparison uses the ground-truth data as guide for training without relying directly on one-to-one mappings between generated and ground-truth data, which enables the network to better produce results that are similar to the ground truth without being direct copies. By definition, this lack of a one-to-one mapping also means this approach is considered an unsupervised learning method. Furthermore, since this approach is unsupervised, all hyperparameters used in this study are preset before training, and all validation of data is done as analysis on objects generated by the network during inference time.

Evaluation

We evaluate our network across several areas. First, we show qualitative results of generated 3D objects from the ShapeNet dataset by [58]. Then, we evaluate the unsupervised learned representations from the discriminator by using them as features for 3D object classification. Next, quantitative results are shown on the popular benchmark ModelNet dataset from [20]. Finally, we apply the network approach




Fig. 4 Experimental training set: Ti–6Al–4V dataset, whose grains were used to train M-GAN on grain shape recognition. Sample is shown to scale in inverse pole figure (IPF) color. The IPF color indicates the orientation for each individual voxel, and these orientations are used in the segmentation process to identify individual grains, which are subsequently used for training

to grain morphologies from real, experimentally gathered materials data. Because grain shapes are stochastic, direct recognition of resultant objects is not straightforward, so shape quality must be validated by other means. Here, shapes are validated using the moment invariant approach described in [66].

We train one M-GAN for each object category using a 512-dimensional random vector that follows a normal distribution with mean 0 and variance 0.2. We compare our generated objects with [27], because [28] used an enhancer network and additional information from 2D rendered images during training. Our network synthesizes high-resolution 3D objects ($64 \times 64 \times 64$) with detailed shape information trained from only 3D input. To ensure that the network is not simply memorizing the training data, we analyze synthesized objects using comparisons to nearest neighbors in discriminator feature space as described by [27]. For this analysis, discriminator feature vectors are captured for both the generated samples and ground-truth data, and nearest neighbors in feature space are determined using minimum ℓ_2 distance, which is far more efficient than direct 3D object comparison. Feature analysis shows that the generated samples were not identical to their nearest neighbors.

To evaluate the unsupervised learned features from our network and assess performance, we use the feature extraction approach discussed in [27] to provide a means of comparison, as there is no established standard. In this approach, features identified in the convolution layers of the discriminator are collected for a variety of different categorically classified objects. This library of features is then integrated into a support vector machine whose classification performance is tested using a new dataset of equivalently classified objects. We train our M-GAN network on six object categories (bed, car, chair, plane, sofa, table) from [27]. Each object category has 25 samples, similar to [27] in the training set from the ShapeNet dataset.

Next, we use the ModelNet dataset from [58] to evaluate the unsupervised features learned by our network. The ModelNet dataset has two categories: ModelNet10 (10 classes) and ModelNet40 (40 classes). ModelNet10 has a total of 3991 training samples and 908 test samples. From this dataset, we use 100 samples from each category, totaling 1000 training samples. ModelNet40 has a total of 9843 training samples and 2468 test samples. From this dataset, we use up to 100 samples from each category when available, totaling 3906 training samples, as some categories had fewer than 100 samples. For ModelNet10 and ModelNet40, we use the entire test dataset provided, though fewer training samples than were used by [27, 28].

When evaluating feature extraction accuracy, to provide a fair comparison, we use the same kernel size = {8, 4, 2} defined by both [27] and [28], as well as the defined stride = {4, 2, 1} by [28]. We calculate features from the second,

third, and fourth layers of the trained discriminator, which are then concatenated after applying max-pooling with the defined kernel size and stride. We then train a linear Support Vector Machine (SVM) on training features and calculate classification accuracy on the test features.

Functional object data from the ShapeNet and ModelNet sets are well suited for network refinement and performance assessment, in part due to their easy visual recognition. Grain shapes, on the other hand, have a stochastic morphological appearance, so they cannot be easily evaluated by visual means or by direct object comparison. For this reason, we use statistical distribution comparison of 3D moment invariants to assess the quality of generated results for single grains. Much like the image moments used in 2D analysis, 3D moment invariants are integration-based descriptors that numerically quantify an object based on the distribution of its solid volume. These types of invariants have been used previously by [67] as shape descriptors for general 3D objects, and in materials science by [66, 68] to describe the shapes of particles such as grains and inclusions. Following the approach of [66], Cartesian moment descriptors μ_{pqr} take the form:

$$\mu_{pqr} = \int \int \int x^p y^q z^r F(r) dr \quad (4)$$

where $F(r)$ is a characteristic function that has a value of one in material regions and a value of zero in void regions. As all ground-truth data are in voxel-based form in this investigation, all integrations herein are calculated as Riemann sums in voxel space rather than as continuous integrals. Following this Cartesian moment form, μ_{000} is directly equal to the volume V of an object, and the centroid (X_c, Y_c, Z_c) of an object can be expressed in the following form:

$$(X_c, Y_c, Z_c) = \left(\frac{\mu_{100}}{V}, \frac{\mu_{010}}{V}, \frac{\mu_{001}}{V} \right) \quad (5)$$

