



Full length article

Statistically conditioned polycrystal generation using denoising diffusion models

 Michael O. Buzzy^a, Andreas E. Robertson^b, Surya R. Kalidindi^{a,b,*}
^a School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

^b Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

ARTICLE INFO

Keywords:

Polycrystalline microstructure generation
 Conditional microstructure generation
 2-Point statistics
 Denoising diffusion models
 Spherical harmonics

ABSTRACT

Synthetic microstructure generation algorithms have emerged as a key tool for enabling large ICME and Materials Informatics efforts. In particular, statistically conditioned generative models allow researchers to systematically explore complex design spaces encountered in microstructure design. In spite of the engineering importance of polycrystalline materials, generative frameworks for these systems remain extremely limited. This stunted development – in comparison to the N-phase microstructure generation problem – occurs because of the complexities inherent to the representation of the polycrystalline orientation fields. For example, these fields exhibit multiple crystal- and sample-level symmetries. In prior work, these difficulties have resulted in instabilities in deep generative models for polycrystalline microstructures. In this work, we propose the use of a Reduced-Order Generalized Spherical Harmonic (ROGSH) basis to address the challenge described above. The proposed approach accounts for the complex sample- and crystal-level symmetries, and produces well behaved and low dimensional representations whose space has a meaningful Euclidean measure. We then demonstrate the ROGSH basis's remarkable ability to produce stable denoising diffusion models by using our recently established Local-Global generative framework to create visually realistic synthetic polycrystalline microstructures. Furthermore, we demonstrate that the generation process can be conditioned on both first- and second-order spatial statistics of the polycrystalline orientation fields.

1. Introduction

Controllable synthetic microstructure generation has been identified as a foundational cornerstone of next generation Materials Informatics and ICME efforts for accelerating materials discovery, design, and manufacturing [1–11]. These are algorithms that facilitate the systematic generation of synthetic microstructures corresponding to targeted features of interest – often microstructure statistics. These tools provide a practical avenue for generating and expanding the material microstructure datasets that support a myriad of important informatics efforts – from building candidate pools for active learning of surrogate models needed to practically drive materials innovation efforts [5,12–15] to expanding design spaces in microstructure sensitive design [10,16–18]. In particular, there have been significant recent advances in developing 1- and 2-point spatial statistics conditioned generative algorithms (i.e., second-order accurate generative algorithms) [10,19–23] because of their extensive utility in PSP (process–structure–property) learning efforts [9,12,13,24–30]. Although polycrystalline microstructures (encountered in most metals and ceramics as well as some polymers) are

particularly important in a wide range of industries, the development of generative algorithms for this class of microstructures has been extremely slow. Limited success in unconditional generation has been reported [31–35]; there do not yet exist any frameworks that achieve 2-point statistics conditioned polycrystalline generation. Primarily, this is because of the subtle complexities of the local state (i.e., crystal lattice orientation) that characterizes polycrystalline microstructures.

Single-phase polycrystalline microstructures are identified by a spatial field that defines the pointwise crystallographic orientation of the solid. While a general orientation field alone is fairly simple to represent (they are often compactly represented using three euler angles), polycrystals display several modes of symmetry that significantly complicate the representation. These symmetries often exist at two length scales: globally, at the sample level, and locally, at each point within the grains (i.e., individual crystals comprising the polycrystal). Sample level symmetries are introduced by manufacturing processes such as rolling [36], while point-wise symmetries are due to symmetries in the underlying atomic lattice [37].¹ In either case, these symmetries

* Correspondence to: George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA.
 E-mail address: surya.kalidindi@me.gatech.edu (S.R. Kalidindi).

¹ In this paper, we are primarily focused on mesoscale microstructures. Therefore, the entire atomic lattice is treated as a point-wise entity in the microstructure.

result in significant redundancies in representation as well as irregular distances in the orientation space [4,36]. The second is especially important; the true distance between two orientations is obfuscated by the presence of symmetric equivalents. Both outcomes present significant challenges to machine learning (ML) paradigms. Without proper mitigation of these effects, ML models are forced to learn the intricate complexities in the representations of orientation and symmetry. Large models and datasets are necessary to achieve even humble results [31], often in direct opposition to the data-scarcity that characterizes many engineering applications.

N-phase generation (where the characteristic symmetries of the salient microstructure features are ignored and each phase is simply labeled in a one-hot encoding scheme) has seen a recent surge of novel techniques [16,19–22,31,38–44] which have achieved second order accuracy and beyond (e.g., [20]). Several efforts have attempted to transfer successful techniques from N-phase generation to polycrystalline generation with only modest results [31,34]. The resulting synthetic microstructures lack realistic features: e.g., synthetic grain boundaries display unnatural smearing [31]. Recent research from Jangid et al. on deep learning methods for polycrystalline systems (specifically, polycrystalline super-resolution [45,46]) has indicated that these limitations arise because of incompatibilities between common polycrystalline orientation representations – namely, Euler angles and, to a lesser extent, quaternions – and accepted deep learning standards – e.g., L^p loss functions. This concept will be discussed in more detail in Section 3.2. They demonstrate that super-resolution performance can be significantly improved by directly designing the learning paradigm to be sensitive to orientation information, instead of forcing the model to learn it on its own. These papers present a variety of specialized techniques, including bespoke loss functions [45] and model architectures for orientation information [46]. While these efforts provide excellent guiding lessons, there is much improvement needed, particularly for complex problems such as polycrystal generation. First, such efforts address symmetries in an adhoc fashion by mapping all orientations to one apriori selected representation. Second, these efforts do not account for sample level symmetries. Finally, a more general approach, which avoids the use of specialized architectures, would be desirable.

In this paper, we propose a Local–Global Decomposition based framework for statistically conditioned polycrystal microstructure generation. Critically, we utilize Generalized Spherical Harmonics (GSH) to account for representational redundancies due to symmetry and stabilize the local diffusion model’s performance. Our framework achieves second order accuracy, and generates visually realistic polycrystalline microstructures with desired 1- and 2-point statistics. The key contributions of this work are divided into two sections.

1. **Reduced GSH Space:** We identify a 3D subspace of the infinite GSH coefficient space suitable for compactly representing spatial data with sample-level and pointwise symmetries for generation problems. This subspace is low dimensional (making it suitable for learning efforts), invertible (allowing transformation between this representation and other common representations such as Euler angles), invariant to salient symmetries, and measurable using standard norms (allowing it to naturally support standard deep learning efforts).
2. **Statistically Conditioned Generation:** Merging this representation with our recently proposed Local–Global Decomposition [20], we demonstrate that we can perform polycrystal generation conditioned on salient microstructure statistics - specifically, 1- and 2-point statistics [28].

2. Notation

To facilitate communication we briefly present the notation adopted throughout the paper. Vector-valued quantities are demarcated in bold,

e.g., \mathbf{a} . Quantities with spatial dependency, such as spatially resolved functions, are demarcated using a subscript for discretely sampled quantities or a spatial dependency for continuous quantities: a_s and $a(\mathbf{x})$, respectively. Components of vector-valued quantities are indexed using a superscript, $a^\beta = \mathbf{a} \cdot \mathbf{e}^\beta$, where \mathbf{e}^β is the β -basis vector. Powers will be denoted outside a parenthesis to avoid confusion – except in the case of common functions such as the exponential – for example, $(m^\alpha)^2$. Finally, summations will always be written explicitly using the summation operator and are never implied by repeated indices.

3. Background

While interest in synthetic microstructure generation has grown significantly in recent years, methodological advancement of polycrystalline microstructure generation (both conditional and unconditional) has been slow in comparison to other prominent microstructure generation efforts, such as N-phase generation [16,19–22,31,38–44,47]. A significant roadblock is the statistical complexity and specificity of the grain structure. Even in relatively morphologically simple materials, like equiaxed aluminum, grains display sharp and highly geometric grain boundaries. It has been well documented in N-phase generation research that such geometric specificity is difficult to achieve without powerful generative models [19,38,39,41,44,48]. This problem is significantly exacerbated in polycrystalline microstructures. Beyond their distinct morphology, grains are also characterized by a crystallographic orientation. The coupling between the structure of the grains and the spatial arrangement of the orientation field makes this problem much more challenging than N-phase generation.

In N-phase generation, deep learning based frameworks are the state of the art for synthesizing microstructures with highly complicated local features [20,31,38–41,49,50]. However, as discussed in the introduction, this success has not translated to polycrystalline generation [31]. The clear phase boundaries that are achieved on the N-phase problem are replaced by murky boundaries and other artifacts. Jangid and colleagues provide an excellent analysis of several important root causes of these artifacts [45,46]. For example, they demonstrate that euclidean distances (such as the L_1 and L_2 -norms) are not well suited for measuring error in polycrystalline microstructures. Because most deep learning paradigms rely on such distance measures, they will naturally struggle generating non-Euclidean measurable local states (such as Euler angles or quaternions). This demonstrates an important point: a successful polycrystal generation framework must be sensitive to the natural representation and distance measures of the orientation space.

Currently, the most mature and widely adopted class of polycrystalline generation methods are space packing methods – the most popular of which is Dream3D [11,35,51,52]. In brief, these methods sample and subsequently, pack a set of individual features (i.e., grains) to form larger microstructures [35,53]. To facilitate this process, they adopt several important simplifications. First, they utilize simplistic representations of the salient microstructure features (e.g., many frameworks utilize an ellipsoid approximation to the grain shape [53]). Second, they ignore any spatial dependency in the orientation field; they largely decouple the placement of grains from the assignment of their orientations. As a result, they are unable to recreate important features such as the spatially varying textures observed in titanium alloys [46] and the morphologically complex microstructures present in 3D printed metals [54]. However, these methods have amassed tremendous popularity and use because they can be conditioned to enforce many important microstructure statistics (such as grain size distributions) during generation [6–9,11,53]. This is an important lesson; next generation deep learning based methods designed to overcome these simplifications must retain statistical conditioning to actualize their practical utility.

3.1. 1- and 2-point statistics

Statistically conditioned microstructure generation frameworks facilitate the synthetic generation of microstructures displaying target important microstructure statistics. N-point statistics are a rigorous statistical framework for quantifying material microstructures derived from statistical physics² [3,57–59]. This theoretical foundation guarantees that similarity in two material's n-point statistics results in similarity in their behavior. As a result, they have been used as powerful feature descriptors in a variety of Materials Informatics efforts (e.g., microstructure quantification [4,13,28,30,60,61] and advanced surrogate modeling for Process–Structure [8,26,62] and Structure–Property [27,63–65] linkages). For many practical applications, 1- and 2-point statistics are sufficient to support highly accurate data-scientific models. These statistics can be efficiently estimated by adopting a voxelized (or pixelized in 2D) representation of an instantiation of a microstructure [4]. In this representation, each voxel contains a summary of the average state within the voxel. The 1-point statistics are computed directly via averaging while the 2-point statistics of a periodic microstructure can be computed efficiently via the following expression.

$$f_r^{\alpha\beta} = \frac{1}{S} \sum_{s=1}^S m_s^\alpha m_{s+r}^\beta \quad (1)$$

$$= \frac{1}{S} \mathcal{F}^{-1} \left[(M_r^\alpha)^* M_r^\beta \right]_r \quad (2)$$

Here, $\mathcal{F}[\cdot]$ is the discrete Fourier transform (DFT) operation (performed using the Fast Fourier Transform algorithm); m_s^α is the value of the α state at the spatial index s and M_r^α is its DFT indexed at frequency index r ; $*$ indicates complex conjugation; S is the total number of voxels.

3.2. Generalized Spherical Harmonics

Generalized Spherical Harmonics (GSH) are an extensively studied function basis for quantifying information in orientation space [4,36]. Congruent to the usage of Fourier expansions, any function in orientation space can be compactly summarized and approximated by identifying the coefficients of its GSH expansion. In crystallography, the GSH basis is typically used to summarize the Orientation Distribution Function (ODF), a popular 1-point statistic for polycrystalline microstructures. The ODF is the probability distribution of orientations present in a given volume – historically, either the entire microstructure or a single pixel/voxel. It is denoted $f(g)dg$, where f is a function which returns the probability of the orientations in the differential dg . Its GSH representation is expressed as

$$f(g) = \sum_{l=0}^{\infty} \sum_{\mu=1}^{M(l)} \sum_{\nu=1}^{N(l)} C_l^{uv} \hat{T}_l^{*uv}(g) \quad (3)$$

where C_l^{uv} are the coefficients which parameterize the GSH expansion, and \hat{T}_l^{*uv} are the symmetrized GSH basis functions. Importantly, the basis functions are symmetrized to naturally incorporate any of the function's salient symmetries [36].

