

Polycrystalline Microstructure Reconstruction Using Markov Random Fields and Histogram Matching[☆]

Iman Javaheri^{a,b}, Veera Sundararaghavan^{b,*}

^a NASA Langley Research Center, Hampton, VA, 23681, USA

^b Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, 48109, USA

ARTICLE INFO

Article history:

Received 16 September 2019

Received in revised form 13 November 2019

Accepted 11 December 2019

Keywords:

Polycrystalline microstructure

Data-driven inference

Markov random field

Histogram matching

Grain size distribution

Orientation distribution function

ABSTRACT

A new numerical method is presented for reconstructing three-dimensional (3D) microstructures from two-dimensional (2D) sections, imaged on orthogonal planes, by exploiting the complete red–green–blue (RGB) color space. The algorithm reconstructs 3D models through sampling voxel neighborhoods to representative 2D micrographs, based upon a Markovian assumption. The sampling is followed by an optimization procedure, ensuring smoothness across the orthogonal sections of the synthesized voxels. Previous 3D Markov random field (MRF) microstructure reconstruction techniques were restricted to traditional grayscale images only. This method now enables the use of the entire RGB spectrum, employing a histogram matching step. This paper examines the algorithm's accurate representation of orientations and morphologies, encompassing a variety of micrographs from electron backscatter diffraction (EBSD) and polarized light microscopy.

Published by Elsevier Ltd.

1. Introduction

Many aspects of microstructures (e.g., size, shape, and connectivity of phases in a multiphase material) are not readily apparent from a two-dimensional (2D) sectional measurement [1]. Yet, the majority of conventional microstructural characterization procedures are limited to only 2D observations [2,3]. Three-dimensional (3D) microstructural information is essential to computational techniques, used for understanding structure–property relationships. As a result, advances in destructive metallographic methods have gained considerable attention for a 3D model generation. The serial sectioning method and its variants are now commonly used to derive 3D microstructural information from 2D sectional images [4–6]. In serial sectioning, layers of controlled thickness are repeatedly removed. This allows for image acquisition of each section, followed by reconstruction of the layers into a 3D representation of the microstructure. The 3D data collected via serial sectioning can then be used as the basis for computer simulations, such as crystal plasticity modeling of mechanical deformation [7,8]. The drawback of these characterization techniques is that they are fairly laborious for

routine engineering applications, due to the time and the effort involved. In response to these limitations several non-destructive experimental techniques for 3D microstructural information acquisition have recently been realized. X-ray based computed tomography (CT) methods are generally deficient in achieving the desired resolution for detailed analyzes of many phenomena in metallic materials [9–11]. High-energy diffraction microscopy (HEDM) techniques can probe the microstructures at or below the grain-scale more accurately [12–15]. However, a lack of general availability for high-energy beam sources restricts their common application in computational modeling development. For these reasons, new numerical techniques that can rapidly synthesize 3D microstructural models have gained significant interests in the materials community.

Advances in microscopy have allowed for data collection on the surfaces of a specimen at a rate of several hundred points per second [16]. As a result, the direct problem of capturing 2D sections using optical or micro-diffraction techniques is considerably more straightforward, compared to the aforementioned 3D acquisition procedures. Current numerical techniques for computer-aided microstructure reconstruction methods, such as geometry-based (e.g., Voronoi tessellations) or feature-based (e.g., simulated annealing), often use 2D sectional information to generate 3D microstructures. Feature-based reconstruction algorithms have long been used to categorize various microstructural snapshots based on common sets of underlying features [17–26]. Some of these features include marginal histograms [17], multi-resolution filter outputs (e.g., Gaussian and wavelet filters)

[☆] No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to <https://doi.org/10.1016/j.cad.2019.102806>.

* Corresponding author.

E-mail addresses: imanajv@umich.edu (I. Javaheri), veeras@umich.edu (V. Sundararaghavan).

Nomenclature	
w	Sampling window size (in lattice units)
G	Number of color levels in each RGB channel
\mathbf{n}	Axis of rotation associated with each node in the fundamental Rodriguez space
θ	Counterclockwise rotation angle associated with each node in the fundamental Rodriguez space
\mathbf{V}	Synthesized 3D microstructure
\mathbf{V}_v	RGB triplet of voxel v in the synthesized 3D microstructure
$\mathbf{V}_v^x, \mathbf{V}_v^y, \mathbf{V}_v^z$	Spatial neighborhoods of voxel v in sections orthogonal to x , y , and z axes
$\mathbf{V}_{v,u}^x, \mathbf{V}_{v,u}^y, \mathbf{V}_{v,u}^z$	RGB triplets of voxel u in respective \mathbf{V}_v^x , \mathbf{V}_v^y , and \mathbf{V}_v^z sections
$\mathbf{S}^x, \mathbf{S}^y, \mathbf{S}^z$	Orthogonal input 2D images corresponding to x , y , and z axes
$\mathbf{S}^{x,w}, \mathbf{S}^{y,w}, \mathbf{S}^{z,w}$	Windows of size w in input 2D micrographs corresponding to x , y , and z axes
$\mathbf{S}_u^{x,w}, \mathbf{S}_u^{y,w}, \mathbf{S}_u^{z,w}$	RGB triplets of pixel u in respective $\mathbf{S}^{x,w}$, $\mathbf{S}^{y,w}$, and $\mathbf{S}^{z,w}$ windows
$\mathbf{S}_v^x, \mathbf{S}_v^y, \mathbf{S}_v^z$	Best-matching windows to neighborhoods of voxel v in respective x , y , and z input 2D images
$\mathbf{S}_{v,u}^x, \mathbf{S}_{v,u}^y, \mathbf{S}_{v,u}^z$	RGB triplets of pixel u in respective \mathbf{S}_v^x , \mathbf{S}_v^y , and \mathbf{S}_v^z windows
$\omega_{v,u}^x, \omega_{v,u}^y, \omega_{v,u}^z$	Gaussian weighting factors associated with pixel u in respective $\mathbf{S}^{x,w}$, $\mathbf{S}^{y,w}$, and $\mathbf{S}^{z,w}$ windows
$E(\cdot)$	Cost function based on a Euclidean distance of matching windows
$H_{2D}(\cdot)$	Reference cumulative probability distribution function of colors for input 2D images
$H_{3D}(\cdot)$	Cumulative probability distribution function of colors for synthesized 3D model
$P(\cdot)$	Apparent 2D grain-size distribution function
$F(\cdot)$	Actual 3D grain-size distribution function

[18,19,27], and point probability functions (e.g., autocorrelation function) [20–25]. Although these methods are efficient at modeling the global features of the microstructure, they usually fail at capturing the local information in the form of per-pixel data, when reconstructing polycrystalline structures [19]. A popular geometry-based tool for 3D reconstructions of microstructures is DREAM.3D [28–30]. DREAM.3D extracts statistical information from images, such as grain morphology as well as orientational features. From such information it generates a full 3D geometrical model, capturing the overall crystalline structure. The super-ellipsoidal grain approach associated with DREAM.3D reconstruction tends to idealize the crystal morphology [31]. Hence, the above-mentioned numerical methods run into difficulties when modeling complexities of microstructures such as non-equilibrium grain structures, non-convex grains, twins, second phases, precipitates, and cell structures [32,33]. Such features

naturally arise from materials processing, playing an important role in their properties and performance. Accordingly, the need to develop robust numerical algorithms for 3D microstructure reconstruction is critical. Such algorithms can provide insight into the microstructure–property relationship of materials; ultimately advancing the development of smarter, multifunctional, and more sustainable structural systems.