Working from a coordinate space originating at the centroid, the non-normalized moment invariants $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ are as follows:

$$\mathcal{O}_1 = \mu_{200} + \mu_{020} + \mu_{002} \quad (6)$$

$$\mathcal{O}_2 = \mu_{200}\mu_{020} + \mu_{200}\mu_{002} + \mu_{020}\mu_{002} - \mu_{110}^2 - \mu_{101}^2 - \mu_{011}^2 \quad (7)$$

$$\mathcal{O}_3 = \mu_{200}\mu_{020}\mu_{002} + 2\mu_{110}\mu_{101}\mu_{011} - \mu_{200}\mu_{011}^2 - \mu_{020}\mu_{101}^2 - \mu_{002}\mu_{110}^2 \quad (8)$$

These can be normalized to object volume to produce the moment invariants $(\Omega_1, \Omega_2, \Omega_3)$:

$$(\Omega_1, \Omega_2, \Omega_3) = \left(\frac{3V^{5/3}}{O_1}, \frac{3V^{10/3}}{O_2}, \frac{V^5}{O_3} \right) \quad (9)$$

The distributions of these moment invariants are used to evaluate the shape quality of the generated grains. Additional analysis of 3D moment invariants can be found in [66].

Results and Discussion

Object-Based Data

For object-based datasets, the network showed strong performance. Examples of generated object output based on ShapeNet training are shown in Fig. 5, and network performance against existing architectures is shown in Table 1. With fewer training samples than [27, 28], we achieved 92.29% accuracy on the ModelNet10 dataset and 85.08% accuracy on the ModelNet40 dataset. If we use a comparable training set size to [27], we achieve a 2.29% improvement on ModelNet10 and a 3.78% improvement on ModelNet40 dataset. Furthermore, [28] uses all available training samples as well as additional rendered 2D images for both datasets, compared to our use of fewer training samples and only 3D input.

Since all benchmarking for M-GAN was done on recognizable objects with well-defined orientations, no rotational augmentation of datasets was used during training with objects or grains, and all results are shown in their as-generated orientation. Analysis of generated objects indicates that the M-GAN network is not generating any rotated objects in cases where rotation effects are distinguishable. For microstructural objects like grains, since they are fundamentally stochastic and have no defined top or bottom, rotation augmentation may prove beneficial as an additional

Table 1 Comparison of object-based performance for the ModelNet10 and ModelNet40 datasets using the feature extraction approach described by [27]


Method (supervised)	ModelNet10 (%)	ModelNet40 (%)
3D ShapeNets [20]	83.54	77.32
VoxNet [69]	92.00	83.00
Geometry image [70]	88.40	83.90
PointNet [71]	77.60	–
GIFT [72]	92.35	83.10
FusionNet [73]	93.11	90.80
Method (unsupervised)	“	“
SPH [74]	79.79	68.23
LFD [75]	79.87	75.47
VConv-DAE [26]	80.50	75.50
3D-GAN [27]	91.00	83.30
3D-GAN (≈ 100 samples)	90.00	81.30
Ours (M-GAN) (100 samples)	92.29	85.08

training tool, especially for anisotropic grain structures, but the influence of rotational augmentation is not explored in this work.

Material Data

Following functional object generation, the network was trained on experimentally gathered Ti–6Al–4V grain morphologies. Representative volumes of grains generated using the M-GAN are shown in Fig. 6. Unlike the benchmark results, which are shown as contour displays, these figures are shown in voxel form, which matches the representation of the ground-truth data. It should be noted that some generation artifacts were present in $64 \times 64 \times 64$ volume containing these grains, with the most common of these artifacts being filled single voxels at random locations in the render

Fig. 5 ShapeNet output: the M-GAN network generates detailed shape information for a diverse range of 3D objects. As the network learns from shape feature distributions, these objects are similar but not identical, such as the different styles of chairs shown in the bottom row. The variation within a class of generated objects that satisfies functional requirements is applied to the generation of grains across different material and processing classes, as shown in Fig. 1






Fig. 6 Visual comparison of grains: sampling of grains from real and generated sets are shown to verify visual similarity. Voxel size is $1.5 \mu\text{m} \times 1.5 \mu\text{m} \times 1.5 \mu\text{m}$. For generated grains, only largest connected

component is shown. Stochastic nature of grains means no directly matching shapes are expected. Shape distributions are compared in Fig. 7

volume. Some examples of these artifacts can be seen in the contour images of the benchmark set shown in Fig. 5. In the case of grain volumes shown in Fig. 6, only the largest connected volume is shown. These representative volumes show shapes that would be reasonably expected of grains in an equiaxed polycrystalline metal like wrought Ti–6Al–4V, but these representative volumes alone are insufficient verification of generated grain quality.