More recently, GSH coefficients have been used to represent the pointwise, spatially varying orientation state of a polycrystalline microstructure (m_s^α in Section 3.1). This has been most useful in computing second-order microstructure statistics via the equations laid out in Section 3.1 [9,28]. In these works, it is standard to limit each voxel to a single orientation by maintaining a voxel lengthscale much smaller than the grain size. Consequently, the discrete local states are the GSH

² In this paradigm, materials are conceptualized as stochastic processes. The stochastic microstructure function (i.e., the microstructure generating process) – not the individual instantiations of possible microstructures – are centralized in a Process–Structure–Property conceptualization. The microstructure process is compactly quantified via its n^{th} order moments – the n-point statistics [3,4,19,55–57].

coefficients, C_l^{uv} , of the expansion of the impulse distribution centered around the single orientation: The integral of $f(g)$ is equal to one only if the domain includes the orientation of interest. For impulse distributions there exists a special relation [36]:

$$C_l^{uv} = (2l + 1) \hat{T}_l^{*uv}(g_o) \quad (4)$$

Representing single orientations as the GSH expansions of impulse distributions comes with several key benefits when compared to other representations such as euler angles or quaternions. **symmetry**: the generalized spherical harmonics can be symmetrized to naturally reflect any salient pointwise (i.e., crystal) or sample symmetries. A subset of the GSH (or linear combinations of subsets) are chosen such that the set of basis functions themselves possess the same symmetries as the physical system in question. This results in one representation for all symmetrically equivalent orientations, in contrast to other representations where a fundamental zone must be defined and enforced throughout any computations. This unique representation is particularly valuable for microstructure generation because it avoids the mixed periodic and mirror boundaries of some fundamental zones [36]. These complex fundamental zone boundaries are particularly detrimental to learning efforts. They make it challenging to ascertain the true distance between two orientations. For example, naively taking the distance within the fundamental zone often neglects that the orientations may be very close to another symmetrically identical orientation. The second benefit of the GSH representation is extremely closely related: **natural L_p distance**. The L_p distance is especially poorly behaved in the Euler angle space for several reasons. First, the aforementioned challenges of symmetry (these also apply to quaternion representations). Second, the Euler angle space possess several degenerate planes in which all orientations within that plane are identical [4]. The L_p distance will be non-zero within those degenerate planes indicating a difference in orientation when there is not one. Finally, the differential volume element in the Euler angle parameterization of orientations is nonuniform. It is defined as $dg = \frac{1}{8\pi^2} \sin(\Phi) d\Phi d\varphi_1 d\varphi_2$. L_p distances assume a Euclidean (uniform) differential element and would assume $dg = d\Phi d\varphi_1 d\varphi_2$ [36].³ Ignoring this difference and measuring the distance between orientations by taking the L_p norm directly on Euler angles introduces distortions into the distance metric. The GSH representation on the other hand naturally accepts L_p distances [4,36]. Since the GSH are the coefficients of a Fourier expansion (weights of functions on SO_3), this representation again lies in a well behaved euclidean space [4,36]. This allows for the direct comparisons of individual orientations (as well as ODFs) through the common L_2 norm [4,36]. In this work, this is critical for avoiding bespoke loss functions and architectures created to cope with the topology of orientation spaces, as is the case for quaternion based neural networks [45,46]. By adopting this representation, we gain access to state-of-the-art ML techniques. The third benefit is **direct averaging**. The entire microstructure's ODF (i.e., the 1-point statistic) can be easily computed by directly numerically averaging the ODF of each voxel. This significantly simplifies conditioning the generation process with 1-point statistics, Section 3.3. Finally, the last benefit of the GSH space is **convexity** [4,36]. This property greatly simplifies downstream tasks necessary for polycrystal generation, because the set of well defined points is easily identifiable. This provides an easy mechanism to check the training performance. Additionally, for the purposes of this work convexity is necessary for inverting back to Euler angles (see Section 4.2). Although other representations provide subsets of these identified properties, the strength of the GSH coefficient representation is that it *uniquely* provides all of them simultaneously.

³ The Euler angles here are taken to be in the Bunge convention.

3.3. Local-Global Decomposition

Historically, modern deep learning frameworks for microstructure generation have excelled at generating highly realistic microstructure features [38,40–42,50] – especially spatially localized features – at the cost of limited statistical controllability. However, the recent Local–Global Decomposition (LGD) framework [20] develops several statistical physics inspired mechanisms to facilitate 1- and 2-point statistical conditioning by guiding the inference mechanism of deep diffusion models. Importantly, this conditioning process is achieved without any changes to the training process and, even, decreases the amount of necessary training data (in the original examples, the models are trained on a single experimental microstructure [20]). Additionally, there is early evidence that the framework helps stabilize the deep learning model [20] facilitating some extrapolation.

The LGD framework globalizes the role of 1- and 2-point statistics, arguing that they define long range patterns. Concurrently, it localizes the role of higher order statistics – and the deep learning algorithm – to defining salient local microstructure features.

$$p(\mathbf{m}_1, \dots, \mathbf{m}_S; \boldsymbol{\mu}, \mathbf{f}_r) = \mathcal{N}(\hat{\mathbf{m}}_1, \dots, \hat{\mathbf{m}}_S; \boldsymbol{\mu}, \mathbf{f}_r) \prod_{i=1}^K p^{\text{cond}}(N_i | \hat{N}_i, N_i^c; \Phi^{(3, \dots)}) \quad (5)$$

Here, $\boldsymbol{\mu}$ and \mathbf{f}_r are the possibly vector valued, constant mean and autocorrelation functions derived from the 1- and 2-point statistics [19]. $\mathcal{N}(\cdot; \cdot)$ is the multivariate Gaussian distribution, entities after the semicolon denote parameterizing statistics. N_i denotes the value of the voxels in the i^{th} spatially compact neighborhood [20]. The hat denotes quantities of the intermediate output – the second order latent approximation. The N_i^c denotes the neighborhood complement – the values of the set of voxels outside of the i^{th} neighborhood.

In practice, a deep diffusion learning model approximates the second term: the localized neighborhood distribution. When the complete framework is utilized, the diffusion model's output is guided by the output of the global model, the first Gaussian term. Robertson and Kalidindi's Multi-Output Gaussian Random Field (NGRF) model [19] is employed as the global model [20]. Importantly, this approximation analytically incorporates the desired 1- and 2-point statistics and can be efficiently sampled in $O(NS \ln S)$ computational complexity (N microstructure states, S spatial voxels). Its worth noting that, to date, this global generative model was only validated on N-phase generation; however, it was able to generate microstructures with large numbers of phases and arbitrary, anisotropic 2-point statistics.

In the complete framework, the diffusion model's inference process is guided in two ways. First, an extension of the denoising function, ensures that the conditional sampling process does not distort the desired 1-point statistics.

$$s_{\theta}^{\text{cond}}(\mathbf{m}, t) = s_{\theta}^{\text{uncond}}(\mathbf{m}, \sigma) - c_t \left(\frac{\mathbf{m}^T \mathbf{1}}{S} - \omega \right) \mathbf{1} \quad (6)$$

For microstructure's with vector valued local states, the correction is applied to each dimension individually. Here, $s_{\theta}(\mathbf{m}, \sigma)$ is the trained diffusion model with trainable parameters θ . $\mathbf{m}^T \mathbf{1}/S$ computes the average of the dimension. ω is the desired average for that phase. c_t is a constant that possibly depends on the step in the diffusion process, t . Second, the 2-point statistics are softly enforced by seeding a truncated diffusion process on the output of the NGRF (i.e., replacing the initial random noise sample with a sample from the NGRF per [20]). Tuning is necessary to identify an optimal subset of the diffusion process that sufficiently introduces the higher order statistics (i.e., clarifies the salient local features) while retaining the desired 2-point statistics [20]. We refer the reader to the original work for an exhaustive analysis of the mechanics of sampling.

It is important to emphasize that it is not necessary to use both guiding processes. In the original work [20], we document that imposing

the 2-point statistics balances the quality of the local features against maintaining the 2-point statistics. Therefore, for systems without long range global patterns, it might be preferable to just use the 1-point conditioning strategy.

3.4. Denoising diffusion probabilistic models

Denoising diffusion models are the natural choice of deep distributional learning model for the proposed polycrystalline generation framework since we intend to utilize the LGD framework to incorporate statistical conditioning. While there exist several different classes of diffusion models [20,66–70], in this work we have chosen to use denoising diffusion probabilistic models (DDPM) introduced by Ho [67] because of their widespread success in a variety of engineering generation problems [71–76]. They have demonstrated state of the art performance and empirically seem less impacted by the blurring artifacts effecting other distributional learning methods [77] and the training instability displayed by score-based models [20,66,69,70], this has been evidenced by the overwhelming preference of DDPM within the machine learning community [32,76]. Training is performed by minimizing the variational bound on the negative log likelihood. This results in the following loss function [67].

$$\text{Loss} = \mathbb{E}[\|s - s_{\theta}^{\text{uncond}}(\sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} z, t)\|^2] \quad (7)$$

Once trained, we iteratively sample the trained model using the following procedure [67].

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} s_{\theta}^{\text{uncond}}(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z} \quad (8)$$

Here, α_t are denoising parameters taken to follow a linear schedule [67]. $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Like before, $s_{\theta}(\mathbf{x}_t, t)$ is the learned denoising function, additionally transformed by the identified correction in our application, Eq. (6).

4. Reduced GSH space

In this section, we develop the foundation of the proposed generative framework; we present a local state representation, m_s^{α} , for polycrystals suitable for the polycrystal generation problem. We begin with the GSH coefficients representation which already overcomes many of the identified problems with other representations. First, differences between orientations in this space are naturally measured using standard distance measures (e.g., the L_2 -norm). Second, it naturally accounts for symmetry without introducing complex boundaries and artifacts. Finally, it provides a sensible notion of averages; the average orientation – i.e., the 1-point statistic – is computed via direct numerical averaging. Unfortunately, for all these benefits, the principal drawback for using the GSH coefficient representation in the context of polycrystalline generation is its high dimensionality. It is well documented that the performance of generative algorithms rapidly deteriorates as the dimensionality of the state space increases [20].

4.1. Selecting the ROGSH basis

We overcome the high dimensionality of the full GSH expansion by restricting the type of ODF that can exist in each pixel or voxel. In general, this high dimensionality is incurred because many terms are required to accurately describe an arbitrary ODF (e.g., [28]). However, for the purposes of microstructure generation, it is standard to assign a single orientation to each pixel or voxel [31,33,34,45,46]. Therefore, we restrict the ODFs to impulse distributions. Considering Eq. (4), the entire expansion associated with an impulse distribution can be computed using only the three Euler angles as input. Therefore, we argue that there must be a reduced set of – at minimum – three GSH coefficients for the restricted set of impulse functions which are equivalent to the Euler angle representation within symmetry. In other words,

an invertible transformation between the two representations exists. As a result, since we are only interested in describing a single orientation, we can exploit these redundancies in the full GSH coefficient set to describe the polycrystal microstructure local state, m_s^g , with only three coefficients from the expansion.

In general, there is not a consistent approach for distilling the ROGSH basis; a bespoke ROGSH basis must be derived for each symmetry group. The selection process begins with the complete GSH basis for that symmetry group. Importantly, the three GSH coefficients comprising the ROGSH basis must be selected carefully to ensure the unique mapping constraint is met. Consider the unsymmetrized GSH set [36] as an informative example to illuminate the challenges in this process.

$$T_l^{mn}(g) = e^{im\varphi_2} P_l^{mn}(\Phi) e^{in\varphi_1} \quad (9)$$

The basis depends on the complex Legendre polynomials $P(\cdot)$, while also containing multiple exponential terms. Therefore, a random selection of these functions is not guaranteed to provide the desired one to one mapping from ROGSH to Euler angles. The selection is further complicated when utilizing the symmetrized GSH basis, as the symmetrized GSH are formed through linear combinations of the unsymmetrized GSH [36]. This results in many possible degeneracies which are symmetry specific and not necessarily obvious. The ROGSH basis functions for hexagonal-triclinic [78] and cubic-triclinic systems – the focuses of the upcoming case studies – are reported in Appendix B. The unique mapping of these bases was verified numerically by thoroughly sampling the fundamental zones, and ensuring that all orientations yielded unique values in the ROGSH space.