It is experimentally observed that sufficiently large slices along a particular direction taken from a 3D polycrystalline microstructure generally ‘look alike’ [26,34–37]. This gives evidence to the notion of stationary probability distribution, which underlies microstructure formations [38,39]. Consequently, the inverse problem of synthesizing 3D polycrystalline structures from 2D sectional images taken along the orthogonal directions is of specific interest to this paper. Here, the 2D input sections are in the form of pixels that contain red–green–blue (RGB) color values representing different constituent phases or orientations. The outcome of the inverse problem is a 3D solid microstructural model made of voxels (i.e., volumetric pixels) that are colored consistently with the 2D orthogonal input images, such that any arbitrary x , y , or z slice of the 3D model is ‘similar’ to its corresponding 2D input micrograph. The measurement of similarities between the slices of the synthesized 3D microstructure and the given 2D sectional images is based on a least-square distance (i.e., Euclidean-norm) cost function. Minimizing the cost function ensures that the local neighborhood of any voxel on a 2D slice taken along the x , y , or z direction through the synthesized 3D microstructure is representative of some neighborhood in the given 2D micrograph, taken along the same direction. This approach is consistent with the Markov random field (MRF) model of texture synthesis [40–42] that recently has been successfully applied for microstructure reconstruction [39,43–46]. The underlying assumption states that the probability density of each pixel within the microstructure can be fully defined, based on a limited number of its surrounding pixels, not necessarily all the pixels in the micrograph [38,39]. To estimate an unknown voxel in the 3D synthesized model, small windows in the 2D micrographs are first identified as being the best-match to the unknown voxel’s neighborhood. Once identified, the matching windows are sampled. The state of the unknown voxel is then determined by the coloring of the center pixels in the best-matching windows. Previous works in 3D microstructure reconstruction using MRF have shown that such an approach can effectively capture microstructural features [38]. However, the results were restricted to grayscale coloring only. This work has laid the foundation for the use of the entire red–green–blue (RGB) color spectrum, through the exploitation of an image histogram matching step (also known as histogram specification). In the image processing community, this step is often performed to achieve consistent intensity (e.g., illumination or contrast) across two datasets (i.e., the reference and the target images) [47,48]. In this paper, the efficacy of this new approach to model 3D features in the polycrystalline structure using orthogonal 2D image datasets is examined. Using the MRF algorithm, it is shown that both global features such as grain size, texture distribution, and grain neighborhood correlations as well as localized information such as precipitates and grain boundary networks closely resemble those of the experimental images.

2. Methodology

The objective here is to first summarize the sampling methodology, used in Reference [38]. The histogram matching extension, which provides the ability to reconstruct color images, is described next. Following this, the computational costs associated with MRF modeling and its relation with sampling window size is discussed in detail.

2.1. Sampling algorithm

In the following discussion, let \mathbf{S}^x , \mathbf{S}^y , and \mathbf{S}^z denote the set of orthogonal slices of the microstructure along their respective x , y , and z directions. The symbol \mathbf{V} indicates the synthesized 3D microstructure. Whereas, \mathbf{V}_v represents the color of voxel v in the solid microstructure. In the previous work, G distinct color levels in the range $\{0, 1, \dots, G - 1\}$ were used to represent the microstructural data. In the current algorithm, three layers of color in the red, green, and blue channels (as a triplet) are allowed for representation of the phase information. The exploitation of RGB scheme allows for significant flexibility in the choice of the microstructural colormap, enabling direct visualization of colors throughout the reconstruction.

The following algorithm presented here is the same as in Ref. [38], except that the averaging step is performed on each color channel separately. For pedagogy, let us consider how one of the color channels in the 3D microstructure is reconstructed. The hypothesis in the sampling approach is that the color assigned to any voxel follows a Markovian property. That is, the probability density function (PDF) of a voxel given the states of its spatial neighborhood is independent of the entire dataset. Thus, here the neighbors are chosen over relatively small user-assigned windows around the voxel v . The vectors denoting the spatial neighborhood of voxel v in the slices orthogonal to the x , y , and z axes are denoted as \mathbf{V}_v^x , \mathbf{V}_v^y , and \mathbf{V}_v^z , respectively (see Fig. 1(a)). Let $\mathbf{S}^{x,w}$, $\mathbf{S}^{y,w}$, and $\mathbf{S}^{z,w}$ denote a window of the same size in the input 2D micrographs. In order to find the coloring of an unknown voxel v given the neighbor voxels in the x -plane, one needs to compute the conditional probability density of \mathbf{V}_v given colorings of x -plane neighbors of v . Explicit construction of such a probability density is often computationally intractable. An alternative approach is to estimate the most likely RGB value of the voxel v by identifying a window $\mathbf{S}^{x,w}$ that is most similar to \mathbf{V}_v^x in the 2D input micrograph. This window is denoted by \mathbf{S}_v^x (see Fig. 1(b)). Similarly, matching windows to the y and the z plane neighborhoods of voxel v are found in the corresponding 2D sectional images (denoted as \mathbf{S}_v^y and \mathbf{S}_v^z , respectively). Each of these matching windows \mathbf{S}_v^x , \mathbf{S}_v^y , and \mathbf{S}_v^z may have different RGB coloring for the center pixel. Thus, an optimization methodology is employed to effectively merge these disparate values and to identify a unique coloring for the voxel v . The optimization approach is described next.

Let the value $\mathbf{V}_{v,u}^x$ denote the RGB coloring of voxel u in the neighborhood \mathbf{V}_v^x . Let the values $\mathbf{S}_{v,u}^x$ and $\mathbf{S}_u^{x,w}$, respectively denote the RGB color of pixel u in the windows \mathbf{S}_v^x and $\mathbf{S}^{x,w}$. Consequently, the 3D microstructure can be synthesized by posing the problem as a Euclidean distance (L^2) minimization of the energy function (also known as the cost function), as formulated below:

$$E(\mathbf{V}) = \sum_{i \in \{x,y,z\}} \sum_v \sum_u \omega_{v,u}^i \|\mathbf{V}_{v,u}^i - \mathbf{S}_{v,u}^i\|^2 \quad (1)$$

Here, $\omega_{v,u}^i$ denotes a per-pixel weighting factor. In order to preserve the short-range correlations of the microstructure as much as possible, a Gaussian weighting is used such that the weights for nearby pixels are taken to be higher than pixels farther away. For instance, if the pixel u along the x -plane is placed at a relative location (y_u, z_u) (in lattice units) with respect to the voxel v , then $\omega_{v,u}^x$ is given as:

$$\omega_{v,u}^x = \frac{\exp\left(-\frac{y_u^2 + z_u^2}{2\sigma^2}\right)}{\sum_j \exp\left(-\frac{y_j^2 + z_j^2}{2\sigma^2}\right)} \quad (2)$$

The summation in the denominator is taken over all the pixels in $\mathbf{S}^{x,w}$. Here, the standard deviation σ , is taken to be $0.16 \times w$, where w refers to the window size [43]. Similarly, $\omega_{v,u}^y$ and $\omega_{v,u}^z$, respectively denoting per-pixel weighting factors in $\mathbf{S}^{y,w}$ and $\mathbf{S}^{z,w}$, can be computed.

