

**CONSTRUCTION OF REPRESENTATIVE 3D
MICROSTRUCTURES FROM COMPLETE AND PARTIAL
STATISTICS**

A Dissertation
Presented to
The Academic Faculty

by

David Michael Turner

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Mechanical Engineering

Georgia Institute of Technology
December 2016

COPYRIGHT © 2016 BY DAVID TURNER

**CONSTRUCTION OF REPRESENTATIVE 3D
MICROSTRUCTURES FROM COMPLETE AND PARTIAL
STATISTICS**

Approved by:

Dr. Surya R. Kalidindi, Advisor
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Karl I. Jacob
School of Mechanical Engineering
Georgia Institute of Technology

Dr. David L. McDowell
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Le Song
College of Computing
Georgia Institute of Technology

Dr. Hamid Garmestani
School of Material Science and
Engineering
Georgia Institute of Technology

Date Approved: November 02, 2016

To my wonderful family, for all their love and support.

ACKNOWLEDGEMENTS

I would like to acknowledge first and foremost my advisor Dr. Surya R. Kalidindi. He has been the principle mentor of my professional life and for this I am eternally grateful. Specifically, I thank him for his kind and generous instruction, professional guidance, patience, and support throughout all these long years. His influence on my understanding of materials science is truly profound.

I acknowledge funding from the Office of Naval Research (ONR) award N00014-11-1-0759 (Dr. William M. Mullins, program manager).

Additionally, I would like to thank my research colleagues throughout the years, listed in no particular order; Dr. Stephen Niezgoda, Dr. David Fullwood, Dr. Brent L. Adams, Dr. Anthony Fast, Dr. Yuksel C. Yabansu, Dr. Mohammed Tahir Abba, Dr. Jordan Weaver, Dr. David Brough, Dr. Ali Khosravani, Dr. Siddiq Qidwai, Dr. Haviva Goldman, Ahmet Cecen, Joshua Shaffer, Dan Satko, Dr. Dipen Patel, Dr. Amanda Levinson, Dr. Shraddha Vacchani, Dr. Hamad Al-Harbi, Dr. Naomi Barth, Dr. Siddhartha Pathak, Dr. Marko Knezevic, Dr. Ayman Salem, Dr. Dejan Stojakovich, Dr. Ko Nishino, Dr. Ulrike Wegst, and Dr. Spiros Mancoridis. I am sure I have left someone out, please forgive me.

I also need to take time to thank my first advisor, Dr. Vassilis Prevelakis. Your guidance throughout my undergraduate career was an invaluable first start at research.

I must also thank my committee who have taken the time to consider this work; Dr. David L. McDowell, Dr. Hamid Garmestani, Dr. Karl I. Jacob, and Dr. Le Song. You cannot imagine how honored I am that you all have agreed to be on my committee.

William Vince has always generously supported and encouraged my education and for this I will always be grateful.

Finally, to my wonderful family who have encouraged and supported me for my entire life. Mom, I hope that I have made you proud, thank you for loving me unconditionally and always believing in me. Dad, thanks for teaching me to love science and buying me my first computer. CB, thanks for being my absolute best friend, for all the help moving and all the rides from Edison every week before I could drive. Aunt Jenny, you are my second mom, don't forget it.

To my wife Lina, thank you for your patience and love. Other than myself, no one has felt the burden of this endeavor more than you. There has never been a moment that I felt you weren't supporting me. I could not have done this without you.

Finally, to Mari, you are part of the original team! To Max and Mateo this will only be history, a story we tell them, but to you it was your life. Thank you for becoming the reason I get up in the morning.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
LIST OF TABLES	vii
LIST OF FIGURES	viii
LIST OF SYMBOLS AND ABBREVIATIONS	xiv
SUMMARY	xvi
CHAPTER 1. INTRODUCTION	1
1.1 Related Work	3
1.2 Contents	11
CHAPTER 2. N-POINT CORRELATIONS	13
2.1 Discretized Microstructure Function	13
2.2 Spatial Statistics of Microstructure	15
2.3 Efficient calculation of 2-point statistics	16
2.4 Boundary Conditions	21
CHAPTER 3. REPRESENTATIVE VOLUME ELEMENTS	25
3.1 Ensemble Averaged Microstructure Statistics	28
3.2 Statistical Volume Element (SVE) Size	31
3.3 Building Weighted Statistical Volume Element Sets	33
3.4 Examples	35
3.4.1 Two-Phase Digitally Created 2D Microstructure	36
3.4.2 2D Polycrystalline Microstructure	40
3.4.3 Large 3D Polycrystalline Microstructure	45
3.5 Conclusions	64
CHAPTER 4. STATISTICAL RECONSTRUCTION	66
4.1 Solid Texture Synthesis for Microstructure Reconstruction	66
4.2 Results and Discussion	77
4.3 Conclusions	88
APPENDIX A. CHORD LENGTH DISTRIBUTIONS	90
4.4 Computational Procedures	93
4.5 Examples	98
4.5.1 Two-Phase 2-D Microstructure	98
4.5.2 Porous 3-D Microstructures	107
4.6 Summary	109
REFERENCES	110

LIST OF TABLES

Table 1	Comparison of scalar error D_s , effective Young's modulus, and interface fraction between the large microstructure data set (labeled as sample) shown in Figure 9 and the corresponding RVE Sets produced in this study (see Figure 11 and Figure 12).	39
Table 2	A comparison of scalar error D_s and the first-order upper and lower bounds for the components of the stiffness tensor and tensile, and shear yield strengths between the entire sample and an optimize SVE set of size 1.	43
Table 3	A comparison of scalar error D_s and the first-order upper and lower bounds for the components of the stiffness tensor and tensile, and shear yield strengths between the entire sample and an optimize SVE set of size 3.	44
Table 4	Specifications of the representative WSVE sets and the traditional RVE. Note the extreme weight given to element E in the suboptimal WSVE set. This is element selected based on 2-point statistics while the others are random. Also note the weights of the optimal WSVE set, they are more equal as the elements have been selected heuristically.	49
Table 5	Statistical error between representative WSVE sets, the traditional RVE and ensemble statistics for the entire sample.	52
Table 6	Representative constitutive model parameters for β -Ti 21S alloy.	53
Table 7	Statistical measures (mean and standard deviation) of distance between the autocorrelations of a 150^3 voxels subvolume extracted from the larger (500^3 voxels) original microstructure and the full original microstructure compared against the distance between the autocorrelations of the original microstructure and its reconstruction. Microstructures numbers correspond to the one shown in Figure 31.	81

LIST OF FIGURES

Figure 1	The image on the left represents a simple 2-phase/binary microstructure. Random vectors are thrown into the data with their heads and tails lying in different spatial bins. The image to the right shows how we could accumulate these vectors within a set of matrices, one for each correlation.	19
Figure 2	Choose a vector length and orientation of interest. Center a copy of the image on top of itself offset by this vector. Multiply each of the overlapping pixels together and add them up and you will have the total count for that particular vector.	19
Figure 3	An example of a microstructure with two local states containing equally sized circles. Yellow example vectors are shown thrown into the microstructure. Not all vectors sampled are shown. A calculation of the 2-point auto-correlation (using the local state red). To read the probability associated with the yellow vector simply read the value at the head after placing the tail on the origin.	20
Figure 4	An example of a microstructure with three local states. The cross-correlation to the right is computed between the red and green local states.	21
Figure 5	An example of the normalization applied to non-periodic 2D 2-pt correlation. The origin is at the center of the plot. Each value is the total count of vectors that can fit into a 256x256 image. In the non-periodic case long and angled vectors can be thrown in less than short vectors so their counts must normalized non-uniformly.	23
Figure 6	A comparison of periodic vs. non-periodic auto-correlation of the image of circles in the top right. Notice the more pronounced effects for longer vectors.	24
Figure 7	A comparison of computation time in seconds as a function of image size between 2-point correlation calculation algorithms. The blue line with $O(n^2)$ growth is a naïve implementation of the convolution. The red and green lines represent the $O(n \log n)$ FFT methods. Notice the increased cost that padding adds for the non-periodic case.	24
Figure 8	(a) Ensemble averaged normalized 2-point particle–particle autocorrelations, $\langle \frac{nn^j}{f_t} - (nf^j)^2 \rangle$, calculated over 100	30

independent 150×150 pixel realizations of non-overlapping disks of diameter 20 pixels with a minimum spacing of 5 pixels. The ensemble average statistics clearly shows the discs are spatially uncorrelated with a coherence length of 25 pixels. (b) Three independent realizations of the microstructure. (c) The normalized 2-point statistics of the realizations shown in (b). Even though the dimensions of each realization are six times larger than the coherence length, the correlation functions for each realization show significant correlations at the longest vectors.

- Figure 9 Digitally created 2-phase microstructure (550×550 pixels) for which an SVE Set is to be created. 32
- Figure 10 (a) Mean scalar error, \bar{D}_S^j , for ensembles of 100 members of different sub-domain sizes. Error bars indicate maximum and minimum observed values of D_S^j for each ensemble. (b) Standard deviation of D_S^j as a function of the sub-domain size. 33
- Figure 11 The 200×200 pixel sub-domain (a) and its associated 2-point statistics (white–white autocorrelation) (b) that comes closest to matching the statistics (c) of the target microstructure (d). The main features of the target 2-point statistics are evident in the RVE statistics, such as the intensity and shape of the central spot. The location of the sub-domain shown in (a) is highlighted in (d). 38
- Figure 12 Members (a–c) and associated statistics (d–f) of the three-member RVE Set. The statistics (white–white autocorrelations) are all plotted with the same color-scale to highlight the differences observed between the members. Individually the statistics of the individual members are a poor match to the target statistics shown in Figure 11c; however, their weighted average (h) is a close match. The location of the individual windows in the target microstructure (g) is highlighted. 39
- Figure 13 The influence of the SVE Set size on the scalar error between the 2-point statistics of the SVE Set and those of the target microstructure, D_S . By adding more members to the SVE Set, the target statistics can be captured to improved accuracy. 40
- Figure 14 (a) EBSD grain map of Fe–Si steel where the fundamental zone of crystal orientation has been binned into 512 bins in the Euler-angle space. Each pixel represents $1 \mu\text{m}$. (b) The corresponding texture of the microstructure shown in (a) as $\langle 100 \rangle$ and $\langle 110 \rangle$ pole figures showing that the sample is only 41

weakly textured.

- Figure 15 Standard deviation of D_{ζ}^j as a function of the sub-domain spatial extent for the microstructure data set shown in Figure 14. 42
- Figure 16 The texture corresponding to the one- (top) and three-member (bottom) SVE Sets, demonstrating that even the very weak textures can be accurately captured by a relatively small SVE set. The maximum intensities in the SVE Sets are somewhat higher due to a much smaller number of grains in these sets. 42
- Figure 17 Sample of the 4300-grain sample of β -stabilized Ti 21S ($1.115 \times 0.516 \times 0.3 \text{ mm}^3$, $1670 \times 770 \times 200$ voxels). The color key corresponds to the crystallographic orientation parallel with the Z-axis $[0\ 0\ 1]$. 46
- Figure 18 Locations and relative sizes of all five members of each representative WSVE set and the traditional RVE in the 4300-grain sample. 48
- Figure 19 A comparison of orientation distributions via pole figures in three crystallographic directions of interest. The color value indicates times random. The samples compared are the full 4300-grain dataset, the traditional RVE, the optimal WSVE set and the suboptimal WSVE set. The pole figures are created using the same level of smoothing for all. 51
- Figure 20 (a) Superposition of a regular mesh on the underlying microstructure is used to create the finite-element model for an approximately 100-grain volume. (b) Description of displacement-based uniaxial loading conditions. The four transverse planes (two are not shown in this 2-D view) are constrained to move as planes. 54
- Figure 21 Global uniaxial stress–strain responses for the suboptimal WSVE set: (a) individual SVEs in the global X-direction; (b) individual SVEs in the global Y-direction; (c) individual SVEs in the global Z-direction; and (d) effective RVE response in the three global directions obtained by weighting each SVE response in respective directions according to Equation 22 with the appropriate weights given in Table 4. 55
- Figure 22 Global uniaxial stress–strain responses for the optimal WSVE set: (a) individual SVEs in the global X-direction; (b) individual SVEs in the global Y-direction; (c) individual SVEs in the global Z-direction; and (d) effective RVE response in the 56

three global directions obtained by weighting each SVE response in respective directions according to Equation 22 with the appropriate weights given in Table 4.

- Figure 23 Effective uniaxial stress–strain responses for the suboptimal WSVE set, optimal WSVE set and the 499-grain RVE under tension in: (a) X-direction, (b) Y-direction and (c) Z-direction. 57
- Figure 24 Comparison of the effective uniaxial stress–strain behaviors as predicted for the optimal WSVE and the random set with equal weights (0.25) in (a) X-direction, (b) Y-direction and (c) Z-direction. The error (defined in the text) is in terms of relative absolute error with respect to the WSVE response. 58
- Figure 25 Frequency plots or histograms by volume of the cumulative shear strain γ for the SVEs of the representative optimal set and their weighted summation, the effective RVE. The 227,700 data points for each SVE are obtained at an applied strain of 1.15% in the X-direction soon after the yield point in the global stress–strain curves (Figure 24). 61
- Figure 26 Comparison of effective histograms by volume of the cumulative shear strain γ for the representative suboptimal and optimal WSVE sets, and the RVE under tension in: (a) X-direction, (b) Y-direction and (c) Z-direction. The data is obtained at an applied strain of 1.15% soon after the yield point in the global stress–strain curves (Figure 21, Figure 22, and Figure 23). 62
- Figure 27 Comparison of the cumulative shear strain data for the optimal WSVE set and random set under loading in (a) X-direction, (b) Y-direction and (c) Z-direction. The fraction of error (defined in the text) indicates the absolute error in terms of fraction of the optimal WSVE set data. The data is obtained at an applied strain of 1.15% soon after the yield point in the global stress–strain curves. 63
- Figure 28 A schematic description of the solid texture reconstruction based on Kopf et al. [17] 67
- Figure 29 A schematic representation of our notation for reconstruction $\{\mathbf{m}\}$ and the set of exemplars. The sample shown in the top left has three orthogonal 2D exemplar scans taken from it. Each voxel in our reconstruction is then assigned multiple best matching neighborhoods taken from the exemplars, where each assignment is treated as a “vote”. The figure conveys how the microstructure value at \mathbf{m}_s has three votes from the three best 80

matching neighborhoods on the three different planes denoted by \tilde{N}_s^i ; these are selected from the set of possible neighborhoods denoted as $\{N^i\}$ from the exemplar set. Furthermore, all of the yellow, green, and red colored voxels in the reconstruction also receive a vote from the same assignment.

- Figure 30 An example 3-D microstructure used in the validation of the algorithm presented in this study. (a) shows the 3-D microstructure volume, (b) shows orientations of the nine slices used in some of the reconstructions as well as a single example slice, and (c) shows an example sub-volume of the same size as our reconstruction. 78
- Figure 32 Selected comparisons of 3-D two point statistics between original microstructures and their reconstructions shown in Figure 31. A subset of the 3-D auto-correlations are shown on three orthogonal slices through the origin of the spatial statistics. 84
- Figure 33 The principle component (PC) embeddings of autocorrelations computed for neighborhoods extracted from the original full microstructure (red points) and its reconstruction (blue points). The percentages next to each plot express the captured overall variance in the ensembles using three PC scores. 87
- Figure 34 This figure details how chords are identified for a particular orientation. In all images the red phase is the phase of interest. The green coloured pixel marks the active pixel in the algorithm, blue marks pixels that have be processed, and gold marks pixels along a chord. In (a) we show how for a given orientation we try every scanline within the image dimensions. Some are not shown here for the sake of brevity. For each scanline, we begin (b) by starting at one edge of the image and marching along the rasterized approximation of an oriented line until (c) we encounter the phase of interest. We then continue moving along the scanline until (d) we encounter a pixel that is not the phase of interest. At this point we have identified a chord as long as the start pixel and end pixel are not touching the edge of the image. In (e) we continue processing the remaining portion of the scanline. 96
- Figure 35 A simple image of identical circles of diameter of approximately 32 pixels placed in a uniformly random manner. Volume fraction of the circular phase is 18.2%. The image dimensions are 2048x2048. 99

Figure 36	Comparison of closed form and calculated results for non-angularly resolved chord length distribution for a simple two dimensional circle with an approximate diameter of 32 pixels.	102
Figure 37	Here we show a comparison between our implementations of chord length distribution calculations using both a scanline and image rotation method. Aside from some small deviations due to differences between Bresenham line discretization and image rotation interpolation the method mostly agree.	102
Figure 38	Full angularly resolved chord length distribution for the image in Figure 35. The polar axis on the figure represents the angular orientation of the chord. The radial axis represents the length of the chord. Each individual distribution (chord angle) has been smoothed with a 5 point/pixel moving window average to eliminate noise resulting from discretization of the circle.	104
Figure 39	Shows a comparison between the analytical and calculated LPFs for the matrix phase in the case of randomly distributed circles of $r=16$ pixels. The calculated LPF derived from our chord length distribution of the matrix phase of the image presented in Figure 35 show excellent agreement.	105
Figure 40	Image of uniform randomly placed ellipses with major axis of 95 pixels and minor axis of 63 pixels. The ellipses are all oriented with major axis at 90 degrees from the horizontal.	106
Figure 41	The angularly resolved chord length distribution of the image shown in figure 7. The 95 pixel major axis of the ellipse is highlighted by the peak of the chord length distribution at the 90 degree orientation.	107
Figure 42	Top left shows a synthetic digitally created 3-D microstructure. The pore network has been coloured to show connected pores with distinguishing colours. In the bottom left, the mean pore chord length in all directions has been plotted on the surface of a sphere. Polar and azimuthal angles correspond to orientations of the chord within the sample. To the right, we show selected chord length distributions for two chord orientations. The angularly resolved chord length distributions show clearly the anisotropic nature of the porous structure.	108

LIST OF SYMBOLS AND ABBREVIATIONS

- $[\phi_1, \Phi, \phi_2]$ Local crystallographic lattice orientation specified by Bunge-Euler angles.
- CLD Chord Length Distribution
- EBSD Electron Backscatter Diffraction
- D_S^j The scalar error between the 2-point statistics of sub-region Ω_j of size S and the ensemble average 2-point statistics.
- n_f^j The 1-point correlation, parametrized by local states \mathbf{n} and sub-region j .
- $n\bar{f}$ The ensemble average 1-point correlation.
- $nn'f_t^j$ The 2-point correlation, parametrized by local states \mathbf{n} and \mathbf{n}' , vector \mathbf{t} , and sub-region j
- $nn'\bar{f}_t$ The ensemble average 2-point correlation.
- $nn'F_k^j$ The discrete Fourier representation of $nn'f_t^j$.
- GPU Graphics Processing Unit
- LPF Lineal Path Function
- RVE Representative Volume Element
- l_c The mean chord length.
- nm_s^j The discretized microstructure realization, where s is the spatial bin, n the local state bin, and j the sub-region of the full microstructure data.
- nM_k^j The discrete Fourier representation of nm_k^j
- $p^i(\mathbf{z})$ The chord length distribution function of local state i . Parameterized by the chord (vector) \mathbf{z} .
- PCA Principle Component Analysis
- SVE Statistical Volume Element
- t_c The coherence length.
- nV^j The volume fraction of local state n over sub-region j

WSVE Weighted Statistical Volume Element

SUMMARY

The principle concern of the material scientist is the connection between microstructure, properties, and processing. Microstructure is characterized via experimental measurements of geometry at the appropriate length scale. This is usually followed by a quantification of microstructure via statistics for which there are a broad base of possibilities including classical stereological measures such as grain size and higher order descriptions like the N-point spatial correlations. Despite advances in 3D characterization of microstructures such as X-ray tomography and serial sectioned SEM, most techniques still capture measurements only in 2D sections. Even when 3D datasets are available they are typically measuring only small volumes leading to uncertainty about their statistical significance. Can we build statistically representative reconstructions of 3D microstructure from the partial information gathered on a collection of 2D cross sections? The proposed work introduces new approaches to these problems for two phase composites with complex anisotropic geometries. Efficient algorithms for the computation of "higher order" statistics, such as N-point correlations and chord length distributions, will be explored. These higher order metrics will form the basis for establishing structure based representative volume elements (RVEs) in both cases where microstructure geometry information is complete and incomplete.

CHAPTER 1. INTRODUCTION

Establishing linkages between the structure of materials, their measured properties, and the processing done to them is the inherent goal of material science. Structure is investigated by first deciding the length scale(s) of interest and then selecting an appropriate experimental technique to measure the internal geometry of the sample of interest. Most experimental techniques measure data in planar cross sections of samples and therefore it is difficult to directly translate into a fully resolved picture of the 3D structure. Characterization in three dimensions is now possible in certain cases utilizing techniques such as X-ray microtomography [1], serial sectioned SEM (using mechanical polishing [2, 3], focused ion beam [4, 5], laser ablation [6, 7], etc.), atom probe [8], and other methods. Suitability of any of these techniques is dependent upon numerous factors; susceptibility of material to investigating radiation, required length scales and resolutions, financial cost, time, etc.

Despite the many advances in microstructure characterization, the number of large true three dimensional spatially resolved datasets that are available to researchers for certain classes of materials is still very low. Currently, materials for which X-ray microtomography is not suitable have been typically analyzed in three dimensions using some form of serial sectioning technique coupled with an image acquisition technology like SEM. Sectioning is usually performed via mechanical polishing or a focused ion beam (FIB) for applicable materials. When FIB is suitable the time spent on sectioning datasets is usually much quicker, easier to implement, and possible to automate fully. However, the volume of material that can be milled away is much smaller than required

for the length scales of interest in a large amount of material systems. This is particularly true for polycrystalline metals where a typical FIB generated dataset size is on the order of 50 cubic microns. When statistically significant volumes of polycrystalline metals are needed in full 3D the only real alternative is mechanically polished serial sectioning with SEM\EBSD. However, the process is much more time consuming and extremely difficult to automate. These types of datasets are rare because of the prohibitive cost in both time and money.

The difficulty in acquiring 3D datasets for a large number of material systems pushed to the forefront the concept of reconstructing statistically representative 3D microstructures from sets of partial spatial correlations. We are dealing with partial statistics simply because the data from which they are derived are not 3D but 2D cross sections taken over large areas of the sample. The motivation for this approach is that we can cover a much larger and thus more statistically significant area of the material with 2D cross sections than we could by investing in direct 3D characterization. If enough 2D cross sections are taken then we should be able to gather enough information to reconstruct the 3D statistics from spatially resolved 2D statistics taken from the cross sections.

The general approach is to take statistics from the 2D cross sections and couple that with the known orientation of the planar section within the sample frame. With all of this information we can reconstruct an approximation to the 3D statistics of the microstructure. We will be missing statistics that are out of the planes that we have not sectioned. However, if enough oblique sections are taken we will be able to approximate the full 3D function. With such an approximation we can then try to derive a structure

digitally that matches the known statistics as closely as possible. This optimization problem begins by guessing a starting point for the microstructure that is as close as possible to the one we are trying to reach of course. We then iterate over this guess making small changes until the structures statistics converge to the target statistics derived from our 2D sections.

1.1 Related Work

Statistical reconstruction has been particularly studied in the metals community because of the difficulty in 3D characterization. The standout success in the field would certainly be the software project Microstructure Builder developed by Rollett et al. [9, 10] . Microstructure builder fits into the general model of statistical reconstruction presented above. The goal of the project is to use SEM\EBSD information gathered from three orthogonal cross sections taken from a polycrystalline sample and derive a full 3D model of the microstructure that is representative. SEM EBSD data contains not only information about the geometry of the grain structures but also crystallographic orientation of each of those grains.

Microstructure builder then calculates a set of statistical distributions that describe the shape and texture (distribution of preferred crystallographic orientations) of the sample. The distributions are derived completely from scans on only the three orthogonal sections. The geometric component of these statistics is the ellipsoid distribution that is used to describe the grain structure. One main assumption of the work is that grains in polycrystalline metals can be approximated via ellipsoids. This assumption is not very reasonable for polycrystalline materials because ellipsoids are not space filling. That is,

while grains could be approximated as ellipsoids and characterized with a distribution of these shapes, this distribution would lack a large amount of information about the structure of interfaces between grains. Additionally, the authors assume that the approximating ellipses are all orientated in the same direction. They assume as well that there are no gradients present in the texture, that is, the ellipsoid distribution is independent of the position in the sample. With these assumptions they formulate a distribution $f(a, b, c)$ where the three parameters are the semi-axis lengths of an ellipsoid and f is the probability of finding a grain with such an approximating ellipsoid. It is obvious how three orthogonal sections could provide complete information to extract this type of 3D statistical distribution. To capture the texture of the sample two statistical distributions are captured from the three sections; the orientation distribution function (ODF) and the misorientation distribution function (MODF). The ODF $f(g)$ gives the probability of finding crystallographic orientation g in the sample, and the MODF $f(\Delta g)$ gives the probability of finding a misorientation Δg between nearest neighboring grains in the sample.

Once these statistical measures are estimated from the orthogonal cross sections the algorithm begins by trying to find a grain structure that matches the ellipsoid distribution. It does this by filling a cubic space with said distribution of ellipsoids meanwhile trying to pack them as tightly as possible. Simulated annealing is used to pack the ellipses. Once this is done the ellipses are converted to space filling grains via a procedure that I will not describe here for sake of brevity.

After a grain structure is produced the final phase of reconstruction is an assignment of orientation to each grain such that the texture of the reconstruction matches

the ODF and MODF of the sample. This process produces a 3D reconstruction that matches the basic grain shape and texture character of the true sample. However, the real drawback in microstructure builder is the statistical measures they chose to capture microstructure detail. The ellipsoid distribution will not capture more complex grain geometries such as colony microstructures with different orientations. Worse yet, they use only the most basic measures of crystallographic orientation and thus their reconstructions will not account for the true spatial correlation between different grains in the structure. These features are extremely important for modeling anisotropic plasticity.

Reconstructions using monte carlo techniques [11-15] directly from correlation functions have been explored extensively. The general approach is to start from an initialization of cells in the reconstruction (nucleation) and then guide the distribution and growth of those cells via minimizing some objective error function. The objective error function in these cases is the classic spatial statistical measure of the two point correlation function. For the case of fiber composites [12], interestingly the authors utilized a geometric description of the fiber composites geometry as cylinder segments with “waviness” connected by node junctions instead of a more traditional image model. This drives home the point that the reconstruction problem is not always a problem in image modelling as often times images are not immediately useful for every application.

Surprisingly, the reconstruction problem has also been explored in the field of computer graphics. This area, called solid texture synthesis [16], endeavors to solve a 3-D reconstruction problem that is essentially identical to the one presented earlier. Computer graphics has adopted the term texture map (or simply texture) to refer to a color image that is applied to 3-D graphical objects to give them a more complex and real

appearance. These textures are most commonly 2-D images that are mapped or interpolated across the surface of the 3-D object. However, when it is necessary to show the internal structure of a 3-D object with texture, a 3-D image or solid texture is then needed. Solid texture synthesis algorithms are being devised to help 3-D animators create realistic 3-D images that specify the appearance of the internal surfaces of objects. The most popular form of these algorithms utilizes what are called 2-D exemplar textures to define the appearance of the solid texture when viewing it from a particular direction. That is, when the solid texture is cut along any plane of a certain orientation it will resemble the 2-D exemplar texture statistically.