4.2. Converting from ROGSH back to Euler angles

Once microstructure generation is completed in the ROGSH space, further analysis and simulation of the generated microstructures necessitates they be converted back into an Euler representation. However, finding an analytical inverse can be challenging. The ROGSH basis often consists of long expressions which contain many trigonometric terms. This makes it difficult to algebraically formulate the inverse, even though it theoretically exists – for example, the inverse for the hexagonal symmetry is derived in the appendix, but we were unable to do the same for the cubic case. In this work, we use a database of known ROGSH–Euler pairs along with a nearest neighbor projection to perform the inverse mapping when an analytical inverse cannot be found. The database is constructed by sampling Euler angles from the fundamental zone [79] of the symmetry in question, then converting those angles to the ROGSH basis. When inverting the ROGSH basis, we simply look up the nearest point in the database and return the associated Euler angles. This is done efficiently with KDTrees [80]. The fundamental zone can be sampled to allow for arbitrary accuracy. In this work we sampled such that the greatest distance between two Euler angles in the database was 1 degree.

An important observation is that using a database for inversion is not equivalent to binning the Euler angle space and treating the structure as an n-phase microstructure, as other works have done in the past [81]. In binning approaches, the fundamental zone is broken into sectors, and any angle within a sector receives the corresponding label of that sector. While the database approach is similar to binning in the sense that it discretizes the Euler angle space and prohibits returning an arbitrary Euler angle, it does so *after* generation. It does not incur the associated dimensionality explosion that binning causes during generation. To achieve the same 1 degree fidelity as this work by binning before generation, tens of thousands of bins would be required. This results in an untenable number of phases for current generative frameworks.

5. Experiments

We demonstrate and test the proposed framework as well as document its strengths and weaknesses by performing two case studies

involving statistically conditioned generation. Here, we maintain a specific focus on generation in a data constrained environment. In both cases, we perform a standard benchmarking exercise [19–21,44,48,82,83]; we aim to synthesize polycrystalline microstructures which are statistically equivalent to a experimentally collected reference. The reference microstructures have the dual purpose of (1) providing exemplary microstructure neighborhoods for training the diffusion models for the LGD generation framework and (2) being a statistical benchmark to quantitatively assess the generative framework's performance. The first case also will serve as an overview of the application of the polycrystalline LGD framework by outlining each step in detail. The complete process is visually summarized in Fig. 1.

5.1. Case study 1: Rolled Ti64

In the first case study, we demonstrate the generation of a rolled Ti64 polycrystalline microstructure. The experimental EBSD used as reference is shown in Fig. 2. Rolled Ti64 is a natural benchmark for polycrystal generation (in fact, it was used in recent works on polycrystal super resolution [45,46]), because it is both challenging to generate and of engineering interest in many applications. This material system poses several challenges for generation because it contains salient microstructure features on multiple length scales. At the grain length scale, it is characterized by a reasonably complex grain morphology (i.e., size and shape distributions of grains). As a result, it serves as a good benchmark to assess the capacity to generate the grain morphologies typically seen in real applications. We leave much more complex grain structures to the next case study. The selected reference microstructure for this case study also exhibits distinct features at a larger length scale because of the spatially heterogeneous texturing of rolled Ti64. It contains characteristic banded, highly textured regions appearing on a lengthscale much larger than that of individual grains. With respect to engineering applications, these bands are extremely important because they are the initiation site of several important engineering phenomena (such as fracture and fatigue phenomena) [84].

The proposed framework is capable of capturing the multiscale characteristic features described above. In Local–Global Decomposition based generation, several generation techniques are combined to construct the complete generative model. Each technique is designed to target features on a specific lengthscale. The following two sections of this case study analyze the performance of these individual methods on polycrystalline generation. Importantly, we see the stabilizing benefit of the ROGSH basis. In addition, this case study outlines the complete usage of the polycrystalline LGD framework in practice.

5.1.1. Global estimate

The Local–Global Decomposition framework begins with a second order approximation of the generating process: a Multi-Output Gaussian Random Field model (GRF). Because previous literature has only benchmarked this model for N-phase generation [19], we briefly address its compatibility with polycrystal generation. Excitingly, the GRF model is still able to generate random field instantiations whose second order statistics match those of the selected reference. A comparison of the experimental microstructure's statistics and the average statistics of 200 samples generated from the GRF are shown in Fig. 3. Note that we adopt the ROGSH microstructure representation throughout generation. The presented statistics are the 2-point statistics of the first ROGSH coefficient field (see Appendix B). The statistics display significant visual similarity; both share the same qualitative characteristics such as the shape of the center peak and vertical bands. The blue plot in Fig. 3 shows the difference between the two statistics (please note the change in scale of the color bar). Quantitatively, the statistics show excellent agreement. The differentiating errors are approximately two orders of magnitude smaller than the original statistics. We note that the error is not randomly distributed, instead the highest errors are localized around the long vertical features ($x = \pm 175$) that quantify the

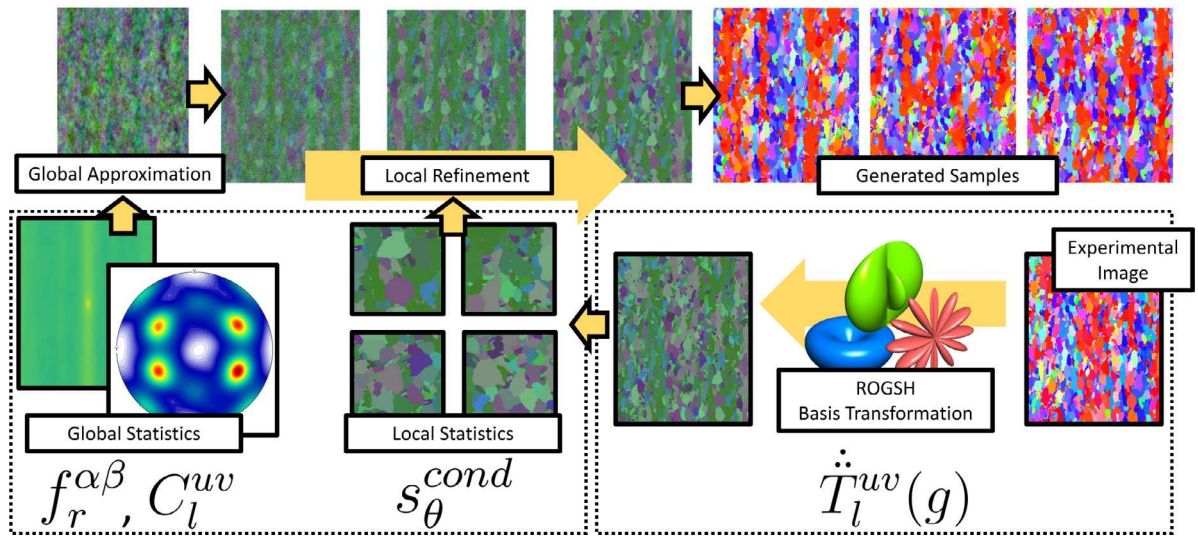


Fig. 1. Visual Overview of the proposed method for generating synthetic polycrystalline microstructures.




Fig. 2. The original sample image shown in the Euler angle space. This sample is rolled Ti64 with hexagonal crystal symmetry (a). Its orientation distribution function (c), and 2-point statistics of the first ROGSH coefficient (b) are also shown.




Fig. 3. A comparison of the 1–1 autocorrelation (i.e., the first ROGSH coefficient, Appendix B) of an experimental Ti64 micro-graph compared to the average of 200 samples from the GRF. The statistics are shown in gray, and their difference is shown in blue using a magnified scale to highlight small discrepancies. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Slices of the sampled statistics generated by the GRF compared to the experimental microstructure. The top plot shows a slice in the X direction, the bottom shows a slice in the Y direction. The experimental microstructure is shown in black, the average GRF statistics in dotted blue, and the samples in solid blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

repeating banding in the microstructure. This is not surprising. As noted in the previous work, this occurs because of the long lengthscale of the bands. These long lengthscale features are underrepresented in an RVE (compared to, for example, the short range features like grains) which introduces higher noise in these statistics. While the nonuniformity of the error is not ideal, the error remains acceptably low throughout and the good statistical performance of the GRF is clear.

In order to better understand the variability in statistics between samples from the GRF, two cross-sections of the statistics are plotted in Fig. 4. These “slices” allow us to plot the statistics of each sample on a single plot, rather than the average statistics which were shown previously. The top plot shows the slice along the X direction. Here we can see that the average output of the GRF aligns almost exactly with the experimental microstructure. This is unsurprising because we adopt no post processing at this stage, so convergence is analytically guaranteed given well sampled statistics. Looking at the variability of the samples, we see good agreement in line with prior work utilizing GRFs [19,20]. Turning our attention to the slice in the Y direction, we see again the slight bias introduced by the poor sampling of the banded features. We also note the slight increase in variability in the samples when compared to the X direction that this under sampling causes. It should be noted that the slice taken in the Y direction is directly across the central ridge seen in Fig. 3. This is precisely where the error is the largest. For other slices in the Y direction the variability is improved and more closely resembles the agreement seen in the X direction.

Fig. 5 shows a sampled output from the GRF. The color scheme is changed to reflect the transition to an ROGSH representation of the polycrystalline microstructure. While Figs. 3 and 4 demonstrated that the GRF structures attain excellent second order statistical agreement with the reference, Fig. 5 demonstrates that the realism of the second order approximation is obviously wanting. Looking closely, the vague impression of grains is given by regions with consistent color (i.e., similar GSH coefficients), but these regions fail to form defined boundaries. In this case, even more than the N-phase case [19], the separation between individual features is blurry at best. A suitable local refinement is necessary to achieve a visually realistic microstructure. Fortunately, the local element of the LGD framework – a ROGSH stabilized diffusion model – provides exactly that.

5.1.2. Local approximation and refinement

The second step of our polycrystalline LGD framework utilizes a DDPM model to locally refine the second-order approximation from the previous step into a finalized microstructure in the reduced GSH space. To train the DDPM model, we cut the original reference into 14,000 128×128 local patches. Patch sizes should be chosen such that there is a clear separation between the global and local length-scales, while also being large enough to accurately quantify the neighborhoods. We found that a patch size large enough to capture 2–3 grains was sufficient. During training we utilized the standard dropout regularization technique to prevent over fitting. In exchange for slightly increased inference times, this highly stabilized the training process, allowing us to avoid much of the painful manual architecture design pursued in our previous work [20]. Further information concerning the model architecture and training can be found in Appendix A. We found for this problem that the ideal number of diffusion steps was 300. This was determined by running the diffusion process for varying numbers of iterations and observing how much the n-point statistics changed between iterations. After 300 steps the statistics started to be perturbed significantly. We found that 300 steps was ideal in most situations.