For quantitative shape evaluation, the distributions of the 3D moment invariants ($\Omega_1, \Omega_2, \Omega_3$) evaluated in this study are shown in Fig. 7. The ground-truth dataset contained 84,215 grains and the M-GAN generated a total of 150,000 grains. For each of the three invariants, less than 0.2% of the grains had values that were either infinite or nonphysical across both the ground truth and generated sets. In observed cases, this resulted from very small grain structures that were linear or planar in nature, which led to extremely large errors in Riemann summation during the calculation of the

invariants. Due to their nonphysicality, these values were omitted from the statistical distribution comparisons.

A comparison of the experimental moment invariant distributions in this investigation compared to those analyzed in [36] shows that ones presented here are of slightly lower average value, which indicates grains being more cube-like in shape. This is expected though, given that the grains in this set are of lower resolution in order to capture a large volume, and lower resolution will mean more cube-like grain features. The means and standard deviations for each of the distributions are shown in the insets of Fig. 7. These values show that the generated data consistently exhibited lower means and larger standard deviations. Comparison of the distribution shapes in Fig. 7 reveals the cause to be that the generated data have a range of moment invariants values in the region close to zero that is not seen in the real dataset.

Analysis of generated grains from these lower magnitude moment invariant values revealed two common issues. The




Fig. 7 Quantitative evaluation of M-GAN grain generation: moment invariants ($\Omega_1, \Omega_2, \Omega_3$) are used to evaluate the similarity of real and generated shape distributions, as grains shapes are stochastic and cor-

rectness is a matter of shape probability rather than object recognition. Histograms are shown for both real (top) and generated (bottom) grain shapes. Distribution means are indicated by vertical lines

first was that many of the generated grains in this region were smaller in volume than the ground truth. In the ground-truth dataset, the smallest grains that could be reliably distinguished from artifacts and noise were on the order of 25 voxels in volume, but many of the generated grains with low moment invariant values had volumes of fewer than 20 voxels. While small volume alone is not necessarily a concern, these smaller grains had a greater negative interaction with the other major issue of generator noise, which was also noticed in the some object-based data as in Fig. 5. Although the generator noise was not displayed in Fig. 6 for visual clarity, this noise was still included in the moment invariant calculation both to minimize quantitative bias and because no steps were included in the architecture or training process to directly mitigate noise. Noise removal with the aid of domain knowledge is a widely accepted practice for microstructural data, but since the goal of this work is direct evaluation of network performance, denoising is considered a post-processing step that is outside the scope of the network. This consideration of noise is critical because some generated grains had very small volumes compared to the ground truth, and many more still had relatively small volumes compared to the $64 \times 64 \times 64$ voxel region in which they were generated. Thus, generator artifacts located far away from the grain itself can cause significant shifts in the centroid location used to establish the coordinate basis for the calculation of the moment invariants, thereby creating distortions in the data. However, beyond these factors, the overall shape and position of the main peaks for both distributions are very similar, indicating that the M-GAN network is producing data that is similar to the ground-truth data without directly replicating it.

Relation to Material Physics

While a quantitative analysis of morphology shows promising results, network performance must be considered in the context of actual material physics. Here, the issue of small grain volumes brings to light another challenge with GAN-based morphology synthesis, namely that of size–shape relationships. Throughout object training, the network is constrained to a $64 \times 64 \times 64$ voxel final output size, regardless of the object morphology being generated. As such, it learns to recognize morphological features independent of the object size, and voxels are assigned whatever physical length scale is necessary to render the object in the available volume. Thus, when considering, for example, cars and guitars, the relevant features on both of these objects are learned in the context of the available rendering scale, even though the physical size of the guitar and car are different.