One of the most effective algorithms in the literature is an optimization based approach presented by Kopf et al. [17]. The algorithm's inputs are a set of three input images that represent the desired appearance of the solid texture when cut anywhere along three orthogonal planes. The algorithm alternates between search and optimization steps until the solid texture converges. The search step is relatively simple. The neighborhood of every pixel in the current iteration's solid texture is examined, and the three best matching neighborhoods from the exemplar input images are found. A neighborhood is defined as a small square region of pixels centered on a certain voxel. Since the algorithm's input is three orthogonal 2D slices we are concerned with the three corresponding orthogonal neighborhoods centered on each voxel. If $N \times N$ represents the size of our neighborhood, then N can range typically between 8 and 32 pixels. The optimization step proceeds by assigning each pixel a value that minimizes the error between all best matching neighborhoods for which the pixel is involved. More

specifically, given the current 3-D reconstruction, S , and a set of 2-D exemplars images, E , this error E_T is defined as

$$E_T(S; E) = \sum_v \sum_{i \in \{x,y,z\}} \|S_{v,i} - E_{v,i}\|^r \quad (1)$$

where v iterates over the entire current 3-D reconstruction, $S_{v,x}$, $S_{v,y}$, and $S_{v,z}$ are vectorized neighborhoods situated around the voxel v in slices orthogonal to the x, y, and z axes. $E_{v,i}$ is the best matching candidate (identified in the previous search phase) for v . The r exponent is used to control sensitivity to outliers. It has been shown that a good value for r is ~ 0.8 [17]. This optimization is solved using an iteratively reweighted least squares (IRLS) method for which a simple closed form expression is derived:

$$S_v = \frac{\sum_{i \in \{x,y,z\}} \sum_{u \in N_i(v)} w_{u,i,v} E_{u,i,v}}{\sum_{i \in \{x,y,z\}} \sum_{u \in N_i(v)} w_{u,i,v}} \quad (2)$$

where $N_i(v)$ represents the different slices forming the neighborhood of v . Here $E_{u,i,v}$ represents the exemplar pixel in the neighborhood $E_{u,i}$ that corresponds to v . The weight $w_{u,i,v}$ is a scalar value for each exemplar neighborhood pixel that can be used to control the influence of a neighborhood to the solution. These weights allow us to drive the solution towards a more global optimum. Equation 14 can be thought of intuitively as setting the value of each voxel in the reconstruction so that it best matches all of the neighborhoods with which it is involved. More specifically, this equation is simply setting the value of the voxel to the average of all the suggested values from the best

matching neighborhoods. If the weighting constants ($w_{u,i,v}$) chosen for each neighborhood's contribution are all equal then the algorithm converges towards a solution in which only local statistics (defined within the selected size of the neighborhood) are matched. The figure below shows a schematic description of the general process of solid texture synthesis.

Kopf et al. [17] present reweighting schemes which incorporate global statistical information to drive the solution towards a compromise between both local and global constraints. Since Kopf et al. are reconstructing color images, the global statistics they are trying to incorporate are simply color histograms. That is, they are trying to produce a 3D reconstruction whose histogram of colors matches that of the 2-D input images. This would be something akin to trying to simply match the volume fraction in the case of a two-phase microstructure. To accomplish this, the weights are selected as the algorithm progresses so that neighborhoods that drive the color histogram away from their target are diminished. Another more recent approach [18] uses a dual approach of position and neighborhood histogram matching. More specifically, a record is kept between iterations of which pixels and neighborhoods from the exemplars are copied to the solution. The optimization weights are adjusted so that the exemplar images are sampled uniformly.

Methods that explicitly model the material or texture with non-parametric Markov Random Fields (NP-MRF) have been known for quite some time [19] and have been among the most successful for reconstructing 2D images from other 2D example images (2D texture synthesis). Before that, MRFs were studied extensively in the context of image texture modelling as a statistical model for local image structure. They have been used for problems like edge detection, image registration, reconstruction, segmentation,

texture classification, and synthesis [20-22]. Recently these methods have been extended to the problem of 3D solid texture synthesis [23]. The MRF approach can be viewed as the application of the Markovian principle to the image modelling domain. More specifically, the probability of any pixel\voxel taking on a certain value depends only on the values at neighboring pixels\voxels. The neighborhood of a pixel\voxel is usually defined a small square region around its location. This Markovian condition is usually called the local conditional probability density function (LCPDF), which may be expressed;

$$L(s) = P(M_s = m_s | M_t = m_t, s \neq t) = P(m_s | m_t, t \in N_s) \quad (3)$$

where M_s is the random variable describing the intensity value at pixel\voxel s and N_s is the neighborhood of s . In NP-MRF, the LCPDF is estimated by constructing a multidimensional histogram from image neighborhoods where each dimension is a position in the neighborhood. This histogram may then be smoothed and normalized into a pdf by Parzen-window density estimation [24] as done by [23]. Using the estimated LCPDF the reconstruction problem can then be treated as a relaxation problem for which there are a number of suitable algorithms; stochastic methods such as Gibbs sampling [21] and Metropolis-Hastings [25] as well as deterministic methods like Iterated Conditional Modes (ICM) [26].

Recently some researchers [27, 28] in the material science field have seemed to take inspiration from the solid texture synthesis community. Liu et al. [27] argue primarily that the microstructure reconstruction problem is analogous to the texture synthesis problem. They argue that modelling microstructure as an MRF should capture

completely all other local spatial statistical measures. Most of their experiments are conducted as a 2D texture synthesis problem using 2D exemplars. There is brief mention and results of extending the algorithms to 3D reconstruction from 2D exemplars. The algorithm is implemented for binary images and is adapted from [29]. It is a surprisingly simple deterministic algorithm which iteratively assigns unknown pixels in the synthesized microstructure from the best matching neighborhoods in the exemplar. Like many solid texture synthesis algorithms it builds a multiscale Gaussian pyramid of both the exemplar image and synthesis solution. It then proceeds from lower resolution to higher, upsampling each time the solution converges or in this case all pixels are assigned. This algorithm performs excellently when 2D images are synthesized from other 2D images because in this case LCPDF of the MRF is completely known. However, in the compressed sensing problem of 3D reconstruction from 2D information only we have an ill-posed problem with many possible valid solutions. More state of the art solid texture synthesis algorithms such as [17, 18] have had more success with utilizing a hybrid approach of trying to match both local and global statistics.

Sundararaghavan [28] seems to have taken the approach of adapting Kopf et al. [17] work to the material science domain. However, there are some interesting differences discussed in the implementation of the underlying algorithms of this work. First, Kopf et al. utilizes a reweighting scheme in the optimization step of their algorithm which down-weights exemplar pixels that drive the solution away from the global color histogram of the exemplar. This can be seen as a form of regularization enforced on the optimization algorithm to help ensure it does not copy the same regions of the exemplar again and again. Or otherwise, so that it does not get stuck in local minima.

Sundararaghavan discusses the weighting in his algorithm as utilizing a simple Gaussian scheme which assigns weights to pixels from the exemplar as a function of distance from the center of the neighborhood. This is done to “to preserve short-range correlations of the microstructure”. This is interesting because there does not seem to be any regularization in the optimization as done in Kopf et al.[17] and others[18]. The presented results are shown for a few simple microstructures using small exemplar images. Numerous artifacts are present in the reconstructions that are not visible in the exemplar images which are most likely the result of the algorithm converging to local minima, a typical problem with solid texture synthesis. None the less, to the author’s knowledge, this is probably the best example of solid texture synthesis algorithms being applied directly to the field of material science to date.

1.2 Contents

This document describes the work of adapting state of the art solid texture synthesis algorithms to the problem of 3D microstructure reconstruction. This effort began with research into statistical measures of microstructure for assessing the accuracy of these reconstructions. Chapter 2 includes detailed discussions on a classical measure of microstructure known as the 2-point correlation function. Specifically, it discusses extensions made to the fast Fourier transform (FFT) convolution algorithm for computing these spatial statistics on microstructure images that allows for avoiding periodic boundary condition assumptions. In chapter 3, we discuss the application of microstructure statistics towards a definition of the representative volume element (RVE). Our weighted sets of statistical volume elements (SVE) are an attempt to redefine the RVE in terms of statistics of microstructure as opposed to simulated performance

bounds. In chapter 5, we discuss the implementation of a solid texture synthesis inspired algorithm for the reconstruction of 3D microstructure images from sets of 2D images. Finally, appendix A discusses another metric of microstructure known as the chord length distribution and its expression in an angularly resolved form. We discuss an algorithm for the efficient computation of these statistics on 2D and 3D microstructure images.

CHAPTER 2. N-POINT CORRELATIONS

In order to reconstruct a statistically representative microstructure we must first define what we mean by “statistically representative”. There have been numerous ways in which materials researchers attempted to capture features of microstructure in a statistical space; including the n -point spatial correlations [30-37], shape moment invariants [38], topological descriptors including homology [39-41], entropic descriptors [42, 43], chord length distributions [44-50], lineal path functions [51], among others. The following chapter of this work is dedicated to n -point correlations. We will define these measures, describe how to compute them efficiently, and show why they can be effective statistical representations for microstructures. If you have additional interest in chord length distributions and how they can be efficiently computed, please consult Appendix A [44].

2.1 Discretized Microstructure Function

In order to characterize microstructure, we must first specify a mathematical description of it. Let the spatial domain of the microstructure be binned into a uniform grid of S cells/pixels/voxels, whose nodes are enumerated by the ordered triplet of indices, one per dimension of the image, $[s_1, s_2, s_3]$. Each component is numbered in a zero based indexing, $s_i = 0, 1, 2 \dots S_i - 1$. We can represent this triplet by vector \mathbf{s} , the vector from the origin to $S_1 - 1, S_2 - 1, S_3 - 1$ as $\mathbf{S} - \mathbf{1}$, and the zero vector as $\mathbf{0}$. Just as we have discretized our spatial domain we may do the same for the local state space represented at each individual spatial bin. The local state space H , is binned into a uniform grid of N discrete local states. Similarly, local state bins will be represented by an ordered vector \mathbf{n} .

The discrete microstructure function over the sub-region Ω^j described here is therefore denoted in a condensed form as ${}^n m_s^j$. It is necessary to introduce a convention for indexing sub-region Ω^j because this work often discusses aggregating over sub-samples of a microstructure dataset.

We define the following constraints on the discretized microstructure function;

$$\sum_{n=0}^{N-1} {}^n m_s^j = 1, \quad {}^n m_s^j \geq 0, \quad \sum_{s=0}^{S-1} {}^n m_s^j = {}^n V^j S \quad (4)$$

where ${}^n V^j$ is the volume fraction of local state \mathbf{n} in the sub-region Ω^j and $S = S_1 S_2 S_3$ (the number of spatial cells). By formalizing the microstructure function in this manner we are expressing the idea that all characterization techniques probe the material over finite volumes (spatial bins), thus the measurement at bin \mathbf{s} is effectively quantifying the local state distribution at the probe volume. More succinctly, we can express the inherent uncertainty with individual measurements using this description.

However, for most of our cases we deal with images that only have one local state represented in any spatial bin. That is, for any spatial bin \mathbf{s} of ${}^n m_s^j$, only one local of state bin \mathbf{n} has a non-zero value. Such microstructure datasets are very common in material science because often times experimentalists simplify their data by deciding the most dominant signal present in a spatial bin will represent that bin's local state entirely. For example, micro-porous datasets can be approximated by binarizing (thresholding) the local state space of X-ray tomography data. For convenience, throughout this work we will most

likely express such two-phase microstructures using a simplified form of this notation, m_s . Here, the local state index and spatial region index have been dropped for brevity.

2.2 Spatial Statistics of Microstructure

The most comprehensive and hierarchical measures of a microstructure are provided by the framework of n-point spatial correlations [30, 37, 52-55]. For example, the 1-point and 2-point statistics corresponding to the discretized microstructure realization ${}^n m_s^j$, on the region Ω^j are defined as;

$${}^n f^j = \frac{1}{S} \sum_{s=0}^{S-1} {}^n m_s^j, \quad (5)$$

$${}^{nn'} f_t^j = \frac{1}{S} \sum_{s=0}^{S-1} {}^n m_s^j {}^{n'} m_{s+t}^j \quad (6)$$

${}^n f^j$ is the probability of finding local state \mathbf{n} within sub region Ω^j , or in other words, the volume fraction of local state \mathbf{n} , ${}^n V^j$. ${}^{nn'} f_t^j$ can be thought of as the joint probability of finding local states \mathbf{n} and \mathbf{n}' separated by the discrete set of vectors \mathbf{t} within sub region Ω^j . That is, if we throw a discrete vector corresponding to \mathbf{t} randomly into our discrete microstructure, ${}^{nn'} f_t^j$ gives the probability of finding local states \mathbf{n} and \mathbf{n}' at the tail and head of that vector respectively. Throughout this work, we will refer to the case when $\mathbf{n} = \mathbf{n}'$ as an auto-correlation and the alternative as a cross-correlation. When computing these quantities assumptions must be made regarding how to treat vectors that extend past the boundary of the data. It is important to note that this definition of the 2-point correlation

differs with respect to some definitions within the literature in one important regard. Here, we are capturing the probability for any arbitrary but specific vector (fixed length and orientation). Some definitions of the 2-point correlation, particularly when characterizing isotropic structures, integrate the angular component of this function and only parameterize by the length. Throughout this work, when discussing 2-point correlations we will always be referring to the full correlation, unless otherwise indicated. In most cases throughout this work we will assume periodic boundary conditions, however, in this section we will discuss methods for calculating these quantities using other boundary assumptions. Note the large difference in the amount of information contained in the 2-point statistics when compared with 1-point statistics. We go from only capturing volume fraction information to a measure that requires more memory storage than the original image itself. While it is possible to extend these definitions to an arbitrary numbers of points (thus the name, N-point correlations), this work does not discuss these measures. The computation of 3-point or greater correlations represents a difficult and open problem in the literature, at least to the author’s knowledge. Additionally, the exponential growth in the size of these measures necessitates finding redundancies or ways to compress this data to make it tractable to compute and store for even modest size microstructures datasets.

2.3 Efficient calculation of 2-point statistics

The most naïve approach for computing the 2-point statistics of a given microstructure image/realization $n m_s^j$ would be a sampling approach. That is, throw as many vectors of the given size and orientation into our microstructure and enumerate the counts for different local states. We would need to have a set of matrices for each possible

pair of local states in the image. These matrices could be used to enumerate counts for any combination of local states found separated by the given vector, and then one has to repeat this computation for all vectors of interest. Figure 1 represents this algorithm, specifically how vectors of different lengths and orientations could be accumulated within a counting matrix by translating the tail of the vector to the center of the matrix and using that as the origin. Such an algorithm would be effective however computationally it is not very efficient. In addition, how many vectors must we sample before the statistics converge? Within a m by n pixel image we would have up to m^2n^2 possible samples (exact number depends on which boundary assumptions that we would want to make). This is of a quadratic computational complexity. For large images, especially those with more than a few local states, this would be computationally inconvenient even on modern hardware.

Fortunately, there is a much more efficient way to compute these correlations if we realize that what we are actually computing is a convolution. First, create a version of our microstructure image where ones represent our local state of interest and all other pixels are zero. Choose a vector length and orientation of interest. Center a copy of the image on top of itself but offset it by this vector. Compute an elementwise multiplication of these matrices and sum them up. This will give us the count for this particular vector. If we repeat this for each vector we will have computed the auto-correlation of this local state of interest. That is, we can convolve this image with itself to compute the correlation. This works to our advantage because convolution is an operation that can be efficiently computed in Fourier space using a simple element multiplication. So with the use of the fast Fourier transform (FFT) and a subsequent inverse FFT we can compute the correlation in a much faster $O(n \log n)$ asymptotic complexity [37, 54, 56]. In order to express this

more formally, the discrete Fourier transform (DFT) of the microstructure function is written as

$${}^n M_k^j = \mathcal{F}({}^n m_s^j) = \sum_{s=0}^{S-1} {}^n m_s^j e^{2\pi i s \cdot k/S}, \quad (7)$$

where $\mathcal{F}(x)$ is the DFT, and k is the index of the terms of the Fourier series. We know that we can evaluate the DFT in $O(n \log n)$ computational complexity using the FFT. It follows from the convolution theorem that we can then perform the convolution in Fourier space, as an element wise product of the series,

$${}^{nn'} F_k^j = \mathcal{F}({}^{nn'} f_t^j) = \frac{1}{S} {}^n M_k^{j*} {}^{n'} M_k^j, \quad (8)$$

where the $*$ symbol denotes the complex conjugate. Note the normalization of statistics by the number of pixels/cells in the image. This is because, for any particular vector, we can place that vector on any pixel in the image by invoking fully periodic boundaries. Throughout this work, we will employ a simple visualization of the 2-point correlation function where the origin is placed at the center of the plot. To read the value ${}^{nn'} f_t^j$ for any particular \mathbf{t} , simply go to its coordinates and read the value. That is, place the tail of the vector at the origin and read the value at the head. An example of these visualizations is shown in Figure 3 and Figure 4.

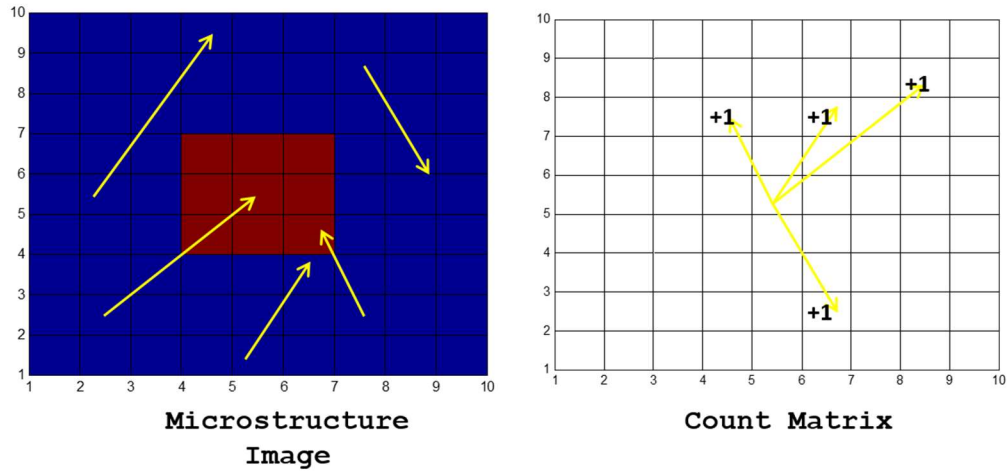


Figure 1: The image on the left represents a simple 2-phase/binary microstructure. Random vectors are thrown into the data with their heads and tails lying in different spatial bins. The image to the right shows how we could accumulate these vectors within a set of matrices, one for each correlation.




Figure 2: Choose a vector length and orientation of interest. Center a copy of the image on top of itself offset by this vector. Multiply each of the overlapping pixels together and add them up and you will have the total count for that particular vector.

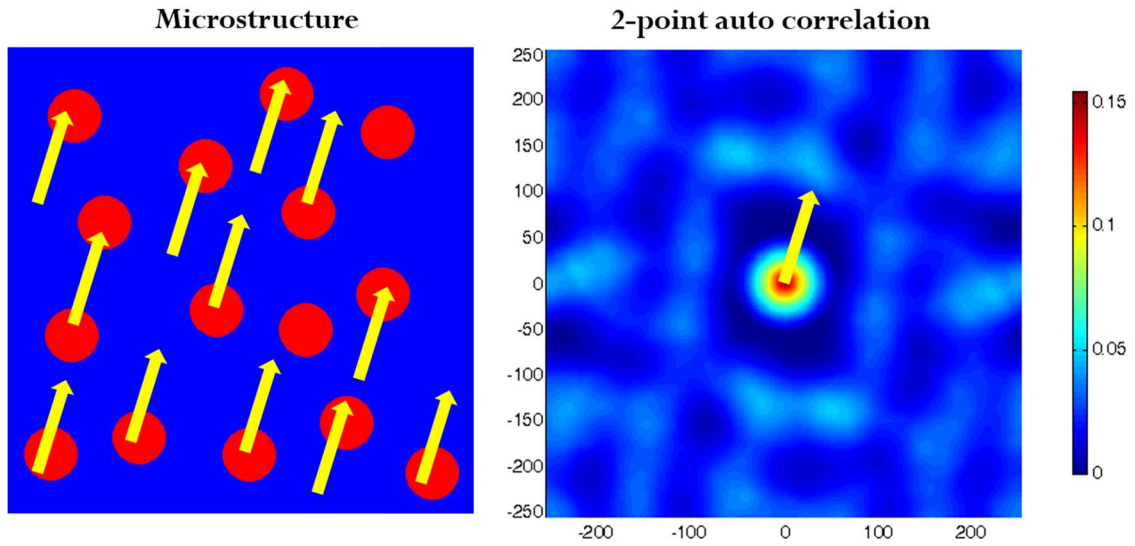


Figure 3: An example of a microstructure with two local states containing equally sized circles. Yellow example vectors are shown thrown into the microstructure. Not all vectors sampled are shown. A calculation of the 2-point auto-correlation (using the local state red). To read the probability associated with the yellow vector simply read the value at the head after placing the tail on the origin.

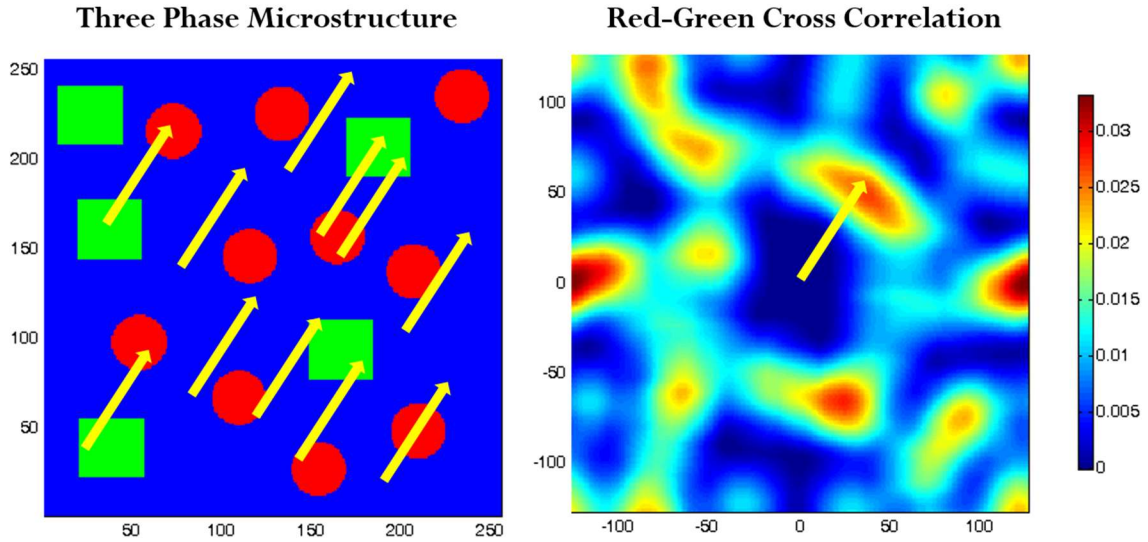


Figure 4: An example of a microstructure with three local states. The cross-correlation to the right is computed between the red and green local states.

2.4 Boundary Conditions

One result of utilizing Fourier methods to compute the 2-point statistics is the implied assumption that the microstructure is a periodic structure. That is, for vectors that extend beyond the boundary of our image it will be as if they wrap around to the other side and encounter whatever is present there in the image. Or to put it another way, it is as if we have padded the image with copies of itself. While typically this is not a problem in a lot of cases, there are alternatives for those that feel this may impact their use of these measures. One approach is to treat the boundary as if it is padded with zeros beyond the range of the image. In this case vectors that extend beyond the boundary will not be counted. However, this requires two considerations. First, normalizing these statistics is not as simple as the periodic case (divide by the number of spatial cells\pixels). All vectors are not sampled equally in this case and thus we need to normalize each count accordingly. The adjusted normalization and statistics is as follows,

$$\aleph(\mathbf{t}) = (S_1 - t_1)(S_2 - t_2)(S_3 - t_3) \quad (9)$$

$$m'_{f_t^j} = \frac{1}{\aleph(\mathbf{t})} \sum_{s=0}^{S-1} n m_s^{j n'} m_{s+t}^j \quad (10)$$

A plot of the normalization factor $\aleph(\mathbf{t})$ is shown in Figure 5. The second consideration that one must make when computing these non-periodic statistics is that we must actually pad the image with zeros to twice its size in all dimensions. This will increase the computational cost of the calculation when using an FFT approach. There may be intelligent ways to mitigate this cost (because the pad is all zeros) but the authors have not explored customizing the FFT for these purposes. A computational time comparison between different methods for computing a 2D 2-point auto-correlation are shown in Figure 7. Finally, because there are less samples for longer vectors than for shorter vectors one must consider whether this effects the accuracy of the statistics themselves.




Figure 5: An example of the normalization applied to non-periodic 2D 2-pt correlation. The origin is at the center of the plot. Each value is the total count of vectors that can fit into a 256×256 image. In the non-periodic case long and angled vectors can be thrown in less than short vectors so their counts must be normalized non-uniformly.




Figure 6: A comparison of periodic vs. non-periodic auto-correlation of the image of circles in the top right. Notice the more pronounced effects for longer vectors.

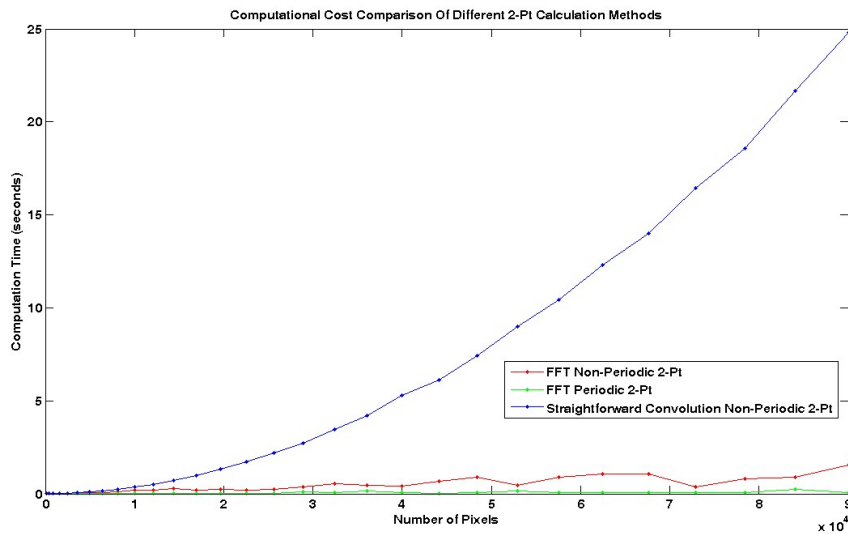


Figure 7: A comparison of computation time in seconds as a function of image size between 2-point correlation calculation algorithms. The blue line with $O(n^2)$ growth is a naïve implementation of the convolution. The red and green lines represent the $O(n \log n)$ FFT methods. Notice the increased cost that padding adds for the non-periodic case.

CHAPTER 3. REPRESENTATIVE VOLUME ELEMENTS

The previous chapters outlined a comprehensive framework for statistics of microstructure that can now be used to define representative volume elements (RVE or SVE). For this work we have elected to utilize 2-point correlations because of their computational efficiency. However, there is no reason that other statistical measures like the chord length distributions could not be used as a basis for defining a statistically representative volume element. The work described in this chapter culminated in the publication of two journal papers [57, 58].

The concept of the RVE in material science most likely dates back to Hill [59], who proposed a working definition of the RVE as;

“a sample that (a) is structurally entirely typical of the whole mixture on average, and (b) contains a sufficient number of inclusions for the apparent overall moduli to be effectively independent of the surface values of traction and displacement, so long as these values are macroscopically uniform.”