Fig. 6 visually summarizes the denoising process of the GRF output. The yellow circles highlight a single example grain which forms from a noisy, but relatively uniform region during the denoising process. The consistent sharpening throughout the process, without excessive distortion, qualitatively demonstrates the successful application of the proposed denoising post-processing. Like in the N-phase case, it is able to correct the local features in the microstructure without significantly perturbing the original samples. Importantly, in addition to clear grains emerging, the long range patterning described earlier is maintained and becomes clearly visible. This is especially clear once we use the identified inversion procedure, Section 4.2, to transition back to Euler-angles and the IPF visualization, as shown in Fig. 7. Inspecting these microstructures we can see that the generative framework is able to successfully reconstruct both sets of important features in the original reference. Locally, clear grain boundaries are present. This result is especially exciting because of the repeated historic difficulty associated with using deep learning to generate clean Euler angle fields

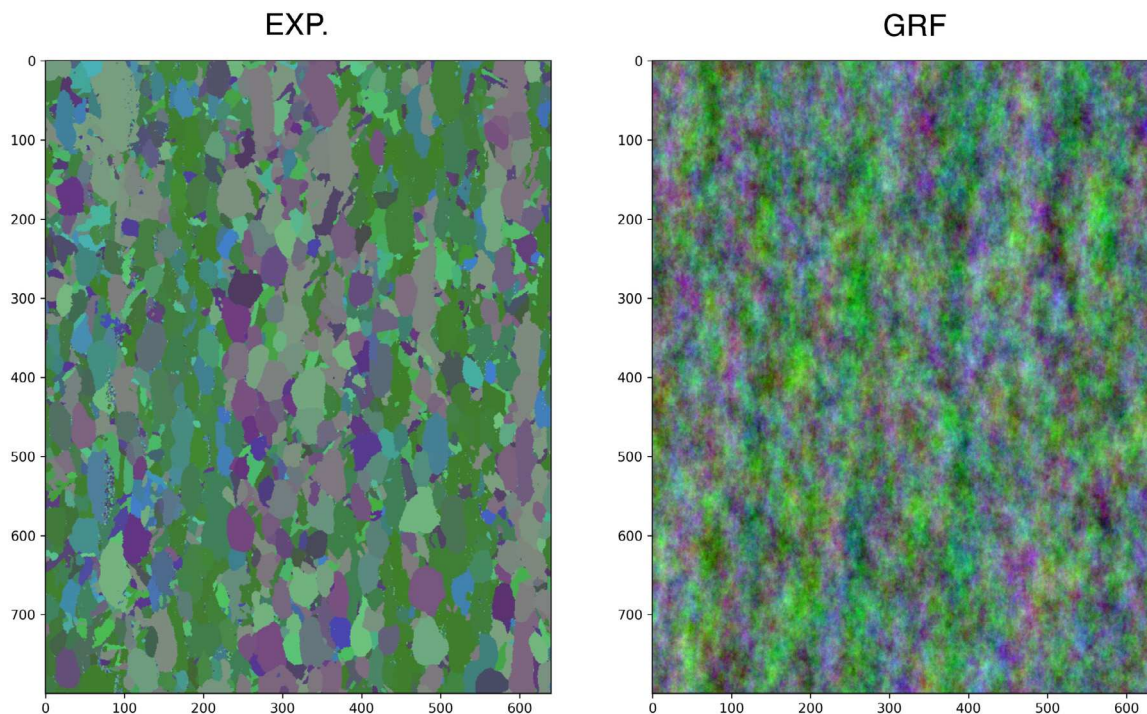


Fig. 5. A microstructure instance in ROGSH space generated by a Gaussian Random Field with its covariance parameterized by the 2-point statistics of the original experimental microstructure. Note: The color change in the image represents the transformation into the ROGSH basis. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 6. Initial noise from the Gaussian Random Field undergoing an iterative transformation into a final polycrystalline microstructure. Inside the yellow circle a purple region is highlighted to show how the final microstructure maintains the patterns from the initial noising field. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(e.g., [31]). In addition, globally, the texturing and orientation preference of the generated microstructures visually agrees with the original reference showing the same alternating vertical bands in the red/blue orientations. We emphasize that the banding occurs on a lengthscale larger than the diffusion model's 128×128 window.

5.1.3. Generative quality

Contrasting salient microstructure statistics between the generated microstructures and the experimental reference, we quantify the observations described above. The original statistics are taken directly from the reference microstructure. They are compared with the statistics of 200 synthetic microstructures generated by conditioning the proposed framework using the reference statistics.

Fig. 8 contrasts the experimental and generated Orientation Distribution Functions (ODFs). The angles were extracted from the generated microstructures, and the ODFs were estimated and visualized in MTEX

using the DeLaVallee Poussin Kernel⁴ [79]. The ODFs are nearly identical, and only vary slightly in low density regions. This agreement, while maintaining visual realism, is very important in the context of previous efforts. Historically, methods have achieved visual realism by using corrective measures to remove artifacts [35]. However, these methods introduce systematic biases in the generated ODFs. In contrast, the direct ODF conditioning in the used DDPM model, Eq. (6), accurately yields the correct ODF despite also acting as a denoising measure for the GRF.

Fig. 9 contrasts the 2-point statistics of the reference and the average 2-point statistics from the generated ensemble of microstructures. These are seen to be in good agreement with each other. The variation within

⁴ MTEX estimated the optimal kernel halfwidth, selecting a kernel bandwidth of 102 and a halfwidth of 2.2 degrees.




Fig. 7. Several samples from the generative process converted back to an Euler representation.

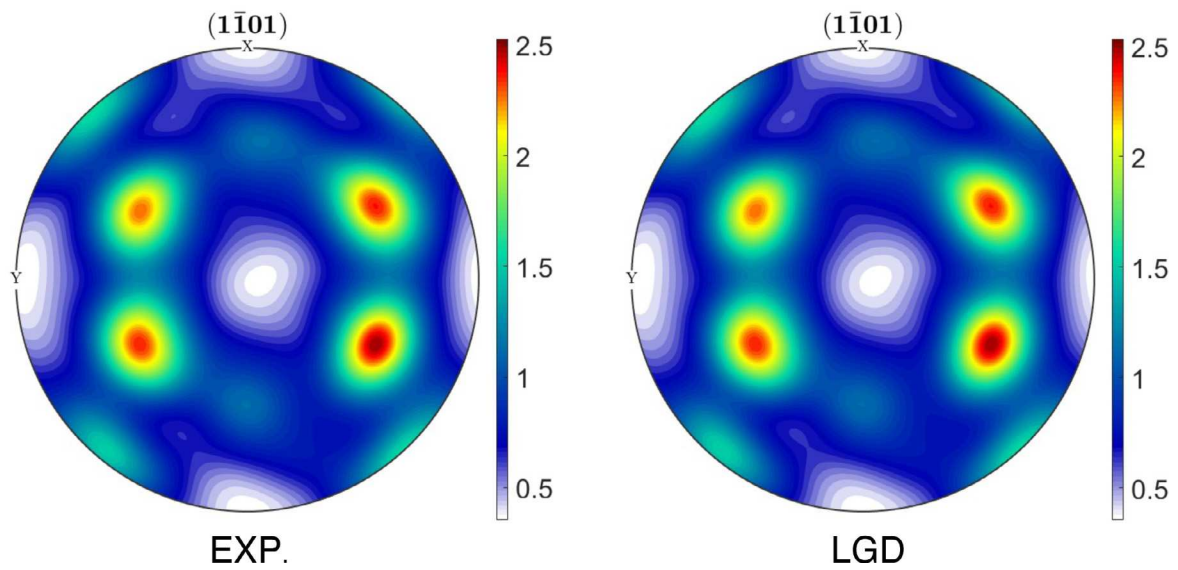


Fig. 8. A comparison between the experimental and generated ODFs for Ti64.




Fig. 9. The errors in 1-1 ROGSH autocorrelation between the generated microstructures and the experimental microstructure.




Fig. 10. Slices of the sampled statistics generated by the LGD compared to the experimental microstructure. The top plot shows a slice in the X direction, the bottom shows a slice in the Y direction. The experimental microstructure is shown in black, the average statistics of the full LGD framework in dotted blue, and the samples in solid blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the distribution of generated microstructures is largely bound within ± 0.003 of the reference. This is nearly two orders of magnitude less than the peak value of 0.125 in the original statistics. Importantly, the agreement demonstrates that the denoising process is not significantly altering the long range statistics produced by the GRF while it corrects the local features. The error is only slightly higher than that observed in the GRF (approx. 0.001). We can analyze the statistical fluctuations in the ensemble by looking at Fig. 10. In the cross sections, we see similar errors to that of the GRF, but now slightly exaggerated. Again, the slice in X direction matches well, and the average output nearly coincides perfectly with the experimental microstructure. Most importantly, the secondary peaks corresponding to the long range layering are clearly present. In the Y direction, the bias associated with the vertical bands has increased slightly. Additionally, the peak of the generated sample is now noticeably wider than the peak of the experimental sample. This increase occurs because the banding features are larger than the chosen neighborhood size of 128×128 in the vertical direction. As a result, the local approximation struggles to learn these features, since it cannot observe them in their entirety. Because of this, grains inside the bands struggle to resolve. This can be seen clearly in Fig. 7 where inside the red bands some grains do not become distinctly separate. This highlights a fundamental boundary to the LGD framework's performance; when there is a pattern that fully saturates the "field of view" of either component model (i.e., the bands being poorly sampled in the GRF, or the bands being too large to be captured in a neighborhood), some error will be introduced. Although a larger neighborhood size which fully contains the salient patterns would likely reduce these artifacts, the selected size is also limited by other practical limitations – such as the availability of training data. Overall, the deviations remain small. The framework achieves excellent first- and second-order agreement.

Finally, contrasting salient grain statistics, we see that synthetic microstructures locally approximate the reference as well. The agreement in the 2-point statistics near zero shown in Figs. 9 and 10 provides an initial indication that the grain structure attains good agreement in the eccentricity of the grains. Additionally, we use chord-length distributions to quantify the distribution of grain sizes present in the microstructures. As can be seen in Fig. 11, the synthetic microstructure's chord-length distribution approximates the reference well. Briefly, again we see the previously identified effect of the neighborhood size:

the generated microstructures are slightly biased towards longer cord lengths. This visually is apparent in the generated microstructures, as there is slight smoothing of the grain boundaries in regions with a high density of highly similar orientations (i.e., the bands) when compared to the original microstructure. Unsurprisingly, we can see that this effect is slightly worse for the CLD in the Y direction.

5.2. Case study 2: Additively manufactured Inconel 625

The first case study showcases the ROGSH stabilized LGD framework's ability to generate polycrystals with both local and global heterogeneity. For this second case study, we explore the full capabilities of the local distribution approximation by generating a microstructure with very rich local features. For this purpose, we have selected a sample of additively manufactured Inconel 625 from the AFRL AM Bench challenge [54]. This sample is an ideal case study to showcase the local approximation, as it has sharp grain boundaries of varying size and shape, and is visually much less predictable than the previous Ti64 sample. Furthermore, we take this opportunity to highlight our ability to deal with microstructures with cubic crystal symmetry by using the more complex cubic ROGSH basis rather than the HCP ROGSH basis used before, Appendix B.

The sample in question is shown in Fig. 12. We can immediately see from the EBSD image on the left that the microstructure displays extremely complex local features, characteristic of additive manufacturing. There exists a wide variety of grain sizes and shapes. The smaller grains take on a plate like appearance, and are visually almost rectangular. The larger grains take on ill defined asymmetric shapes with very rough and noisy grain boundaries. Furthermore, the larger grains are also partially bifurcated by smaller grains. Turning our attention to more global features, we can see that the inconel microstructure lacks the same global ordering as the Ti64, as there are no long range patterns such as banding present. For this reason, we will generate inconel microstructures using only the local approximation, and forgo the global approximation using the GRF. This highlights a key benefit of the LGD framework in that we are approximating the ordering across different length scales quasi-independently. Therefore when performing generation we are able to modularize our approach, and incorporate information from each approximation one at a time in a successive

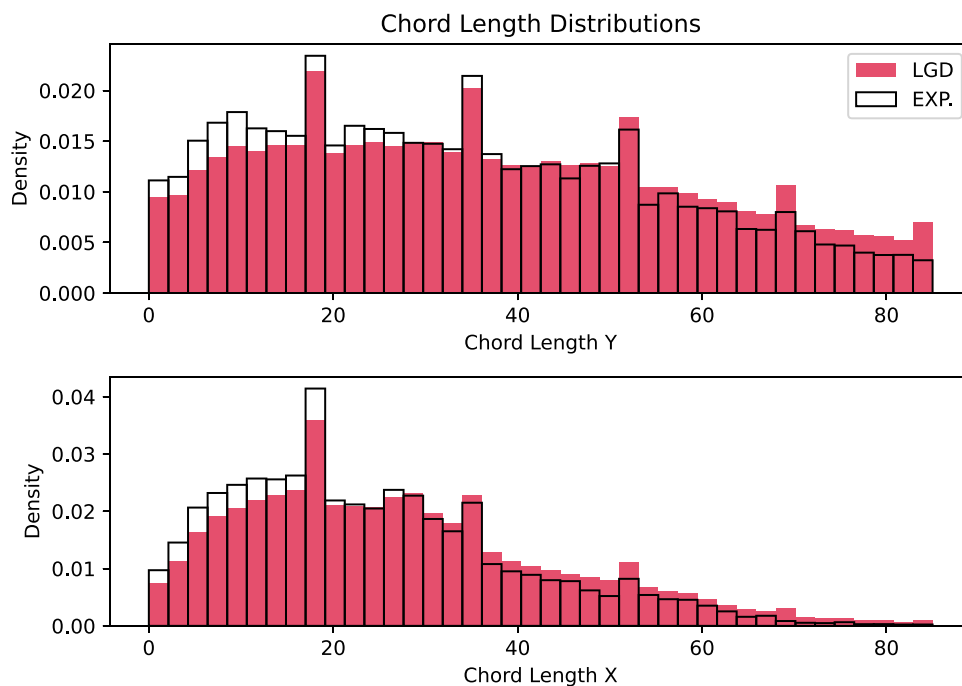


Fig. 11. Experimental (black outline) compared to generated (red) chord length distributions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