In many scenarios, size-independent morphological recognition is an asset that offers greater versatility. However, for materials data, it has some undesired consequences. In

the previously discussed case of low-volume grains, the network generated objects smaller than the experimental resolution of the dataset, even going as far as to produce an instance of a grain with single-voxel volume. While these results are inaccurate with respect to the ground-truth data, this is primarily due to voxel-based resolution limitations rather than physical inaccuracies. Even for the most extreme case of a single voxel grain, if we apply the ground-truth length scale of $1.5 \mu\text{m} \times 1.5 \mu\text{m} \times 1.5 \mu\text{m}$ per voxel, it is certainly possible for a grain with a diameter of $1.5 \mu\text{m}$ to exist in a variety of real materials. However, regardless of material, that grain will possess morphological constraints that are not represented by a single-voxel rendering. This tells us the network has challenges with distinguishing feature constraints associated with physics from those associated with resolution limitations. This is a challenging problem to address in EBSD images, where data are gathered using a rastering electron beam, and therefore will always be pixel- or voxel-based. Currently, the most promising outlook for resolving this issue is an automated means of enhancing 3D images with resolution artifacts, such as super-resolution, a survey of which is presented by [76]. With sufficient resolution enhancement, small-scale grains could be presented to the network at a volume scale where physical features can be readily recognized beyond resolution limitations. It may be possible to remove some of these resolution limitations using either networks that produce non-voxel output, or networks that allow for continuous resolution scaling, such as Infinity GAN by [77]. However, achieving meaningful results on materials data with either of these approaches would require detailed 3D information across a variety of different length scales, which is nontrivial to obtain.

The other challenge that must be addressed for crystalline materials is grain connectivity relationships. For both the object and material data, all learning was performed with independent renderings of single objects. Recognition of individual objects regardless of context can be beneficial; however, in polycrystals, the shape of each grain depends on physical interaction with its surrounding neighbors. The network presented here has a good understanding of grain morphology, but cannot infer how different grain shapes are related to one another in the bulk material. It is possible that graph- or label-based approaches may offer a means to clarify these 3D connectivity relationships, but to date, the available resources for achieving this are limited.

The presented M-GAN approach demonstrates how fundamental materials concepts, such as grain morphology, can be learned by network-based approaches in cases where an abundance of data is not readily available. Grain morphology is addressed herein as it is a principal component in understanding the relationships between microstructure and material properties. Network-based approaches such as M-GAN enable the generation of material features (e.g.,

grain morphologies) in a generalizable context. This broad generative capacity expands our exploration of the processing–structure and structure–property relationships that are critical to materials discovery and development. For example, the M-GAN approach is directly applicable to the generation of bulk synthetic polycrystalline microstructures from constituent components. M-GAN can readily be used to generate collections of grain morphologies based on a distribution of interest, and these morphologies can be iteratively packed and assigned crystallographic orientations to form a bulk polycrystalline solid. Network-based approaches like M-GAN also serve as a foundation to explore morphology-dependent relationships, such as the formation of microtextured regions in titanium alloys [78–80], or the formation of thermomechanically dependent grain structures, such as columnar grains from directional melt solidification in additive manufactured metals [81–83].

Conclusions

We present M-GAN, a generative adversarial network, to recognize and synthesize 3D grain morphologies in crystalline microstructures. All objects are generated from latent space vectors without any supplemental 2D input.

Network capabilities are demonstrated on the ModelNet benchmark datasets and on an experimental 3D material dataset of Ti–6Al–4V. On benchmark datasets, the M-GAN yielded more reliable discriminator classification with exposure to fewer objects, highlighting the quality of its discriminator feature recognition. On experimental data, the M-GAN network produced results that were morphologically similar to ground truth without being a direct replication. The effects of generator noise produced some variation in the moment invariant distribution compared to experimental data, particularly at low value invariants, and most often with grains of smaller volume. These challenges with smaller grains indicate a limitation of the network in distinguishing grain shape features from spatial resolution artifacts.

This approach lays a foundation for the use of network based approaches to understand key relationships between grain structure and properties of crystalline materials. In future work, this approach could also be extended to relate grain shape to neighbor connectivity in the bulk material structure.

Acknowledgements All authors gratefully acknowledge financial support from the National Science Foundation (Award: 1934641) as part of the HDR IDEAS Institute. In addition, Manjunath acknowledges support from the NSF SSI Award 1664172 that is supporting the core BisQue software development infrastructure for sharing data and analysis modules. The authors sincerely thank J. C. Stinville, A. Polonsky, and J. Wendorf for collection of the 3D Ti–6Al–4V datasets. The MRL Shared Experimental Facilities are supported by the MRSEC Program

of the NSF under Award No. DMR 1720256; a member of the NSF-funded Materials Research Facilities Network (www.mrfn.org). On behalf of all authors, the corresponding author states that there is no conflict of interest associated with this publication.