Part (a) speaks to the idea that our volume element should look like the microstructure in an average sense. That is, if we sample many volumes they should not look entirely different than our sub-volume. Our volume should be large enough that any inherent heterogeneity in the structure of the sample is represented within it. Part (b), which in Hill’s definition is the stronger requirement, makes the claim that the RVE should be large enough such that calculated (simulated) effective properties should be independent of boundary effects (i.e. the number of boundary elements is small relative to the number

of total elements for finite element method (FEM)). Many authors determine the RVE based solely on the convergence of effective properties for both Neuman and Dirichlet boundary conditions (for a review of important results and studies on the application of Hill's RVE see [60, 61]). It is important to note that using this approach the RVE is a function of both the effective property under consideration and the contrast in local property values between the constitutive microstructure components. In general, the RVE for highly nonlinear properties or high contrast composites (porous materials) will be substantially larger than for low contrast composites (phase separated metallic alloys) or linear properties.

It is often the case that, from a practical point of view, the RVE defined above is simply too large. Thus, an ensemble of smaller regions, often termed statistical volume elements (SVEs), is used to approximate effective properties. Properties calculated from the members of the ensemble are termed apparent properties and will exhibit a variance that is a function of SVE spatial size. Effective (RVE) properties can be bounded by apparent (SVE) property calculations [61-63]. For example, Huet [63] showed that for linear elastic materials the effective stiffness is always bounded from below by the harmonic average of moduli under Neuman boundary conditions and bounded from above by the arithmetic average of moduli under Dirichlet boundary conditions. Kanit et al. [64] constructed confidence intervals on the value of the effective property as a function of members in the SVE ensemble, while other authors determine ensemble size from convergence studies on the effective property as new members are added. In general, with smaller SVEs, more members are required in the ensemble, and the upper and lower bounds of effective properties are wider. This approach is reasonable when

defect-insensitive homogenized properties are sought for macroscale design. However, such approaches have not enjoyed much success in applications involving defect-sensitive macroscale properties or performance criteria that depend strongly on the spatial distributions of specific microscale features.

However, the work presented in this chapter attempts to more formally define part (a) of Hill's criteria, in hopes that in doing so we will arrive at a more robust and effective definition of the RVE. That is, using our previously defined concepts of N-point correlations as a measure of microstructure statistics, can we define an RVE that is statistically representative of the sample microstructure. There are at least two reasons that this approach might be more effective than the convergence of effective properties criteria described above. First, we can define an RVE that is entirely defined by structure, regardless of property of interest. It is known that structure-property relationships are many-to-one [65] and [66], and thus just because we have found volume that has similar effective properties does not mean it is truly representative of the structure. Effective properties emerge at different length scales and there is no reason to assume an RVE for one will be an RVE for another. However, if we represent the structure accurately, it should follow that many properties are captured accurately as well.

Second, what does one do when the necessary RVE size using Hill's criteria is not computationally tractable for simulation. In this case, using the concept of microstructure spatial statistics, we can try to optimize the selection of an ensemble of smaller SVEs rather than take a random sample, which is typically done. The goal of such an optimization would be to find a set of SVEs that best approximate the statistics of the full sample microstructure. In doing so, researchers could optimize for the best RVE under a

give computational budget. This can be thought of as a form of intelligent sub-sampling. In this work we refer to such optimized ensembles as weighted statistical volume element sets (WSVE sets).

3.1 Ensemble Averaged Microstructure Statistics

In chapter 2, we discussed the definition of a discretized microstructure function and their related N-point correlations. These two components form the basis for defining a representative volume element in a statistical sense. To begin with constructing a statistical RVE we must first estimate the statistics of the microstructure we wish to represent. We may think of these sample statistics as the target or goal of our approach. The most-straightforward way to estimate these statistics would be to take an ensemble average of a large number of characterized regions of the microstructure in question. Or, more typically, given a large characterized volume (for which simulation is impractical), select a sub-volume size and randomly extract as many as needed from the full sample until our statistics converge. The ensemble-averaged 1-point and 2-point statistics over J such sub regions\ volumes ($\Omega^1, \Omega^2 \dots, \Omega^J$) are simply defined as

$${}^n f_t^\Omega \approx {}^n \bar{f}_t = \frac{1}{J} \sum_{j=1}^J {}^n f_t^j = \langle {}^n f_t^j \rangle \quad (11)$$

$${}^{nn'} f_t^\Omega \approx {}^{nn'} \bar{f}_t = \frac{1}{J} \sum_{j=1}^J {}^{nn'} f_t^j = \langle {}^{nn'} f_t^j \rangle \quad (12)$$

where ${}^n f_t^\Omega$ and ${}^{nn'} f_t^\Omega$ are the material sample statistics. These equations express a simple ensemble average of the statistics as an estimate of the true sample statistics. The

natural origin of the spatial statistics allows us to do this. If we did a similar average over the microstructure realizations themselves the result would be effectively meaningless. However, one important consideration to make is that all of this assumes that the SVEs are drawn from a statistically homogenous microstructure. If the sample contains any structural gradients, then these assumptions do not apply and this work is not applicable.

Using the spatial statistics, we can define the idea of coherence length, which is the length scale beyond which the spatial distributions of local states are completely uncorrelated. Or mathematically, the coherence length t_c

$$\langle nn' f_t^j - n f^j \cdot n' f^j \rangle \leq \epsilon \quad \forall \|t\| \geq t_c \quad (13)$$

where ϵ denotes some sufficiently small tolerance. This allows us to define a minimum length scale of interest based on our SVE ensemble. However, this does not mean that any two members of our SVEs will have similar statistics beyond this length scale (see Figure 8)




Figure 8: (a) Ensemble averaged normalized 2-point particle-particle autocorrelations, $\langle (n_t^{nn} f_t^j - (n f^j)^2) \rangle$, calculated over 100 independent 150×150 pixel realizations of non-overlapping disks of diameter 20 pixels with a minimum spacing of 5 pixels. The ensemble average statistics clearly shows the discs are spatially uncorrelated with a coherence length of 25 pixels. (b) Three independent realizations of the microstructure. (c) The normalized 2-point statistics of the realizations shown in (b). Even though the dimensions of each realization are six times larger than the coherence length, the correlation functions for each realization show significant correlations at the longest vectors.

Finally, for our purposes, it is important to define a scalar distance metric between two different sets of microstructure statistics. In our case, we will be computing the distance between the estimate sample statistics (ensemble average) and a sub-domain Ω^j of our WSVE set. We will define this distance as

$$D_S^j = \frac{1}{S} \sqrt{\sum_{t=0}^{S-1} \sum_{n=0}^{N-1} \sum_{n'=0}^{S-1} (nn' f_t^j - nn' \bar{f}_t)^2} \quad (14)$$

where S defines the spatial extent of the sub-domain as in Equation 4. This definition is dependent on this size so it is denoted as a subscript. As S approaches the size of the full sample the distance will converge to zero. However, our original goal is to sub-sample the data, thus a researcher must choose an effective size S to both minimize D_S^j while still considering the practicality of a WSVE set with elements of size S .

3.2 Statistical Volume Element (SVE) Size

Selecting the size of each element of an SVE ensemble is necessary step as the size remains fixed for all members. One simple approach is to only consider the computational budget that the researcher has and let that determine the SVE size. That is, if the researcher's goal is to perform FEM analysis on each member of the SVE set, then they could set the size to the maximum for which they could run such simulations. However, let us assume for the sake of discussion that our hypothetical researcher wishes to optimize both their compute budget and their results. A more sophisticated approach would then be to look at the value of D_S^j as a function of S . More specifically, can we see a point where changing the size of the SVE element does not improve the error (on average) by an amount that would justify the extra computational cost.

More formally, let us extract for each SVE size in question, a large number of sub-volumes\regions from our full sample and compute their individual 2-point statistics. We can then compute their ensemble average statistics using Equation 12. Using this as our

sample target statistics $nn' \bar{f}_t$ in Equation 14, we can then compute a scalar error for each individual sub-volume D_S^j . The mean error \bar{D}_S^j and standard deviation error σ_{D_S} will decrease as the SVE size approaches the full sample size. We can look at this convergence to see at which size S this value seems to stop sharply decreasing in order to get an idea where to define the SVE size. For example, in Figure 9 we show an example of a digitally created microstructure with two distinct local states. In Figure 10, we show the relationship between SVE size and the scalar error D_S^j of individual sub-volumes. We can see a plateau in both the mean and variance at around 200x200 pixels. This means that there is probably not much need to extend the size of the SVE beyond this range for this particular sample. For the rest of this chapter we will refer to the selected SVE size as \tilde{S} .




Figure 9: Digitally created 2-phase microstructure (550×550 pixels) for which an SVE Set is to be created.

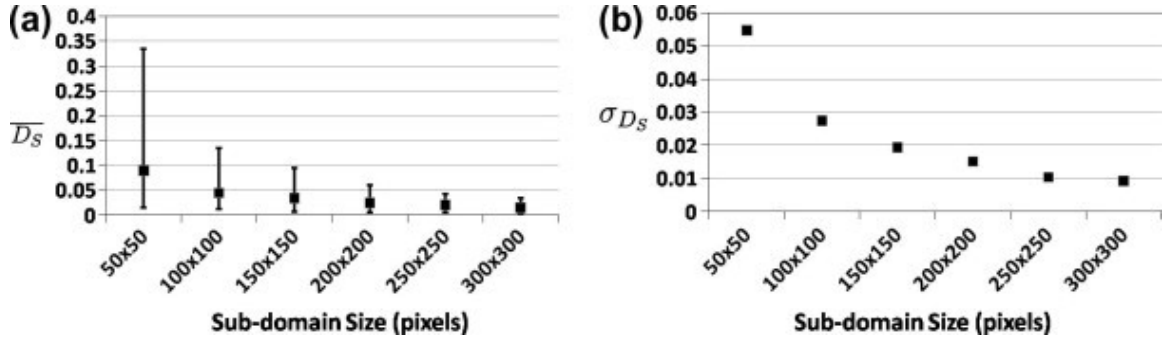


Figure 10: (a) Mean scalar error, \bar{D}_S^j , for ensembles of 100 members of different sub-domain sizes. Error bars indicate maximum and minimum observed values of D_S^j for each ensemble. (b) Standard deviation of D_S^j as a function of the sub-domain size.

3.3 Building Weighted Statistical Volume Element Sets

Now that we have determined the size of our individual SVE sub-regions we can construct an optimized set to best match the target sample statistics. Our ability to precisely match the ensemble-averaged statistics of the sample depends on the number of elements in the SVE Set. In the limiting case, when we extract a very large number of SVEs of size Ω_S^j from the original sample and consider them all as equally weighted elements of the RVE set, we can recover precisely the statistics of the entire sample as an SVE ensemble. It is worth reminding ourselves that this is what was precisely done in establishing the plot in Figure 10 (in this plot 100 elements were used for each selected domain size).

Instead of creating a very large SVE set (large in terms of number of SVEs, not SVE size), we can optimize each selected SVE element such that it drives us closer to the sample statistics. Doing this would hopefully allow us to create a much smaller SVE set that is still statistically representative. In addition, because we are using statistics, we can allow the weights assigned to each SVE be non-equal. Doing this allows us to build a

weighted SVE (WSVE) set. Let N denote the number of volume elements in the WSVE set. Let S denote the number of pixels\voxels dictating the size of the spatial domain of the SVE. Let the ensemble $(\Omega^1, \Omega^2 \dots, \Omega^J)$ denote the set of potential volume elements (each of size S) that could be included in the WSVE set (presumably these have been extracted from the sample). For sake of discussion, let us denote the simple concatenation of all individual 2-point correlations $^{nn'} f_t^j$ as a very high-dimensional point f_2^j . Here the subscript 2 denotes the 2-point correlation. We can visualize the two-point statistics for the ensemble $(f_2^1, f_2^2, \dots, f_2^J)$ as a set of points in a very high-dimensional (on the order $O(H^2S)$, where H is the number of distinct local states) space. For a given WSVE set size, N , the search for the best possible elements of the WSVE set reduces to a search in the convex hull defined by $(f_2^1, f_2^2, \dots, f_2^J)$ for points, lines and hyperplanes that are closest (by Equation 14) to the ensemble average $^{nn'} \bar{f}_t$. For example, if N is selected to be 1, then we are looking for the specific f_2^j that is closest to ensemble average. If N is selected to be 2, we are searching for two specific points, say f_2^j and f_2^k , such that a bounded line between them passes closest to the ensemble average. With N equal to 3, we are searching for a bounded triangular surface defined by a triplet of (f_2^i, f_2^j, f_2^k) that passes closest to the ensemble average. For arbitrary sizes of N , we are looking for the N individual statistics that form the closest hyperplane to the ensemble average. For any WSVE set size greater than 1, we can easily compute the corresponding set of weights for each member by considering the convex combination of their statistics that comes closest to ensemble statistics.

While the brute force (try all combinations in the ensemble) computations for small values of N are manageable, as the size N of the WSVE set increases, the number of possible combinations increases exponentially. We have therefore developed and implemented new heuristics that are particularly suited for larger values of N (typically larger than 2 or 3). The search begins by selecting a random point f_2^i . It then identifies the closest line to the ensemble average using f_2^i and any other point in the ensemble. After finding the closest line, passing through f_2^i and f_2^j , it identifies the convex combination of f_2^i and f_2^j that lies closest to the ensemble average; let us denote this point \hat{f}_2 . At this stage, we have established the first two members of our WSVE (Ω^i and Ω^j). To add another member, we simply search again for the closest line between \hat{f}_2 and any other point in our ensemble, other than those already in the WSVE set. This procedure continues until we have found N members of the WSVE set. Since the establishment of the WSVE by this procedure is critically dependent on the selection of the first point in the ensemble, we repeat the entire WSVE selection process described above several times, each time with a different (but random) selection of the first point in the WSVE set. If there are J volume elements in ensemble, this can lead to a maximum of J trials. Finally, we compare all the constructed WSVE sets and pick the one which provides a convex combination that is closest to the target statistics. It should be noted that although these heuristics do not guarantee the most optimal WSVE set, we have successfully constructed acceptable ones for numerous case studies papers [57, 58].

3.4 Examples

The following sections outlines two example case studies in which the methods described in this chapter were used.

3.4.1 Two-Phase Digitally Created 2D Microstructure

Utilizing, the microstructure shown in Figure 9, we performed a case study showing the usefulness of these methods for defining an RVE. Figure 11 and Figure 12 depict SVE sets of size 1 and 3 respectively. The corresponding difference between the 2-point statistics of the produced SVE Set (as a weighted combination of the instantiations shown) and the ensemble-averaged statistics of the sample, defined by Equation 12, as well as the calculated properties for the sample and RVE, are summarized in Table 1. To validate our concept of the SVE set we have chosen two properties of interest; the macroscale Young's modulus and the interface fraction. These specific attributes were selected as representative examples of a macroscale property and a microstructure metric that depend on the 2-point statistics in the microstructure. The macroscale Young's modulus was computed using the second-order homogenization theory described and validated in the following works; [65, 67, 68]. In this work, the stiff (white) phase and the soft (black) phase were assumed to possess Young's moduli of 70 and 45 GPa, respectively, along with a Poisson's ratio of 0.3 (same for both phases). The interface fraction (as an approximation to interface area per unit volume) is simply the number of interface pixels over the total number of pixels in the SVE.

If we examine results shown in Table 1, we can see that the optimal SVE set of size 1 captures the macroscale module with less than 1% error and the interface fraction with 4%. If we utilize and SVE set size of 3 we can reach less than 1% error and less than

0.1% error with a set size of 5. Also note the decreasing value of D_s shown in Table 1 and Figure 13. An important point to realize from these results is that for different properties of interest a different SVE size \tilde{S} and set size N might be needed to accurately capture the effective properties. We can see this with the macroscale Young's modulus when compared to the interface fraction. It is likely that any arbitrary sub-region of size 100x100 pixels on this sample would capture the Young's modulus to a 1% error. However, this is not true for the interface fraction. This is why constructing SVE sets using a statistical basis is more robust than defining them based on certain specific properties and then hoping that they are representative for others.

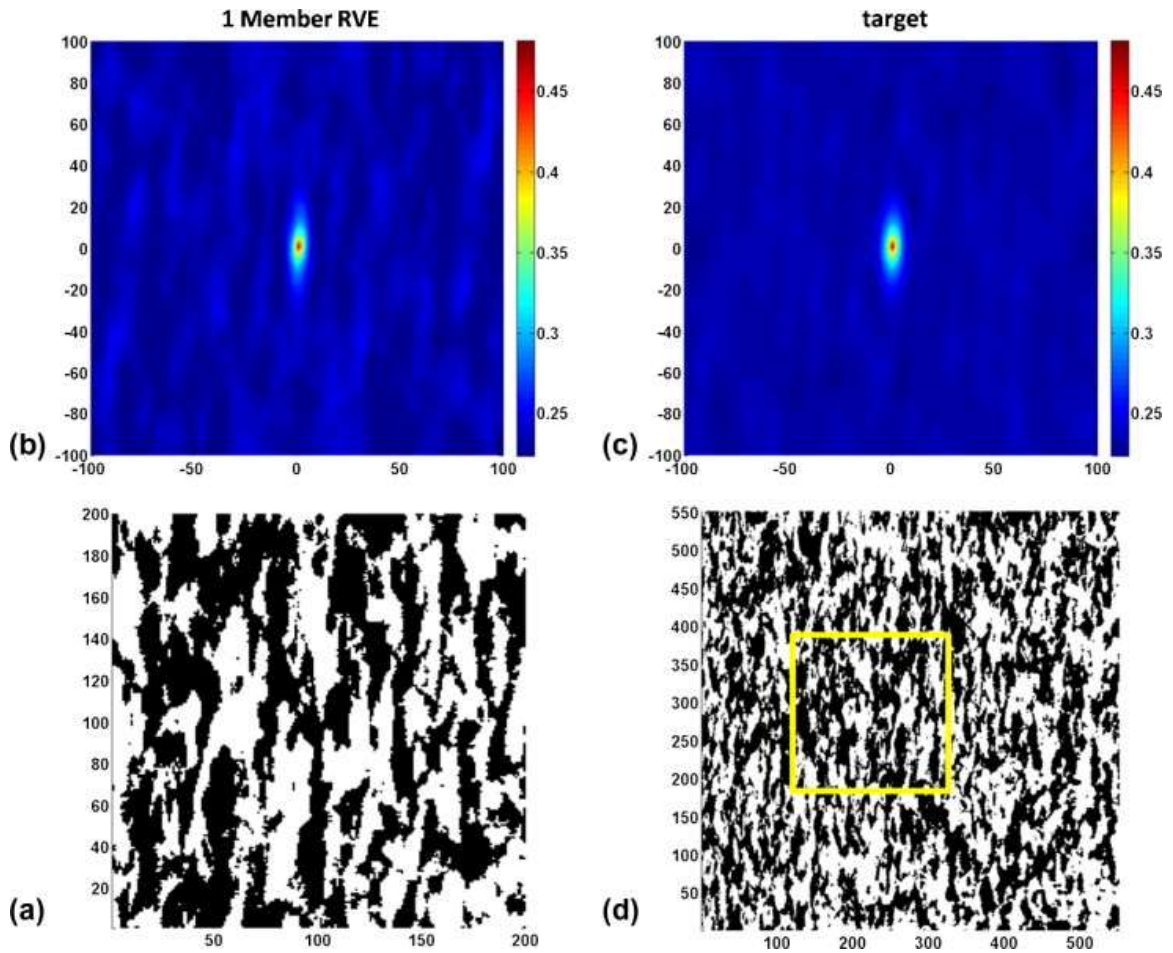


Figure 11: The 200×200 pixel sub-domain (a) and its associated 2-point statistics (white–white autocorrelation) (b) that comes closest to matching the statistics (c) of the target microstructure (d). The main features of the target 2-point statistics are evident in the RVE statistics, such as the intensity and shape of the central spot. The location of the sub-domain shown in (a) is highlighted in (d).




Figure 12: Members (a–c) and associated statistics (d–f) of the three-member RVE Set. The statistics (white–white autocorrelations) are all plotted with the same color-scale to highlight the differences observed between the members. Individually the statistics of the individual members are a poor match to the target statistics shown in Figure 11c; however, their weighted average (h) is a close match. The location of the individual windows in the target microstructure (g) is highlighted.

Table 1: Comparison of scalar error D_S , effective Young’s modulus, and interface fraction between the large microstructure data set (labeled as sample) shown in Figure 9 and the corresponding RVE Sets produced in this study (see Figure 11 and Figure 12).

	SAMPLE	SVE SET SIZE 1	SVE SET SIZE 2	SVE SET SIZE 5
D_S		0.0247	0.0111	0.0078
EFFECTIVE MODULUS (GPA)	59.7	60.1	60	59.7
% ERROR		0.67	0.5	0
INTERFACE FRACTION	0.125	0.1295	0.1259	0.1251
% ERROR		3.6	0.72	0.08

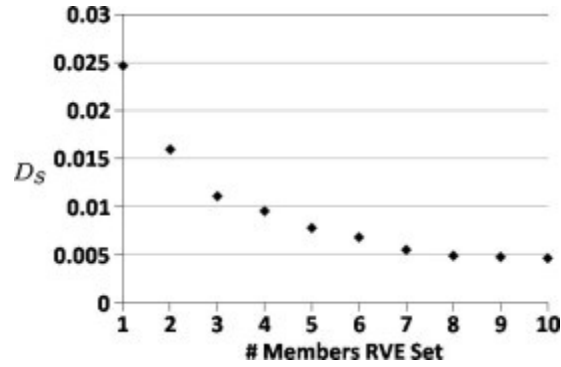


Figure 13: The influence of the SVE Set size on the scalar error between the 2-point statistics of the SVE Set and those of the target microstructure, D_s . By adding more members to the SVE Set, the target statistics can be captured to improved accuracy.

3.4.2 2D Polycrystalline Microstructure

Here we demonstrate the construction of an SVE set for a large ($1936 \mu\text{m} \times 1956 \mu\text{m}$ scan with 946,704 measurement points) polycrystalline microstructure obtained from a lightly rolled Fe–Si steel sample, characterized by electron backscattered diffraction (see Figure 14). The microstructure local state description for each spatial bin is a triplet of three Bunge-Euler angles $[\phi_1, \Phi, \phi_2]$ [69] specifying the local crystallographic lattice orientation. In order to create a manageable number of discrete local states we have binned the orientation space into (cubic-triclinic fundamental zone) into 512 equal volume bins [70]. The complete set of possible 2-point correlations includes $512^2 = 262,144$ correlations – far too many to enumerate directly. Fortunately, this set contains a large number of interrelations [37] that reduce the set of independent correlations to just 511. For this case study, we have fixed the value of \mathbf{n} at the most populous bin (#460 corresponding to $\phi_1 = 321^\circ$, $\Phi = 85^\circ$, $\phi_2 = 5^\circ$) and let \mathbf{n}' vary across the full local state space (511 bins).

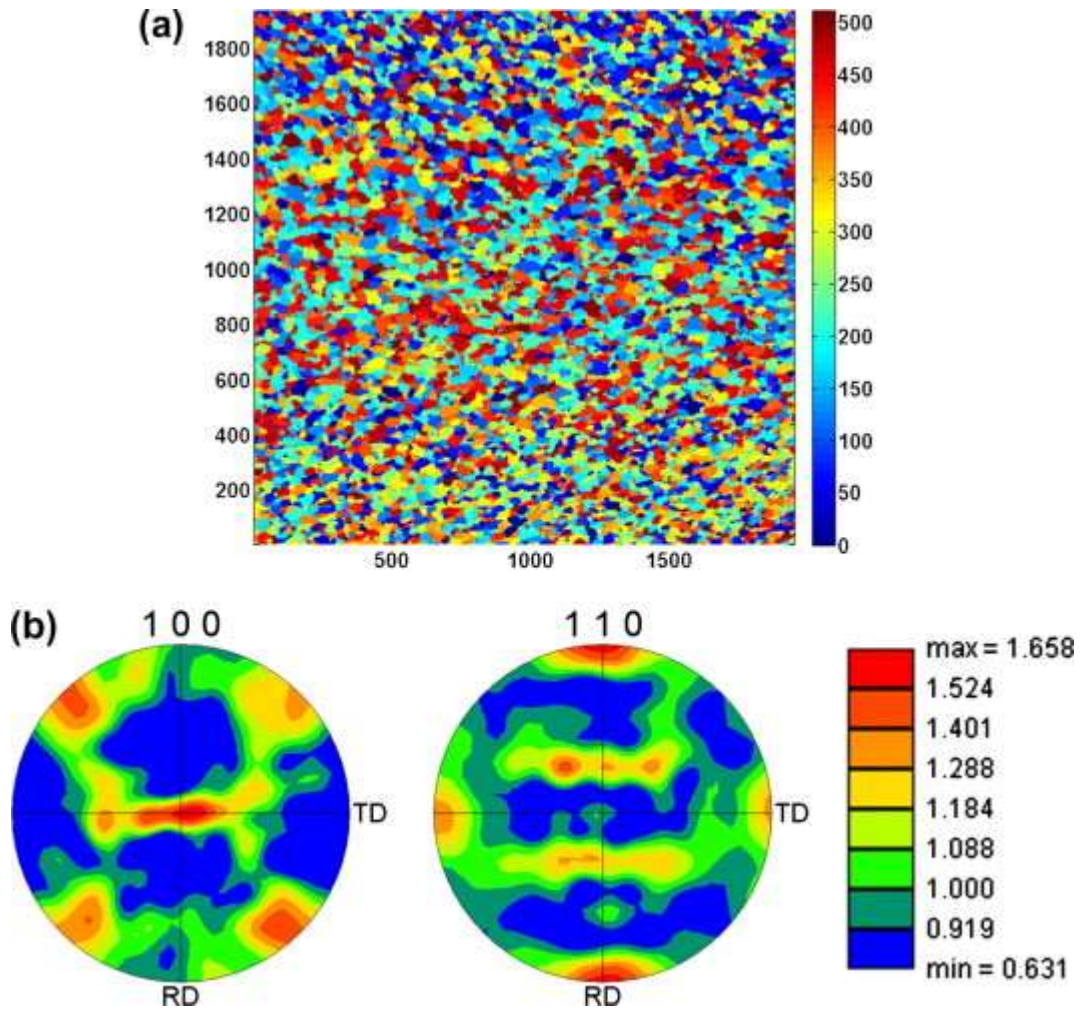


Figure 14: (a) EBSD grain map of Fe–Si steel where the fundamental zone of crystal orientation has been binned into 512 bins in the Euler-angle space. Each pixel represents 1 μm . (b) The corresponding texture of the microstructure shown in (a) as $\langle 100 \rangle$ and $\langle 110 \rangle$ pole figures showing that the sample is only weakly textured.




Figure 15: Standard deviation of D_S^j as a function of the sub-domain spatial extent for the microstructure data set shown in Figure 14.




Figure 16: The texture corresponding to the one- (top) and three-member (bottom) SVE Sets, demonstrating that even the very weak textures can be accurately captured by a relatively small SVE set. The maximum intensities in the SVE Sets are somewhat higher due to a much smaller number of grains in these sets.

Table 2: A comparison of scalar error D_s and the first-order upper and lower bounds for the components of the stiffness tensor and tensile, and shear yield strengths between the entire sample and an optimize SVE set of size 1.

	Entire sample		SVE Set size 1			
D_s			0.0247			
Property	Upper bound	Lower bound	Upper bound	Lower bound	% Error UB	% Error LB
Elastic						
C1111	281.74	261.45	283.91	263.55	0.77	0.8
C2222	281.76	261.47	280.64	260.51	0.4	0.37
C3333	274.83	255.07	274.83	255.1	0	0.01
C1212	86.17	71.43	85.64	70.95	0.61	0.67
C1313	93.1	78.36	91.45	76.59	1.77	2.25
C2323	93.08	78.34	94.72	80.18	1.77	2.36
C1122	121.96	90.88	121.38	90.37	0.47	0.56
C1133	128.62	97.55	126.99	95.98	1.26	1.61
C2233	128.6	91.15	130.15	90.69	1.21	0.51
Yield						
Sigma11	2.93	2.23	2.97	2.24	1.28	0.18
Sigma22	2.93	2.23	2.9	2.22	1.19	0.26
Sigma33	2.85	2.18	2.84	2.19	0.31	0.34
Tau12	1.59	1.3	1.58	1.29	0.45	0.57
Tau13	1.66	1.33	1.64	1.32	1.07	0.84
Tau23	1.66	1.33	1.67	1.33	0.66	0.41

Table 3: A comparison of scalar error D_s and the first-order upper and lower bounds for the components of the stiffness tensor and tensile, and shear yield strengths between the entire sample and an optimize SVE set of size 3.