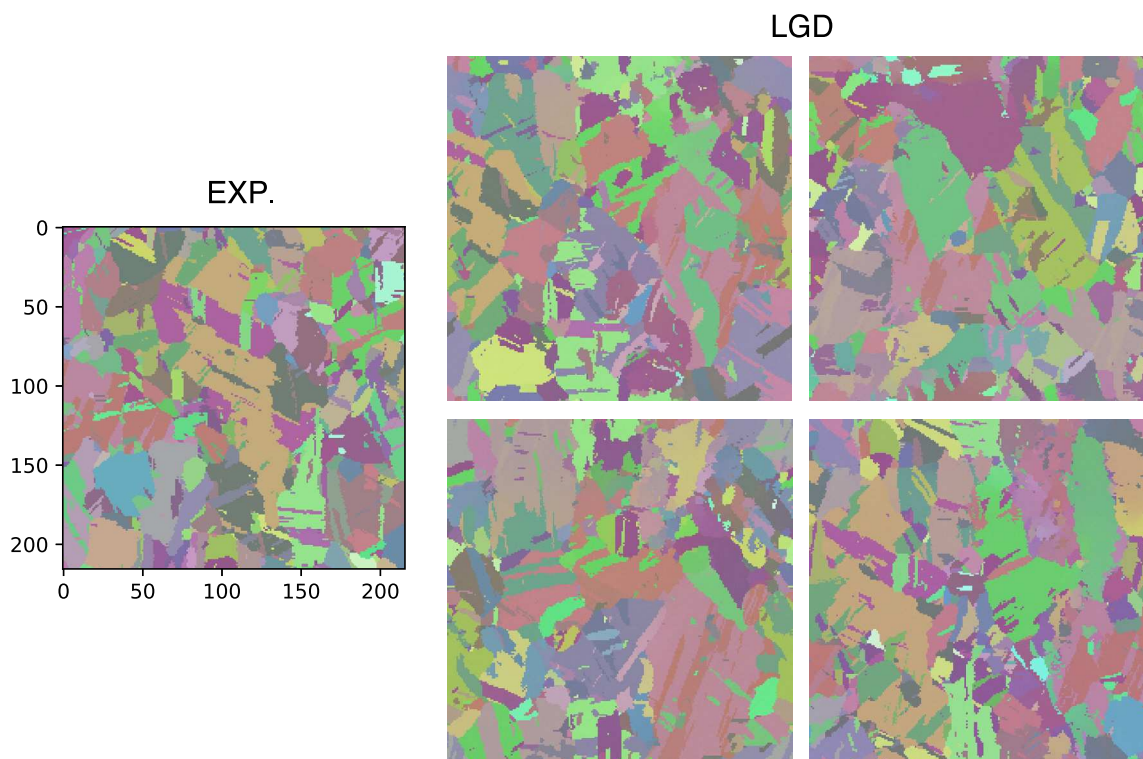


Fig. 12. Experimental and Generated Inconel samples shown in the Cubic ROGSH Space.

manner. This gives the user the freedom to generate microstructures using whatever subset of information is relevant to the microstructure or the problem at hand.

The local approximation method was applied to the inconel sample following a similar protocol to the Ti64 sample. We aimed to pick a patch size that captures 2–3 grains on average, which in this case

was 64×64 pixels. We took a total of 128 images of the AM Inconel from the AFRL AM Bench challenge which resulted in a dataset of approximately 16000 local patches. We can see sample generated microstructures from the trained model in Fig. 12. Qualitatively the generated microstructures are very similar to the original, displaying both large and small grains. The smaller grains have plate like structure

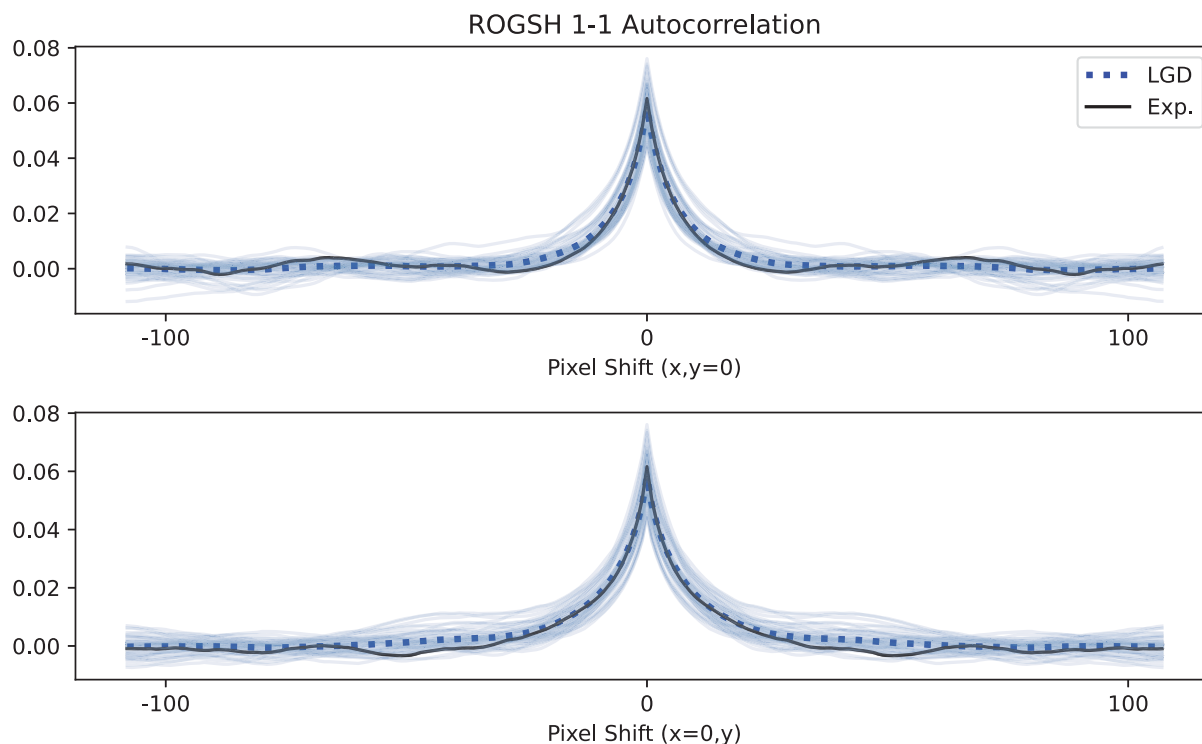


Fig. 13. Slices of the sampled statistics generated by the LGD compared to the experimental microstructure. The top plot shows a slice in the X direction, the bottom shows a slice in the Y direction. The experimental microstructure is shown in black, the average statistics of the LGD in dotted blue, and the individual samples in solid blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

similar to those in the original, and the larger grains take on asymmetric amorphous shapes with rough grain boundaries. Furthermore, we see the same bifurcation of larger grains by smaller grains that was present in the original microstructure.

Quantitatively, we briefly compare the same statistics as Case Study 1. They display strong agreement with the experimental reference. Here, we will highlight 2-point statistics and chord length distributions. The remainder of the comparisons can be found in [Appendix C](#). As in Case Study 1, the generated statistics are averaged over 200 generated microstructures.

[Fig. 13](#) compares cross-sections of the experimental and generated 2-point statistics. Looking at the experimental statistics we can confirm our initial observation that these microstructures contain no long range ordering. This is clear from the nearly flat tails of the 2-point statistics. The values of the statistics within these flat tails hovers around the square of the peak value, further indicating that they correspond to random spatial patterns [4]. Making note of this is key as it validates our decision to forgo the global approximation method in this case study, and utilizing only the learned local statistics initialized on random noise. Furthermore, we can see that the generated microstructures agree on average with the experimental microstructure, and that the variance between individual samples is far tighter than in the previous case study. We believe the improved agreement is primarily due to the lack of global ordering in this microstructure. As previously discussed, long range patterns become more difficult to generate because we observe fewer examples of them when compared to local features. In the case of the inconel samples, we can clearly see the obtainable fidelity of these methods when trained on properly sampled spatial patterns.

Turning our attention to chord length distributions we see similar results to that of the previous case study. [Fig. 14](#) compared the generated samples to the experimental reference. The agreement is strong, with the exception of a slight bias towards larger grains. While the same phenomena was observed in the previous case study, the causes seem to be different. Previously this shift seemed to be due to the

merging and smoothing some grains in particularly challenging regions, but these artifacts are not observed in the generated inconel samples. Instead, we believe the shift to be from a reduction of noise in the generated samples when compared to the original. The experimental data contains small single pixel artifacts which artificially reduce the length of chords. The regularization in the learned local approximations prevents it from learning these small, truly random, perturbations resulting in cleaner overall microstructures.

5.3. Comparison against other methods and representations

We provide a comparison against alternative generation approaches as a final showcase of the proposed framework. Our goal in this comparison is two-fold. First, we demonstrate the strengths of higher-order generation with respect to the current first-order approaches in use today. We do this through a direct comparison against the popular polycrystalline generation/reconstruction software Dream3D⁵ [85]. Second, we validate our prior claims concerning the strengths of the ROGSH basis representation – that the representation is ideally suited for learning tasks by, for example, providing sensible notions of averages and distances while accounting for symmetry. We utilize the Euler angles directly in the LGD framework (EULER-LGD), and compare the output to the ROGSH generations (ROGSH-LGD). Quaternions were excluded from this comparison because training diffusion models on quaternion representations is a non-trivial task requiring highly specialized models and techniques (e.g., [86]). We leave the exploration of quaternions for polycrystalline generation as a possible opportunity for future work. However, we would like to reiterate that a key benefit of the ROGSH basis is that it does not require such bespoke models, and can be

⁵ For the Dream3D structure the raw EBSD data was fed into the software, and the example protocols were followed for extracting the first-order microstructure statistics. The standard generation procedure was then executed using the computed statistics.




Fig. 14. Experimental (black outline) compared to generated (red) chord length distributions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

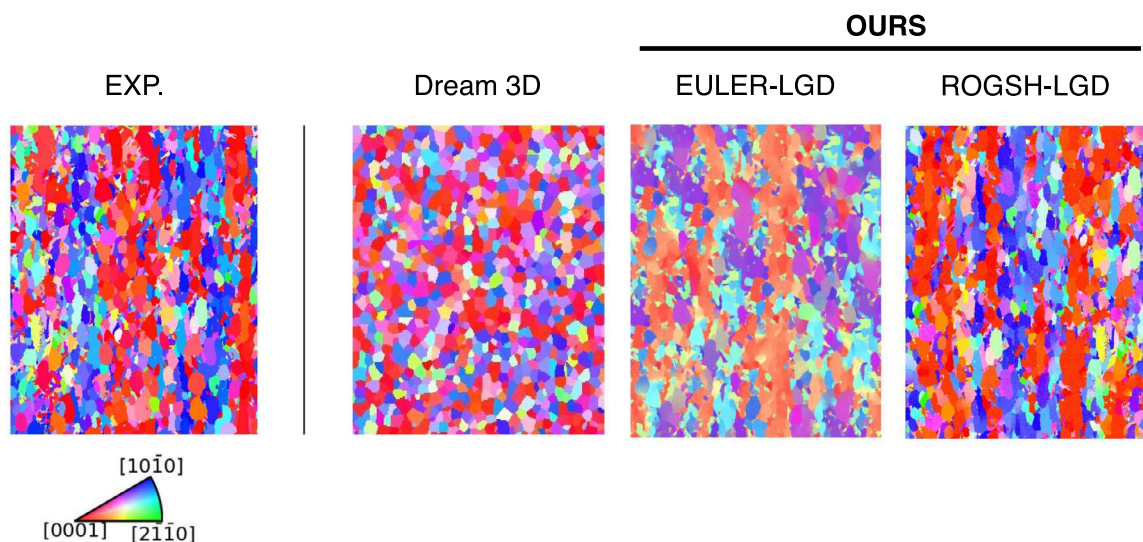


Fig. 15. A comparison of various generation approaches. The original microstructure is shown on the left followed by a synthetic generation from Dream3D, then the LGD framework using Euler angles, and finally the LGD framework using the ROGSH basis.

readily combined with any future state of the art generative modeling techniques without modification.

The generated structures in Fig. 15 contrast the three techniques against the experimental reference from Case Study 1. Comparing them qualitatively, we can immediately see the strengths of the proposed ROGSH-LGD method. The Dream3D microstructure contains the expected outcome for a first-order generation. Primarily, the long range vertical banding in the original microstructure is absent. Additionally, the grains display simple isotropic geometric forms which are representative of the complex, elongated grains in the original morphology. We also see a greater homogeneity in the grain size. The original microstructure contains both large and small, fleck like grains. The Dream3D generation grains visually seem biased towards the average

grain size. Nevertheless, the ODF (the most important first-order statistic) seems to match reasonably well as the coloring of the image is similar to the original (we will later numerically confirm this as well).