The optimization of the energy function $E(\mathbf{V})$, is carried out in two steps: (i) the searching step, and (ii) the expectation step. In the first step, the energy function is minimized with respect to \mathbf{S}_v^i . Here, the assumption is that the most likely sample from the conditional probability distribution of the center voxel in the 3D model along a particular direction is the center pixel of the best-matching window in an experimentally-obtained 2D slice on the corresponding plane. Hence, the best-matching neighborhood of voxel v along the x -plane is selected by solving the following minimization problem:

$$\mathbf{S}_v^x = \arg \min_{\mathbf{S}^{x,w}} \sum_u \omega_{v,u}^x \|\mathbf{V}_{v,u}^x - \mathbf{S}_u^{x,w}\|^2 \quad (3)$$

This is an exhaustive search that compares all the windows in the input 2D micrograph to the corresponding x -slice neighborhood of voxel v and identifies a window that leads to a minimum weighted squared distance.

Thus, for each voxel v , a set of three best-matching neighborhoods within the experimental images are identified. Generally the center pixel values in each of these neighborhoods are composed of different RGB colorings. Yet, a unique value of v needs to be found by weighting colors pertaining to location v in not only the matching windows of voxel v but also its surroundings. This is done in the second step of the optimization procedure. Here, the optimal color of voxel v is computed by setting the derivative of the energy function to zero, with respect to \mathbf{V}_v . This leads to the following weighted-average expression for the color of voxel v :

$$\mathbf{V}_v = \left(\sum_{i \in \{x,y,z\}} \sum_u \omega_{u,v}^i \mathbf{S}_{u,v}^i \right) / \left(\sum_{i \in \{x,y,z\}} \sum_u \omega_{u,v}^i \right) \quad (4)$$

Note that the subscripts u and v are switched in the above expression, compared to Eq. (1). This implies that the optimal color of the voxel v is the weighted-average of the colors at locations corresponding to voxel v in the best-matching windows of voxels u (i.e. \mathbf{S}_u^i) found within the synthesized 3D solid microstructure. The color channels are averaged independently in the equation above. Since \mathbf{V}_v is continuously changing after each step, the set of closest input neighborhoods \mathbf{S}_v^i varies accordingly. Hence, the above two steps are repeated until convergence; that is until the set \mathbf{S}_v^i remains unchanged.

2.2. Histogram matching algorithm

The primary issue with the above equations is the assumption that the color space is continuous. Such an assumption allows for the partial derivatives of the energy function to be obtained. However, the color space is typically discrete and range bound. Consequently, the averaging performed in Eq. (4), always tends to shrink the color levels. For example the color level 0 always tends to increase, since it is repeatedly averaged with all the color levels that are greater or equal to 0. Similarly the maximum color level G of each RGB channel (typically equals to 255) has a tendency to decrease, due to it being averaged with other color levels that are always smaller or equal to G . However, the general assumption in MRF reconstruction remains, that the two datasets (i.e., input images and solid synthesized model) have the same range of color levels. To solve this issue, histogram matching is applied after each iteration, allowing for the color levels to be appropriately stretched back to values from 0 to G at each channel.

The stretching process is based on the cumulative distribution function (CDF) of colors in the input micrographs. That means, the




Fig. 1. (a) The neighborhoods of the voxel v in the slices orthogonal to the x , y , and z axes are shown. (b) The windows in the input 2D micrograph shown in dotted lines are denoted by S_v^i for $i = x, y$, and z . These windows closely resemble the neighborhoods of the voxel, v .

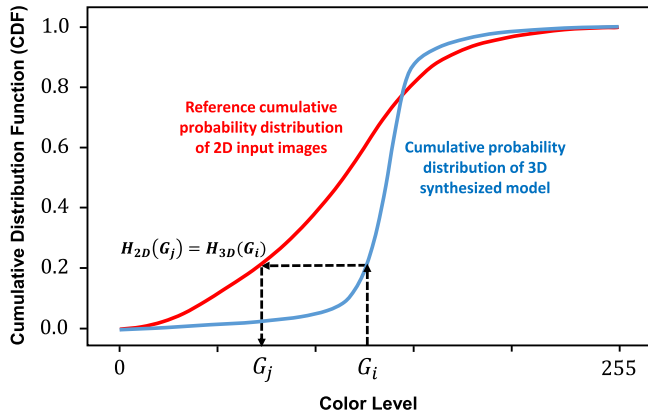


Fig. 2. Histogram matching algorithm: the color G_i in the synthesized 3D image is switched for each channel independently. The new color G_j has the same cumulative distribution function as in the 2D images.

2D micrographs are considered as the reference images, whereas the 3D solid model's color histogram is to be matched following each iteration. Here, all three orthogonal images are lumped together in order to generate the reference cumulative probability distribution. In current implementation, the histogram matching algorithm is applied separately, for each color channel. Given the reference (2D) and the target (3D) images, the color histograms are obtained by binning the pixels into G discrete intervals. Subsequently, CDFs of the two color histograms H_{2D} and H_{3D} are computed and normalized. For each voxel (with RGB color value \mathbf{V}_v) in the 3D synthesized model, the color level X in the 2D image is found such that $H_{2D}(X) = H_{3D}(\mathbf{V}_v)$. Accordingly, the color level associated with \mathbf{V}_v (in the 3D model) is replaced with color level X (from the 2D images). The histogram matching scheme is depicted in Fig. 2.

The application of histogram matching using an RGB image is illustrated in Fig. 3. The first reconstruction (shown on the top of Fig. 3) is based on the algorithm presented in Ref. [38]. Note, that phase information (e.g., grain size, shape, orientation, boundary networks, etc.) is lost within the first few iterations. Though with the current implementation (using histogram matching) the volume fraction of each color level is preserved with the input images. The result of which is the reconstruction of a representative 3D microstructure. In Fig. 3, the same 2D micrograph is considered as the reference image for both reconstruction schemes. The snapshots of the 3D solid models at iterations 0, 7, 21, and 35 are shown, respectively.

Although the general approach presented in this work is not tied to a specific input resolution, all the input images are re-sampled to 128×128 pixels prior to MRF reconstruction. As a starting condition for 3D reconstruction, a random RGB color from the input 2D images is assigned to each voxel v . The reconstruction process is then carried out in a multi-resolution (or multi-grid) fashion [38]: starting with a coarse voxel mesh while progressively interpolating the results to a finer mesh once the coarser 3D image has converged to a local optimum. The multi-resolution approach drastically increases the convergence rate to an optimum cost function. As such, three resolution levels (32^3 , 64^3 , and 128^3) are used for the following three numerical examples. Respective window sizes of 7, 9, and 11 are chosen in the three resolution levels. The pseudocode in Algorithm 1 summarizes the sampling and reconstruction schemes:

Algorithm 1: Pseudocode for 3D MRF reconstruction of size N^3 .