	Entire sample		RVE Set size 3			
D_s			0.0111			
Property	Upper bound	Lower bound	Upper bound	Lower bound	% Error UB	% Error LB
Elastic						
C1111	281.74	261.45	282.96	262.65	0.43	0.46
C2222	281.76	261.47	281.64	261.42	0.04	0.02
C3333	274.83	255.07	274.83	255.08	0	0.01
C1212	86.17	71.43	85.62	70.93	0.64	0.69
C1313	93.1	78.36	92.43	77.64	0.72	0.91
C2323	93.08	78.34	93.75	79.09	0.72	0.96
C1122	121.96	90.88	121.38	90.34	0.47	0.6
C1133	128.62	97.55	127.95	96.91	0.52	0.65
C2233	128.6	91.15	129.23	90.62	0.49	0.58
Yield						
Sigma11	2.93	2.23	2.95	2.24	0.84	0.37
Sigma22	2.93	2.23	2.92	2.23	0.3	0.08
Sigma33	2.85	2.18	2.84	2.19	0.15	0.11
Tau12	1.59	1.3	1.58	1.29	0.51	0.37
Tau13	1.66	1.33	1.65	1.33	0.6	0.35
Tau23	1.66	1.33	1.66	1.33	0.2	0.16

Figure 15 shows our analysis for determining the size of the SVE for this case study. Like before, we plot the standard deviation of the scalar error σ_{D_s} as a function of SVE size S . In this case we selected a size of $600 \mu\text{m} \times 600 \mu\text{m}$ because the variance seemed to plateau at this point with diminishing returns beyond this size. The textures of the one- and three-member SVE sets are shown in Figure 16 as $\langle 100 \rangle$ and $\langle 110 \rangle$ pole figures, demonstrating that, even though the sample is only very weakly textured, the orientation distribution is reasonably well captured in the RVE set of size 3. To validate the properties of these SVE sets we looked at a range of elastic properties (predicted using the elementary first-order upper and lower bounds [59, 71-73]) and yield properties (Hill's anisotropic yield parameters [74], predicted using crystal plasticity theories [75]) and calculated them for the SVE sets and the full sample shown in Figure 14. The results are summarized in Table 2 and Table 3. For a single-member SVE Set, the first-order bounds on all of the selected macroscale elastic and yield properties were captured to within 3%. However, it was observed that a three-member RVE set captured the same properties to within 1%.

3.4.3 Large 3D Polycrystalline Microstructure

Throughout this chapter we have been referring to volume elements. However, the case studies above dealt with 2D samples only so this has been a bit of misnomer. We will now focus on a case study involving a very rare and large 3D polycrystalline dataset. This work has been published as well. [58].

The dataset in this study is a large (4300 grain, $1.115 \times 0.516 \times 0.3 \text{ mm}^3$, $1670 \times 770 \times 200$ voxels) 3D sample of bcc β -stabilized Ti 21S. The microstructure was

reconstructed using serial sectioning, optical microscopy with periodic electron backscatter diffraction (EBSD). In addition, image processing and segmentation were required to create the final sample. The complete methodology has been published in a paper by Rowenhorst et al. [76]. For each grain, an average crystallographic orientation is determined from values measured using EBSD. In Figure 17, we show the volumetric dataset, with each color signifying its orientation parallel to the Z direction. The depicted Z direction is parallel to the sectioning direction. To reduce memory and computational requirements in the subsequent analyses and modeling, the dataset is compressed by sampling every third voxel in the XY-plane of the reconstruction, and every second voxel in the Z-direction.




Figure 17: Sample of the 4300-grain sample of β -stabilized Ti 21S ($1.115 \times 0.516 \times 0.3$ mm³, $1670 \times 770 \times 200$ voxels). The color key corresponds to the crystallographic orientation parallel with the Z-axis $[0\ 0\ 1]$.

In order to test our methodologies, we constructed three RVE sets. First, a sub-optimal WSVE set was built. We call this set sub-optimal because it contains four

arbitrarily selected SVE's of similar size and number of grains (approximately 100), with a fifth one selected so as to make the two-point statistics of the five-member set match the average ensemble statistics as closely as possible. We call the first four SVE's the random WSVE and the four SVEs with optimal fifth SVE the sub-optimal WSVE. Second, for the optimal case we use the algorithms described in section 3.3 to construct the full five member WSVE set. In addition, we chose larger SVE sizes (approximately 200 grains) for this set. The size difference and the level of optimization in the selection procedure were intentionally selected to evaluate vastly different examples of WSVEs that could be produced using the statistical microstructure-based representation methodology described in this paper. Finally, we have what we call a traditional single volume element. This traditional RVE has as many as five times the number of grains as our smallest WSVE sets. Based on prior modeling experience with this material, this size is expected to adequately represent both effective and local material behavior in the context of initial yield (small deformations). In Figure 18, we illustrate the locations and relative sizes of this RVE and the WSVE sets described. In Table 4, we show the corresponding weights for WSVE members as well as their dimensions and number of grains.

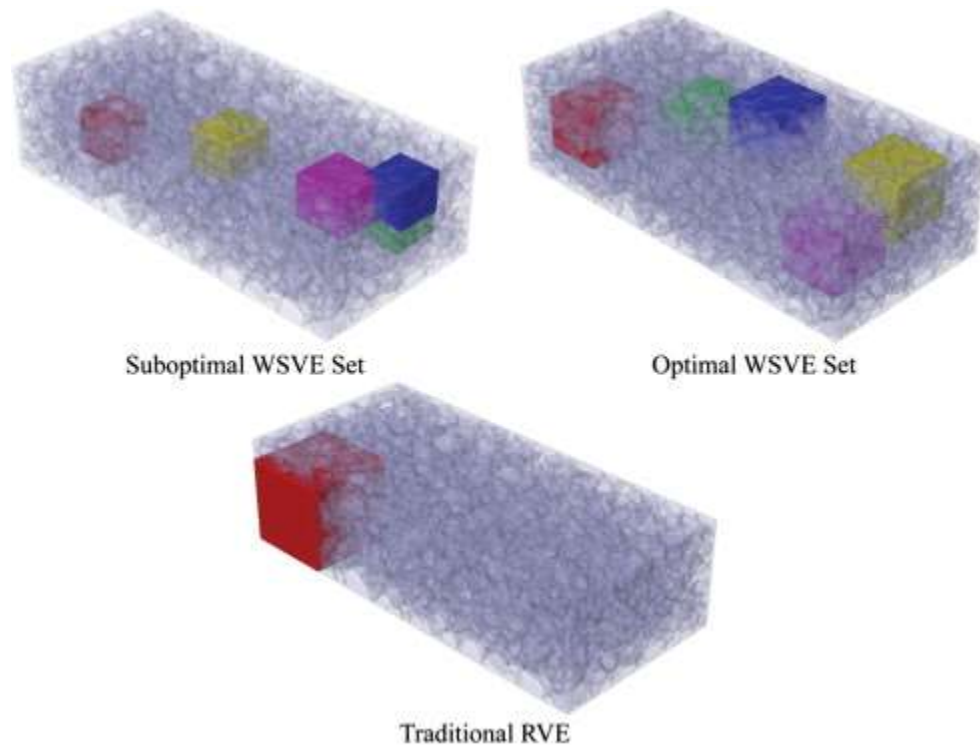


Figure 18: Locations and relative sizes of all five members of each representative WSVE set and the traditional RVE in the 4300-grain sample.

Table 4: Specifications of the representative WSVE sets and the traditional RVE. Note the extreme weight given to element E in the suboptimal WSVE set. This is element selected based on 2-point statistics while the others are random. Also note the weights of the optimal WSVE set, they are more equal as the elements have been selected heuristically.

SVE	Grains	Dimensions (μm)	Elements	Degrees of freedom (approx.)	Weights (W_i)
<i>Suboptimal WSVE Set</i>					
A	99	132 × 148 × 113	185,592	668,131	0.1732
B	101	120 × 132 × 125	166,320	598,752	0.0701
C	87	148 × 144 × 118	213,120	767,232	0.0630
D	97	152 × 140 × 112	202,160	727,776	0.0820
E	92	120 × 132 × 112	150,480	541,728	0.6117
Average	~95	134 × 139 × 116	183,534	660,723	–
<i>Optimal WSVE Set</i>					
A	224	168 × 184 × 157	409,584	1,474,502	0.2282
B	217				0.1462
C	198				0.2177
D	192				0.1944
E	171				0.2135
Average	~200	168 × 184 × 157	409,584	1,474,502	–
<i>Traditional RVE</i>					
–	499	243 × 241 × 244	1,210,484	4,357,742	–

3.4.3.1 1-point and 2-point Statistics

To compare the structural representation of our sample and these different RVE sets we computed pole figures and have included them in Figure 19. Pole figures basically represent a form of 1-pt statistics, that is, volume fractions of different orientations found in the sample. As is typical, the values shown in the pole figures depicted in Figure 19 indicate times random. Since we have used 2-point correlations as the statistical basis for constructing our WSVE sets we should expect that they accurately capture the 1-point statistics of the sample. This is because N-point correlations are a hierarchical measure and thus the 2-point correlations contain the 1-point correlations

entirely. This is exactly what we show in Figure 19. The first distinctive feature to note is that while the texture of the sample is very random with a maximum magnitude near 2.2 times random, the suboptimal WSVE set exhibits a stronger texture with maximum near 4.5 times random. The traditional RVE is around 3 and the optimal WSVE set is around 2.2. These distributions are expected, since the total volume of the optimal WSVE set is larger than the traditional RVE and similarly the traditional RVE is larger than the suboptimal WSVE set. That is, while the single traditional RVE has a larger spatial range than any single element of the optimal WSVE set, the total number of crystallographic grains is larger in the full optimized set of five weighted SVEs. This of course leads to a closer statistical match when comparing the samples via a one-point statistical representation. In addition, we examined the scalar error D_S between the 2-point statistics of the sample and our WSVE sets. These results are shown in Table 5. Note that as expected the completely random WSVE set has the highest error with the sample, followed by the sub-optimal WSVE set, then the optimal WSVE set, and finally the traditional single RVE. Interestingly, though close, the optimal WSVE set does not have lower error than the traditional RVE despite its larger combined volume of the WSVE. This is most likely due to the larger spatial extent of the traditional RVE and thus its greater ability to capture longer range correlations. Still, the suboptimal WSVE set has a significantly larger statistical error, whereas the statistical error of the random set (i.e. SVEs A–D of the suboptimal set with equal weights) is significantly worse than all. This result shows that the WSVE sets constructed via matching the two-point statistics are indeed more accurate than a simple random selection.




Figure 19: A comparison of orientation distributions via pole figures in three crystallographic directions of interest. The color value indicates times random. The samples compared are the full 4300-grain dataset, the traditional RVE, the optimal WSVE set and the suboptimal WSVE set. The pole figures are created using the same level of smoothing for all.

Table 5: Statistical error between representative WSVE sets, the traditional RVE and ensemble statistics for the entire sample.

Sample	<u>D (Statistical Error)</u>
Suboptimal WSVE Set	10.7588
Random WSVE Set	17.8115
Optimal WSVE Set	5.0826
Traditional RVE	3.0477

3.4.3.2 Global and Local Mechanical Behaviour

Similarly, to our previous two examples, we have also conducted simulated experiments to gauge how accurately our WSVE sets are capturing specific properties of interest. These simulations were conducted by our collaborators Qidwai et al. at the Naval Research Laboratory (NRL). We would like to acknowledge their specific and significant contribution to this section. The mechanical simulations discussed here were calculated using commercial finite element (FE) software ABAQUS® [77]. All SVEs and RVEs were converted to FE meshes by using a one-to-one correspondence between voxel centroids and the centroids of eight-node brick elements (ABAQUS designation: C3D8). An example of one of these SVE in mesh form is shown in Figure 20. Global displacement based loading conditions are applied such that uniaxial tensile behavior is obtained effectively. The displacements are applied to attain up to 2.5% uniaxial strain, enough to cause the global response of the SVE to evolve beyond the initiation of plastic flow. The rate-dependent crystal plasticity constitutive model described by Asaro (1983) and Huang (1991) on the basis of resolved shear stress, is used to describe the material behavior. The interested reader can find details of the model in Ref. [78]. The user-material subroutine UMAT developed by Huang [79] based on this model is employed in

the simulations. Representative material parameters available in the literature for the β -Ti 21S are used and shown in Table 6. The elastic constants are taken from Ledbetter et al. [80], whereas the constants related to plastic slip are incorporated from Deka et al. [81]. The crystal plasticity constitutive description and its implementation in FE models described here is fully consistent with prior efforts in literature [75] and [82].

Table 6: Representative constitutive model parameters for β -Ti 21S alloy.

Material parameter	Value
<i>Family of slip systems (48 independent)</i>	$s = \langle 111 \rangle \quad n = \{110\}$ $s = \langle 111 \rangle \quad n = \{112\}$ $s = \langle 111 \rangle \quad n = \{123\}$
<i>Elastic moduli</i>	$C_{11} = 97.7 \text{ GPa}, C_{12} = 82.7 \text{ GPa}, C_{44} = 37.5 \text{ GPa}$
<i>Shearing rate parameters</i>	$m = 50, \quad g_o = 0.0023 \text{ s}^{-1}$
<i>Hardening moduli parameters</i>	$h_{o1} = 1.5 \text{ GPa}, h_{o2} = 1.5 \text{ GPa}, h_{o3} = 1.5 \text{ GPa}$ $t_{o1} = t_{o2} = t_{o3} = 200 \text{ MPa}, t_{s1} = t_{s2} = t_{s3} = 500 \text{ MPa}$ $q1 = q2 = q3 = 1^*$

Figure 21 contains the global uniaxial stress–strain responses associated with the suboptimal WSVE set in four graphs. The first three (a, b and c) show the uniaxial behavior of the individual SVEs in the set for loading in the three global directions, X, Y and Z, respectively, whereas the last one shows the effective RVE response in the three global directions defined as:

$$RVE \text{ Response} = \sum_i^N \text{weight}_i \times WSVE \text{ Response}_i \quad (15)$$




(a) Finite Element Mesh

Uniaxial Tension

— undeformed

⋯ deformed



(b) Displacement Boundary Conditions

Figure 20: (a) Superposition of a regular mesh on the underlying microstructure is used to create the finite-element model for an approximately 100-grain volume. (b) Description

of displacement-based uniaxial loading conditions. The four transverse planes (two are not shown in this 2-D view) are constrained to move as planes.

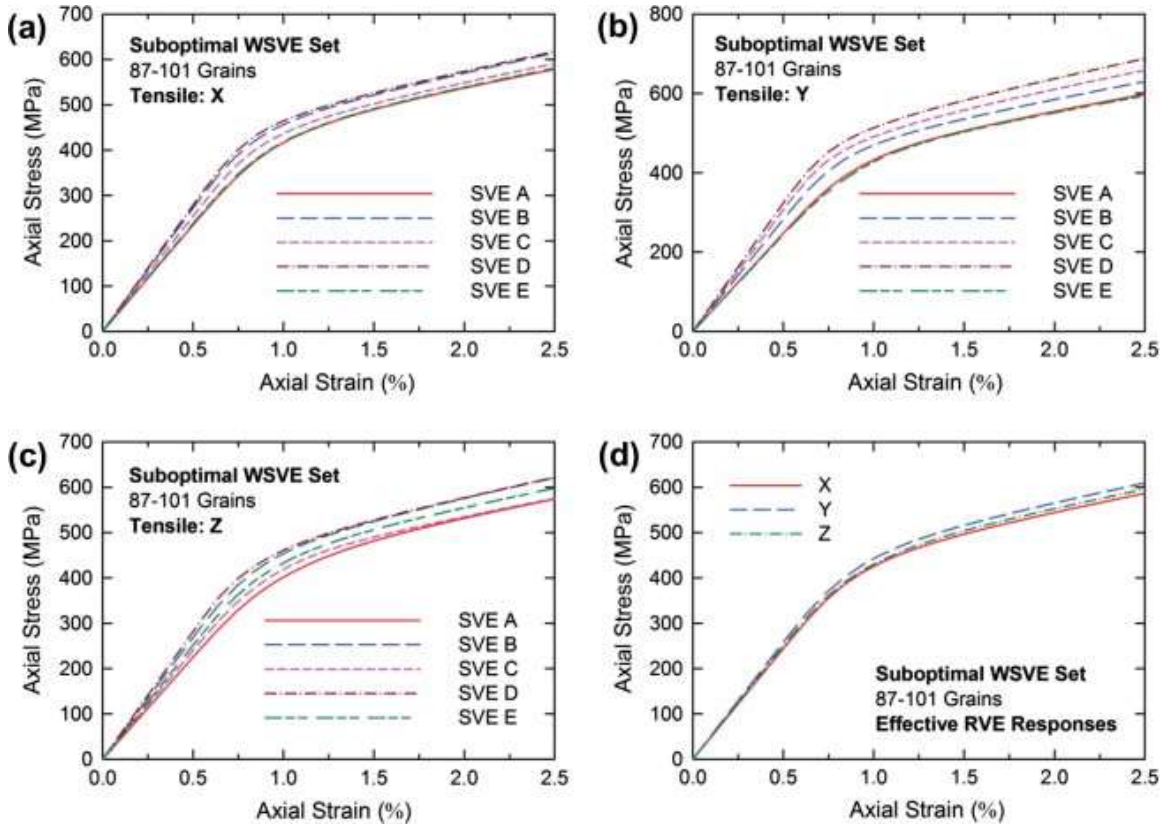


Figure 21: Global uniaxial stress–strain responses for the suboptimal WSVE set: (a) individual SVEs in the global X-direction; (b) individual SVEs in the global Y-direction; (c) individual SVEs in the global Z-direction; and (d) effective RVE response in the three global directions obtained by weighting each SVE response in respective directions according to Equation 22 with the appropriate weights given in Table 4.

The uniaxial responses for each SVE show a considerable amount of variance for all loading directions. However, when combined by weighted sum (using Equation 22) the converge to a narrow band.

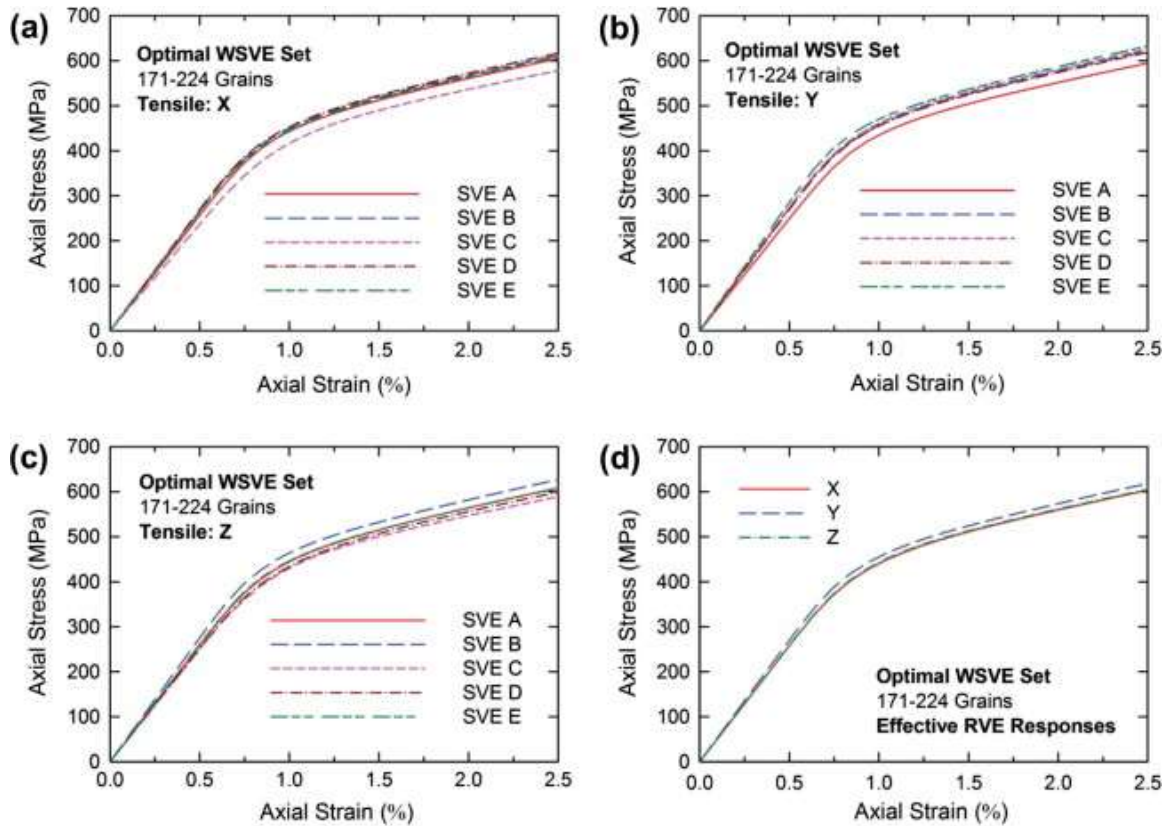


Figure 22: Global uniaxial stress–strain responses for the optimal WSVE set: (a) individual SVEs in the global X-direction; (b) individual SVEs in the global Y-direction; (c) individual SVEs in the global Z-direction; and (d) effective RVE response in the three global directions obtained by weighting each SVE response in respective directions according to Equation 22 with the appropriate weights given in Table 4.

In Figure 22, the same data is shown for the optimal WSVE set. However, in this case the individual SVE responses are much less scattered. In each loading directions there is a single SVE that diverges more from the others. For tension in the X-direction, it is the third SVE (C), and the first (A) and second (B) SVEs, respectively, for the Y- and Z-directions. Again, the effective RVE responses in the three directions converge into a tighter band, as shown in Figure 22(d).

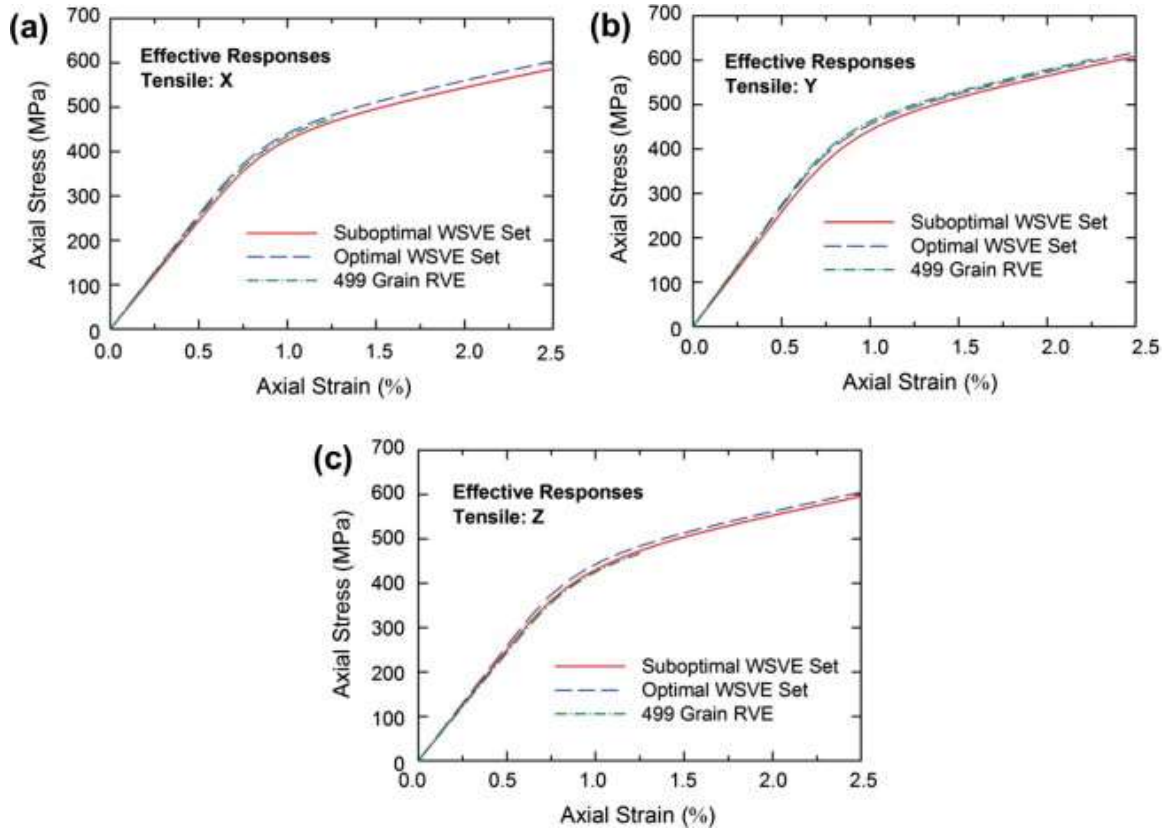


Figure 23: Effective uniaxial stress–strain responses for the suboptimal WSVE set, optimal WSVE set and the 499-grain RVE under tension in: (a) X-direction, (b) Y-direction and (c) Z-direction.

In Figure 23, we have our first comparison of effective uniaxial stress–strain responses for the suboptimal WSVE set, the optimal WSVE set, and the traditional 499 grain RVE. Due to computational cost considerations, the responses for the 499-grain RVE are obtained up to an applied strain of 1.125% except for the tensile loading in the Y-direction where an applied strain of 2.25% is achieved. The comparisons show that for each loading case, the effective responses match quite well and the differences are minimal.

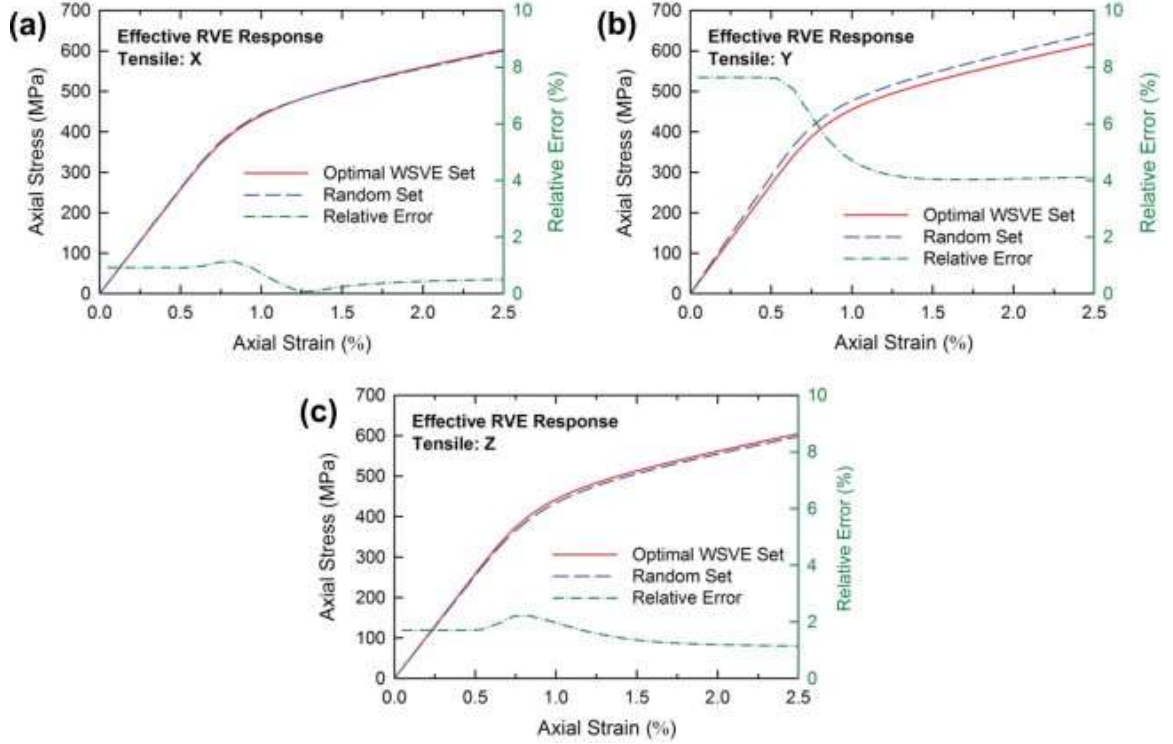


Figure 24: Comparison of the effective uniaxial stress–strain behaviors as predicted for the optimal WSVE and the random set with equal weights (0.25) in (a) X-direction, (b) Y-direction and (c) Z-direction. The error (defined in the text) is in terms of relative absolute error with respect to the WSVE response.