In contrast, the EULER-LGD qualitatively displays the opposite strengths and weaknesses. Inspecting this microstructure we can clearly see a shift in grain orientations as the entire structure has a “washed out” appearance. This indicates that the generated microstructure will have very little statistical agreement in both 1- and 2-order statistics since the generated orientations are fundamentally shifted. Furthermore, in comparison to the ROGSH-LGD, we see an increase in blurring artifacts as well as grains which are ill formed. This is most evident in the gray grains observable throughout. Despite the apparent statistical inaccuracies, the inclusion of 2-point statistics has resulted in some capture of the vertical banding phenomenon from the

Table 1
Metrics comparing the performance of the various generations. The most desirable result is bolded for each row.

	Dream 3D	Euler-LGD	ROGSH-LGD
ODF (MSE Error)	0.050	0.319	3.240 · 10⁻⁶
2-Point statistics (MSE Error)	0.203	0.312	5.759 · 10⁻⁶
Generation time (Min)	1.8	6.7	2.1
# Trainable parameters (Millions)	N/A	230	12

original microstructure. When comparing the ROGSH-LGD structures to the original we can see that they are nearly visually indistinguishable except for the small amount of smoothing discussed earlier.

To facilitate a quantitative comparison, we compute the ODF and 2-Point statistics via a 43 term expansion in to the generalized spherical harmonics [4] and compute the MSE errors between the original and generated structures. This comparison provides a comprehensive statistical comparison; the 1- and 2-point statistics computed on the typical GSH expansion have been shown to be a super-set of many important polycrystalline descriptors (ODF, Misorientation Distribution Function, Grain size Distributions, etc.) [4,36]. Please note that this is **not** the ROGSH basis. While statistics computed on the ROGSH basis clearly contain some of this information, evidenced by the successful efforts in this paper, no guarantees have yet been shown. Therefore we will use the GSH in the typical way for the comparison here [28].

Table 1 shows the summarized results of this comparison. First looking at the statistical accuracy, we can see the ROGSH-LGD approach hits the target statistics with orders of magnitude more accuracy than the other approaches. The Dream3D microstructures have acceptable accuracy in their first order statistics, and, as expected, do not match with respect to second-order statistics. The EULER-LGD structure is highly inaccurate with respect to all statistical measures. This is due to having incorrect predictions of the orientations as discussed earlier. Beyond statistical accuracy there are also practical implications which result for the different techniques.

We found that to get even visually reasonable results the EULER-LGD method required a much larger diffusion model (230 Million parameters compared to 12 Million). This is a clear demonstration of the sum total of the benefits of the ROGSH representation discussed in Sections 3.2 and 4.1. Without the ROGSH basis the model is now forced to learn internal representations of orientations as well as effect of symmetry and mirror/periodic boundaries in the fundamental zone. Not only did this hurt the model performance as these internal representations are sub-par, but the increased parameter count results in a much more difficult training process as well as a slower generation time.⁶ This increased computational cost can become very limiting for downstream tasks, such as dataset generation or structure optimization, where many tens of thousands of microstructures may need to be generated. By using the ROGSH basis the model can focus only on capturing high-order spatial information, thus performing generation tasks more accurately and efficiently.

6. Conclusions

In conclusion, this paper presents a novel approach to statistically conditioned polycrystalline microstructure generation through a Local-Global Decomposition-based framework. By leveraging the ROGSH basis, we address representational redundancies caused by symmetry and stabilize the training of a deep diffusion generative model. This allows our framework to successfully produce visually realistic polycrystalline microstructures while preserving desired 1- and 2-point statistics.

The case studies shown in this work highlight the tremendous fidelity diffusion models offer as local approximations within the LGD

framework. We believe that the microstructures shown here are near the upper limit of complexity that one would face when desiring to perform generation. This suggests that diffusion models may be flexible enough to model any spatial phenomena typical in materials science. However, the use of these models (and any future generative models which have yet to be developed) is predicated on identifying a suitable data representation. This avoids the construction of bespoke models made to address peculiarities unique to the domain, and allow for the direct application of the most sophisticated techniques from the machine learning and generative AI communities. This observation is one of the key messages of this work.

This work also lays the foundation for several important future efforts. Importantly, in this work, we are evaluating the generation performance of our model for the application of “statistical reconstruction”: generating a microstructure with the same statistics as an example reference. A major area for future development is exploring the extrapolation stability of these methods (i.e., changing the global conditioning statistics, while using the same local model). This would allow for the generation of microstructures with unseen 2-point statistics, and allow for exploration to novel microstructures. The extrapolation abilities for the LGD framework were investigated in our prior work [20], where we outline the necessity of picking lower order statistics which introduce viable neighborhoods. Due to the more complex lower order statistics in polycrystalline generation (ODFs rather than volume fractions), choosing compatible statistics becomes a much more challenging task. Furthermore, we observed that there is an extremely strong tie between the neighborhood distribution and the generated orientation distribution function, making it possible that the stability degrades if the target ODF is perturbed significantly away from the training data. Defining this range of suitable ODFs is a significant open problem.

Another future effort of high interest is the generation of 3D structures from 2D micrographs. The method proposed in this work can be applied to 3D generation when 3D data is present, but performing 3D generation with only 2D data remains an open problem. Recent work has shown that 2D generative models can be used in an iterative fashion on individual planes of a 3D volume [87]. This approach, while promising, has yet to be applied to polycrystalline materials due to (1) the lack of a suitable generative model (which this work directly has addressed), and (2) challenges with ensuring the generated structures are physically reasonable. From the authors experience, ensuring compatibility between the statistics from each direction is critical to get a physically sensible structure, and the difficulty of ensuring compatibility increases with the sophistication of the microstructure. Most 2D to 3D generation efforts have focused on two phase systems [31,87] where issues with compatibility seem to only arise in edge cases. For polycrystalline materials the complex local state makes these issues much more pervasive. If suitable data is available (orthogonal images along each axis), then this compatibility is guaranteed from the physical system. However this is typically not the case, and enforcing this compatibility without three orthogonal images is another high value open problem.

Code availability

Code will be made available upon request.

⁶ All tests were run on an RTX 6000 with 24 GB of VRAM.

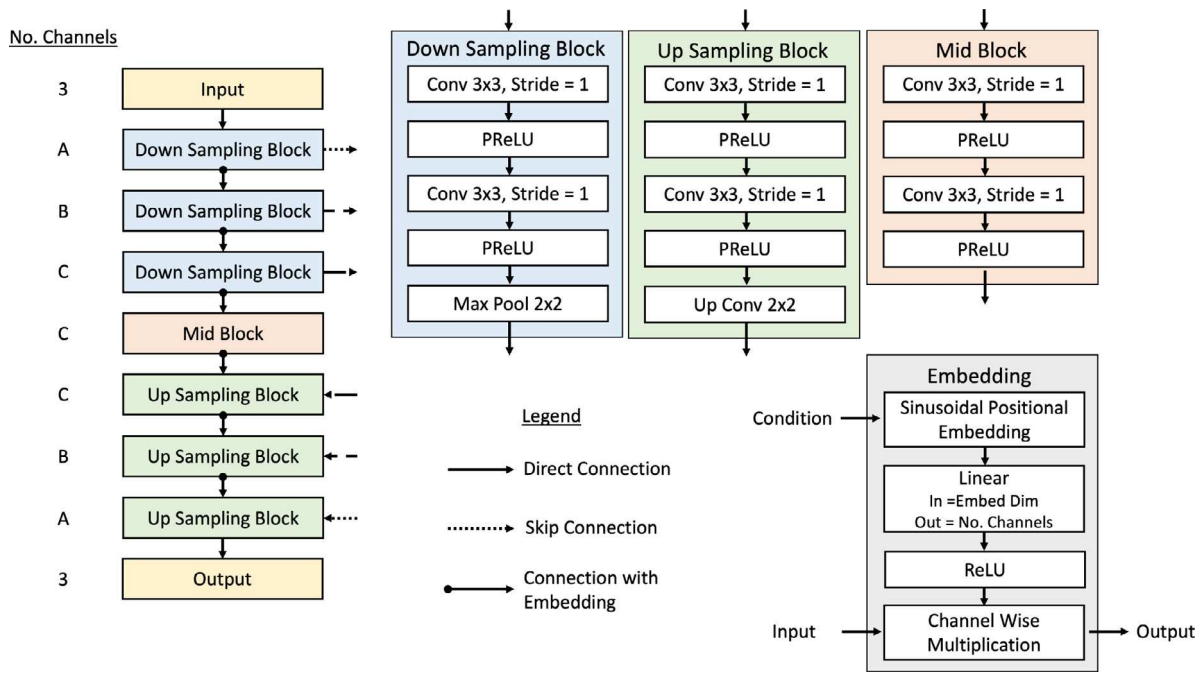


Fig. 16. The UNet model architecture used throughout this work. The overall block diagram is shown on the left. Direct connections between blocks are shown with solid arrows. Skip connections are shown with dashed arrows where arrows with the same pattern represent the same connection. Fourier time embedding is used, and is represented by arrows with a circular tail. The number of channels per block vary depending on the case study, and are denoted by A, B, and C.

CRediT authorship contribution statement

Michael O. Buzzy: Conceptualization, Methodology, Software, Data curation, Investigation, Writing – original draft, Writing – review & editing. **Andreas E. Robertson:** Conceptualization, Methodology, Software, Writing – review & editing. **Surya R. Kalidindi:** Conceptualization, Supervision, Funding acquisition, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

M.O. Buzzy would like to acknowledge Dr. Aditya Venkatraman, Cameron Gwynn, and Alexandra Nguyen for their helpful discussions and input. He would also like to thank Natalie Mata for her unwavering support. In addition, M.O. Buzzy acknowledges support from the National Science Foundation, USA under the grant NSF DMREF 2119640. A.E. Robertson and S.R. Kalidindi thank the National Science Foundation, USA for their support under NSF 2027105. Additionally, A.E. Robertson would like to acknowledge the continued support of the Jack Kent Cooke Foundation.

Appendix A. Model architecture

The UNet shown in Fig. 16 shows the model architecture used in the work. While case studies 1 and 2 both used the same overall architecture, different numbers of channels were used in each block. The more complex local neighborhoods shown in case study 2 required a larger network. The number of channels in each block are denoted by A,B,C, and their values are shown in Table 2.

Table 2
The values of A, B, and C for Case Study 1 and 2.

	Case Study 1	Case Study 2
A	64	128
B	128	256
C	128	512

Appendix B. Symmetrized ROGSH basis

All Euler angles are taken to be in the Bunge convention [36].

B.1. Hexagonal-triclinic ROGSH basis

The following are the basis functions for a sample with hexagonal crystal symmetry and triclinic sample symmetry. A direct inverse exists, and is unique within symmetry (there are six unique solutions to these equations but all are symmetrically equivalent). Note that these are symmetrized bases, so the indices are transformed from the prior unsymmetrized discussion.

$$\dot{T}_2^{0,1} = \frac{1}{2}(3 \cos^2 \Phi - 1) \Rightarrow \Phi = \cos^{-1} \sqrt{\frac{2\dot{T}_2^{0,1}}{3}} \tag{10}$$

$$Re(\dot{T}_2^{-1,1}) = \frac{\sqrt{6}}{4} \sin \varphi_1 \sin 2\Phi \Rightarrow \varphi_1 = \sin^{-1} \frac{4}{\sqrt{6} \sin 2\Phi} Re(\dot{T}_2^{-1,1}) \tag{11}$$

$$\dot{T}_6^{0,2} = -\frac{\sqrt{462}}{32} \sin^6 \Phi \cos 6\varphi_2 \Rightarrow \varphi_2 = \frac{1}{6} \cos^{-1} \left(-\frac{32\dot{T}_6^{0,2}}{\sqrt{462} \sin^6 \Phi} \right) \tag{12}$$

The first and third equations also form a reduced-order set of the surface spherical harmonics [36] (within a proportionality constant), and can be used for problems with a dependence on only two angles.

B.2. Cubic-triclinic ROGSH basis

The following are the basis functions for a sample with cubic crystal symmetry and triclinic sample symmetry. A direct inverse may exist,




Fig. 17. A comparison between the experimental and generated ODFs for AM Inconel.