References

1. National Research Council (2008) Integrated computational materials engineering: a transformational discipline for improved competitiveness and national security. The National Academies Press, Washington, DC. <https://www.nap.edu/catalog/12199/integrated-computational-materials-engineering-a-transformational-discipline-for-improved-competitiveness>
2. National Science and Technology Council (US) (2011) Materials genome initiative for global competitiveness. Executive Office of the President, National Science and Technology Council
3. Lu XG (2015) Remarks on the recent progress of materials genome initiative. *Sci Bull* 60(22):1966–1968. <https://doi.org/10.1007/s11434-015-0937-2>
4. M. I. I. M2I (2016) ICMEG: integrated computational materials engineering expert group. <https://www.m2i.nl/portfolio-items/icmeg-integrated-computational-materials-engineering-expert-group/>
5. Jain A, Persson KA, Ceder G (2016) Research update: the materials genome initiative: data sharing and the impact of collaborative ab initio databases. *APL Mater* 4(5):053102. <https://doi.org/10.1063/1.4944683>
6. Draxl C, Scheffler M (2018) NOMAD: the FAIR concept for big-data-driven materials science. *MRS Bul* 43(9):676–682
7. Pennekamp J, Glebke R, Henze M, Meisen T, Quix C, Hai R, Gleim L, Niemietz P, Rudack M, Knape S, Epple A, Trauth D, Vroomen U, Bergs T, Brecher C, Buhrig-Polaczek A, Jarke M, Wehrle K (2019) Towards an infrastructure enabling the internet of production. In: Proceedings of the 2019 IEEE international conference on industrial cyber physical systems, ICPS 2019
8. Jensen DJ, Lauridsen EM, Margulies L, Poulsen HF, Schmidt S, Sørensen HO, Vaughan GB (2006) X-ray microscopy in four dimensions. *Mater Today* 9(1–2):18–25
9. Ludwig W, Reischig P, King A, Herbig M, Lauridsen EM, Johnson G, Marrow TJ (2009) Three-dimensional grain mapping by x-ray diffraction contrast tomography and the use of Friedel pairs in diffraction data analysis. *Rev Sci Instrum* 80(3):033905. <https://doi.org/10.1063/1.3100200>
10. Poulsen H (2012) An introduction to three-dimensional X-ray diffraction microscopy. *J Appl Crystal* 45(6):1084–1097
11. Echlin MP, Mottura A, Torbet CJ, Pollock TM (2012) A new TriBeam system for three-dimensional multimodal materials analysis. *Rev Sci Instrum* 83(2):023701. <https://doi.org/10.1063/1.3680111>
12. Miller MP, Pagan DC, Beaudoin AJ, Nygren KE, Shadle DJ (2020) Understanding micromechanical material behavior using synchrotron X-rays and in situ loading. *Metall Mater Trans A* 51(9):4360–4376. <https://doi.org/10.1007/s11661-020-05888-w>
13. Hata S, Honda T, Saito H, Mitsuhara M, Petersen TC, Murayama M (2020) Electron tomography: an imaging method for materials deformation dynamics. *Curr Opin Solid State Mater Sci* 24(4):100850
14. Rowenhorst DJ, Nguyen L, Murphy-Leonard AD, Fonda RW (2020) Characterization of microstructure in additively manufactured 316L using automated serial sectioning. *Curr Opin Solid State Mater Sci* 24(3):100819
15. Chapman MG, Shah MN, Donegan SP, Scott JM, Shade PA, Menasche D, Uchic MD (2021) AFRL additive manufacturing modeling series: challenge 4, 3D reconstruction of an IN625