In Figure 24, we have a direct comparison between the random selection WSVE set (A – D of sub-optimal WSVE with equal weights, 0.25) and the optimal WSVE set in terms of effective uniaxial stress–strain behaviors. We compute the relative error using;

$$relative\ error = \frac{|\sigma_{owsve} - \sigma_{rsve}|}{\sigma_{owsve}} \times 100 \quad (16)$$

where σ_{owsve} and σ_{rsve} are the optimal WSVE set and random SVE set axial stresses, respectively. The effective responses are comparable with the relative error not exceeding 1%, 8% and 2% for tensile loading in the X-, Y- and Z-directions. Note that the largest

error occurs in the elastic regime for each loading case, which is related to the prediction of stiffness values by each model.

All of our comparisons so far have been looking at effective global responses, now let us turn to errors in local response. To perform these comparisons, we chose to use histograms of cumulative shear strain γ , accumulated on all slip systems up to a specific time at a material point. For each FE, we obtain the cumulative shear strain at the centroid through interpolation. To minimize the effect of boundary conditions we discard values from FE's that are within 10% of the overall distance to the surface in the global direction. This criteria is based on the calculated boundary-to-boundary variations in mechanical responses of polycrystalline RVEs for four different types of boundary conditions in [83]. In Figure 25, we show the histograms of cumulative shear strain for each individual SVE in the optimal WSVE set and the effective RVE response (when combined in a weighted sum). Following this, we compare the cumulative shear strain histograms for the suboptimal WSVE set, optimal WSVE set, and traditional 499 grain RVE in Figure 26. We see good overall agreement between these local scale responses.

In our final comparison, shown in , we look at the cumulative shear strain data for the optimal WSVE set and random set. The histograms in Figure 27 show the responses for both. In addition, we show the relative fractional error between the optimized and random WSVE sets. The error is defined similarly to Equation 16;

$$relative\ error = \frac{|VC_{owsve} - VC_{rsve}|}{VC_{owsve}} \times 100 \quad (17)$$

where VC_{owsve} and VC_{rsve} are the volumetric counts for the optimal WSVE and random WSVE sets respectively. The error in local responses are higher than their global counterparts with the minimum to maximum fraction of error values of 0.08–0.30, 0.03–0.35 and 0.0–0.21 for tensile loadings in the X-, Y- and Z-directions, respectively.

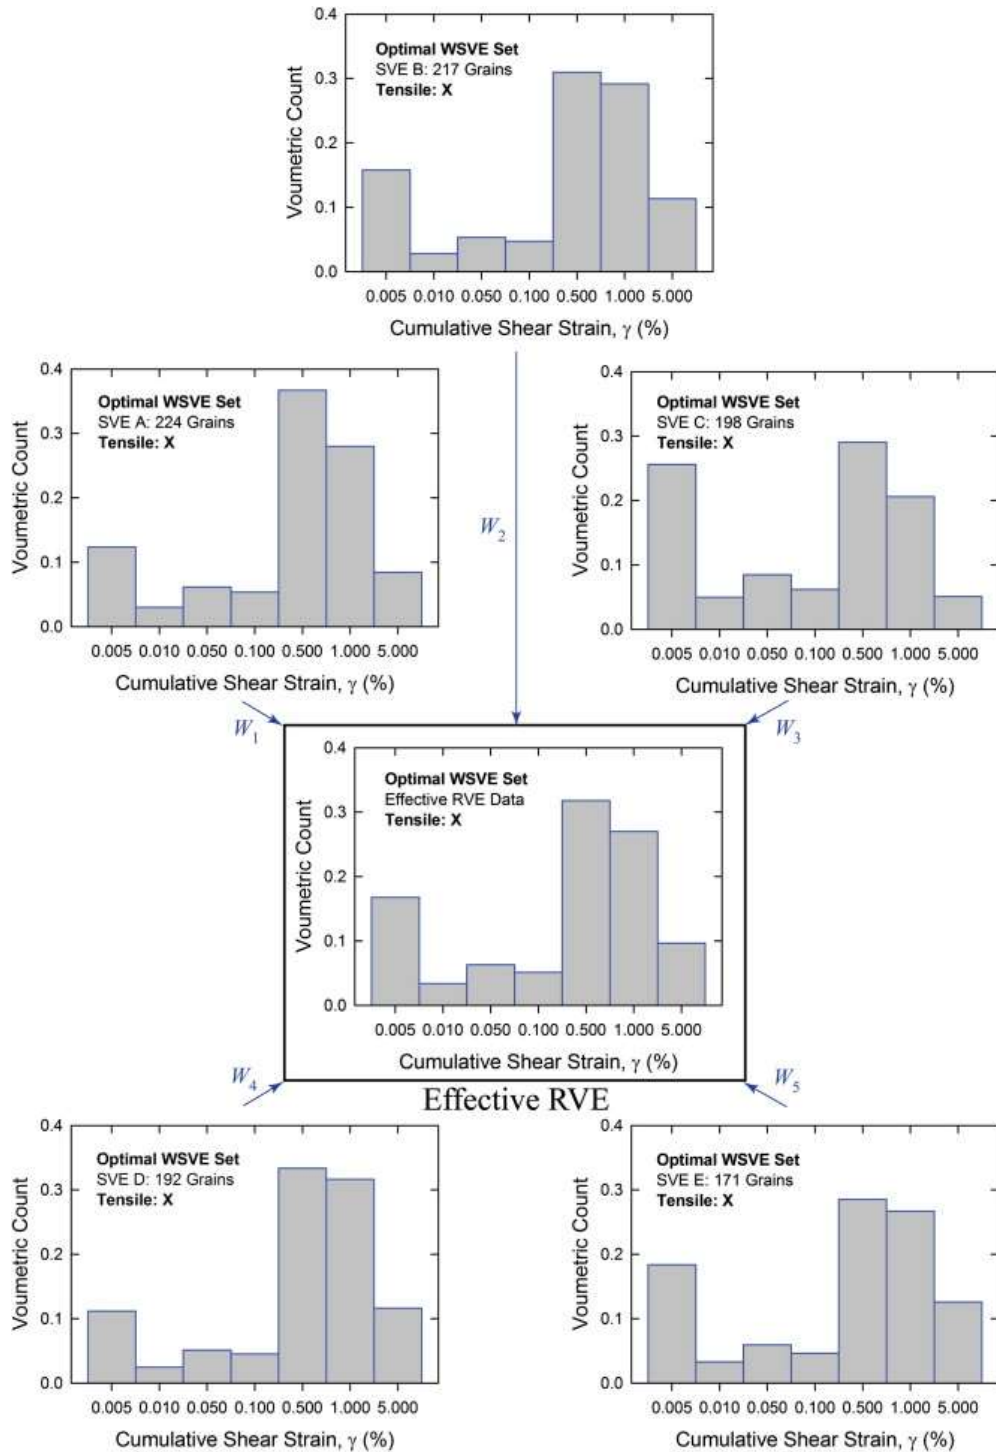


Figure 25: Frequency plots or histograms by volume of the cumulative shear strain γ for the SVEs of the representative optimal set and their weighted summation, the effective RVE. The 227,700 data points for each SVE are obtained at an applied strain of 1.15% in the X-direction soon after the yield point in the global stress–strain curves (Figure 24).

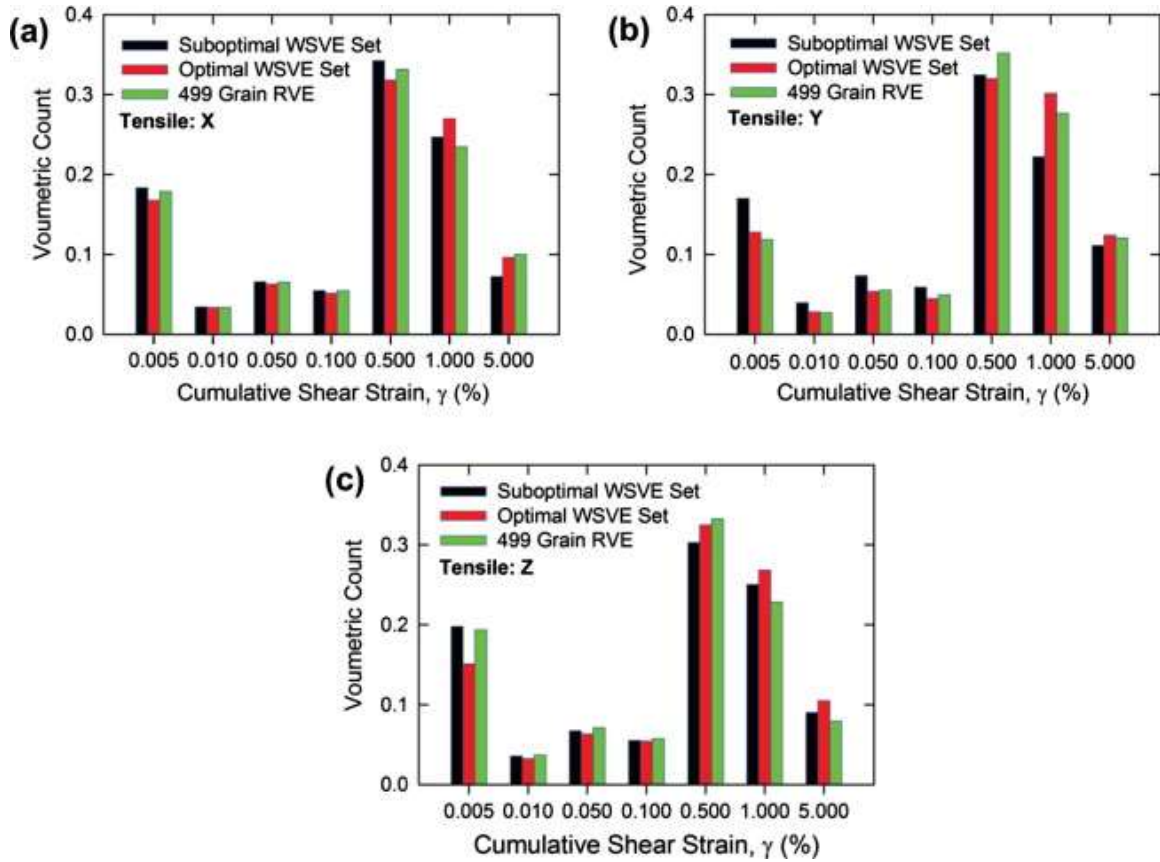


Figure 26: Comparison of effective histograms by volume of the cumulative shear strain γ for the representative suboptimal and optimal WSVE sets, and the RVE under tension in: (a) X-direction, (b) Y-direction and (c) Z-direction. The data is obtained at an applied strain of 1.15% soon after the yield point in the global stress–strain curves (Figure 21, Figure 22, and Figure 23).

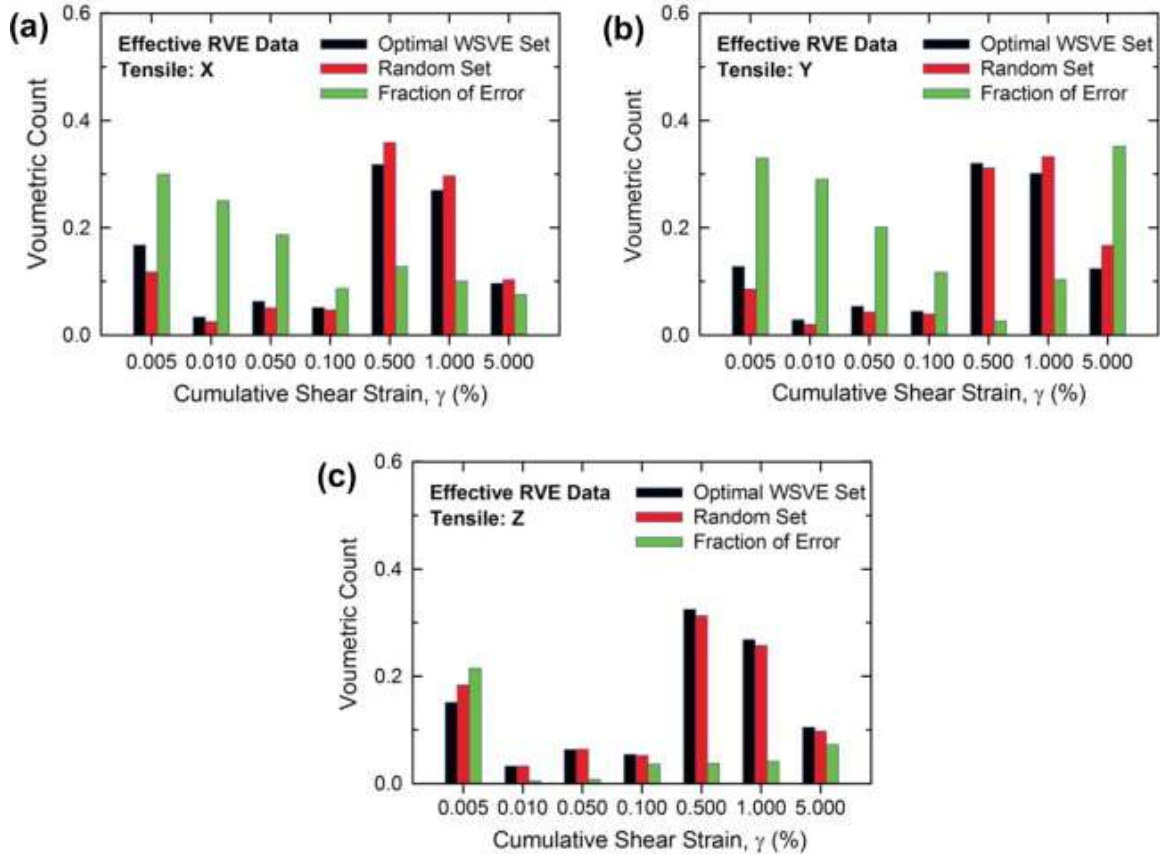


Figure 27: Comparison of the cumulative shear strain data for the optimal WSVE set and random set under loading in (a) X-direction, (b) Y-direction and (c) Z-direction. The fraction of error (defined in the text) indicates the absolute error in terms of fraction of the optimal WSVE set data. The data is obtained at an applied strain of 1.15% soon after the yield point in the global stress–strain curves.

We have shown that the suboptimal WSVE set and optimal WSVE capture similar global and local responses to the more traditional RVE. It is important to note that in the suboptimal WSVE set, each SVE contain approximately 100 grains while in the optimal WSVE there are approximately 200 grains each. Despite this, the local and global response is captured to a similar accuracy. This shows that the size and number of elements in the WSVE set are not paramount, provided all microstructural features pertinent to the material behavior under study are accounted for in a weighted sense

within the collection of selected SVEs. This is the inherent advantage of building our RVE based on microstructure statistics and not specific properties of interest. We can incrementally improve WSVE set by adding structurally representative SVEs. This is much more computationally efficient than building a traditional RVE by running simulations with larger and larger volumes until the convergence of our simulated property irrespective of boundary conditions. More to this point, the average volume of the sub-optimal SVE is eight times smaller than the traditional RVE. This means that the degrees of freedom in our computational models for the suboptimal WSVE set and traditional RVE are 660,723 (an average) and 4,357,742 respectively. So, even with these much smaller models, we are able to capture global and local response because we have chosen them intelligently using microstructure statistics. On average the simulation clock time for the suboptimal and optimal SVEs with 56 SGI Altix processors at 2 GB per processor was 1.5 and 3 days, respectively, whereas for the traditional RVE it was 14 days with 156 processors. (Note: the simulation clock time for the RVE is based on an applied strain of 2.25% compared to the 2.5% for the SVEs.) By using structure based WSVE sets we can realize enormous computational and memory cost savings when building an RVE.

3.5 Conclusions

The methods shown here in this chapter have been called representative or statistical volume element methods based on structure. However, the lack of large 3D datasets for which to apply these methods created an impetus to examine the problem from a different perspective. What can be done when the full 3D 2-point statistics are not known a priori. That is, what if we cannot afford to characterize our material using 3D

techniques and we are left with only 2D data. The next chapter explores this problem of statistical reconstruction of 3D microstructure from partial 2D data.

CHAPTER 4. STATISTICAL RECONSTRUCTION

We conclude our work by arriving at the problem stated in the introduction; can we construct a statistically representative 3D microstructure using only 2-D samples (exemplars) collected on oblique sections. The work presented in this chapter culminated in a single journal publication by Turner et al. [84].

4.1 Solid Texture Synthesis for Microstructure Reconstruction

To attempt an answer to this question we have explored the application of solid texture synthesis approaches [17, 18] from computer graphics in the context of material science. The inputs to the microstructure reconstruction problem of interest here are the 2-D exemplar images that represent the desired appearance of the microstructure when the solid 3-D volume is cut along a given plane. Although we will use three orthogonal sections to illustrate our methodology, we have ensured that the algorithm presented here is not restricted to this specific choice (specific examples will be presented later). Also, in this initial foray, we have restricted our attention to two-phase or gray scale images; however, we strongly believe that the basic algorithm presented here can be easily extended to multiphase composites.

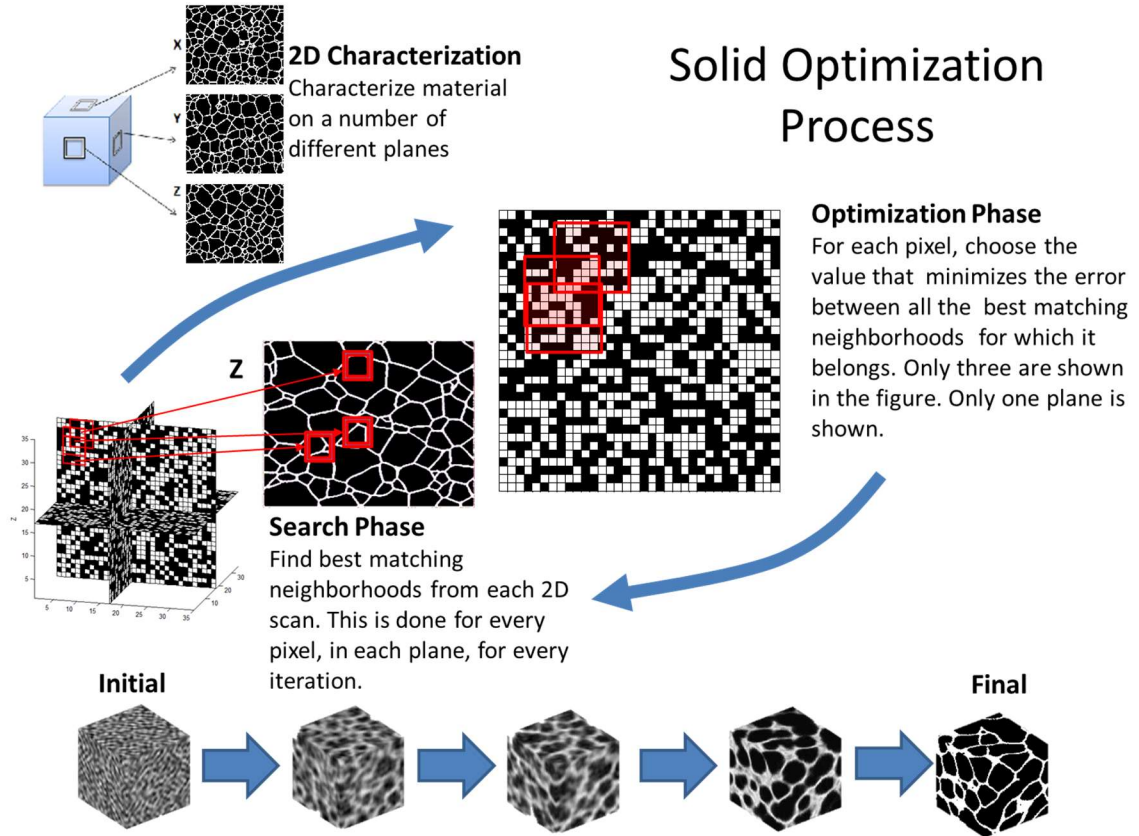


Figure 28: A schematic description of the solid texture reconstruction based on Kopf et al. [17]

The reconstruction of interest will be pursued in a cubical volume that is uniformly tessellated into cuboids (this restriction can be relaxed as needed in future efforts). Let the array $\{m\}$ denote the 3-D solid reconstruction (i.e., the reconstructed microstructure), where its components m_s [66] denote the volume fraction of a selected phase in the spatial bin (can also be interpreted as a pixel or a voxel) enumerated by the index s . Since our interest in this work is restricted to two-phase microstructures, we only need to work with the volume fraction of any one selected phase (the volume fraction of the other phase is identically equal to $(1 - m_s)$). Note also that in most practical

situations we are only interested in the final reconstructions being expressed as eigen microstructures, where m_s is only allowed to take values of either zero or one (i.e., each spatial bin is allowed to be occupied by only one of the phases present in the microstructure). Furthermore, it should be recognized that the index s itself can be expressed as an array of indices to help visualize easily the 3-D microstructure volume being reconstructed (i.e., $s = (s_1, s_2, s_3)$). Let S denote the set of all possible values that can be assigned to s in the reconstruction being pursued (i.e., $s \in S$).

The algorithm begins with an initial guess of the reconstructed microstructure, $\{m\}$. A simple random field can work as an initial guess but like most large scale optimization problems, the closer one starts to a satisfactory local optimum the better. Depending on the exemplar images, this may be difficult or easy. From the initial guess, the algorithm alternates between two steps called the search and the optimization steps, until the reconstruction converges (stops changing significantly). This process is shown schematically in Figure 1. The goal of the search step is to simply identify the best matching 2-D neighborhoods (from the available 2-D exemplars) for every spatial bin in the reconstruction. For this purpose a neighborhood is defined as a small square region of spatial bins of the size $n \times n$, where n may be selected typically in the range between 8 and 32. Each neighborhood array, N , then captures the values of the microstructure (i.e., m) assigned to the selected $n \times n$ spatial cells in a contiguous square area. Since we do not want to mix up neighborhoods from the exemplars collected from the differently oriented sections in the original sample, we shall use the superscript i to refer to each distinct orientation in the input set of 2-D exemplars. In other words, every potential neighborhood identified from all available 2-D exemplars collected on a specified cut

into the sample would be denoted as N^i , and would constitute the set \mathcal{N}^i (i.e., $N^i \in \mathcal{N}^i$). As a specific example, if one were to use as input 2-D exemplars collected from three orthogonal sections into the sample, then \mathcal{N}^1 , \mathcal{N}^2 , and \mathcal{N}^3 would denote the complete sets of all distinct neighborhoods (of the selected size $n \times n$) that can be extracted from the input 2-D exemplars available from each of the three orthogonal sections, respectively. It should be noted that establishing these neighborhood array sets needs to be done only once for each 3-D solid reconstruction. In the search step of the 3-D reconstruction, the current neighborhood array in each selected 2-D orientation centered around every spatial bin s (denoted as N_s^i) is compared with the elements of the corresponding neighborhood array set \mathcal{N}^i to identify the one that is closest (using a simple Euclidean distance measure). Let \tilde{N}_s^i denote the identified closest neighborhood array (note that this has to be an element of \mathcal{N}^i). The search phase dominates the computational complexity of the entire algorithm, which is mainly impacted by the size of \mathcal{N}^i and the choice of n . In the examples presented here, this was accomplished utilizing approximate nearest neighbor search (ANN) available from the library called FLANN [85]. In this protocol, the 2-D neighborhood arrays of the exemplars are flattened and stored in an ensemble of randomized KD-trees. Search is performed in parallel across multiple cores on a single machine.

An important consequence of the search step is that each identification of \tilde{N}_s^i produces a set of n^2 recommendations for the elements of the reconstructed microstructure. This is because each neighborhood consists of n^2 spatial bins. In other words, the selection of \tilde{N}_s^i not only constitutes a specific recommendation (or a vote) for what should be placed in the spatial bin s (i.e., the value of m_s) but also in its

neighboring spatial bins. Seen from a different perspective, each spatial bin s receives n^2 recommendations for m_s from a consideration of one set of distinctly oriented exemplars (since each spatial bin participates in n^2 neighborhoods in the search step for each orientation i). If there are I distinct orientations included in the reconstruction, then each spatial bin receives a total of $n^2 I$ recommendations for m_s at the end of the search step. The schematic in Figure 29 shows this schematically. Put another way, the recommendations (or votes) for each spatial bin in the reconstruction will come from all of the best matching neighborhoods that it is involved in, not simply the ones centered on it.

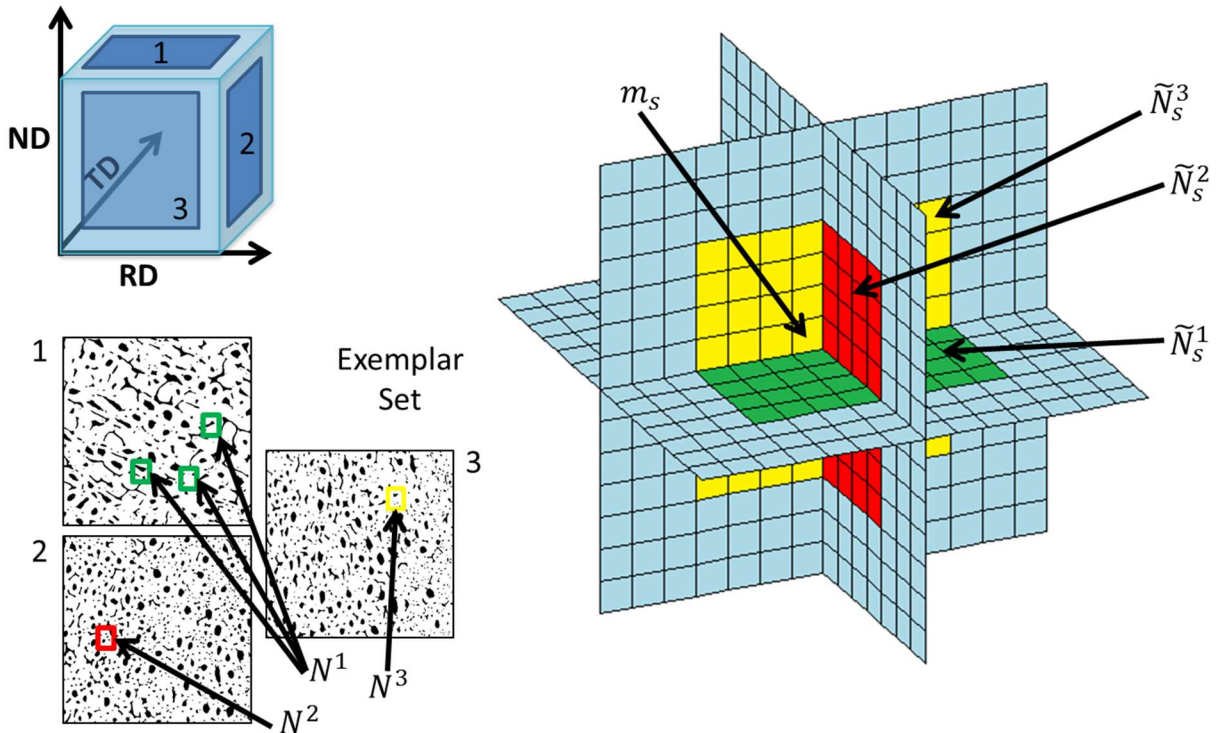


Figure 29: A schematic representation of our notation for reconstruction $\{m\}$ and the set of exemplars. The sample shown in the top left has three orthogonal 2D exemplar scans

taken from it. Each voxel in our reconstruction is then assigned multiple best matching neighborhoods taken from the exemplars, where each assignment is treated as a “vote”. The figure conveys how the microstructure value at \mathbf{m}_s has three votes from the three best matching neighborhoods on the three different planes denoted by \tilde{N}_s^i ; these are selected from the set of possible neighborhoods denoted as $\{N^i\}$ from the exemplar set. Furthermore, all of the yellow, green, and red colored voxels in the reconstruction also receive a vote from the same assignment.