Fig. 18. A comparison between the experimental and generated 1-1 ROGSH autocorrelation for AM Inconel.

but could not be found. It is therefore inverted with a nearest neighbor projection using KDTrees to within 1 degree of accuracy (on the order of experimental noise). Note that these are symmetrized bases, so the indices are transformed from the prior unsymmetrized discussion.

$$\begin{aligned} \text{Re}(\hat{T}_4^{-4,1}) = & \frac{\sqrt{30}}{192} \left((14 (\cos(\Phi) - 1)^2 \cos(4.0\varphi_1) \right. \\ & + (\cos(\Phi) + 1)^2 \cos(4.0\varphi_1 + 4.0\varphi_2)) \\ & \left. (\cos(\Phi) + 1)^2 + (\cos(\Phi) - 1)^4 \cos(4.0\varphi_1 - 4.0\varphi_2) \right) \end{aligned} \quad (13)$$

$$\hat{T}_4^{0,1} = \frac{\sqrt{21}}{48} (5 \sin^4(\Phi) \cos(4\varphi_2) + 35 \cos^4(\Phi) - 30 \cos^2(\Phi) + 3.0) \quad (14)$$

$$\begin{aligned} \hat{T}_{12}^{0,2} = & \frac{\sqrt{166305594}}{40304640} (1025 (\cos(\Phi) - 1)^6 (\cos(\Phi) + 1)^6 \cos(12\varphi_2) \\ & + 66 (\cos(\Phi) - 1)^4 (\cos(\Phi) + 1)^4 \cdot (161.0 \sin^4(\Phi) \\ & - 280.0 \sin^2(\Phi) + 120.0) \cos(8\varphi_2) + 99 (\cos(\Phi) - 1)^2 (\cos(\Phi) + 1)^2 \\ & (7429 \cos^8(\Phi) - 9044 \cos^6(\Phi) + 3230 \cos^4(\Phi) - 340 \cos^2(\Phi) + 5) \cos(4\varphi_2) \\ & + 1352078 \cos^{12}(\Phi) - 3879876 \cos^{10}(\Phi) + 4157010 \cos^8(\Phi) - 2042040 \cos^6(\Phi) \\ & + 450450 \cos^4(\Phi) - 36036 \cos^2(\Phi) + 462) \end{aligned} \quad (15)$$

The second and third equations also form a reduced-order set of the surface spherical harmonics [36] (within a proportionality constant), and can be used for problems with a dependence on only two angles.

Appendix C. Comparison of inconel statistics

Here we show the remainder of the comparisons for AM inconel. Looking at the ODFs in Fig. 17 we can see they match nearly exactly. Fig. 18 shows good agreement for the 1-1 autocorrelation consistent with the slices shown in case study 2. We also see further confirmation that the inconel samples have no long range ordering, as the 2-point statistics have a noisy random pattern away from the peak. The differences between the generated and experimental 2-point statistics is also generally randomly distributed and has little spatial dependence. This corroborates our conclusions that errors are introduced via spatial patterns which exceed the length-scale of a generative approach - analytical or learned. Since the inconel has no such patterns the error is therefore minimal.

References

- [1] N. Brodnik, C. Muir, N. Tulshibagwale, J. Rossin, M. Echlin, C. Hamel, S. Kramer, T. Pollock, J. Kiser, C. Smith, S. Daly, Perspective: Machine learning in experimental solid mechanics, *J. Mech. Phys. Solids* 173 (2023) 105231, <http://dx.doi.org/10.1016/j.jmps.2023.105231>.
- [2] D. Dimiduk, E. Holm, S. Niezgoda, Perspective on the impact of machine learning, deep learning, and artificial intelligence on materials, processes, and structures engineering, *Integr. Mater. Manuf. Innov.* 7 (2018) 157-172, <http://dx.doi.org/10.1007/s40192-018-0117-8>.
- [3] S. Torquato, *Random Heterogeneous Materials*, Springer, New York, NY, 2002.
- [4] B. Adams, S. Kalidindi, D. Fullwood, *Microstructure Sensitive Design for Performance Optimization*, Butterworth-Heinemann, Waltham, MA, 2013.

- [5] A. Marshall, S. Kalidindi, Autonomous development of a machine-learning model for the plastic response of two-phase composites from micromechanical finite element models, *JOM* 73 (2021) 2085–2095, <http://dx.doi.org/10.1007/s11837-021-04696-w>.
- [6] F. Roters, P. Eisenlohr, L. Hantcherli, D. Tjahjanto, T. Bieler, D. Raabe, Overview of constitutive laws, kinematics, homogenization and multiscale methods in crystal plasticity finite element modeling: Theory, experiments, applications, *Acta Mater.* 58 (2010) 1152.
- [7] M. Diehl, M. Groeber, C. Haase, D. Molodov, F. Roters, D. Raabe, Identifying structure-property relationships through DREAM.3D representative volume elements and DAMASK crystal plasticity simulations: An integrated computational materials engineering approach, *JOM* 69 (2017) 848–855, <http://dx.doi.org/10.1007/s11837-017-2303-0>.
- [8] S. Hashemi, S.R. Kalidindi, Gaussian process autoregression models for the evolution of polycrystalline microstructures subjected to arbitrary stretching tensors, *Int. J. Plast.* 162 (2023) 103532, <http://dx.doi.org/10.1016/j.ijplas.2023.103532>.
- [9] S. Hashemi, S. Kalidindi, A machine learning framework for the temporal evolution of microstructure during static recrystallization of polycrystalline materials simulated by cellular automaton, *Comput. Mater. Sci.* 188 (2021) 110132, <http://dx.doi.org/10.1016/j.commatsci.2020.110132>.
- [10] Y. Gao, Y. Liu, Reliability-based topology optimization with stochastic heterogeneous microstructure properties, *Mater. Des.* (2021) <http://dx.doi.org/10.1016/j.matdes.2021.109713>.
- [11] S. Krishnamoorthi, R. Bandyopadhyay, M.D. Sangid, A microstructure-based fatigue model for additively manufactured Ti-6Al-4V, including the role of prior β boundaries, *Int. J. Plast.* 163 (2023) 103569, <http://dx.doi.org/10.1016/j.ijplas.2023.103569>.
- [12] A. Generale, S. Kalidindi, Reduced-order models for microstructure-sensitive effective thermal conductivity of woven ceramic matrix composites with residual porosity, *Compos. Struct.* 274 (2021) 114399, <http://dx.doi.org/10.1016/j.compstruct.2021.114399>.
- [13] M.C. Barry, J.R. Gissinger, M. Chandross, K.E. Wise, S.R. Kalidindi, S. Kumar, Voxelized atomic structure framework for materials design and discovery, *Comput. Mater. Sci.* 230 (2023) 112431, <http://dx.doi.org/10.1016/j.commatsci.2023.112431>.
- [14] D. Khatamsaz, B. Vela, P. Singh, D.D. Johnson, D. Allaire, R. Arróyave, Multi-objective materials Bayesian optimization with active learning of design constraints: Design of ductile refractory multi-principal-element alloys, *Acta Mater.* 236 (2022) 118133, <http://dx.doi.org/10.1016/j.actamat.2022.118133>.
- [15] N. Wilson, D. Wilhelm, X. Qian, R. Arróyave, X. Qian, Batch active learning for accelerating the development of interatomic potentials, *Comput. Mater. Sci.* 208 (2022) 111330, <http://dx.doi.org/10.1016/j.commatsci.2022.111330>.
- [16] N.N. Vlassis, W. Sun, Denoising diffusion algorithm for inverse design of microstructures with fine-tuned nonlinear material properties, *Comput. Methods Appl. Mech. Engrg.* 413 (2023) 116126, <http://dx.doi.org/10.1016/j.cma.2023.116126>.
- [17] J. Jung, J.I. Yoon, H.K. Park, H. Jo, H.S. Kim, Microstructure design using machine learning generated low dimensional and continuous design space, *Materialia* 11 (2020) 100690, <http://dx.doi.org/10.1016/j.mtla.2020.100690>.
- [18] A. Generale, A. Robertson, C. Kelly, S. Kalidindi, Inverse stochastic microstructure design, 2023, <http://dx.doi.org/10.2139/ssrn.4590691>, SSRN: Preprint.
- [19] A. Robertson, S. Kalidindi, Efficient generation of N-field microstructures from 2-point statistics using multi-output Gaussian random fields, *Acta Mater.* 232 (2022) 117927, <http://dx.doi.org/10.1016/j.actamat.2022.117927>.
- [20] A. Robertson, C. Kelly, M. Buzzy, S. Kalidindi, Local-global decompositions for conditional microstructure generation, 2023, <http://dx.doi.org/10.2139/ssrn.4388214>, URL <https://ssrn.com/abstract=4388214>.
- [21] Y. Gao, Y. Jiao, Y. Liu, Ultra-efficient reconstruction of 3D microstructure and distribution of properties of random heterogeneous materials containing multiple phases, *Acta Mater.* 204 (2021) 116526, <http://dx.doi.org/10.1016/j.actamat.2020.116526>.
- [22] P. Seibert, M. Ambati, A. Rabloff, M. Kastner, Reconstructing random heterogeneous media through differentiable optimization, *Comput. Mater. Sci.* 196 (2021) 110455, <http://dx.doi.org/10.1016/j.commatsci.2021.110455>.
- [23] P. Seibert, A. Rabloff, M. Ambati, M. Kastner, Descriptor-based reconstruction of three-dimensional microstructures through gradient-based optimization, *Acta Mater.* 227 (2022) 117667, <http://dx.doi.org/10.1016/j.actamat.2022.117667>.
- [24] S. Torquato, Effective stiffness tensor of composite media: 1. Exact series expansions, *J. Mech. Phys. Solids* 45 (1997) 1421–1448.
- [25] M. Safdari, M. Baniassadi, H. Garmestani, M. Al-Haik, A modified strong-contrast expansion for estimating the effective thermal conductivity of multiphase heterogeneous materials, *J. Appl. Phys.* 112 (2012) 114318.
- [26] D.M.d. Zapiain, J. Stewart, R. Dingreville, Accelerating phase field based microstructure evolution predictions via surrogate models trained by machine learning methods, *NPJ Comput. Mater.* 3 (2021) 1–11, <http://dx.doi.org/10.1038/s41524-020-00471-8>.
- [27] N. Paulson, M. Priddy, D. McDowell, S. Kalidindi, Reduced-order microstructure-sensitive protocols to rank-order the transition fatigue resistance of polycrystalline microstructures, *Int. J. Fatigue* 119 (2019) 1, <http://dx.doi.org/10.1016/j.ijfatigue.2018.09.011>.
- [28] N. Paulson, M. Priddy, D. McDowell, S. Kalidindi, Reduced-order structure-property linkages for polycrystalline microstructures based on 2-point statistics, *Acta Mater.* 129 (2017) 428, <http://dx.doi.org/10.1016/j.actamat.2017.03.009>.
- [29] P.R. Kaundinya, K. Choudhary, S.R. Kalidindi, Machine learning approaches for feature engineering of the crystal structure: Application to the prediction of the formation energy of cubic compounds, 2021, <http://dx.doi.org/10.48550/arXiv.2105.11319>, ArXiv.
- [30] P. Altschuh, Y.C. Yabansu, J. Hötzer, M. Selzer, B. Nestler, S.R. Kalidindi, Data science approaches for microstructure quantification and feature identification in porous membranes, *J. Membr. Sci.* 540 (2017) 88–97, <http://dx.doi.org/10.1016/j.memsci.2017.06.020>.
- [31] S. Kench, S. Cooper, Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion, *Nat. Mach. Intell.* 3 (2021) 299–305, <http://dx.doi.org/10.1038/s42256-021-00322-1>.
- [32] P. Zelaia, J. Cheng, J. Mayeur, A. Ziabari, M. Kirka, Digital polycrystalline microstructure generation using diffusion probabilistic models, 2023, <http://dx.doi.org/10.2139/ssrn.4419461>, SSRN.
- [33] A. Senthilnathan, P. Acar, M.D. Graef, Markov random field based microstructure reconstruction using the principal image moments, *Mater. Charact.* 178 (2021) 111281, <http://dx.doi.org/10.1016/j.matchar.2021.111281>.
- [34] I. Javaher, V. Sundararaghavan, Polycrystalline microstructure reconstruction using Markov random fields and histogram matching, *Comput. Aided Des.* 120 (2020) 102806, <http://dx.doi.org/10.1016/j.cad.2019.102806>.
- [35] M. Groeber, S. Ghosh, M. Uchic, D. Dimiduk, A framework for automated analysis and simulation of 3D polycrystalline microstructures. Part 2: Synthetic microstructure generation, *Acta Mater.* 56 (2008) 1274–1287, <http://dx.doi.org/10.1016/j.actamat.2007.11.040>.
- [36] H.-J. Bunge, *Texture Analysis in Materials Science: Mathematic Methods*, Butterworth & Co., Berlin, 1982.
- [37] C. Hammond, *The Basics of Crystallography and Diffraction*, vol. 21, International Union of Crystallography texts on crystallography, 2015.
- [38] L. Mosser, O. Dubrule, M. Blunt, Stochastic reconstruction of oolitic limestone by generative adversarial networks, *Transp. Porous Med.* 125 (2018) 81–103, <http://dx.doi.org/10.1007/s11242-018-1039-9>.
- [39] C. Dureth, P. Seibert, D. Rucker, S. Handford, M. Kastner, M. Gude, Conditional diffusion-based microstructure reconstruction, 2022, ArXiv.
- [40] J. Tang, X. Geng, D. Li, Y. Shi, J. Tong, H. Xiao, F. Peng, Machine learned-based microstructure prediction during laser sintering of alumina, *Sci. Rep.* 11 (2021) 10724, <http://dx.doi.org/10.1038/s41598-021-89816-x>.
- [41] A. Bhaduri, A. Gupta, A. Olivier, L. Graham-Brady, An efficient optimization based microstructure reconstruction approach with multiple loss functions, *Comput. Mater. Sci.* 199 (2021) 110709, <http://dx.doi.org/10.1016/j.commatsci.2021.110709>.
- [42] R. Bostanabad, A. Bui, W. Xie, D. Apley, W. Chen, Stochastic microstructure characterization and reconstruction via supervised learning, *Acta Mater.* 103 (2016) 89–102, <http://dx.doi.org/10.1016/j.actamat.2015.09.044>.
- [43] X. Liu, V. Shapiro, Random heterogeneous materials via texture synthesis, *Comput. Mater. Sci.* 99 (2015) 177–189, <http://dx.doi.org/10.1016/j.commatsci.2014.12.017>.
- [44] S. Cheng, Y. Jiao, Y. Ren, Data-driven learning of 3-point correlation functions as microstructure representations, *Acta Mater.* 229 (2022) 117800, <http://dx.doi.org/10.1016/j.actamat.2022.117800>.
- [45] D. Jangid, N. Brodnik, M. Goebel, A. Khan, S. Majeti, M. Echlin, S. Daly, T. Pollock, B. Manjunath, Adaptable physics-based super-resolution for electron backscatter diffraction maps, *NPJ: Comput. Mater.* 8 (2022) 255, <http://dx.doi.org/10.1038/s41524-022-00924-2>.
- [46] D. Jangid, N. Brodnik, M. Echlin, T. Pollock, S. Daly, B. Manjunath, Q-RBSA: High-resolution 3D EBSD map generation using an efficient quaternion transformer network, 2023, <http://dx.doi.org/10.48550/arXiv.2303.10722>, ArXiv.
- [47] P. Seibert, M. Husert, M. Wollner, K. Kalina, M. Kastner, Fast reconstruction of microstructures with ellipsoidal inclusions using analytic descriptors, 2023, <http://dx.doi.org/10.48550/arXiv.2306.08316>, ArXiv.
- [48] Y. Jiao, F. Stillinger, S. Torquato, A superior descriptor of random textures and its predictive capacity, *Proc. Natl. Acad. Sci. USA* 106 (2009) 17634–17639, <http://dx.doi.org/10.1073/pnas.0905919106>.
- [49] T. Hsu, W.K. Epting, H. Kim, H.W. Abernathy, G.A. Hackett, A.D. Rollett, P.A. Salvador, E.A. Holm, Microstructure generation via generative adversarial network for heterogeneous, topologically complex 3D materials, *JOM* 73 (2021) 90–102, <http://dx.doi.org/10.1007/s11837-020-04484-y>.
- [50] O. Ogoke, K. Johnson, M. Glinsky, C. Laursen, S. Kramer, A. Farimani, Deep-learned generators of porosity distributions produced during additive manufacturing, 2022, ArXiv: Submitted to Additive Manufacturing.
- [51] R. Quey, P. Dawson, F. Barbe, Large-scale 3D random polycrystals for the finite element method: Generation, meshing and remeshing, *Comput. Methods Appl. Mech. Engrg.* 200 (17–20) (2011) 1729–1745.
- [52] M. Prasad, N. Vajragupta, A. Hartmaier, Kanapy: A python package for generating complex synthetic polycrystalline microstructures, *J. Open Source Softw.* 4 (2019) 1732, <http://dx.doi.org/10.21105/joss.01732>.