Require: Orthogonal 2D images of size N^2 as well as the sampling window size for each level
Ensure: Optimum cost function given in Eq. (1)
 Resize 2D images to L^2 , where $L = N/4$
 Initialize a random 3D microstructure of size L^3 , where $L = N/4$
 Initialize level index: $level \leftarrow 1$
while $level < 4$ **do**
 while $level$ unchanged **do**
 Use Eq. (3) to identify the best matching windows in orthogonal 2D images (*Searching Step*)
 Use Eq. (4) to compute the expected voxel channel values (*Expectation Step*)
 for each channel **do**
 Compute CDFs of 2D images (H_{2D}) and 3D solid model (H_{3D})
 Find color level X in the 2D images such that $H_{2D}(X) = H_{3D}(\mathbf{V}_v)$
 Replace voxel channel color with color level X (*Histogram Matching Step*)
 end
 end
 Update image size: $L \leftarrow 2L$
 Resize 3D synthesized model to L^3 and 2D images to L^2
 Update level index: $level \leftarrow level + 1$
end

2.3. Computational cost vs. window size

Although there are several free parameters throughout the reconstruction strategy, window size is chosen to be the only




Fig. 3. Effect of histogram matching: without histogram matching (as illustrated on the first row), the color space continuously shrinks and phase information is lost after first few iterations; contrarily, with histogram matching (as shown on the bottom), the color space remains consistent with the input reference image throughout the simulation.

adjustable variable for simplicity. Window size plays an important role in the MRF reconstruction model. At window sizes much smaller than the correlation lengths, false matches lead to inaccurate reconstructions. Contrarily, at very large window sizes, not enough matching windows can be identified. Additionally, high window sizes require extensive computational power for both sampling and reconstruction. Between these two extremes, there is an ideal window size that needs to be realized through numerical trials. Normally, the window size should be larger than the average microstructural features, especially at low-resolution levels. This allows the algorithm to capture global attributes (e.g., grain boundary network, inter-connectivity, size distribution, etc.) in the polycrystalline microstructure at early iterations while reconstructing finer details (e.g., precipitate, twin, etc.) at later resolution levels. Fig. 4 depicts the effect of window size (within the highest resolution level, 128^3) on the quality of synthesized 3D images. Note that only odd values are allowed for the window size, so that the sampled neighborhood is symmetric around the center voxel.

The computational cost to obtain 3D solid models is directly related to the voxel resolution. In other words, higher voxel resolutions (128^3 vs. 64^3) require more computational power to perform searching and expectation steps. Table 1 provides a breakdown of the computational cost as a function of the sampling window size associated with the highest resolution level. The processing times here correspond to the solid models shown in Fig. 4. Note, simulation times generally vary drastically depending on iteration number as well as the features being sampled. As shown in Table 1, most of the computational burden is related to identifying the best-matching windows in the orthogonal input images (i.e., searching step). This is expected as the searching stage consists of an exhaustive approach that compares all the windows in the input micrographs, in order to identify a neighborhood that leads to a minimum weighted squared distance, formulated in Eq. (3). Furthermore, note that the computational burden associated with performing the histogram matching step is typically related to the 3D voxel resolution, and it is independent of the sampling window size. Given this reason, histogram matching is excluded from Table 1. At the highest grid resolution, the typical time required to perform histogram matching is minimal (usually fewer than 5 s) when compared to expectation and searching steps.

Table 1

Breakdown of computational cost associated with highest resolution ($128 \times 128 \times 128$) level of the MRF reconstruction. For reference, all reconstructions are executed using a single processor and 4 GB of memory.

	Relative time (%) Searching step	Relative time (%) Expectation step	Average time (min)
Window Size 3	65.05	34.95	1.69
Window Size 5	71.94	28.06	5.53
Window Size 7	74.31	25.69	11.09
Window Size 9	78.80	21.20	20.37
Window Size 11	81.31	18.69	32.25

In light of the above discussion regarding computational burden caused by different window sizes, it is expected that mapping sample images to lower resolutions would significantly improve the overall computational efficiency of the current algorithm. Note, the reconstruction scheme presented in this work is not tied to a specific input resolution. Thus, any 2D image resolution can be reconstructed, if needed. The implication is that the user should specify a sufficient pixel resolution that would capture the phenomena of interest.

3. Results and discussion

In the example below, the MRF algorithm is examined by comparing grain sizes for a previously studied sample in Reference [38] using the Saltykov method. Following this, the reconstruction procedure is applied to an orthogonal set of EBSD images and polarized light microscope data of an Aluminum alloy. Thereafter, the restrictions of current MRF microstructural modeling is highlighted.

3.1. Example 1: Validation of 3D grain size

Here, a polycrystal grayscale micrograph is chosen. The image is reconstructed in the RGB color space using the present algorithm. In this example, the ability of MRF to accurately model 3D grain size distribution as predicted by the stereological formula is studied. Such validation has not been done previously for 3D MRF algorithms. Grain size distribution is essential for simulating mechanical properties in polycrystals via Hall–Petch models [49]. The Saltykov approximation method provides a closed-form expression for the 3D grain sizes, given a 2D observation. The




Fig. 4. Effect of window size: none of the sections in 3D images look identical to the sample image, but the image generated with window size of 11 has statistical correlation function similar to that of input image.




Fig. 5. An experimental 2D polycrystalline microstructure (as shown on the left) is used as an exemplar for 3D reconstruction. The 3D MRF reconstructed model is depicted on the right.

assumption is that the grains are approximately equiaxed, which is reasonable given that only one micrograph is used for all three orthogonal sections. The 3D reconstruction is shown in Fig. 5.

The Saltykov formula is used to unfold the 2D apparent grain-size distribution into the actual 3D PDF ($F(R)$), as shown in Fig. 6(a). The results are then compared against the 3D grain size distribution of the MRF reconstruction. The assumption in the Saltykov formula is that the reconstructed volume and the 2D images are of unit volume and unit area, respectively. The derivation is based on Fig. 6(b), where $z = \sqrt{R^2 - r^2}$ corresponds to a random cut location across a sphere of radius R with r being its apparent 2D grain size. Since the probability of z follows a uniform distribution (i.e., $P(z) = \frac{1}{R}$), the probability of r can be computed in a closed-form by the change of variables, as follows:

$$P_r(r) = P(z) \left| \frac{dz}{dr} \right| = \frac{r}{R\sqrt{R^2 - r^2}} \quad (5)$$

Integrating this expression leads to:

$$P(r_1 < r < r_2) = \int_{r_1}^{r_2} P_r(r) dr = \frac{1}{R} \left(\sqrt{R^2 - r_1^2} - \sqrt{R^2 - r_2^2} \right) \quad (6)$$

where R corresponds to the 3D grain size, r_1 and r_2 are lower and upper limits of the interval for 2D grain size, and $P(r_1 < r < r_2)$ is the probability to cut the section within the defined interval.

The probability that a plane cuts a sphere of radius R is $2R$ (per unit length). The likelihood of cutting such spheres not only depends on their numbers but also on their sizes, since larger spheres are more likely to intersect a random section plane, as shown in Fig. 6(c). Also, the likelihood of a sphere of radius R in the unit volume is $F(R)\Delta$. Here $\Delta = (R_{max} - R_{min})/k$ denotes the size of each bin, where k equals the total number of bins. Hence, the probability a 2D section of radius r can be numerically written by taking a product of the equation above with $2RF(R)\Delta$ and by summing it over all bins (Saltykov formula):

$$P(r_1 < r < r_2) = 2\Delta \sum_i \left(\sqrt{R_i^2 - r_1^2} - \sqrt{R_i^2 - r_2^2} \right) F(R_i) \quad (7)$$

For the Saltykov formula, the same number of bins ($k = 24$) in the 2D and 3D case is used such that the above set of equations (one equation for each 2D bin) can be exactly inverted to obtain $F(R)$ [50]. In this example, the 2D grain statistics are captured in the following manner. The incomplete grains on the outer-edge adjacent to the borders of the image are removed. For every inner grain, the maximum possible diameter at all its boundary points is computed in pixel units. The PDF of the raw diameter data is then calculated over 24 bins using a log-normal fit. The result of the Saltykov algorithm on this data is plotted as solid points in Fig. 7. To obtain the 3D grain size distribution, edge detection is first performed on the solid MRF volume, followed by segmentation of grains. In the segmentation algorithm, grains are described by the number of connected faces (six-fold connectivity) of a similar color. To obtain the grain size distribution, the same process is followed as that used in the 2D image. That is the external grains are first eliminated, and for each inner grain, the maximum diameter for every boundary voxel is obtained. As illustrated in Fig. 7, a close agreement of the Saltykov estimation and the sampled grain size distribution from the MRF reconstruction is observed. Deviations from Saltykov estimates are seen at the tails of the distribution, where the 3D reconstruction marginally predicts a higher probability of small and large grains.