The goal of the optimization step is to consider all the multiple recommendations and arrive at a single update for the current iteration of the reconstruction, i.e., a single update for $\{m\}$. As described earlier, the multiple recommendations for each spatial bin in the reconstruction are a natural consequence of the tremendous coupling between the neighborhoods associated with neighboring spatial bins. Let \tilde{S}_s denote the set of all spatial bins in the neighborhood of s whose closest neighborhood matches impact the final update of the m_s in the current iteration (an example of this is shown in Figure 29 to include all of the yellow, green, and red colored voxels for a selected voxel in the reconstruction). As mentioned earlier, \tilde{S}_s should comprise of n^2I elements. Further, note that $s \in \tilde{S}_s \subset S$. Let $m_{s,\tilde{s}}$ denote the recommendation for the value of m_s from each member of \tilde{S}_s (enumerated by index \tilde{s}). Note that the values of $m_{s,\tilde{s}}$ come directly from the closest neighborhoods \tilde{N}_s^i identified in the search step (using the neighborhoods extracted from the input 2-D exemplars images). The goal of the optimization step then is to update the value of m_s in such a way that it minimizes the overall discrepancy for all neighborhoods in the reconstruction in relation to their respective closest matching neighborhoods identified in the search step, while accounting for the tremendous coupling that exists between the neighborhoods as described earlier. For this purpose, we can define an overall measure of error as

$$e = \sum_{s \in \mathcal{S}} \sum_{\tilde{s} \in \tilde{\mathcal{S}}_s} \|m_s - m_{s,\tilde{s}}\| \quad (18)$$

where $\| \cdot \|$ represents the L2-norm. The minimization of this error is solved using an iteratively reweighted least squares (IRLS) method for which a simple closed form expression is derived as

$$m_s = \frac{\sum_{\tilde{s} \in \tilde{\mathcal{S}}_s} w_{s,\tilde{s}} m_{s,\tilde{s}}}{\sum_{\tilde{s} \in \tilde{\mathcal{S}}_s} w_{s,\tilde{s}}} \quad (19)$$

where $w_{s,\tilde{s}}$ denotes a value (i.e., weight) assigned to control the influence of a neighborhood to the solution. As a specific example, one might decrease the values of $w_{s,\tilde{s}}$ with increasing values of the physical distance between spatial bins s and \tilde{s} , if one desires to reduce the influence of the neighborhood in the reconstruction.

Equation 19 essentially embodies a weighted recommendation for the update at the end of each iteration of the reconstruction. Consequently, a better use of the weights might be to use them to drive the reconstruction towards some desired global requirements for the overall statistics in the reconstruction. Since $w_{s,\tilde{s}}$ can be varied for each choice of s and \tilde{s} , it can be readily interpreted as a weighting factor for the specific 2-point statistic corresponding to $t = (s - \tilde{s})$. By letting the weights change dynamically from one iteration to another, one can devise clever ways to drive the overall reconstruction to capture closely any selected subset of 2-point statistics of interest in the reconstruction. However, if the weights are all set equal to each other, it essentially implies equal consideration of all the 2-point statistics defined by $t = (s - \tilde{s})$ for the

reconstruction. Finally, rather than updating the value at m_s directly using Equation 19, we find the nearest value within the set of votes ($m_{s,\tilde{s}}$ for all $\tilde{s} \in \tilde{\mathcal{S}}_s$) and copy this value directly. This “discrete optimization” as described by [18], is recommended to reduce blending of values and ensures that only values actually present in the exemplars will be present in the reconstruction.

The two step algorithm described above is implemented in the successive hierarchy of a Gaussian pyramid. That is, the exemplars are down-sampled (i.e., coarser spatial resolution) a number of times (two times in practice) to build a hierarchy of images. The reconstruction begins at the coarsest spatial resolution. When it converges at this level, it is then up-sampled and used as the starting guess for the algorithm at next level. This process is repeated until we reach the highest level and the original resolution of the exemplars. This strategy dramatically reduces the number of iterations needed at the higher spatial resolutions, where nearest neighbor search is very expensive. Additionally, it indirectly allows capture of longer range features or spatial statistics in the iterations at the coarser resolutions.

We implemented our algorithm in a hybrid of MATLAB and C++ codes. Most of the time critical portions of the code including the crucial search and optimization phases are implemented using multi-core parallelism with highly optimized C++ for performance needs. This allowed us to use a higher number of larger exemplar images than any prior reported work in the literature.

As mentioned earlier, reweighting schemes where the weights are adjusted dynamically from one iteration to another are an important component of the

reconstruction algorithm presented here. Kopf et al. [17] have presented reweighting schemes that drive the reconstruction to capture closely the overall ensemble averaged color histograms (gathered from all of the input 2-D exemplars). They accomplish this by adjusting the weights at each iteration such that the specific N^i helping drive the solution to the desired global statistics are up-weighted. This technique has the tendency to cause the algorithm to sample more uniformly from the available exemplars (note that there was no criterion for uniform sampling of the exemplars in the basic two-step algorithm described earlier). In fact, in more recent work [18], reweighting schemes have been developed that explicitly select the weights based on how often a particular pixel from the exemplar has been used in the reconstruction. This strategy encourages a more uniform utilization of the input exemplar images within the reconstruction.

In the present work, we implemented a reweighting scheme based on the exemplar position histogram [18], $H(p)$, that describes the number of times each pixel p from the exemplar is copied and used in the current reconstruction normalized by the number of voxels in the reconstruction. In other words, each time Equation 19 is used for a spatial bin and we select the closest value to copy from our recommendation set to the reconstruction, we keep a count in our histogram $H(p)$. This information is then used to adjust the weights for the next iteration using

$$w_{s,\bar{s}} = \frac{w_{s,\bar{s}}}{1 + \max(0, H(L(m_{s,\bar{s}})) - \emptyset)}, \quad (20)$$

where the function $L(m_{s,\bar{s}})$ gives the pixel location within the exemplar from where the value $m_{s,\bar{s}}$ was copied. The constant \emptyset is simply the number of times that a pixel from the

exemplar is expected to appear in the reconstruction if the image was sampled uniformly, once again normalized by the size of reconstruction. By defining the weights using this equation we will down-weight pixels from the exemplars that have been overused. While this helps alleviate some under-sampling of the exemplars it does not eliminate it [18]. In fact, using a weighting scheme only during the optimization phase of the algorithm can lead to the algorithm not converging to a reasonable solution for certain microstructures. This is particularly a problem in those microstructures where one of the two phases has a low volume fraction (below 15%). In these cases, the algorithm will quickly converge to a local optimum in which the entire reconstruction is the high volume fraction phase. At this point the algorithm will not be able to down-weight the pixels from the exemplars enough because they will all be originating from the same regions of the exemplar. This is because the search phase of the algorithm does not take these weights into account and simply finds the neighborhoods in the exemplar that best match the now uniform reconstruction.

To counteract this we adapt the recommendations presented in [18] and utilize an additional weighting scheme in the search phase. The authors of Ref. [18] redefine the Euclidean distance metric utilized for calculating neighborhood similarity to be scaled by a weight defined based on the neighborhood histogram. This weight is defined in an analogous way to the position histogram weight discussed above. The goal is to down-weight neighborhoods that have been overused in the reconstruction. We adopted a slightly different approach in the present work to improve performance. Instead of redefining the neighborhoods similarity metric, we simply added the normalized neighborhood histogram as an additional component to the high dimensional vector

representing the neighborhoods pixel values. The neighborhood histogram then becomes an additive factor to the neighborhood similarity instead of a multiplicative factor. This choice was made because it allowed efficient implementation of the neighborhood search using approximate nearest neighbor search. Even though we need to rebuild the KD-trees each time the neighborhood histograms change, this extra cost in time is dwarfed because of the large number of query points (all neighborhoods in the 3-D reconstruction) when compared with the number of points in our tree (all neighborhoods in our 2-D exemplars).

Like most non-convex optimization problems, reconstructing a 3-D solid image from only 2-D exemplars suffers from the problem of multiple solutions or local optima. The success of these algorithms in navigating this immense search space is highly dependent on a number of factors. Certain exemplar inputs will converge to reasonable solutions while others will not. There is no known general solution to these problems and thus heuristics must be derived. A general guiding principal is that we should try to constrain the algorithm to a search space of reasonable solutions as much as possible so that it does not converge to poor but locally optimal solution. The reweighting schemes discussed above are examples of general constraints that can be applied to any reconstruction from exemplars. However, sometimes it is possible to constrain a solution even further for particular types of structures and achieve better results. For example, imagine a granular microstructure with two phases, one being grain boundary and the other being grain interior. In a typical polycrystalline structure there will not be many grains which percolate the entire sample from one side to the other. We can enforce such a constraint on our reconstructions if it is not being met by the simple reweighting

approaches discussed earlier. While not implemented in this work, we suggest this idea for future studies.

4.2 Results and Discussion

In an effort to test the performance of our algorithm, we selected a class of large 3-D two-phase microstructures. These were assembled from multiple sources. Some were digitally generated using various algorithms [86, 87], while others were previously collected using X-ray microCT techniques on natural materials (ivory and macadamia nutshells) [88]. Each of the microstructures comprised of 500^3 or more voxels. The exact details of the protocols used in acquiring these datasets are not very relevant to the present study. Our main interest is in using these datasets to test our algorithm by starting with a limited set of 2-D images extracted from the 3-D microstructures, conducting the reconstructions using the algorithm described above, and then comparing the reconstructions to the original 3-D microstructure datasets (used to represent the ground truth).

Final reconstruction volume size was selected as 150^3 voxels for all the reconstructions performed in this work. We needed to keep the reconstruction size to be small enough to allow the computations to be performed within available computational resources, but be large enough to ensure that the reconstruction captures a sufficient number of the features seen in the microstructure. Since the average feature size in the microstructures used in this study was around 15 pixels, the selected size of 150^3 voxels was deemed a good compromise for this present study.

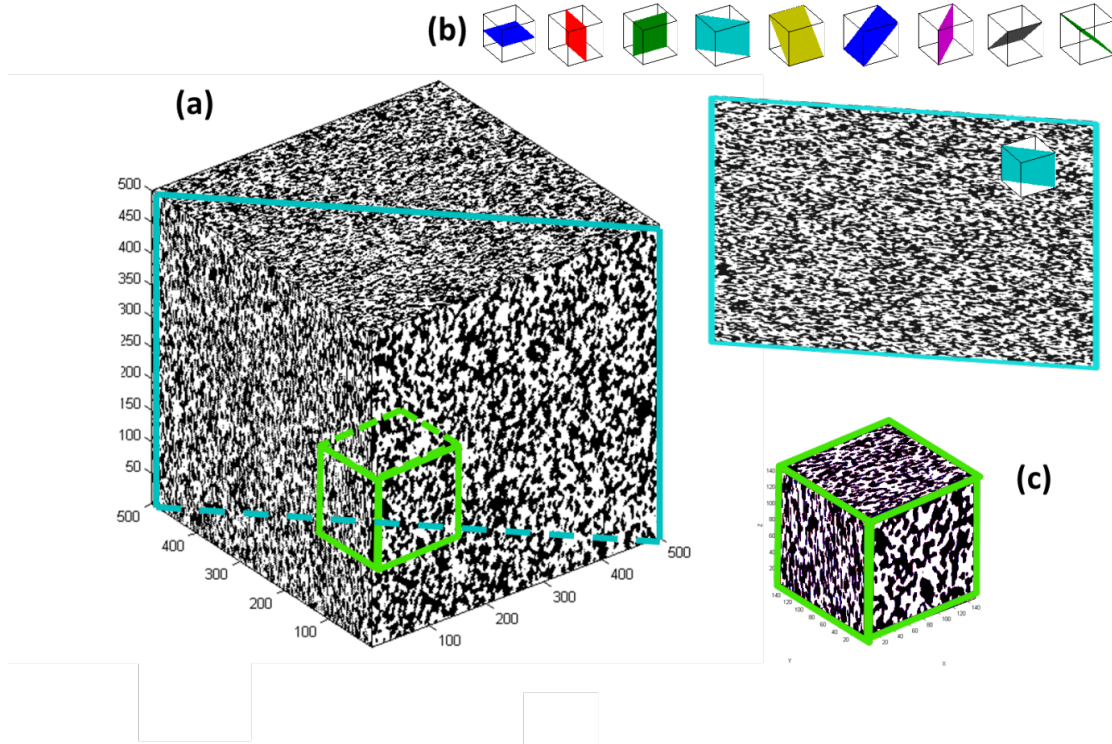


Figure 30: An example 3-D microstructure used in the validation of the algorithm presented in this study. (a) shows the 3-D microstructure volume, (b) shows orientations of the nine slices used in some of the reconstructions as well as a single example slice, and (c) shows an example sub-volume of the same size as our reconstruction.

The first step is to extract the 2-D exemplars that will serve as the input to the reconstruction algorithm. For this purpose, the 3-D microstructures were “digitally” sectioned along nine oblique planes as shown in Figure 30. The nine planes selected are defined by Miller indices (100) , (010) , (001) , (110) , (011) , (101) , $(\bar{1}01)$, $(\bar{1}10)$, and $(0\bar{1}1)$. However, we carried out the reconstructions twice for each microstructure studied - once with only three orthogonal exemplars from (100) , (010) , and (001) sections, and the second time with all nine oblique sections. The goal of doing the reconstructions twice in this manner was to explore the effect of the number of oblique sections on the accuracy of the reconstructions. Neighborhood sizes of 11×11 pixels

were used (unless otherwise noted) at all stages (i.e., resolutions) of the reconstruction. Full reconstruction took around four hours for cases with nine oblique sections, but only about 1.5 hours for the cases with three oblique sections. All reconstructions were executed using a single Dell C8220 compute node with two Intel E5 8-core processors and 32GB of memory.

Our results show that the two-step algorithm presented in this paper using position and neighborhood histogram re-weighting is a viable approach for reconstruction of a variety of microstructures. A comparison of the reconstructions against the original microstructures is shown in Figure 31 for five selected example microstructures. Several others were also performed as a part of this study. The five examples shown in Figure 31 were specifically selected to illustrate the capabilities as well as the shortcomings of the algorithm presented in this work.

The comparison of the reconstructions with the original microstructure shown in Figure 31 is highly qualitative. In order to conduct a quantitative comparison, we need to compute and compare salient statistical measures of the original and reconstructed microstructures. As described in chapter 2 and 4, the most comprehensive and hierarchical measures of a microstructure are provided by the framework of n-point spatial correlations [30, 37, 52-55]. We will limit our attention in this work to the non-redundant 2-point spatial correlations for the present study computed [37, 54] and denoted as f_t .




Figure 31: Comparisons of reconstructions with the original microstructures for several selected examples. The ground truth is shown as a 150^3 voxel subsample taken randomly from a larger 500^3 voxel sample. The reconstructions are of the same size as the 150^3 voxel subsamples from the original microstructures.

Visualization of f_t presents a significant challenge as it represents a 3-D array of values. Figure 32 presents a comparison of the autocorrelations computed for the original

microstructure (500^3 voxels) with those of the reconstructions for selected examples from Figure 31. These comparisons are presented in the form of three orthogonal sections in f_t corresponding to the orientations of the three oblique sections used in both the 3-slice reconstructions as well as the 9-slice reconstructions. A careful comparison of these autocorrelation plots indicates that the reconstructions do capture the salient statistical measures to a very large extent. It is important to keep in mind that the correlations shown in Figure 32 for the original microstructure used the full 500^3 voxels, while the correlations for the reconstruction are based on the 150^3 voxels. Therefore, one should not expect a close match between these autocorrelations.

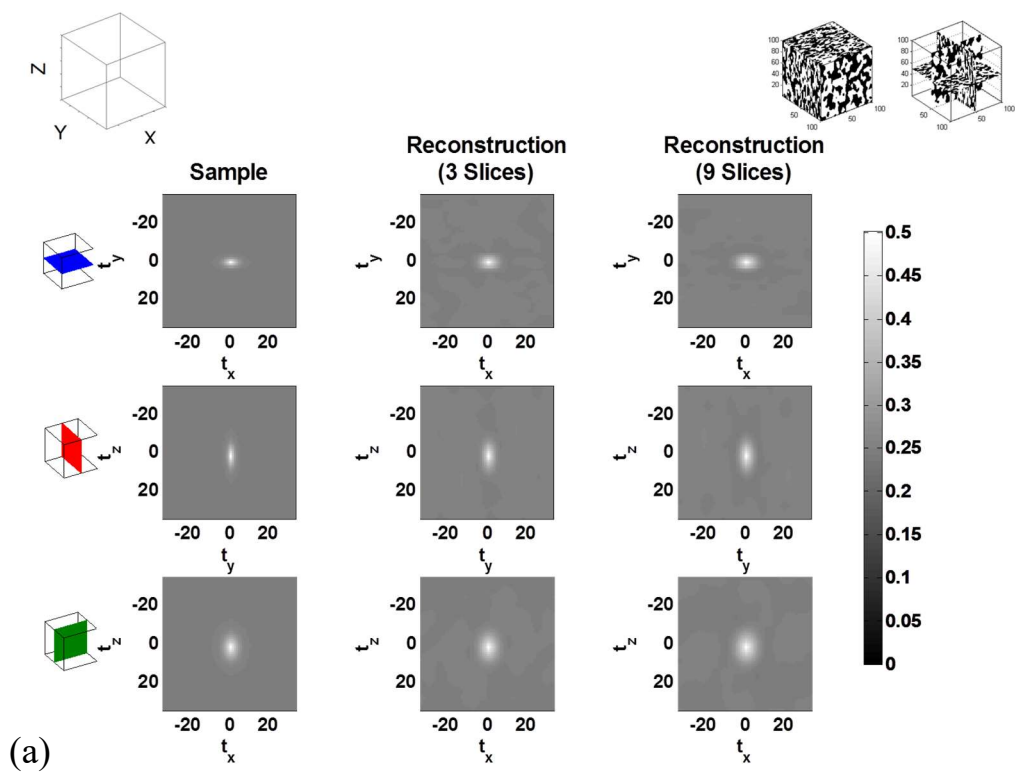
In order to arrive at a more quantitative comparison between the original and the reconstructed microstructures, a large ensemble of 600 subvolumes of size 150^3 voxels were extracted from the original microstructure. Euclidean distance measures were then computed between the autocorrelations of each of the members of the ensemble and that of the entire original sample (i.e., the ensemble average). These distance measures then provide an estimate of the inherent approximation involved in random sampling of the original larger volume by the smaller subvolume size used for the reconstruction. The mean distance measure and the standard deviation in the ensemble are presented in Table 7 for each of the five microstructures shown in Figure 31, and are compared with the distance measure between the autocorrelations of the original and reconstructed microstructures shown in Figure 32.

Table 7: Statistical measures (mean and standard deviation) of distance between the autocorrelations of a 150^3 voxels subvolume extracted from the larger (500^3 voxels) original microstructure and the full original microstructure compared against the distance

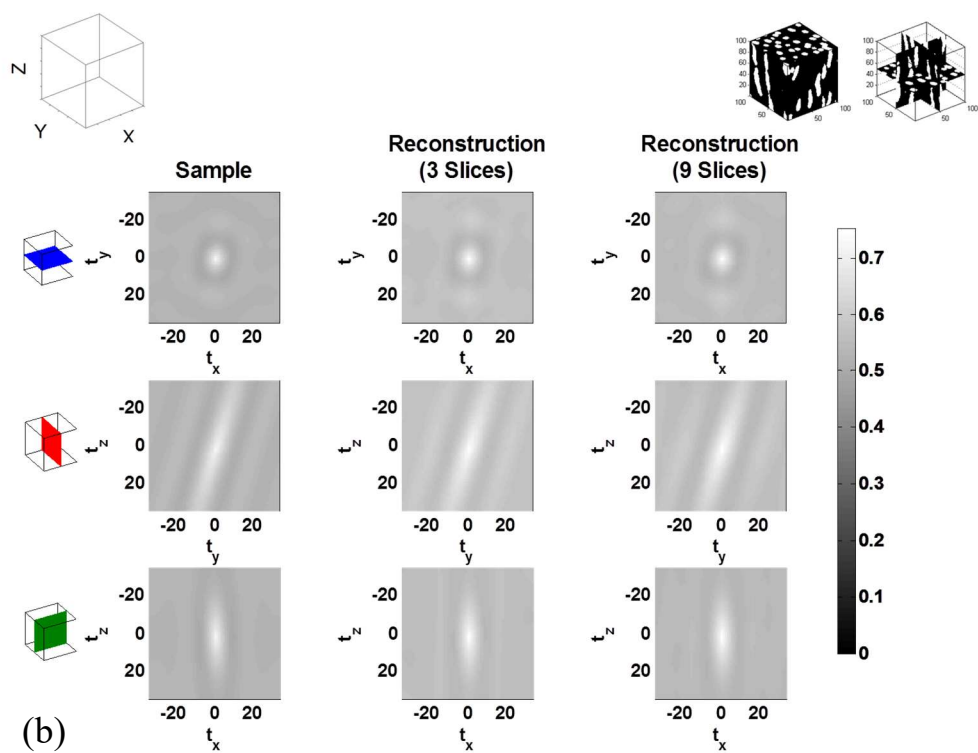
between the autocorrelations of the original microstructure and its reconstruction. Microstructures numbers correspond to the one shown in Figure 31.

Micro-structure	Mean Distance of Subvolumes from the Original	Standard Deviation in Distance of Subvolumes from the Original	Distance Between 3-Slice Reconstruction and the Original	Distance Between 9-Slice Reconstruction and the Original
1	2.124999E-05	1.775622E-05	2.702421E-05	1.389484E-04
2	8.914490E-04	1.195266E-03	1.462192E-03	5.438459E-04
3	2.728381E-05	2.631370E-05	3.936405E-05	4.987727E-05
4	8.186787E-05	5.296932E-05	1.527004E-04	1.871140E-04
5	2.116913E-05	2.170245E-05	1.967602E-03	5.765046E-03

All of the distance measures shown in Table 7 can be interpreted as mean squared errors. Keeping in mind the actual values of f_t are about 0.5 (see Figure 32), these error values confirm that the reconstructions are indeed very good. Columns 2 and 3 provide an estimate of the error and its standard deviation that can be attributed to the fact that reconstructions are significantly smaller in volume compared to the original microstructure. Table 7 indicates that these errors are comparable to the errors observed between the original and the reconstructed microstructures.



(a)



(b)

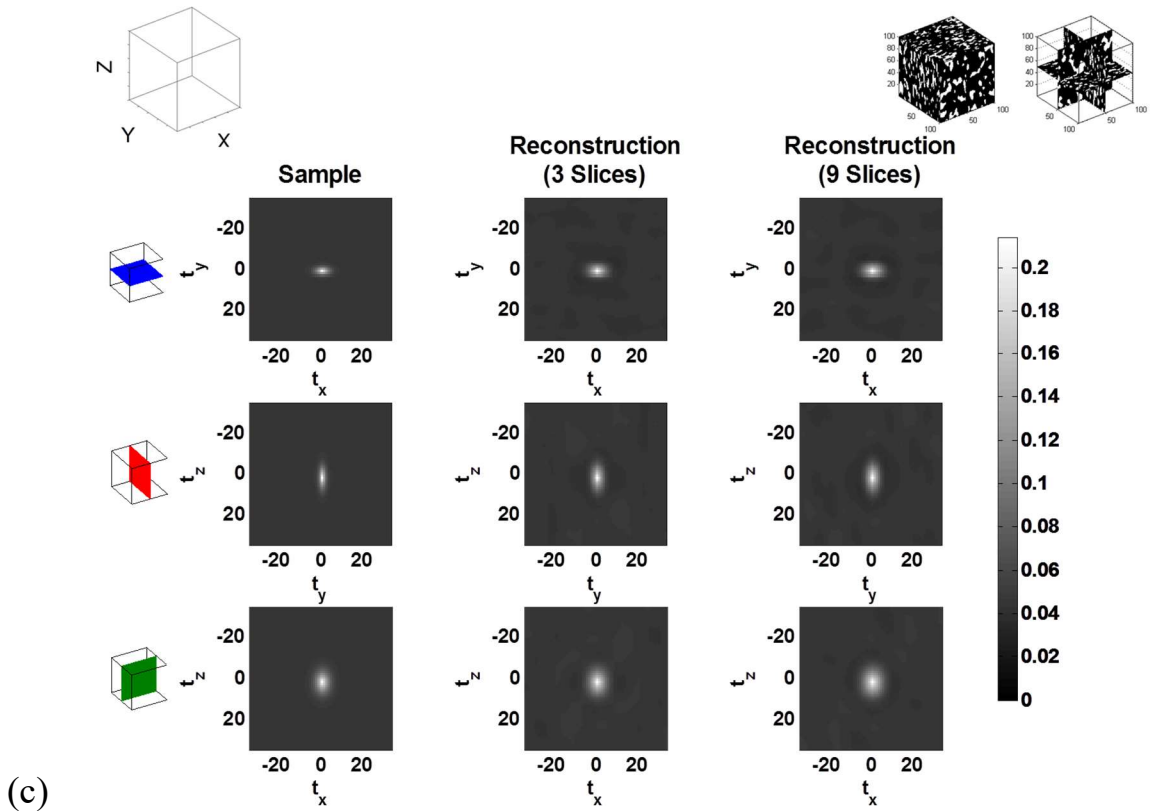


Figure 32: Selected comparisons of 3-D two point statistics between original microstructures and their reconstructions shown in Figure 31. A subset of the 3-D auto-correlations are shown on three orthogonal slices through the origin of the spatial statistics.

Table 7 indicates that the 9-slice reconstruction was significantly better than the 3-slice reconstruction only for Microstructure #2. For the rest of the microstructures, the 9-slice reconstruction was either no better or worse than the 3-slice reconstruction. This is because only Microstructure #2 showed a significant anisotropy (see the autocorrelations shown in Figure 32). In the rest of the microstructures, the anisotropy was essentially confined to short vectors and was adequately captured in the three orthogonal slices. In other words, there isn't much new information on the additional slices included in the 9-

slice reconstruction compared to the information already present in the 3-slice reconstruction for these microstructures (all except Microstructure #2). For some of these cases (e.g., Microstructure #1) the 9-slice reconstruction performed worse because the computations were needlessly penalized by the dramatically increased higher dimensional optimization space. It is therefore important to use the additional orientations (i.e. larger number of distinct oblique sections) only for microstructures that exhibit strong anisotropy that is revealed in these additional orientations.

Table 7 also indicates that the reconstructions for Microstructures #1, #2, and #3 are significantly better than for Microstructures #4 and #5. We believe that the higher errors for the later indicate the limitations of the current algorithm. It is noted that Microstructures #4 and #5 exhibit a certain non-Markovian nature, and further work is needed to constrain solutions during the iterative optimization to achieve better results. In addition, Microstructure #5 needed to be down-sampled by a factor of two before exemplar extraction and reconstruction, which led to smaller input images. The down-sampling was necessary so that we could capture longer range features in the small neighborhoods we used for the present study. This structure was not a synthetic microstructure so it was impossible to generate a higher resolution sample. We believe that this highlights the importance of having high resolution images for reconstructing complex geometries.