- [53] S. Mandal, J. Lao, S. Donegan, A. Rollett, Generation of statistically representative synthetic three-dimensional microstructures, *Scr. Mater.* 146 (2018) 128–132, <http://dx.doi.org/10.1016/j.scriptamat.2017.11.034>.
- [54] M.G. Chapman, M.N. Shah, S.P. Donegan, J.M. Scott, P.A. Shade, D. Menasche, M.D. Uchic, 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 (2021) 129–141.
- [55] D. Fullwood, S. Niezgoda, B. Adams, S. Kalidindi, Microstructure sensitive design for performance optimization, *Prog. Mater. Sci.* 55 (2010) 477–562, <http://dx.doi.org/10.1016/j.pmatsci.2009.08.002>.
- [56] S. Niezgoda, Y. Yabansu, S. Kalidindi, Understanding and visualizing microstructure and microstructure variance as a stochastic process, *Acta Mater.* 59 (2011) 6387–6400, <http://dx.doi.org/10.1016/j.actamat.2011.06.051>.
- [57] W. Brown Jr., Solid mixture permittivities, *J. Chem. Phys.* 23 (1955) 1514–1517.
- [58] E. Kroner, Bounds for effective elastic moduli of disordered materials, *J. Mech. Phys. Solids* 25 (1977) 137–155.
- [59] D. Fullwood, B. Adams, S. Kalidindi, A strong contrast homogenization formulation for multi-phase anisotropic materials, *J. Mech. Phys. Solids* 56 (2008) 2287–2297.
- [60] T. Fast, O. Wodo, B. Ganapathysubramanian, S. Kalidindi, Microstructure taxonomy based on spatial correlations: Application to microstructure coarsening, *Acta Mater.* 108 (2016) 176, <http://dx.doi.org/10.1016/j.actamat.2016.01.046>.
- [61] A. Robertson, S. Kalidindi, Digital representation and quantification of discrete dislocation structures, *JOM* 73 (2021) 2143–2158, <http://dx.doi.org/10.1007/s11837-021-04669-z>.
- [62] Y.C. Yabansu, P. Steinmetz, J. Hötzer, S.R. Kalidindi, B. Nestler, Extraction of reduced-order process-structure linkages from phase-field simulations, *Acta Mater.* 124 (2017) 182–194, <http://dx.doi.org/10.1016/j.actamat.2016.10.071>.
- [63] M. Yuan, S. Paradiso, B. Meredig, S. Niezgoda, Machine learning-based reduced order polycrystalline modelling for ICME applications, *Integr. Mater. Manuf. Innov.* 7 (2018) 214–230, <http://dx.doi.org/10.1007/s40192-018-0123-x>.
- [64] A. Castillo, A. Venkatraman, S. Kalidindi, Mechanical responses of primary- α Ti grains in polycrystalline samples: Part II – Bayesian estimation of the crystal-level elastic-plastic mechanical properties from spherical-indentation measurements, *Integr. Mater. Manuf. Innov.* 10 (2021) 99–114, <http://dx.doi.org/10.1007/s40192-021-00204-9>.
- [65] J. Rossin, P. Leser, K. Pusch, C. Frey, S. Vogel, A. Saville, C. Torbet, A. Clarke, S. Daly, T. Pollock, Single crystal elastic constants of additively manufactured components determined by resonant ultrasound spectroscopy, *Mater. Charact.* 192 (2022) 112244, <http://dx.doi.org/10.1016/j.matchar.2022.112244>.
- [66] Y. Song, J. Sohl-Dickstein, D.P. Kingma, A. Kumar, S. Ermon, B. Poole, Score-based generative modeling through stochastic differential equations, in: *International Congress for Learning Representation*, 2021, pp. 1–36.
- [67] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, 2020, CoRR abs/2006.11239. [arXiv:2006.11239](https://arxiv.org/abs/2006.11239), URL <https://arxiv.org/abs/2006.11239>.
- [68] C. Luo, Understanding diffusion models: A unified perspective, 2022, <http://dx.doi.org/10.48550/arXiv.2208.11970>, ArXiv.
- [69] Y. Song, S. Ermon, Generative modeling by estimating gradients of the data distribution, in: *NeurIPS 2019*, 2019.
- [70] Y. Song, D. Kingma, How to train your energy based models, 2021, pp. 1–22, ArXiv 1.
- [71] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, D. Fleet, Video diffusion models, 2022, <http://dx.doi.org/10.48550/arXiv.2204.03458>, ArXiv.
- [72] W. Harvey, S. Naderiparizi, V. Masrani, C. Wilbach, F. Wood, Flexible diffusion modeling of long videos, 2022, <http://dx.doi.org/10.48550/arXiv.2205.11495>, ArXiv.
- [73] N. Anand, T. Achim, Protein structure and sequence generation with equivariant denoising diffusion probabilistic models, 2022, <http://dx.doi.org/10.48550/arXiv.2205.15019>, ArXiv.
- [74] E. Hoogeboom, V. Satorras, C. Vignac, M. Welling, Equivariant diffusion for molecule generation in 3D, in: *International Conference on Machine Learning*, 2022, <http://dx.doi.org/10.48550/arXiv.2205.15019>.
- [75] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, M. Norouzi, Palette: Image-to-image diffusion model, in: *ACM SIGGRAPH Conference Proceedings*, 2022, <http://dx.doi.org/10.1145/3528233.3530757>.
- [76] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, M. Chen, Hierarchical text-conditional image generation with CLIP latents, 2022, <http://dx.doi.org/10.48550/arXiv.2204.06125>, ArXiv.
- [77] D.P. Kingma, P. Dhariwal, Glow: Generative flow with invertible 1x1 convolutions, in: *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [78] A. Muth, A. Venkatraman, R. John, A. Pilchak, S.R. Kalidindi, D.L. McDowell, Neighborhood spatial correlations and machine learning classification of fatigue hot-spots in Ti–6Al–4V, *Mech. Mater.* 182 (2023) 104679.
- [79] F. Bachmann, R. Hielscher, H. Schaeben, Texture analysis with MTEX–Free and open source software toolbox, *Solid State Phenomena* 160 (2010) 63–68.
- [80] J.L. Bentley, Multidimensional binary search trees used for associative searching, *Commun. ACM* 18 (9) (1975) 509–517.
- [81] D. Fullwood, S. Niezgoda, S. Kalidindi, Microstructure reconstruction from 2-point statistics using phase recovery algorithms, *Acta Mater.* 56 (2008) 942–948, <http://dx.doi.org/10.1016/j.actamat.2007.10.044>.
- [82] C. Yeong, S. Torquato, Reconstructing random media, *Phys. Rev. E* 57 (1998) 495–506.
- [83] Y. Jiao, F. Stillinger, S. Torquato, Modeling heterogeneous materials via two-point correlations. II. Algorithmic details and applications, *Phys. Rev. E* 77 (2008) 031135, <http://dx.doi.org/10.1103/PhysRevE.77.031135>.
- [84] Z. Xu, A. Liu, X. Wang, Influence of macrozones on the fatigue cracking behavior and fracture mechanisms of rolled Ti–6Al–4V alloy, *Mater. Sci. Eng. A* 824 (2021) 141824.
- [85] M. Groeber, M. Jackson, Dream3D: A digital representation environment for the analysis of microstructure in 3D, *Integr. Mater. Manuf. Innov.* 3 (2014) 56–72, <http://dx.doi.org/10.1186/2193-9772-3-5>.
- [86] A. Leach, S.M. Schmon, M.T. Degiacomi, C.G. Willcocks, Denoising diffusion probabilistic models on so(3) for rotational alignment, in: *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, 2022.
- [87] K.-H. Lee, G.J. Yun, Multi-plane denoising diffusion-based dimensionality expansion for 2D-to-3D reconstruction of microstructures with harmonized sampling, 2023, [arXiv:2308.14035](https://arxiv.org/abs/2308.14035).