- high-energy diffraction microscopy sample using multi-modal serial sectioning. *Integr Mater Manuf Innov* 10(2):129–141. <https://doi.org/10.1007/s40192-021-00212-9>
16. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. In: Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ (eds) *Advances in neural information processing systems*, vol 27. Curran Associates Inc, New York, pp 2672–2680
 17. Chaudhuri S, Kalogerakis E, Guibas L, Koltun V (2011) Probabilistic reasoning for assembly-based 3D modeling. *ACM Trans Graph* 30(4):1
 18. Funkhouser T, Kazhdan M, Shilane P, Min P, Kiefer W, Tal A, Rusinkiewicz S, Dobkin D (2004) Modeling by example. *ACM Trans Graph* 23(3):652–663. <https://doi.org/10.1145/1015706.1015775>
 19. Kalogerakis E, Chaudhuri S, Koller D, Koltun V (2012) A probabilistic model for component-based shape synthesis. *ACM Trans Graph (TOG)* 31(4):1–11
 20. Wu Z, Song S, Khosla A, Fisher Y, Zhang L, Tang X, Xiao J (2015) 3D ShapeNets: a deep representation for volumetric shapes. Technical Report. <http://3dshapenets.cs.princeton.edu>
 21. Choy CB, Xu D, Gwak J, Chen K, Savarese S (2016) 3d-r2n2: a unified approach for single and multi-view 3d object reconstruction. In: *European conference on computer vision*. Springer, pp 628–644
 22. Zhang X, Zhang Z, Zhang C, Tenenbaum JB, Freeman WT, Wu J (2018) Learning to reconstruct shapes from unseen classes. In: *Proceedings of the 32nd international conference on neural information processing systems*, series NIPS'18. Curran Associates Inc., Red Hook, pp 2263–2274
 23. Wu J, Wang Y, Xue T, Sun X, Freeman B, Tenenbaum J (2017) Marrnet: 3d shape reconstruction via 2.5d sketches. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) *Advances in neural information processing systems*, vol 30. Curran Associates Inc., New York, pp 540–550
 24. Noguchi A, Harada T (2019) Rgb-d-gan: unsupervised 3d representation learning from natural image datasets via rgb-d image synthesis. [arXiv:1909.12573](https://arxiv.org/abs/1909.12573)
 25. Nguyen-Phuoc T, Li C, Theis L, Richardt C, Yang YL (2019) Hologan: unsupervised learning of 3d representations from natural images. In: *Proceedings of the IEEE international conference on computer vision*, pp 7588–7597
 26. Sharma A, Grau O, Fritz M (2016) Vconv-dae: deep volumetric shape learning without object labels. In: *European conference on computer vision*. Springer, pp 236–250
 27. Wu J, Zhang C, Xue T, Freeman B, Tenenbaum J (2016) Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In: Lee DD, Sugiyama M, Luxburg UV, Guyon I, Garnett R (eds) *Advances in neural information processing systems*, vol 29. Curran Associates Inc., New York, pp 82–90
 28. Zhu J, Xie J, Fang Y (2018) Learning adversarial 3d model generation with 2d image enhancer. <https://www.aaii.org/ocs/index.php/AAAI/AAAI18/paper/view/16064>
 29. Coster M, Arnould X, Chermant J-L, Moataz AE, Chartier T (2005) A microstructural model by space tessellation for a sintered ceramic: cerine. *Image Anal Stereol* 24(2):105
 30. Johnson W, Mehl R (1939) Reaction kinetics in processes of nucleation and growth. *Trans Metall Soc AIME* 135:416–442
 31. Avrami M (1939) Kinetics of phase change. I general theory. *J Chem Phys* 7(12):1103–1112
 32. Kolmogorov AN (1937) A statistical theory for the recrystallization of metals. *Izvestiya Rossiiskoi Akademii Nauk, Seriya Matematicheskaya*
 33. Nosonovsky M, Zhang X, Esche SK (2009) Related content Scaling of Monte Carlo simulations of grain growth in metals. *Modell Simul Mater Sci Eng* 17:1–13
