

MIT Open Access Articles

*An Algorithm to Generate Synthetic 3D
Microstructures from 2D Exemplars*

The MIT Faculty has made this article openly available. ***Please share***
how this access benefits you. Your story matters.

As Published: <https://doi.org/10.1007/s11837-019-03825-w>

Publisher: Springer US

Persistent URL: <https://hdl.handle.net/1721.1/131916>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



An Algorithm to Generate Synthetic 3D Microstructures from 2D Exemplars

Cite this article as: Tristan N. Ashton, Donna Post Guillen and William Harris, An Algorithm to Generate Synthetic 3D Microstructures from 2D Exemplars, JOM <https://doi.org/10.1007/s11837-019-03825-w>

This Author Accepted Manuscript is a PDF file of an unedited peer-reviewed manuscript that has been accepted for publication but has not been copyedited or corrected. The official version of record that is published in the journal is kept up to date and so may therefore differ from this version.

Terms of use and reuse: academic research for non-commercial purposes, see here for full terms. <https://www.springer.com/aam-terms-v1>

Author accepted manuscript

An Algorithm to Generate Synthetic 3D Microstructures from 2D Exemplars

Tristan N. Ashton¹, Donna Post Guillen¹, and William Harris²

¹Idaho National Laboratory, 2351 N. Boulevard Ave., Idaho Falls, ID 83415

²Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139

Abstract

The inverse problem of constructing 3D microstructures from 2D data is an area of active research within the materials science community. This paper presents the implementation of a robust, computationally efficient algorithm, Hierarchical Algorithm for the Reconstruction of Exemplars (HARE), written in Python to reconstruct 3D features in a given microstructure from up to three orthogonal 2D exemplars using nearest neighbor matching to reproduce feature qualities, such as shape, size, and distribution. HARE's feature sampling implements histogram reweighting to avoid both over- and undersampling. A neighborhood voting scheme allows each pixel to provisionally affect its neighbors according to its weight. The algorithm is presently configured for two-phase materials and is being extended to accommodate multiple phases. HARE is a convenient and robust base from which to generate statistically representative synthetic microstructures for use in multi-scale modeling or machine learning applications to support advanced manufacturing and materials discovery.

1 Introduction

Recent years have seen an accelerating expansion of computational resources available to the materials science research community, due in large part to advances in parallel computing architectures and the resulting increases in processing speeds. These advances have prompted renewed focus in the area of predictive computational modeling of materials, in which novel materials are designed digitally to produce a prescribed functionality and performance. This approach reduces both the time and resources necessary for material design, and opens the frontier for multiscale modeling to connect emergent properties with nano-, micro-, and macroscale features. Key materials processes occur at microstructural scales and are inherently three-dimensional in nature¹. Richard Feynman in his celebrated 1959 lecture, "There's Plenty of Room at the Bottom," imagined the possibilities that could be achieved by manipulating and controlling matter at such small scales². However, the development of such predictive frameworks is challenging for at least two reasons. First, as a classic inverse problem, the space of possible material microstructures to explore to achieve the desired properties is extremely large, which makes the location of a global optimum difficult without specialized searching algorithms. Secondly, predictive models must be verified through comparison to experimental measurements of fabricated materials. This requires that as-fabricated materials be digitally reconstructed in some manner so that the candidate predictive model can be performed on them. The method described herein is much less resource intensive than obtaining 3D material microstructures by means of electron backscatter diffraction with successive serial sectioning³. The first challenge has been addressed with a wide variety of techniques⁴, such as modified Markov Chain Monte Carlo sampling and similar Monte Carlo methods⁵⁻⁷, and statistical correlation functions^{6,8-11}, largely in the context of molecular modeling and computational chemistry, and is not explored here. This paper focuses on solutions to the second challenge, namely, digital reconstruction of three-dimensional microstructures.