The comparisons discussed thus far were focused on the autocorrelations evaluated on the entire reconstructed microstructure. However, our interest is also in comparing information between the original microstructure and the reconstruction at the scale of smaller neighborhoods. Such a comparison would constitute a higher level of evaluation

compared to what is presented in Figure 32 and Table 7. In order to perform such a comparison, we extracted ensembles of small neighborhoods (of specified size) from both the reconstructions as well as the original full sample. The autocorrelations of each of these neighborhoods was then computed. However, this is still a high dimensional representation and makes it difficult to quantitatively compare the distributions within the ensembles and between the ensembles. As a specific illustration of the challenge, we selected the neighborhood sizes as 30^3 voxels and 75^3 voxels. The number of autocorrelations computed for each of these neighborhoods would be 30^3 and 75^3 , respectively. In other words, each neighborhood is now represented by a vector of size 30^3 and 75^3 , respectively, for the two different neighborhood sizes selected for the study.

Following approaches outlined in earlier work [55, 86, 87, 89, 90] we employed principal component analyses (PCA) to obtain objective low-dimensional representations of autocorrelations of the neighborhoods for this study. For each pair of ensembles extracted from the original image and its reconstruction, one PCA was performed. The top three principle components are then plotted as 3-D scatter plots in Figure 33 for the two different neighborhood sizes as well as for 3- and 9-slice reconstructions. In each plot, the red points denote the neighborhoods extracted from the original full sample and the blue points denote the neighborhoods from the reconstruction. For most of the structures examined, there is a large overlap of these point clouds indicating that the neighborhoods extracted from both the original sample and the reconstruction have similar spatial statistics. The differences are somewhat significant for Microstructure #4 and most significant for Microstructure #5, consistent with the observations we made earlier from Figure 32 and Table 7. Also, the differences are larger for the larger

neighborhood size. This essentially indicates that the algorithm needs additional development and tuning to capture the slightly longer range spatial correlations.

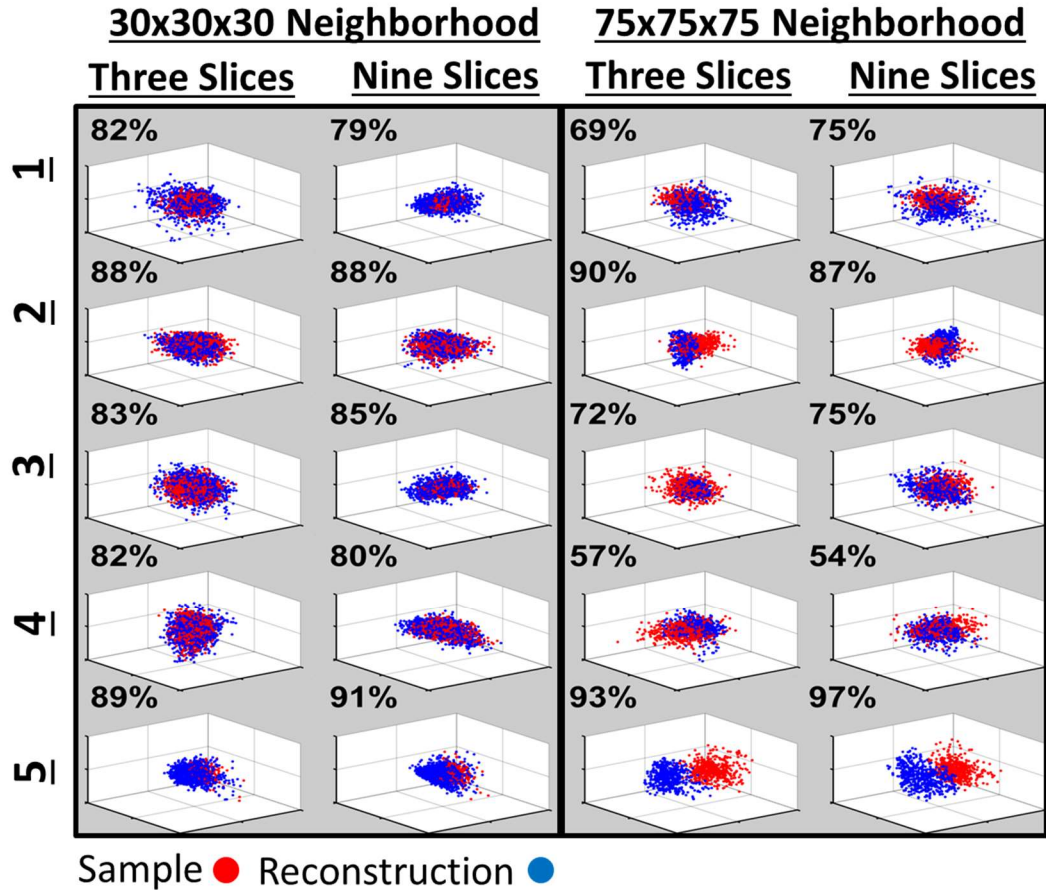


Figure 33: The principle component (PC) embeddings of autocorrelations computed for neighborhoods extracted from the original full microstructure (red points) and its reconstruction (blue points). The percentages next to each plot express the captured overall variance in the ensembles using three PC scores.

Although the examples presented here demonstrate tremendous promise, they also point to much needed further development. The reconstructions are extremely sensitive to the heuristics employed (including the many hyper parameters and the choices made in

the algorithms). These include neighborhood size, the number of Gaussian layers, the input exemplar size, the exemplar phase volume fractions, among several others, and can all have dramatic effects on the convergence. In particular, microstructures with low volume fractions (less than 10%) are found to be particularly challenging for this algorithm, as it has a tendency to converge to solutions of entirely the high volume fraction phase. Although, this can be largely alleviated with histogram weighting as explained in this work, it is by no means a perfect solution.

4.3 Conclusions

We have demonstrated the effectiveness of an algorithm inspired by the field of solid texture synthesis for reconstruction of several material microstructures. This algorithm provides a new approach for material scientists to acquire statistically representative 3-D microstructure samples, when only measurements on 2-D sections are possible or tractable. Moreover, in cases where large 3-D scans are not possible experimentally, researchers may be able to acquire more statistically representative ensembles of samples from 2-D data using the approach described in this work. Although three orthogonal sections were sufficient for reconstructing most of the microstructures explored in this study, it was observed that additional orientations (i.e., new oblique sections) would be necessary in cases where the microstructure shows strong anisotropy in new directions or planes (not contained in the orthogonal sections). Additionally, certain structures, those that do not satisfy a Markovian property, may not be suitable for reconstruction using even large numbers of oblique sections. To our knowledge, this work represents the first implementation of a solid texture synthesis algorithm using position and neighborhood histogram reweighting for the problem of microstructure

reconstruction. In addition, it is the first example of more than three oblique sections being used to synthesize a solid texture or reconstruct a microstructure.

APPENDIX A. CHORD LENGTH DISTRIBUTIONS

While 2-point correlations represent a powerful statistical tool for quantifying microstructure, their adoption has been limited in some cases. This could be because of their difficult to interpret nature or the computational cost\difficulty of calculating them. Instead, simple stereological estimators such as intercept lengths are mostly employed due to their a) easy interpretation, b) straightforward manual calculation using widely available microscopy software tools such as ImageJ [27], and c) strong connection to properties relevant to design (for example, grain shape and size distributions in polycrystalline metals are known to strongly influence their mechanical strength [28-32], while pore size and shape distributions are known to have a dominant effect on the mechanical and transport properties in porous solids [13, 33]). Typically, only mean or average feature sizes are calculated by application of the broadly adopted ASTM E-112 “Standard Test Methods for Determining Average Grain Size” protocols [34], which involve counting intercepts of randomly oriented lines on two-dimensional sections or counting numbers of grains per area. The obvious weakness of such approaches is that they are not applicable to structures that have large structural anisotropies or where the size distribution of features includes distinct populations of several large and small features. The limited number of features counted also means that rare features corresponding to the extreme tails of the size distributions are unlikely to be counted. The manual or semi-automated nature of the analysis can also be problematic in that great care must be taken to ensure that the results are independent of who is performing the analysis, as large variances can result. For example, ASTM E-112 gives

reproducibility at ± 0.5 G which corresponds to $\approx \pm 9\mu\text{m}$ uncertainty in the average diameter for material with $45\mu\text{m}$ average grain diameter.

In this chapter, we demonstrate a computationally efficient scheme to compute anisotropic microstructure chord length distributions, which capture the complete size distribution of microstructure features in two- or three-dimensional material datasets. This approach retains the main benefits of the average intercept techniques in that the results are quick to compute, and simple to interpret and visualize, while providing significantly more detailed information beyond the average feature size. Turner et al. [91] is the culmination of the work presented in this chapter.

Chord length distributions (CLDs) [45, 48-50, 92-96] have been successfully utilized in prior literature as rigorous measures of the size and shape distributions of the important microscale constituents in the material system. A chord is defined as any line segment in the microstructure whose interior points are all occupied by the specific local state of interest (e.g., a pore or a distinct phase) and the end points abut against other local states. In other words, chords cannot be extended in either direction and still remain completely in the local state region of interest. Formally, the CLD, $p^i(\mathbf{z})$, denotes the probability density associated with finding a chord \mathbf{z} (including a specified magnitude and a specified direction) within the local state i in a given microstructure dataset. When the chords are identified by both their length and direction, the resulting distributions are termed as angularly resolved CLDs in this work.

Another important microstructure metric that is closely related to the chord length distribution is the lineal path function (LPF) [30, 51]. Formally, the LPF, $L^i(\mathbf{z})$, denotes

the probability that a line segment \mathbf{z} (including a specified magnitude and a specified direction) thrown randomly into a given microstructure dataset lies entirely within local state i . In this definition, the line segment does not have to be a chord with end points at the interfaces as defined earlier. As before, when the line segments are identified by both their length and direction, the resulting function is termed as the angularly resolved LPF in this work.

Prior work [37] has established a relationship between the CLD and the LPF. This relationship can be expressed as

$$p^i(\mathbf{z}) = \frac{l_c}{V^i} \frac{d^2 L^i(\mathbf{z})}{dz^2}, \quad l_c = \int_0^\infty zp^i(\mathbf{z})dz \quad (21)$$

where z is the magnitude of \mathbf{z} , V^i is the volume fraction of local state i , and l_c is the mean chord length along the direction of \mathbf{z} . Both the CLD and the LPF are commonly estimated by sampling, and accurate estimation requires identifying a very large number of chords in the microstructure or throwing a very large number of line segments into the microstructure. Because of the high computational cost involved in these estimations, most prior applications [30, 45, 48-50, 92-98] have been largely limited to isotropic structures, where the CLD and the LPF were assumed to be independent of the direction of the chord or the line segment, respectively.

It should also be noted that CLDs and LPFs have been utilized successfully in reconstructions of microstructures [99]. More germane to the present discussion, Talukdar et al. [100] utilize angularly resolved CLDs in their microstructure reconstructions. However, they employ only four selected directions in a two-

dimensional (2-D) microstructure dataset in computing the angularly resolved CLDs. This is presumably because of the prohibitively high computational cost involved in the currently used methods for estimating these distributions. In fact, Talukdar et al. [100] utilize a sampling approach for the calculation of the CLDs by translating a line successively to every pixel in the image. This approach has the unique advantage that it does sample all chords present in the given dataset. Singh et al. [51] use a similar approach for computing the lineal path function for high resolution 2D micrographs. They compute angularly resolved lineal path functions in three directions; zero degrees, ninety degrees, and forty-five degrees. However, certain optimizations are needed to make the approach computationally practical for use in large 2-D and 3-D microstructure datasets, while including a larger number of directions.

The main purpose of this paper is to explore an efficient algorithm for computation of angularly resolved CLDs and LPFs. The approach shown in this paper is referred as the Scan Line Method, and is an extension of the simple idea of translating a line successively to every pixel in the dataset (utilized earlier by Talukdar et al. [100] and Singh et al. [51]). We present in this paper several new ideas for optimization of the computations involved in the Scan Line Method for estimating angularly resolved CLDs in large 2-D and 3-D microstructure datasets. Furthermore, we present numerical procedures for extracting LPFs from the computed CLDs in this approach. The protocols developed in this work are demonstrated on selected example 2-D and 3-D microstructure datasets. We show clearly that angularly resolved CLDs and LPFs can be computed efficiently for large 3D microstructures for large numbers of directions.

4.4 Computational Procedures

A line scan approach, such as the one employed by Talukdar et al. [100], might prove to be computationally efficient if implemented in a highly optimized environment. As an illustrative example, consider a 2-D binary image represented as a square matrix I . Furthermore, let us restrict our initial attention to only the chords in this image that have an orientation of 0° , i.e., those that are completely horizontal. We can precisely identify and compute the distribution of lengths of these chords by simply scanning through the rows of the matrix I (assuming these correspond to pixels in the horizontal direction), and keeping track of the starts and ends of chords as we encounter them. For instance, we can scan along each pixel in a row of the image and record when we encounter a phase value of interest. After this, we will continue scanning forward until we encounter a pixel with a value of another phase. At this point, we will have identified a chord of a certain length and it can be recorded for calculation of a distribution. This approach works just as well for vertical chords as we can simply look at the columns of the image. One final note is that we must discard chords that originate or touch the boundary of the image because we cannot be sure these are in fact chords. This is because pixel\voxels at the edge of our image may not be interface pixels\voxels. The criteria of non-edge chords not being counted can be removed when one wishes to calculate lineal path functions from chord length distributions because lineal path functions do not require that line segments span interfaces like chords.

A scanning algorithm for orthogonal lines is extremely efficient both computationally and in terms of memory. It will require on the order of $O(n)$ number of operations, where n is the number of pixels in the image. The memory requirements are even better as it requires only keeping track of a distribution of chord lengths, where the

zero chord length and the longest possible chord lengths define the range. So we have a memory requirement upper bound of $O(n)$ because the computation can be done in place. However, we must be able to adapt this algorithm for chords that are not perfectly aligned along the reference axes of our image. Fortunately, the field of computer graphics has already developed a fast and simple algorithm for determining arbitrarily oriented scan lines through an image. This novel idea is described below.

Bresenham's line drawing algorithm [101] is efficient algorithm for drawing an approximation of a line between two points in an n -dimensional image. Given the start and endpoint of the line in the image, Bresenham's algorithm determines the pixels lying on the line between these two points using only integer addition, subtraction, and bit shifting. Using only integer arithmetic is not only faster, but more accurate because it avoids the common drawbacks of accumulated floating point error. With Bresenham's algorithm, we can determine the pixels that correspond to a scan line through the image of any orientation. Once we have determined the pixels for a given scan line, we can then simply translate this line such that it sweeps through the entire image (see Figure 34).

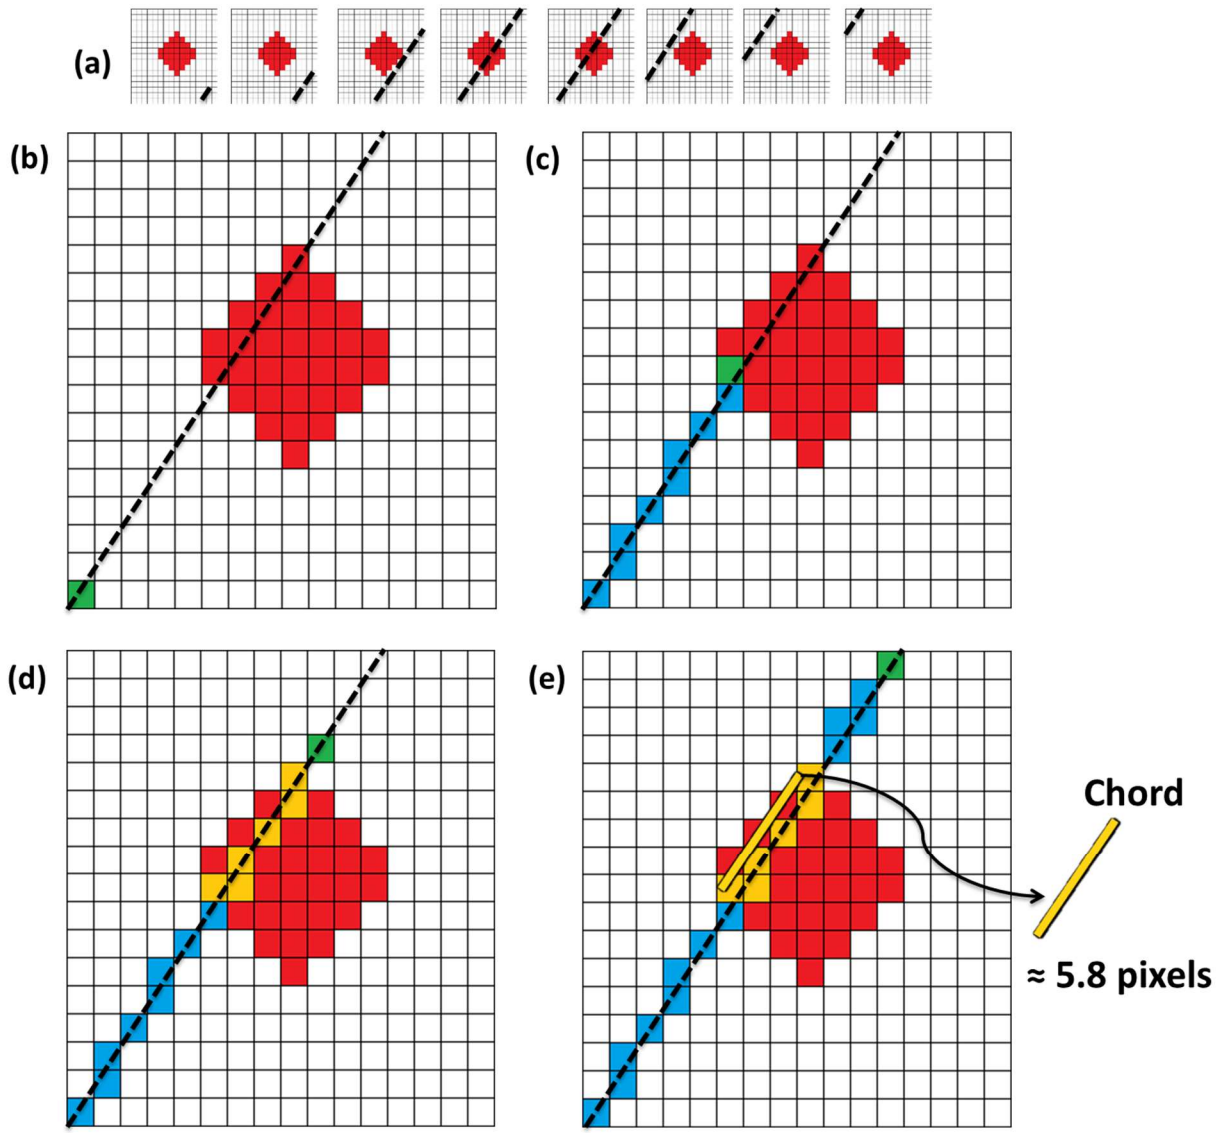


Figure 34: This figure details how chords are identified for a particular orientation. In all images the red phase is the phase of interest. The green coloured pixel marks the active pixel in the algorithm, blue marks pixels that have be processed, and gold marks pixels along a chord. In (a) we show how for a given orientation we try every scanline within the image dimensions. Some are not shown here for the sake of brevity. For each scanline, we begin (b) by starting at one edge of the image and marching along the rasterized approximation of an oriented line until (c) we encounter the phase of interest. We then continue moving along the scanline until (d) we encounter a pixel that is not the phase of interest. At this point we have identified a chord as long as the start pixel and end pixel are not touching the edge of the image. In (e) we continue processing the remaining portion of the scanline.

Using translated Bresenham lines we can scan through an image in any arbitrary direction and identify all chords in the local state regions of interest. This scan converted image will allow us to calculate the distribution of chord lengths in that direction using a process that is almost identical to the case of vertical or horizontal chords described earlier. The main difference is that we will have to account for the conversion between number of pixels along a Bresenham line and the length of the line. The length of a chord can be approximated by calculating the Euclidean distance between the chords start and end pixels.

This scan line approach is roughly equivalent to the rotation of a discretized image. Computationally, the rotation of an image is worse than the scan line method described above since it requires a copy of the image to be made. However, because 2D image rotation is such a common image operation in computer graphics, a tremendous amount of work has been done to develop optimized routines for this purpose. We have thus implemented two algorithms for 2D images, one uses the above described scanline approach and the other using simple image rotation with bilinear interpolation. Bilinear interpolation is probably the most common interpolation method used for image registration to a new coordinate system. Given a point to register it find the four nearest pixels and computes weighted average of their values [102]. Our results for 2-D images are shown and compared for both methods. However, such a simple approach of rotating the image was not effective for 3D images for two principal reasons. First, there are no real efficient implementations of 3D image rotation algorithms known to the authors. Memory requirements of copying a large 3D image could be prohibitive; therefore, an in-place algorithm like the scan line method is more attractive in the authors' opinion.

The same scanline algorithm outlined above for 2-D images can be extended to 3-D images as well. We have developed a parallelized implementation of this algorithm that computes a full angularly resolved CLD. Performance of this algorithm running on an 8 core 2.4 GHz Intel Xeon machine for a 500^3 voxel dataset with azimuthal and polar angular resolution of 5° was approximately 20 minutes. This algorithm may be significantly improved if it was developed for modern graphics processing units (GPUs) because of its massively parallel nature. That is, each computation for a specific scan line orientation is completely independent of the others.

It is relative easy to compute the lineal path function directly from a chord length distribution. Such an algorithm is described by Singh et al. [51] and simply requires counting all the line segments that could possibly fall along each chord. That is, if we have a chord of length 10 pixels\voxels. Then we can fit one line of length 10, two of length 9, three of length 8, and so on. The only thing left after enumerating all such line segments is to normalize by the total number of possible line segments for each length which is a function of the image dimensions and line length. One important note, as mentioned before, is that chords that touch the boundary of the image should be added when constructing lineal path functions from chord length distributions.

4.5 Examples

To demonstrate our algorithms, we have selected a set of example microstructures (both 2-D and 3-D samples).

4.5.1 Two-Phase 2-D Microstructure




Figure 35: A simple image of identical circles of diameter of approximately 32 pixels placed in a uniformly random manner. Volume fraction of the circular phase is 18.2%. The image dimensions are 2048x2048.

As a first demonstration, we consider a simple 2-D digital microstructure (shown in Figure 35) with a uniformly random distribution of circles each with a diameter of approximately 32 pixels. Volume fraction of the circular phase is 18.2 percent. The diameter is approximate because of the nature of representing a circle on a raster. The main advantage of selecting such a simple example microstructure is that we have a closed form analytical solution for both the CLD and LPF. For many simple convex shapes/bodies, closed form expressions for the non-angularly resolved CLD have been derived and studied [46, 50, 103, 104]. For a simple circle, we have the following expression [46] for the non-angularly resolved CLD:

$$p(z) = \frac{z}{4R\sqrt{R^2 - \left(\frac{z}{2}\right)^2}} \quad (22)$$

where R is the radius of the circle, which is 16 pixels for our test image described above. A comparison of the integrated values recovered using our algorithm (rotated 2D images) with the closed form Equation 22 is presented in Figure 36. We can see that the calculation is accurate aside from the noise present due to the discretized nature of the rasterized circle. That is, our digitized image of a circle is only an approximation to a true circle made from a discrete set of Bresenham lines. Therefore, certain chords are overrepresented and others are underrepresented. Nowhere is this more obvious than in the regime of very short chords (around 1 to 3 pixels). Several authors [42, 51] discuss this problem and its causes extensively, and describe a few very simple mitigation strategies. These simple strategies generally involve discarding certain chords by defining a minimum acceptable chord length criterion. The strictest of these criteria is that in order for a chord to be counted it must at least traverse a distance of one voxel unit size in all dimensions. That is, we only accept chords that traverse at least a single resolution distance in both x and y for 2D and x , y , and z for 3D. We have not implemented this criterion in the calculations shown in Figure 36, so that we can illustrate the inherent inaccuracies of approximating chords and shapes on a raster. In Figure 37, we show a comparison between results computed using the scanline method and the image rotation methods described above which show general agreement. We would like to note that the rotation method seemed to result in a slightly less noisy result when compared to the scanline approach. We conjecture that this is a result of the bilinear interpolation and thresholding of the image when it is rotated. Our hypothesis would be that scan line

method captures the discretized nature of the circle at different angles more accurately while the image rotation method has a tendency to average out these differences at each orientation. Future work is needed to verify this hypothesis however.




Figure 36: Comparison of closed form and calculated results for non-angularly resolved chord length distribution for a simple two dimensional circle with an approximate diameter of 32 pixels.




Figure 37: Here we show a comparison between our implementations of chord length distribution calculations using both a scanline and image rotation method. Aside from some small deviations due to differences between Bresenham line discretization and image rotation interpolation the method mostly agree.

If we do not integrate over the angular dependence of the chord length distribution, we arrive at the results shown in Figure 38. This figure depicts the chord length distribution for each angle sampled, where each exhibits inherent noise as shown in Figure 37. However, based on how the circle is rasterized, certain chord orientations will exhibit more noise than others. To alleviate this noise, we have performed a five point/pixel moving window average on each distribution (for each sampled angle). These results show, as we expect, that there is no angular dependence on the chord length distribution for a circle. The computation of the full angularly resolved full distribution for a 500x500 pixel image took less than a second on a 2.00 GHz Pentium 4 PC.

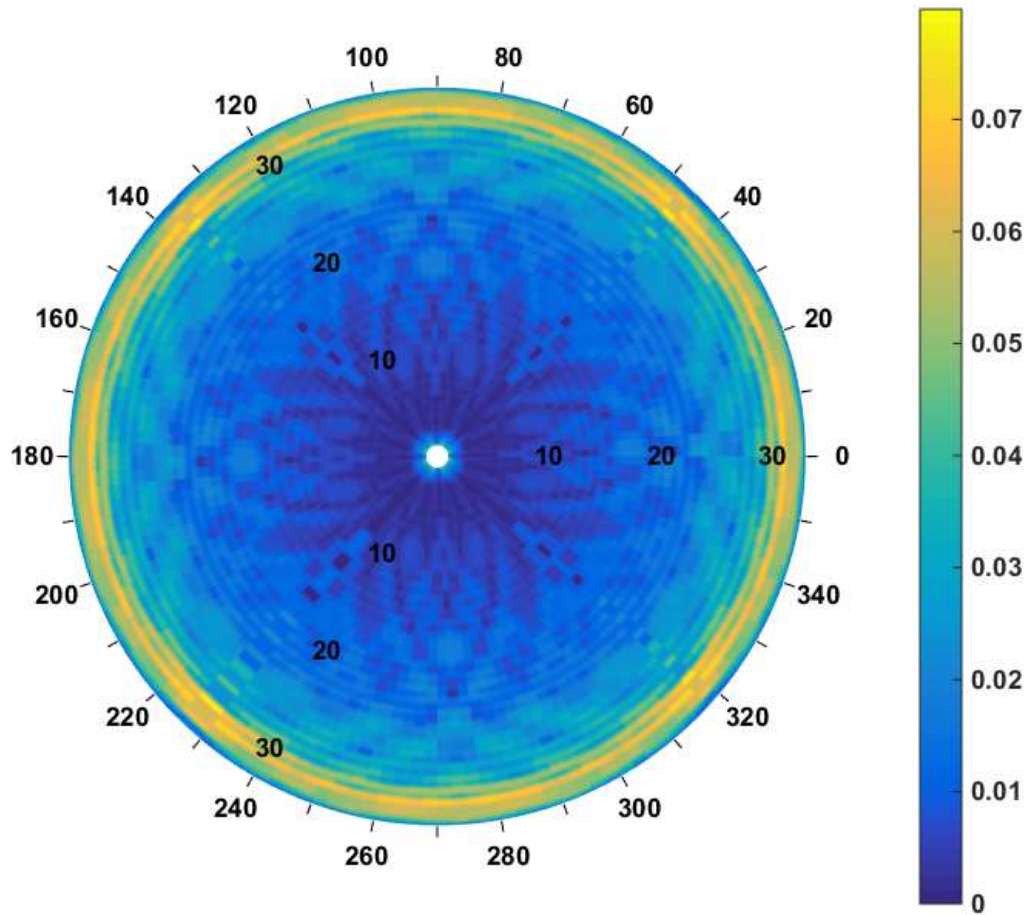


Figure 38: Full angularly resolved chord length distribution for the image in Figure 35. The polar axis on the figure represents the angular orientation of the chord. The radial axis represents the length of the chord. Each individual distribution (chord angle) has been smoothed with a 5 point/pixel moving window average to eliminate noise resulting from discretization of the circle.