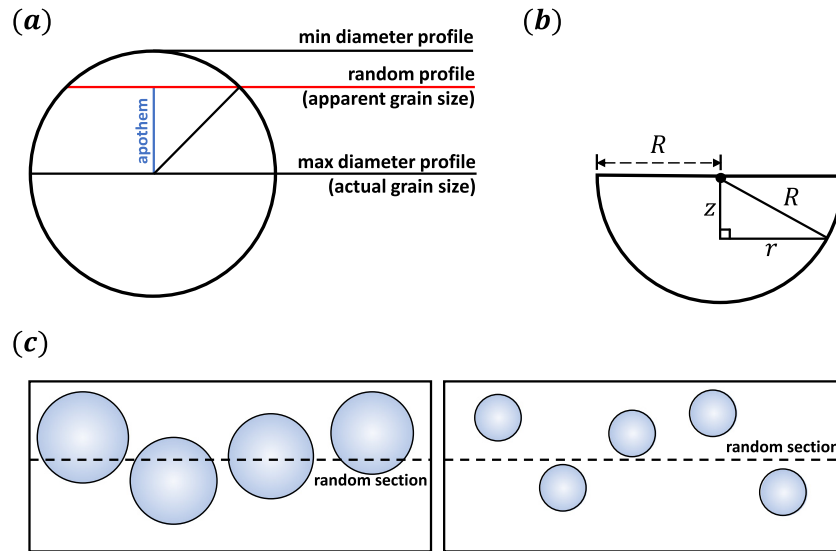


Fig. 6. Schematic explaining the Saltykov formula: (a) Grains are assumed as spheres with a random section seen in the 2D image. (b) Probability of a given 3D grain size R given a 2D section radius r can be obtained by a change of variables. (c) In the case where multiple 3D grain sizes exist, the Saltykov estimate takes into account the probability that a plane can cut a sphere of radius R . Hence, the spheres of larger radii are more likely to intersect the section plane.

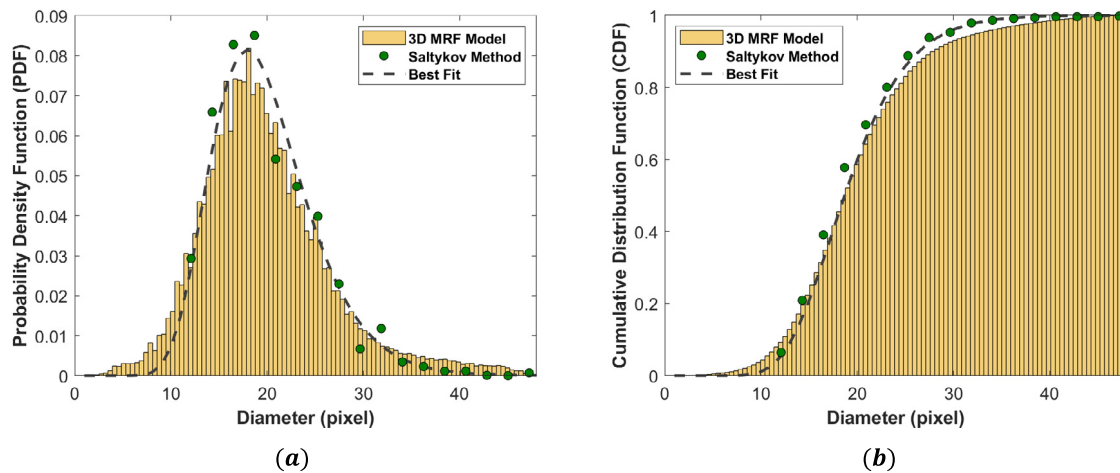


Fig. 7. (a) Comparison of the grain size distribution (PDF) of the 3D microstructure along with the Saltykov prediction. (b) Comparison of cumulative distribution function (CDF) of 3D grain sizes with Saltykov prediction.

For the 3D model presented here, one reference image is supplied as an exemplar for all three orthogonal directions. Therefore, the model is assumed to be an equiaxed microstructure. This implies that slices taken from the diagonal directions should exhibit similar statistics as orthogonal planes in the solid 3D model. Hence, for comparison, the grain size distributions of these sections are studied. The process is carried out for total of 15 orthogonal planes (five in each direction) and 18 oblique sections (three in every face-diagonal direction) across the 3D microstructure. A number of these sectional images are shown in Fig. 8.

To acquire the sectional grain statistics, same process as described for 2D input image, is employed. The partial grains adjacent to the borders are eliminated, and for each inner grain, the maximum diameter at every boundary node is calculated using pixel units. The probability density of diameters for both orthogonal and diagonal sections are then computed over 50 bins. As illustrated in Fig. 9, a close trend between the two data is observed.

3.2. Example 2: 3D reconstruction based on EBSD data

The following EBSD images, shown in Fig. 11, correspond to an Al-Li alloy after forging to 2.5-inch thickness. Here, three orthogonal sections are characterized and used for MRF reconstruction. Experimentally, EBSD data is obtained using a scanning electron microscope (SEM) equipped with detectors normal to each plane. To ensure a consistent MRF sampling across all EBSD sections, the orientations need to be remapped such that all the images are observed in the same sample coordinate system.

Here, the raw Euler angle data is mapped to nodes in a discretized orientation distribution function (ODF). The ODF is a one-point probability function. It describes the volume fractions of crystals as a function of orientation. The orientation is represented here using an axis-angle format initially proposed by Rodrigues. The format is based on the unique association of orientation with an axis of rotation, \mathbf{n} , and a counterclockwise rotation angle, θ about \mathbf{n} [51,52]. The Rodrigues' parametrization, \mathbf{r} , is scaled as $\mathbf{r} = \mathbf{n} \cdot \tan(\frac{\theta}{2})$. A proper rotation, \mathbf{R} , relates the




Fig. 8. Several sectional images of the 3D reconstructed model along diagonal (shown on the left) and orthogonal (shown on the right) directions are illustrated. The grain sizes and shapes along these sections nearly resemble the 2D reference image.