 34. Quey R, Dawson P, Barbe F (2011) Large-scale 3d random polycrystals for the finite element method: generation, meshing and remeshing. *Comput Methods Appl Mech Eng* 200(17–20):1729–1745
 35. Groeber MA, Jackson MA (2014) DREAM.3D: a digital representation environment for the analysis of microstructure in 3D. *Integr Mater Manuf Innov* 3(1):56–72
 36. Callahan PG, Groeber M, De Graef M (2016) Towards a quantitative comparison between experimental and synthetic grain structures. *Acta Mater* 111:242–252
 37. Bapst V, Keck T, Grabska-Barwińska A, Donner C, Cubuk E, Schoenholz S, Obika A, Nelson A, Back T, Hassabis D, Kohli P (2020) Unveiling the predictive power of static structure in glassy systems. *Nat Phys* 16:448–454
 38. Cecen A, Dai H, Yabansu YC, Kalidindi SR, Song L (2018) Material structure-property linkages using three-dimensional convolutional neural networks. *Acta Mater* 146:76–84
 39. Bostanabad R, Bui AT, Xie W, Apley DW, Chen W (2016) Stochastic microstructure characterization and reconstruction via supervised learning. *Acta Mater* 103:89–102
 40. Li X, Yang Z, Brinson LC, Choudhary A, Agrawal A, Chen W (2018) A deep adversarial learning methodology for designing microstructural material systems. In: *ASME 2018 international design engineering technical conferences and computers and information in engineering*
 41. DeCost BL, Hecht MD, Francis T, Webler BA, Picard YN, Holm EA (2017) UHCSDB: ultrahigh carbon steel micrograph database. *Integr Mater Manuf Innov* 6(2):197–205
 42. DeCost BL, Francis T, Holm EA (2017) Exploring the microstructure manifold: image texture representations applied to ultrahigh carbon steel microstructures. *Acta Mater* 133:30–40
 43. Iyer A, Dey B, Dasgupta A, Chen W, Chakraborty A (2019) A conditional generative model for predicting material microstructures from processing methods. Technical Report
 44. Fokina D, Muravleva E, Ovchinnikov G, Oseledets I (2020) Microstructure synthesis using style-based generative adversarial networks. *Phys Rev E* 101(4):043308
 45. Cang R, Xu Y, Chen S, Liu Y, Jiao Y, Yi Ren M (2017) Microstructure representation and reconstruction of heterogeneous materials via deep belief network for computational material design. *ASME J Mech Des* 139(7):071404. <https://doi.org/10.1115/1.4036649>
 46. Torquato S (2002) *Random heterogeneous materials*, vol 16. Series interdisciplinary applied mathematics. Springer, New York. <https://doi.org/10.1007/978-1-4757-6355-3>
 47. Liu R, Kumar A, Chen Z, Agrawal A, Sundararaghavan V, Choudhary A (2015) A predictive machine learning approach for microstructure optimization and materials design. *Nat Sci Rep* 5(1):1–12
 48. Hsu T, Epting WK, Kim H, Abernathy HW, Hackett GA, Rollett AD, Salvador PA, Holm EA (2021) Microstructure generation via generative adversarial network for heterogeneous, topologically complex 3D materials. *JOM* 73(1):90–102. <https://doi.org/10.1007/s11837-020-04484-y>
 49. Bragg WH, Bragg WL (1913) The reflection of X-rays by crystals. *Proc R Soc Lond Ser A Contain Pap Math Phys Charac* 88(605):428–438
 50. Humphreys FJ (1999) Quantitative metallography by electron backscattered diffraction. *J Microsc* 195(3):170–185
 51. Wilkinson AJ, Britton TB (2012) Strains, planes, and EBSD in materials science. *Mater Today* 15(9):366–376
 52. Chen YH, Park SU, Wei D, Newstadt G, Jackson MA, Simmons JP, De Graef M, Hero AO (2015) A dictionary approach to electron backscatter diffraction indexing. *Microsc Microanal* 21(3):739–752. <https://doi.org/10.1017/S1431927615000756>