To confirm our calculation of LPFs from our chord length distribution, similar to [51], we used the following expression for the LPF of the matrix phase for a set of uniformly distributed equal size circles;

$$L^{matrix}(z) = V^{matrix} e^{-\frac{2(1-V^{matrix})z}{R\pi V^{matrix}}} \quad (23)$$

The V^{matrix} is the volume fraction of the matrix phase, 81.8 percent in our case. The radius R of our circles is 16 pixels as described in Figure 35. The comparison of the calculated and analytical LPF is shown in Figure 39. The results show excellent agreement between the two functions.




Figure 39: Shows a comparison between the analytical and calculated LPFs for the matrix phase in the case of randomly distributed circles of $r=16$ pixels. The calculated LPF derived from our chord length distribution of the matrix phase of the image presented in Figure 35 show excellent agreement.

To highlight the anisotropic nature of the angularly resolved chord length distributions we have also selected another simple synthetic 2D microstructure to analyze with our method. In this case, we have an image of uniform randomly placed ellipses with major axis of 95 pixels and minor axis of 63 pixels. The ellipses are all oriented with

major axis at 90 degrees from the horizontal. This microstructure is shown in Figure 40 and its corresponding angularly resolved chord length distribution is shown in Figure 41. Though analytical solutions for isotropic chord length distributions for ellipses have been derived [47], they are not closed form expressions. We have chosen for sake of brevity to limit our experiment to the results shown in Figure 41.




Figure 40: Image of uniform randomly placed ellipses with major axis of 95 pixels and minor axis of 63 pixels. The ellipses are all oriented with major axis at 90 degrees from the horizontal.




Figure 41: The angularly resolved chord length distribution of the image shown in figure 7. The 95 pixel major axis of the ellipse is highlighted by the peak of the chord length distribution at the 90 degree orientation.

4.5.2 Porous 3-D Microstructures

To illustrate the capabilities of the extended approach described in this chapter for 3-D microstructures, we have selected a digitally simulated 3D porous microstructure. This high resolution sample consists of 500x500x500 cubic voxels. It has been segmented into two distinct phases: pore and matrix. To quantify the shape of the porous structure, we ran our chord length algorithm on the pore phase of this dataset. We used scanlines at azimuthal and polar angular increments of every 5 degrees.

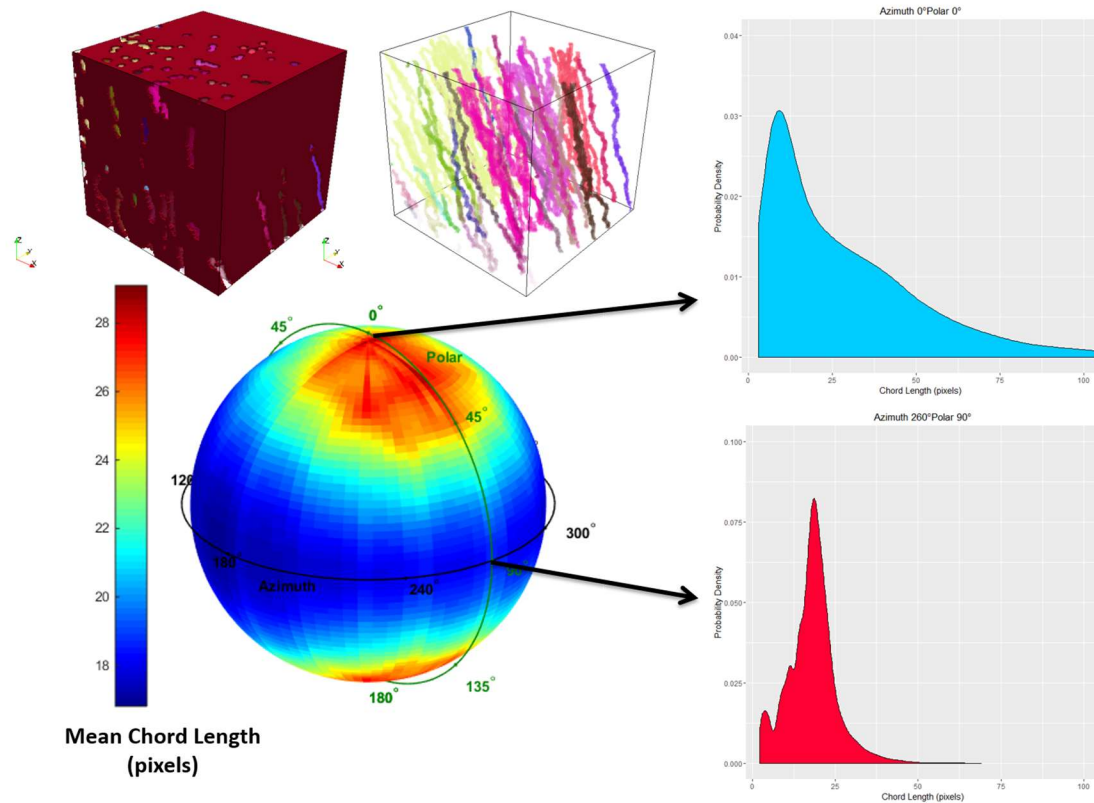


Figure 42: Top left shows a synthetic digitally created 3-D microstructure. The pore network has been coloured to show connected pores with distinguishing colours. In the bottom left, the mean pore chord length in all directions has been plotted on the surface of a sphere. Polar and azimuthal angles correspond to orientations of the chord within the sample. To the right, we show selected chord length distributions for two chord orientations. The angularly resolved chord length distributions show clearly the anisotropic nature of the porous structure.

Analysis of these results shows a difference of approximately 11 voxels between the maximum and minimum mean chord lengths for sampled directions. These anisotropic features of the pores in this sample are illustrated in Figure 42. It is important to remember that for each chosen orientation of scanlines through the sample we calculate a full chord length distribution. It is difficult to render all of this information into a single figure so we have elected to show an illustrative subset of these

distributions. Typical calculations of a non-angularly resolved chord length distribution would obviously not capture these features. Calculation of these distributions on the entire 500x500x500 voxel dataset was performed using a dual core Intel® i7-2460M CPU at 2.8 GHz in roughly 20 minutes.

4.6 Summary

Despite the known complex anisotropic nature of materials, non-angularly resolved implementations of the chord length distribution and lineal path function are frequent in the literature. We have described an efficient algorithm for the calculation of the angularly resolved chord length distributions in both 2D and 3D digital images. To the author's knowledge, this work presents the first calculations of these angularly resolved distributions for high resolution 3D datasets in thousands of directions at once. Moreover, we believe the algorithm presented can be extended to leverage general purpose computing on graphics processing units (GPGPU) to reach an additional order of magnitude or more performance increase. These results should provide researchers wishing to quantify the structure of materials with a new and more efficient tool to do so.

REFERENCES

1. Salvo, L., et al., *X-ray micro-tomography an attractive characterisation technique in materials science*. Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms, 2003. **200**: p. 273-286.
2. Spowart, J., H. Mullens, and B. Puchala, *Collecting and analyzing microstructures in three dimensions: A fully automated approach*. JOM Journal of the Minerals, Metals and Materials Society, 2003. **55**(10): p. 35-37.
3. Spowart, J.E., *Automated serial sectioning for 3-D analysis of microstructures*. Scripta Materialia, 2006. **55**(1): p. 5-10.
4. V. G. M. SIVEL, J.V.D.B., W. R. WANG, H. MOHDADI, F. D. TICHELAAR, P. F. A. ALKEMADE, H. W. ZANDBERGEN,, *Application of the dual-beam FIB/SEM to metals research*. Journal of Microscopy, 2004. **214**(3): p. 237-245.
5. Giannuzzi, L.A. and F.A. Stevie, *Introduction to focused ion beams : instrumentation, theory, techniques, and practice*. 2005, New York: Springer.
6. Echlin, M. and T. Pollock, *Femtosecond Laser Serial Sectioning: A New Tomographic Technique*. Ann Arbor, 2008. **1001**: p. 48109-2136.
7. Echlin, M.P., et al., *A new TriBeam system for three-dimensional multimodal materials analysis*. Review of Scientific Instruments, 2012. **83**(2): p. 023701-023701-6.
8. Midgley, P.A. and M. Weyland, *3D electron microscopy in the physical sciences: the development of Z-contrast and EFTEM tomography*. Ultramicroscopy, 2003. **96**(3-4): p. 413-431.
9. Rollett, A.D., et al., *Modeling Polycrystalline Microstructures in 3D*. AIP Conference Proceedings, 2004. **712**(1): p. 71-77.
10. Rollett, A.D., et al., *Three-Dimensional Characterization of Microstructure by Electron Back-Scatter Diffraction*. Annual Review of Materials Research, 2007. **37**(1): p. 627-658.
11. Garmestani, H., et al., *Semi-inverse Monte Carlo reconstruction of two-phase heterogeneous material using two-point functions*. International journal of theoretical and applied multiscale mechanics, 2009. **1**(2): p. 134-149.
12. Li, D., et al., *3D reconstruction of carbon nanotube composite microstructure using correlation functions*. Journal of Computational and Theoretical Nanoscience, 2010. **7**(8): p. 1462-1468.

13. Baniassadi, M., et al., *Three-dimensional reconstruction and homogenization of heterogeneous materials using statistical correlation functions and FEM*. Computational Materials Science, 2012. **51**(1): p. 372-379.
14. Tabei, S., et al., *Microstructure reconstruction and homogenization of porous Ni-YSZ composites for temperature dependent properties*. Journal of Power Sources, 2013. **235**: p. 74-80.
15. Sheidaei, A., et al., *3-D microstructure reconstruction of polymer nano-composite using FIB-SEM and statistical correlation function*. Composites Science and Technology, 2013. **80**: p. 47-54.
16. Pietroni, N., et al., *Solid-Texture Synthesis: A Survey*. Computer Graphics and Applications, IEEE, 2010. **30**(4): p. 74-89.
17. Kopf, J., et al., *Solid texture synthesis from 2D exemplars*. ACM Trans. Graph., 2007. **26**(3): p. 2.
18. Chen, J. and B. Wang, *High quality solid texture synthesis using position and index histogram matching*. The Visual Computer, 2010. **26**(4): p. 253-262.
19. Paget, R. and I.D. Longstaff, *Texture synthesis via a noncausal nonparametric multiscale Markov random field*. Image Processing, IEEE Transactions on, 1998. **7**(6): p. 925-931.
20. Besag, J., *Spatial interaction and the statistical analysis of lattice systems*. Journal of the Royal Statistical Society. Series B (Methodological), 1974: p. 192-236.
21. Geman, S. and D. Geman, *Stochastic relaxation, Gibbsdistributions, and the Bayesian restoration of images*. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6**:721-741. Geman7216IEEE Transactions on Pattern Analysis and Machine Intelligence1984, 1984.
22. Sinha, A. and S. Gupta, *A fast nonparametric noncausal MRF-based texture synthesis scheme using a novel FKDE algorithm*. Image Processing, IEEE Transactions on, 2010. **19**(3): p. 561-572.
23. Urs, R.D., J.P. Da Costa, and C. Germain, *Maximum-Likelihood Based Synthesis of Volumetric Textures From a 2D Sample*. Image Processing, IEEE Transactions on, 2014. **23**(4): p. 1820-1830.
24. Silverman, B.W., *Density estimation for statistics and data analysis*. Vol. 26. 1986: CRC press.
25. Metropolis, N., et al., *Equation of state calculations by fast computing machines*. The journal of chemical physics, 1953. **21**(6): p. 1087-1092.

26. Besag, J., *On the statistical analysis of dirty pictures*. Journal of the Royal Statistical Society. Series B (Methodological), 1986: p. 259-302.
27. Liu, X. and V. Shapiro, *Random heterogeneous materials via texture synthesis*. Computational Materials Science, 2015. **99**: p. 177-189.
28. Sundararaghavan, V., *Reconstruction of three-dimensional anisotropic microstructures from two-dimensional micrographs imaged on orthogonal planes*. Integrating Materials and Manufacturing Innovation, 2014. **3**(1): p. 19.
29. Wei, L.-Y. and M. Levoy. *Fast texture synthesis using tree-structured vector quantization*. in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. 2000. ACM Press/Addison-Wesley Publishing Co.
30. Torquato, S., *Random Heterogeneous Materials*. 2002, New York: Springer-Verlag.
31. Adams, B.L., H. Garmestani, and G. Saheli, *Microstructure design of a two phase composite using two-point correlation functions*. Journal of Computer-Aided Materials Design, 2004. **11**: p. 103–115.
32. Tewari, A., et al., *Quantitative characterization of spatial clustering in three-dimensional microstructures using two-point correlation functions*. Acta Materialia, 2004. **52**(2): p. 307-319.
33. Gokhale, A.M., *Experimental measurements and interpretation of microstructural N-point correlation functions*. Microscopy and Microanalysis, 2004. **10**(SUPPL. 2): p. 736-737.
34. Gokhale, A.M., A. Tewari, and H. Garmestani, *Constraints on microstructural two-point correlation functions*. Scripta Materialia, 2005. **53**: p. 989-993.
35. Huang, M., *The n-point orientation correlation function and its application*. International Journal of Solids and Structures, 2005. **42**(5-6): p. 1425-1441.
36. Etingof, P.I., D.D. Sam, and B.L. Adams, *Tensorial representation of two-point correlation functions for polycrystalline microstructure by harmonic polynomials*. Philosophical Magazine A: Physics of Condensed Matter, Defects and Mechanical Properties, 1995. **72**(1): p. 199.
37. Niezgodá, S.R., D.T. Fullwood, and S.R. Kalidindi, *Delineation of the space of 2-point correlations in a composite material system*. Acta Materialia, 2008. **56**(18): p. 5285-5292.
38. MacSleyne, J.P., J.P. Simmons, and M. De Graef, *On the use of 2-D moment invariants for the automated classification of particle shapes*. Acta Materialia, 2008. **56**(3): p. 427-437.

39. Kaczynski, T., K. Mischaikow, and M. Mrozek, *Computational homology*. Vol. 157. 2006: Springer Science & Business Media.
40. Gerrard, D.D., et al., *Computational Homology, Connectedness, and Structure-Property Relations*. Computers, Materials, & Continua, 2010. **15**(2): p. 129-152.
41. Niezgoda, S.R., et al., *Applications of the phase-coded generalized hough transform to feature detection, analysis, and segmentation of digital microstructures*. Computers, Materials, & Continua, 2010. **14**(2): p. 79-98.
42. Piasecki, R. *Microstructure reconstruction using entropic descriptors*. in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 2011. The Royal Society.
43. Balzani, D., et al., *Construction of two- and three-dimensional statistically similar RVEs for coupled micro-macro simulations*. Computational Mechanics, 2014. **54**(5): p. 1269-1284.
44. Turner, D.M., S.R. Niezgoda, and S.R. Kalidindi, *Estimation of the angularly resolved chord length distribution for digital microstructures*. in preparation, 2012.
45. Gille, W., D. Enke, and F. Janowski, *Pore Size Distribution and Chord Length Distribution of Porous VYCOR Glass (PVG)*. Journal of Porous Materials, 2002. **9**(3): p. 221-230.
46. Jokisch, D., et al., *Chord distributions across 3D digital images of a human thoracic vertebra*. Medical physics, 2001. **28**(7): p. 1493-1504.
47. Kellerer, A.M., *Chord-length distributions and related quantities for spheroids*. Radiation research, 1984. **98**(3): p. 425-437.
48. Rajon, D., et al., *Voxel effects within digital images of trabecular bone and their consequences on chord-length distribution measurements* Physics in Medicine and Biology, 2002. **47**(10).
49. Shah, A.P., et al., *A Comparison of Skeletal Chord-Length Distributions in the Adult Male*. Health Physics, 2005. **89**(3): p. 199-215
10.1097/01.HP.0000164653.55582.fd.
50. Torquato, S. and B. Lu, *Chord-length distribution function for two-phase random media*. Physical Review E, 1993. **47**(4): p. 2950.
51. Singh, H., et al., *Image based computations of lineal path probability distributions for microstructure representation*. Materials Science and Engineering: A, 2008. **474**(1-2): p. 104-111.

52. Brown, W.F., *Solid Mixture Permittivities*. The Journal of Chemical Physics, 1955. **23**(8): p. 1514-1517.
53. Fullwood, D.T., et al., *Microstructure sensitive design for performance optimization*. Progress in Materials Science, 2010. **55**(6): p. 477-562.
54. Fullwood, D.T., S.R. Niezgoda, and S.R. Kalidindi, *Microstructure reconstructions from 2-point statistics using phase-recovery algorithms*. Acta Materialia, 2008. **56**(5): p. 942-948.
55. Kalidindi, S.R., S.R. Niezgoda, and A.A. Salem, *Microstructure informatics using higher-order statistics and efficient data-mining protocols*. JOM, 2011. **63**(4): p. 34-41.
56. Fullwood, D.T., et al., *Gradient-based microstructure reconstructions from distributions using fast Fourier transforms*. Materials Science and Engineering a-Structural Materials Properties Microstructure and Processing, 2008. **494**(1-2): p. 68-72.
57. Niezgoda, S.R., et al., *Optimized structure based representative volume element sets reflecting the ensemble-averaged 2-point statistics*. Acta Materialia, 2010. **58**(13): p. 4432-4445.
58. Qidwai, S.M., et al., *Estimating the response of polycrystalline materials using sets of weighted statistical volume elements*. Acta Materialia, 2012. **60**(13-14): p. 5284-5299.
59. Hill, R., *Elastic properties of reinforced solids: Some theoretical principles*. Journal of the Mechanics and Physics of Solids, 1963. **11**(5): p. 357-372.
60. Ostoja-Starzewski, M., *Microstructural randomness and scaling in mechanics of materials*. 2007: CRC Press.
61. Ostoja-Starzewski, M., *Material spatial randomness: From statistical to representative volume element*. Probabilistic Engineering Mechanics, 2006. **21**(2): p. 112-132.
62. Hazanov, S. and C. Huet, *Order relationships for boundary conditions effect in heterogeneous bodies smaller than the representative volume*. Journal of the Mechanics and Physics of Solids, 1994. **42**(12): p. 1995-2011.
63. Huet, C., *Application of variational concepts to size effects in elastic heterogeneous bodies*. Journal of the Mechanics and Physics of Solids, 1990. **38**(6): p. 813-841.
64. Kanit, T., et al., *Determination of the size of the representative volume element for random composites: Statistical and numerical approach*. International Journal of Solids and Structures, 2003. **40**(13-14): p. 3647-3679.

65. Kalidindi, S.R., et al., *Elastic properties closures using second-order homogenization theories: Case studies in composites of two isotropic constituents*. Acta Materialia, 2006. **54**(11): p. 3117-3126.
66. Adams, B.L., X. Gao, and S.R. Kalidindi, *Finite approximations to the second-order properties closure in single phase polycrystals*. Acta Materialia, 2005. **53**(13): p. 3563-3577.
67. Binci, M., D. Fullwood, and S.R. Kalidindi, *A new spectral framework for establishing localization relationships for elastic behavior of composites and their calibration to finite-element models*. Acta Materialia, 2008. **56**(10): p. 2272-2282.
68. Kalidindi, S.R., G. Landi, and D.T. Fullwood, *Spectral representation of higher-order localization relationships for elastic behavior of polycrystalline cubic materials*. Acta Materialia, 2008. **56**(15): p. 3843-3853.
69. Bunge, H., *Texture Analysis in Materials Science*. Butterworths, 1982.
70. Gao, X., C.P. Przybyla, and B.L. Adams, *Methodology for Recovering and Analyzing Two-Point Pair Correlation Functions in Polycrystalline Materials*. Metallurgical and Materials Transactions A, 2006. **37**(8): p. 2379-2387.
71. Proust, G. and S.R. Kalidindi, *Procedures for construction of anisotropic elastic-plastic property closures for face-centered cubic polycrystals using first-order bounding relations*. Journal of the Mechanics and Physics of Solids, 2006. **54**(8): p. 1744-1762.
72. Paul, B., *Prediction of elastic constants of multiphase materials*. Trans. Metall. Soc. AIME, 1960(218): p. 36-41.
73. Hill, R., *The elastic behavior of a crystalline aggregate*. Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 1952. **65**: p. 349-54.
74. Hill, R., *A Theory of the Yielding and Plastic Flow of Anisotropic Metals*. Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 1948. **193**(1033): p. 281-297.
75. Kalidindi, S.R., C.A. Bronkhorst, and L. Anand, *CRYSTALLOGRAPHIC TEXTURE EVOLUTION IN BULK DEFORMATION PROCESSING OF FCC METALS*. Journal of the Mechanics and Physics of Solids, 1992. **40**(3): p. 537-569.
76. Rowenhorst, D.J., A.C. Lewis, and G. Spanos, *Three-dimensional analysis of grain topology and interface curvature in a β -titanium alloy*. Acta Materialia, 2010. **58**(16): p. 5511-5519.
77. ABAQUS, *Simulia*. 2009, Dassault Systemes: Providence, RI.

78. Qidwai, M.A.S., A.C. Lewis, and A.B. Geltmacher, *Using image-based computational modeling to study microstructure–yield correlations in metals*. Acta Materialia, 2009. **57**(14): p. 4233-4247.
79. Huang, Y., *A user-material subroutine incorporating single crystal plasticity in the ABAQUS finite element program*. 1991: Harvard Univ.
80. Ledbetter, M., et al., *Elastic constants of body-centered-cubic titanium monocrystals*. Journal of Applied Physics, 2004. **95**(9): p. 4642-4644.
81. Deka, D., et al., *Crystal plasticity modeling of deformation and creep in polycrystalline Ti-6242*. Metallurgical and Materials Transactions A: Physical Metallurgy and Materials Science, 2006. **37**(5): p. 1371-1388.
82. Kalidindi, S.R., A. Bhattacharyya, and R.D. Doherty, *Detailed analyses of grain-scale plastic deformation in columnar polycrystalline aluminium using orientation image mapping and crystal plasticity models*. Proceedings of the Royal Society of London Series a-Mathematical Physical and Engineering Sciences, 2004. **460**(2047): p. 1935-1956.
83. Barbe, F., S. Forest, and G. Cailletaud, *Intergranular and intragranular behavior of polycrystalline aggregates. Part 2: Results*. International Journal of Plasticity, 2001. **17**(4): p. 537-563.
84. Turner, D.M. and S.R. Kalidindi, *Statistical construction of 3-D microstructures from 2-D exemplars collected on oblique sections*. Acta Materialia, 2016. **102**: p. 136-148.
85. Muja, M. and D.G. Lowe. *Flann, fast library for approximate nearest neighbors*. in *International Conference on Computer Vision Theory and Applications (VISAPP'09)*. 2009.
86. Niezgod, S.R., A.K. Kanjarla, and S.R. Kalidindi, *Novel microstructure quantification framework for databasing, visualization, and analysis of microstructure data*. Integrating Materials and Manufacturing Innovation, 2013. **2**:3.
87. Niezgod, S.R., Y.C. Yabansu, and S.R. Kalidindi, *Understanding and Visualizing Microstructure and Microstructure Variance as a Stochastic Process*. Acta Materialia, 2011. **59**: p. 6387-6400.
88. Betz, O., et al., *Imaging applications of Synchrotron x-ray micro-tomography in biological morphology and biomaterial science. I. General aspects of the technique and its advantages in the analysis of arthropod structures*. Journal of Microscopy, 2007 **227**(1): p. 51-71.

89. Gupta, A., et al., *Structure-Property Linkages for Non-Metallic Inclusions/Steel Composite System using a Data Science Approach*. Acta Materialia, 2015. **91**: p. 239–254.
90. CeCen, A., et al., *A Data-driven Approach to Establishing Microstructure-Property Relationships in Porous Transport Layers of Polymer Electrolyte Fuel Cells*. Journal of Power Sources, 2014. **245**: p. 144-153.
91. David, M.T., R.N. Stephen, and R.K. Surya, *Efficient computation of the angularly resolved chord length distributions and lineal path functions in large microstructure datasets*. Modelling and Simulation in Materials Science and Engineering, 2016. **24**(7): p. 075002.
92. Darley, P.J., *Measurement of linear path length distributions in bone and bone marrow using a scanning technique*. Journal Name: pp 509-26 of Proceedings on the Symposium on Microdosimetry, Ispra, Italy, November 13-15, 1967. Ebert, H. G. (ed.). Brussels, Euratom, 1968.; Other Information: See EUR--3747; CONF-671109. Orig. Receipt Date: 31-DEC-68, 1968: p. Medium: X.
93. Krag, M.H., et al., *Morphometry of the Thoracic and Lumbar Spine Related to Transpedicular Screw Placement for Surgical Spinal Fixation*. Spine, 1988. **13**(1): p. 27-32.
94. Legrand, E., et al., *Trabecular Bone Microarchitecture, Bone Mineral Density, and Vertebral Fractures in Male Osteoporosis*. Journal of Bone and Mineral Research, 2000. **15**(1): p. 13-19.
95. Libouban, H., et al., *Comparison of Histomorphometric Descriptors of Bone Architecture with Dual-Energy X-ray Absorptiometry for Assessing Bone Loss in the Orchidectomized Rat*. Osteoporosis International, 2002. **13**(5): p. 422-428.
96. Spiers, F.W. and A.H. Beddoe, *'Radial' scanning of trabecular bone: consideration of the probability distributions of path lengths through cavities and trabeculae (dosimetry)*. Physics in Medicine and Biology, 1977. **22**(4): p. 670.
97. Manwart, C., S. Torquato, and R. Hilfer, *Stochastic reconstruction of sandstones*. Physical Review E, 2000. **62**(1): p. 893.
98. Boyard, N., et al., *Study of the porous network developed during curing of thermoset blends containing low molar weight saturated polyester*. Polymer, 2005. **46**(3): p. 661-669.
99. Roberts, A.P., *Statistical reconstruction of three-dimensional porous media from two-dimensional images*. Physical Review E, 1997. **56**(3): p. 3203.
100. Talukdar, M.S., et al., *Stochastic Reconstruction of Chalk from 2D Images*. Transport in Porous Media, 2002. **48**(1): p. 101-123.

101. Bresenham, J., *Algorithm for Computer Control of a Digital Plotter*. IBM Systems Journal, 1965. **4**(1): p. 25-30.
102. Parker, J.A., R.V. Kenyon, and D.E. Troxel, *Comparison of interpolating methods for image resampling*. IEEE Transactions on medical imaging, 1983. **2**(1): p. 31-39.
103. Kellerer, A.M., *Considerations on the random traversal of convex bodies and solutions for general cylinders*. Radiation Research, 1971. **47**(2): p. 359-376.
104. Coleman, R., *Random paths through convex bodies*. Journal of Applied Probability, 1969: p. 430-441.