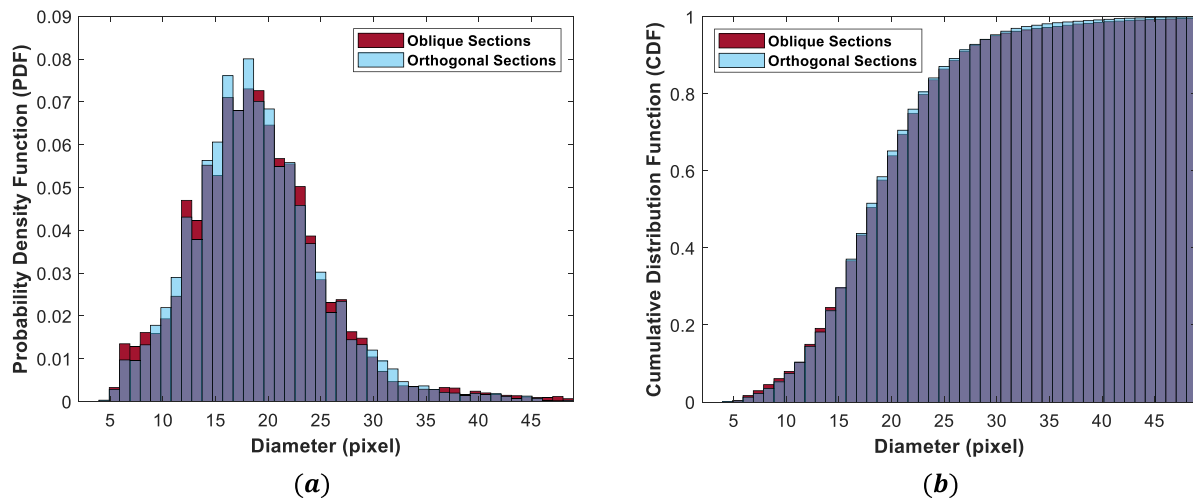


Fig. 9. (a) Comparison of probability density distributions of the 3D grain sizes along oblique and orthogonal sections. (b) Comparison of cumulative distribution function of the 3D microstructure across oblique and orthogonal directions.

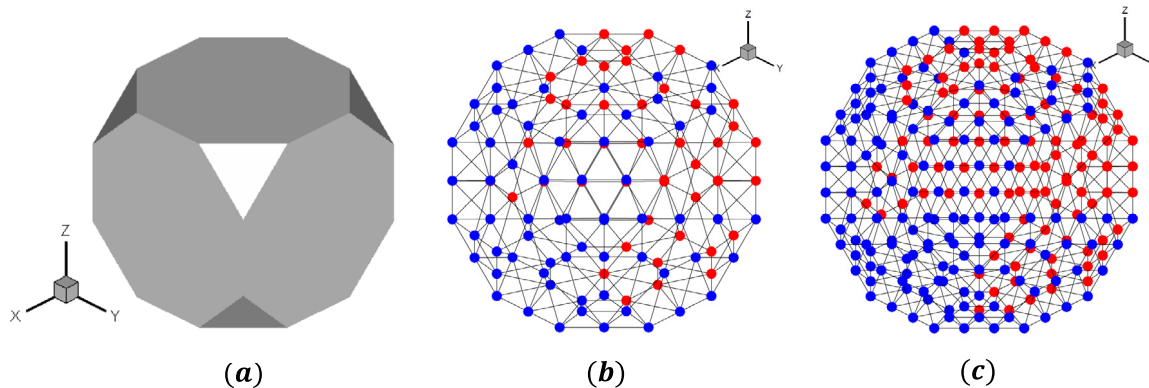


Fig. 10. (a) Rodrigues fundamental region for cubic symmetry. (b) A coarse discretization with independent nodes colored blue. (c) A finer discretization. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

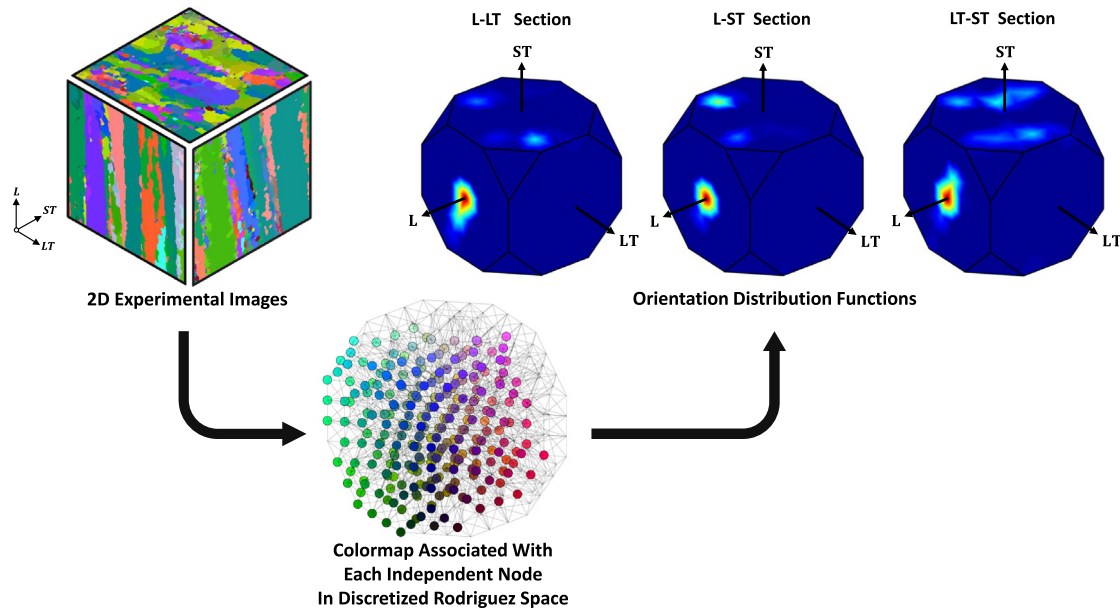


Fig. 11. The colormap shows the RGB values of individual nodes in the face-centered cubic (FCC) fundamental region. The pixels RGB values in the raw 2D orthogonal EBSD images are mapped to the closest node (in Euclidean norm sense) using the colormap. The orientation distribution functions of all three input images are similar, after adjusting the relative alignments of the detectors with respect to the sample.

sample orientation to the crystal orientation. Here, the example material (Al–Li alloy) lattice is of cubic symmetry. The 3D orientation space can be reduced to a small subset called the fundamental region that accounts for the cubic symmetry. The resulting Rodrigues fundamental region is shown in Fig. 10(a). For numerical analysis, the fundamental region is discretized using a finite-element mesh. The symmetry of the ODF is an additional constraint that must be properly considered. Due to cubic symmetries, several of the nodes in the grid are equivalent. Hence, the ODF is represented using a smaller set of independent nodes, as shown in Fig. 10(b), (c).

Each EBSD data point is mapped to the closest independent node in the ODF. If the orientation is close to a non-independent node, the node that is the symmetry equivalent to the node is chosen. The 2D EBSD images are then colored according to the nodal numbers. A simple choice for the colormap is to apply the Rodrigues vector itself as an RGB triplet, after normalization of each Rodrigues vector component to a range of [0, 255]. Fig. 11 shows the orthogonal set of images from the rolled sample; grains are elongated in the rolling direction and shortened in the transverse directions. The grains are colored based on a colormap in Fig. 11, depicting the color of each independent node in the ODF. Here, three 700×700 EBSD sections are chosen and resampled to 128×128 images for the reconstruction. The ODFs for all three sections are comparable after correcting for the relative orientation of the camera with respect to the sample, as shown in Fig. 11. The 3D reconstruction from the MRF algorithm is shown in Fig. 12(b). A number of different sections of the reconstructed model are shown in Fig. 12(c).