53. Lenthe WC, Singh S, Graef MD (2019) A spherical harmonic transform approach to the indexing of electron back-scattered diffraction patterns. *Ultramicroscopy* 207:112841
54. Ding Z, Pascal E, De Graef M (2020) Indexing of electron back-scatter diffraction patterns using a convolutional neural network. *Acta Mater* 199:370–382
55. Kaufmann K, Lane H, Liu X, Vecchio KS (2021) Efficient few-shot machine learning for classification of EBSD patterns. *Sci Rep* 11(1):8172. <https://doi.org/10.1038/s41598-021-87557-5>
56. Karras T, Laine S, Aila T (2019) A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 4401–4410
57. Lin Z, Khetan A, Fanti G, Oh S (2018) Pacgan: the power of two samples in generative adversarial networks. In: *Advances in neural information processing systems*, pp 1498–1507
58. Chang AX, Funkhouser T, Guibas L, Hanrahan P, Huang Q, Li Z, Savarese S, Savva M, Song S, Su H, et al (2015) Shapenet: an information-rich 3d model repository. [arXiv:1512.03012](https://arxiv.org/abs/1512.03012)
59. Hémery S, Naït-Ali A, Guéguen M, Wendorf J, Polonsky AT, Echlin MP, Stinville JC, Pollock TM, Villechaise P (2019) A 3D analysis of the onset of slip activity in relation to the degree of micro-texture in Ti-6Al-4V. *Acta Mater* 181:36–48
60. Echlin MP, Straw M, Randolph S, Filevich J, Pollock TM (2015) The TriBeam system: femtosecond laser ablation in situ SEM. *Mater Charact* 100:1–12. <https://doi.org/10.1016/j.matchar.2014.10.023>
61. Echlin MP, Stinville JC, Miller VM, Lenthe WC, Pollock TM (2016) Incipient slip and long range plastic strain localization in microtextured Ti-6Al-4V titanium. *Acta Mater* 114:164–175
62. Chatterjee K, Echlin MP, Kasemer M, Callahan PG, Pollock TM, Dawson P (2018) Prediction of tensile stiffness and strength of Ti-6Al-4V using instantiated volume elements and crystal plasticity. *Acta Mater* 157:21–32
63. Kingma DP, Ba J (2014) Adam: a method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
64. Arjovsky M, Chintala S, Bottou L (2017) Wasserstein gan. [arXiv:1701.07875](https://arxiv.org/abs/1701.07875)
65. Gulrajani I, Ahmed F, Arjovsky M, Dumoulin V, Courville AC (2017) Improved training of wasserstein gans. [arXiv:1704.00028](https://arxiv.org/abs/1704.00028)
66. MacSleynne JP, Simmons JP, Graef MD (2008) On the use of moment invariants for the automated analysis of 3D particle shapes. *Modell Simul Mater Sci Eng* 16(4):045008
67. Kazhdan M, Funkhouser T, Rusinkiewicz S (2003) Rotation invariant spherical Harmonic representation of 3D shape descriptors. Technical Report
68. Trenkle A, Syha M, Rheinheimer W, Callahan PG, Nguyen L, Ludwig W, Lenthe W, Echlin MP, Pollock TM, Weygand D, Graef M, Hoffmann MJ, Gumbsch P (2020) Nondestructive evaluation of 3D microstructure evolution in strontium titanate. *J Appl Crystallogr* 53(2):349–359
69. Maturana D, Scherer S (2015) Voxnet: a 3d convolutional neural network for real-time object recognition. In: *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp 922–928
70. Sinha A, Bai J, Ramani K (2016) Deep learning 3d shape surfaces using geometry images. In: *Leibe B, Matas J, Sebe N, Welling M (eds) Computer vision – ECCV 2016*. Springer, Cham. https://doi.org/10.1007/978-3-319-46466-4_14
71. Qi CR, Su H, Mo K, Guibas LJ (2017) Pointnet: deep learning on point sets for 3d classification and segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 652–660
72. Bai S, Bai X, Zhou Z, Zhang Z, Jan Latecki L (2016) Gift: a real-time and scalable 3d shape search engine. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 5023–5032
73. Hegde V, Zadeh R (2016) Fusionnet: 3d object classification using multiple data representations. [arXiv:1607.05695](https://arxiv.org/abs/1607.05695)
74. Kobbelt L, Schröder P, Kazhdan M, Funkhouser T, Rusinkiewicz S (2003) Rotation invariant spherical harmonic representation of 3d shape descriptors. In: *Proceedings of the 2003 Eurographics*, vol. 43
75. Chen DY, Tian XP, Shen YT, Ouhyoung M (2003) On visual similarity based 3d model retrieval. In: *Computer graphics forum*, vol 22, no. 3. Wiley, pp. 223–232
76. Wang Z, Chen J, Hoi SCH (2020) Deep learning for image super-resolution: a survey. Technical Report
77. Lin CH, Lee HY, Cheng YC, Tulyakov S, Yang MH (2021) InfinityGAN: towards infinite-resolution image synthesis. [arXiv:2104.03963](https://arxiv.org/abs/2104.03963)
78. Gey N, Bocher P, Uta E, Germain L, Humbert M (2012) Texture and microtexture variations in a near- α titanium forged disk of bimodal microstructure. *Acta Mater* 60(6–7):2647–2655
79. Pilchak AL, Szczepanski CJ, Shaffer JA, Salem AA, Semiatin SL (2013) Characterization of microstructure, texture, and microtexture in near-alpha titanium mill products. *Metall Mater Trans A* 44(11):4881–4890. <https://doi.org/10.1007/s11661-013-1804-x>
80. Semiatin SL (2020) An overview of the thermomechanical processing of α/β titanium alloys: current status and future research opportunities. *Metall Mater Trans A* 51(6):2593–2625. <https://doi.org/10.1007/s11661-020-05625-3>
81. Dehoff RR, Kirka MM, List FA, Unocic KA, Sames WJ (2014) Crystallographic texture engineering through novel melt strategies via electron beam melting: Inconel 718. *Mater Sci Technol* 31(8):939–944. <https://doi.org/10.1179/1743284714Y.0000006097>
82. Polonsky AT, Raghavan N, Echlin MP, Kirka MM, Dehoff RR, Pollock TM (2020) 3D characterization of the columnar-to-equiaxed transition in additively manufactured inconel 718. *Miner Met Mater Ser*. https://doi.org/10.1007/978-3-030-51834-9_97
83. Murray SP, Pusch KM, Polonsky AT, Torbet CJ, Seward GGE, Zhou N, Forsik SAJ, Nandwana P, Kirka MM, Dehoff RR, Slye WE, Pollock TM (2020) A defect-resistant Co-Ni superalloy for 3D printing. *Nat Commun* 11(1):1–11