In this example, the averaging algorithm tends to smoothen the noise, and it supplies an overall smoother reconstruction compared to the input 2D images which consist of a significant amount of noise. Since the averaging step in the MRF algorithm introduces deviation of colors from the original 2D colors, an additional step is employed to remap the colors to the independent nodes of the ODF through the colormap shown in Fig. 11. The top four closest colors in the ODF that closely match the voxel color are found first. If it is a match within a threshold error, the voxel color remains unaltered. Otherwise, the closest colors (in Euclidean norm sense) to the colormap are identified. Of the

colors identified as best matching, one randomly picked to recolor the voxel. It is found that the resulting ODF and pole figures of the MRF reconstructed microstructure matches exceptionally well to the measured ODF, as shown in Fig. 13.

In addition, a new analysis is performed, where the reconstruction algorithm is run sequentially, using random slices of previous 3D model as the input exemplars for the next synthesis. Such an analysis exemplifies the ability of 3D MRF algorithm to preserve texture statistics, generate distinct equivalent microstructures, and most importantly demonstrates how errors can be accumulated after every reconstruction. An example for sequential reconstruction is shown in Fig. 14. These sequences reveal that our 3D reconstruction strategy succeeds in preserving the original exemplars' orientation distribution. Despite all, due to the averaging step, at some point deviations in statistics are expected to creep into reconstructions. This only happens after a significantly larger number of reconstructions.

3.3. Example 3: Reconstruction of polarized light micrograph

The third example is an optical micrograph of an Aluminum alloy from Reference [53]. In this example, the microstructure is first etched with Barker's reagent. Next, an electric field is applied to produce a thick oxide layer over the Aluminum grains. When exposed to the cross-polarized light, interference in these oxide layers produces colors that depend on the grain orientation. For reconstruction purposes, the image is preprocessed such that all grain boundaries are accentuated to demarcate the grains clearly. The 2D microstructure is depicted in Fig. 15(a). Only one sectional representation is taken, and the other sections are assumed to be identical (similar to Example 1). The microstructure is then reconstructed using MRF algorithm, and several sections of the 3D synthesized model in Fig. 15(b) is indicated in Fig. 15(c).

The goal here is to examine the nearest neighbor–grain correlation, which is a critical descriptor relating mechanical properties [54]. For this purpose, the grains are classified in terms of four predominant colors observed in the image: purple (P), blue (B), orange (O), yellow (Y). The 2D and 3D microstructures with grains marked by this classification are shown in Fig. 16 (left). For all yellow grains, the number of blue neighbors, as well as all the




Fig. 12. (a) Experimental 2D polycrystalline microstructures imaged in orthogonal planes. (b) 3D MRF reconstructed model. (c) Sectional images of the reconstructed microstructure.




Fig. 13. Comparison of the ODFs and various pole figures of the 2D sample images against the synthesized 3D microstructure.

other neighbors (purple, orange, and yellow), are counted. The same process is repeated for the remaining colors. Similar data is gathered from five slices in each direction (total of 15 slices) across the 3D microstructure. The neighbor statistics of the colors are compared in Fig. 16 (right). In this figure, the x -axis denotes the color of the chosen grain, while the color bars indicate the probability of occurrence of a yellow, orange, purple, or blue neighbor for that specific grain. There are two bars for each neighbor color, one for the 2D microstructure (located on the left hand side) and another for the 3D reconstructed model (located on the right hand side) as indicated by the legend. The range bar on the tip of the 3D bars shows the maximum and minimum values of the neighbor color percentages along the x , y , and z planes. The numbers on the bars indicate the relative percentage difference between the 2D and 3D statistics. The maximum error between the 2D and 3D neighbor statistics is only 3.32% which is associated with having a purple grain neighboring another purple grain. This example indicates that the MRF algorithm is

exceptional in maintaining the nearest grain neighborhoods in the reconstructed image, as is also observed in the two-point short-range correlation study in Reference [38].

3.4. Example 4: Limitation of Markov random field modeling

The presumption behind MRF reconstruction methodology is that given surface images are representative (i.e., having similar grain morphology and textural information) to every slice in their respective directions. This requires the user to carefully select volumetric regions in which one single image can adequately describe the entire dataset along that particular axis. As shown in previous examples, this approach is efficient in capturing a wide variety of microstructural features. However, there are samples for which above requirement is insufficient. Fig. 17 demonstrates a synthesized 3D microstructure with a non-uniform morphology. Here, grains are densely populated at $z = 128$ and scarcely generated on $z = 0$ slice. Due to the nature of the given microstructure,




Fig. 14. Texture sequence analysis is shown for MRF algorithm. Each reconstruction uses the random slices of the previous 3D model (shown on its left) as its input exemplars. Here, the orientation distributions of each reconstruction exceptionally match with the original exemplars' ODF. (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) Experimental 2D polycrystalline microstructure. (b) 3D MRF reconstructed model. (c) Sectional images of the reconstructed microstructure.

no single image can thoroughly capture the varying morphology along the z direction. As a result, to reconstruct an equivalent 3D model using MRF, one needs to take into account the statistics of additional orthogonal or oblique sections in their exemplars [46].

4. Conclusions

A numerical method for reconstructing diverse 3D multi-color microstructures from 2D micrographs imaged on orthogonal planes is presented. The algorithm reconstructs 3D solid models by matching orthogonal sections of the synthesized microstructure to the representative 2D micrographs through an iterative minimization procedure. The principal contribution in this paper

is the extension of the algorithm to true color (RGB) images via implementing a histogram matching step after every iteration. Such an extension is essential for accurate representation of descriptors such as orientation distribution functions. In this paper, the new algorithm is applied for a wide variety of morphologies, particularly equiaxed and elongated grains. Through the reconstruction of optical and EBSD datasets, it is shown that the MRF approach is able to precisely capture crucial features of polycrystalline microstructure, such as grain size, orientation distribution, and grain neighborhood correlation. The method can be useful to rapidly create a library of 3D microstructures from 2D micrographs for modeling and simulation purposes.




Fig. 16. (left) Four predominant colors (purple (P), blue (B), orange (O), yellow (Y)) as marked on the 2D and 3D microstructures. (right) Neighbor statistics of the colors; the x-axis denotes the color of the chosen grain, and the bars indicate the probability of occurrence of colored neighbor grains. There are two bars for each neighbor color, one for the 2D image and another for the 3D reconstructed microstructure. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 17. This example is specifically selected to illustrate the shortcomings of the MRF algorithm, as presented in this paper. The varying morphology of the given microstructure along the z direction, cannot be captured by only one single image orthogonal to that axis. As a result, the MRF algorithm cannot be used to properly reconstruct such a 3D model. Here, the slices on the right hand side correspond to $z = 0, 64,$ and $128,$ respectively.

Acknowledgments

The authors would like to acknowledge the Air Force Office of Scientific Research, USA Materials for Extreme Environments Program (Grant No. FA9550-18-1-0091) as well as the National Science Foundation, USA Graduate Research Fellowship Program (Grant No. DGE 1256260) for financial support. Gratitude is expressed to Tracy Berman and John Allison from Materials Science and Engineering Department at the University of Michigan for providing EBSD data.

References

- [1] DeHoff R. *J Microsc* 1983;131(3):259–63.
- [2] Spanos G. *Scr Mater* 2006;1(55):3.
- [3] Jensen DJ. Three-dimensional orientation imaging. In: Schwartz AJ, Kumar M, Adams BL, editors. *Electron backscatter diffraction in materials science*. Boston, MA: Springer US; 2000, p. 91–104.
- [4] Saylor DM, Fridy J, El-Dasher BS, Jung K-Y, Rollett AD. *Metall Mater Trans A* 2004;35(7):1969–79.
- [5] Groeber M, Ghosh S, Uchic MD, Dimiduk DM. *Acta Mater* 2008;56(6):1257–73.
- [6] Groeber M, Ghosh S, Uchic MD, Dimiduk DM. *Acta Mater* 2008;56(6):1274–87.
- [7] Zhang C, Li H, Eisenlohr P, Liu W, Boehlert C, Crimp M, Bieler T. *Int J Plast* 2015;69:21–35.

- [8] Yaghoobi M, Ganesan S, Sundar S, Lakshmanan A, Rudraraju S, Allison JE, Sundararaghavan V. *Comput Mater Sci* 2019;169: 109078.
- [9] Buffiere J-Y, Maire E, Verdu C, Cloetens P, Pateyron M, Peix G, Baruchel J. *Mater Sci Eng A* 1997;234:633–5.
- [10] Pokharel R, Lind J, Kanjarla AK, Lebensohn RA, Li SF, Kenesei P, Suter RM, Rollett AD. *Annu Rev Condens Matter Phys* 2014;5(1):317–46.
- [11] Kotula PG, Rohrer GS, Marsh MP. *MRS Bull* 2014;39(4):361–5.
- [12] Lienert U, Li SF, Hefferan CM, Lind J, Suter RM, Bernier JV, Barton NR, Brandes MC, Mills MJ, Miller MP, Jakobsen B, Pantleon W. *JOM* 2011;63(7):70–7.
- [13] Brown DW, Balogh L, Byler D, Hefferan CM, Hunter JF, Kenesei P, Li SF, Lind J, Niezgoda SR, Suter RM. *Mechanical stress evaluation by neutrons and synchrotron radiation VII. Materials science forum*, vol. 777, Trans Tech Publications; 2014, p. 112–7.
- [14] Bingert JF, Suter RM, Lind J, Li SF, Pokharel R, Trujillo CP. *Dynamic behavior of materials*, Vol. 1. Springer; 2014, p. 397–403.
- [15] Eisenlohr P, Shanthraj P, Kieft BRV, Bieler TR, Liu W, Xu R. *JOM* 2017;69(6):1100–5.
- [16] Shivelev A, Shade P, Pilchak A, Tiley J, Kerns R. *J Microsc* 2011;244(2):181–6.
- [17] Heeger DJ, Bergen JR. *Proceedings of the 22nd annual conference on computer graphics and interactive techniques*. Citeseer; 1995, p. 229–38.
- [18] De Bonet JS. *Proceedings of the 24th annual conference on computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co.; 1997, p. 361–8.
- [19] Simoncelli EP, Portilla J. *Proc 5th IEEE int'l conf on image proc*, Vol. 1. 1998.
- [20] Sundararaghavan V, Zabarar N. *Comput Mater Sci* 2005;32(2):223–39.
- [21] Torquato S. *Random heterogeneous materials: microstructure and macroscopic properties*, Vol. 16. Springer Science & Business Media; 2013.
- [22] Yeong C, Torquato S. *Phys Rev E* 1998;57(1):495.
- [23] Jiao Y, Padilla E, Chawla N. *Acta Mater* 2013;61(9):3370–7.
- [24] Fullwood DT, Niezgoda SR, Kalidindi SR. *Acta Mater* 2008;56(5):942–8.
- [25] Xu H, Dikin DA, Burkhart C, Chen W. *Comput Mater Sci* 2014;85:206–16.
- [26] Ashton TN, Guillen DP, Harris WH. *JOM* 2019;1–10.
- [27] Soyarslan C, Bargmann S, Pradas M, Weissmüller J. *Acta Mater* 2018;149:326–40.
- [28] Rollett A, Saylor D, Fridy J, El-Dasher B, Brahme A, Lee S-B, Cornwell C, Noack R. *AIP conference proceedings*, Vol. 712. AIP; 2004, p. 71–7.
- [29] Rollett AD, Lee S-B, Campman R, Rohrer G. *Annu Rev Mater Res* 2007;37:627–58.
- [30] Groeber MA, Jackson MA. *Integr Mater Manuf Innov* 2014;3(1):5.
- [31] Tu X, Shahba A, Shen J, Ghosh S. *Int J Plast* 2019;115:268–92.
- [32] Brahme A, Alvi M, Saylor D, Fridy J, Rollett A. *Scr Mater* 2006;55(1):75–80.
- [33] Falco S, Jiang J, De Cola F, Petrinic N. *Comput Mater Sci* 2017;136:20–8.
- [34] Bargmann S, Klusemann B, Markmann J, Schnabel JE, Schneider K, Soyarslan C, Wilmers J. *Prog Mater Sci* 2018;96:322–84.
- [35] DebRoy T, Wei H, Zuback J, Mukherjee T, Elmer J, Milewski J, Beese AM, Wilson-Heid A, De A, Zhang W. *Prog Mater Sci* 2018;92:112–224.
- [36] Bostanabad R, Zhang Y, Li X, Kearney T, Brinson LC, Apley DW, Liu WK, Chen W. *Prog Mater Sci* 2018;95:1–41.

- [37] Fullwood DT, Niezgoda SR, Adams BL, Kalidindi SR. *Prog Mater Sci* 2010;55(6):477–562.
- [38] Sundararaghavan V. *Integr Mater Manuf Innov* 2014;3(1):19.
- [39] Kumar A, Nguyen L, DeGraef M, Sundararaghavan V. *Modelling Simulation Mater Sci Eng* 2016;24(3). 035015.
- [40] Efros AA, Leung TK. *Proceedings of the seventh IEEE international conference on computer vision*, Vol. 2, IEEE, 1999, pp. 1033–1038.
- [41] Kwatra V, Essa I, Bobick A, Kwatra N. *ACM transactions on graphics (ToG)*, Vol. 24. ACM; 2005, p. 795–802.
- [42] Kopf J, Fu C-W, Cohen-Or D, Deussen O, Lischinski D, Wong T-T. *ACM transactions on graphics (TOG)*, Vol. 26. ACM; 2007, p. 2.
- [43] Acar P, Sundararaghavan V. *Modelling Simulation Mater Sci Eng* 2016;24(7). 075005.
- [44] Bostanabad R, Bui AT, Xie W, Apley DW, Chen W. *Acta Mater* 2016;103:89–102.
- [45] Liu X, Shapiro V. *Comput Mater Sci* 2015;99:177–89.
- [46] Turner DM, Kalidindi SR. *Acta Mater* 2016;102:136–48.
- [47] Coltuc D, Bolon P, Chassery J-M. *IEEE Trans Image Process* 2006;15(5):1143–52.
- [48] Shapira D, Avidan S, Hel-Or Y. 2013 IEEE international conference on image processing. IEEE; 2013, p. 2269–73.
- [49] Sun S, Sundararaghavan V. *Acta Mater* 2012;60(13):5233–44.
- [50] Lopez-Sanchez MA, Llana-Fúnez S. *J Struct Geol* 2016;93:149–61.
- [51] Sundararaghavan V, Zabarar N. *Acta Mater* 2005;53(4):1015–27.
- [52] Acar P, Sundararaghavan V. *AIAA J* 2016;54(5):1751–9.
- [53] Qusted T. *Micrograph 712 available under creative commons license*. 2003, https://www.doitpoms.ac.uk/miclib/micrograph_record.php?id=712, [Online; accessed 19-July-2019].
- [54] Sundararaghavan V, Kumar A, Sun S. *Acta Mater* 2015;93:12–23.