As described by Turner et al.¹², the primary aim of predictive materials modeling is the development of processing-microstructure-property-performance relationships to address the inverse problem. For example, consider a binary composite in which a matrix phase contains spherical inclusions of a precipitate phase B of varying diameter. One could completely specify this microstructure with two ordered configuration-space vectors of the form \vec{R}_i where $i = 1, 2, 3, \dots, n$ specifying the location of the n^{th} inclusion particle, and the vector \vec{D}_i stores

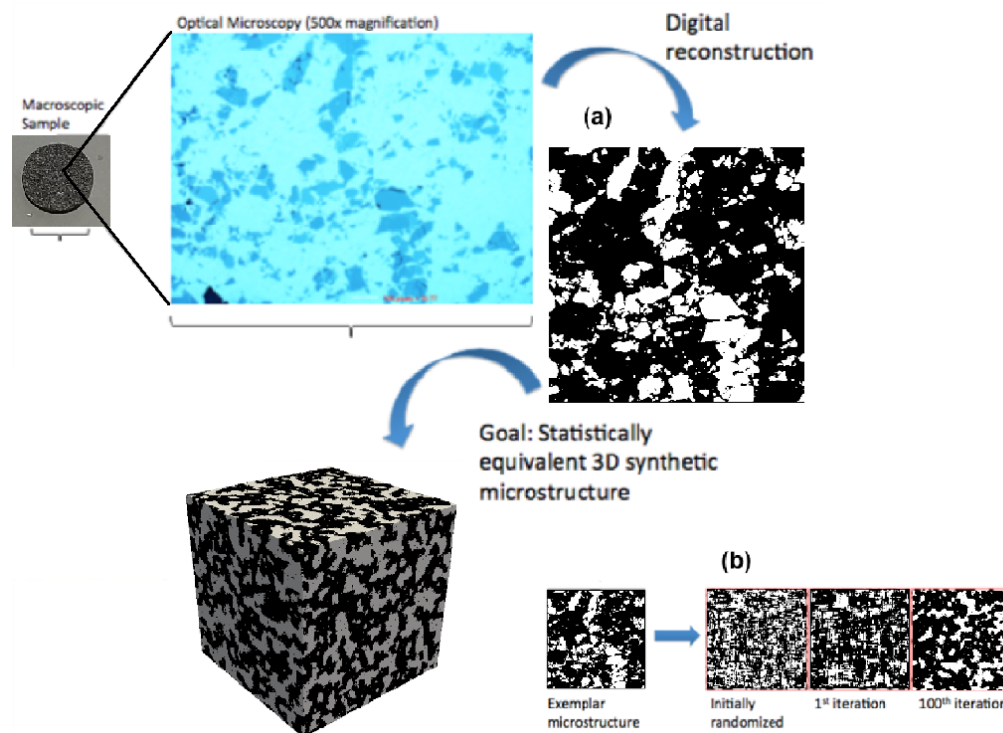


Figure 1: (a) HARE generates 3D microstructures from segmented 2D exemplars while preserving important characteristics, such as volume fraction of a given phase. An optical microscopy image of a metal matrix composite material was successfully reconstructed into 3D. (b) A proof-of-concept two-dimensional reconstruction of an exemplar, showing the algorithm's progress across 100 iterations.

their diameters. Finally, boundary conditions (i.e., periodic in all directions) must be specified, assuming that the configuration-space vectors above are meant to constitute a representative sample of the material. Note that this information could equivalently, though more compactly, be stored within a single higher-dimensional vector for actual computational applications, but that might obscure the present discussion. With this information, the microstructure is known deterministically.

Such an exact specification of a microstructure is of limited utility in achieving microstructure-property relationships because such results are not generalizable. Suppose the aforementioned composite is measured to have an unexpectedly high electrical conductivity. Is that due to the absolute positions of the inclusions or their relative spacing? The distribution of diameters? The length of the longest continuous path through the matrix? These questions cannot be answered readily by studying a single microstructure. Turner et al.¹² and others have instead supported the probabilistic or statistical approach to microstructure characterization, in which the bulk properties of a material are assumed to be uniquely and completely specified by the configurational statistics of the microstructure^{12,13}. This has led to numerous efforts to build both 2D and 3D solid textures from 2D inputs^{14–17}, some of which have emerged as leaders in the rapidly-growing materials development field^{15,16}.

This statistical approach is preferable for a number of reasons. First, if two materials are known to be statistically equivalent relative to one descriptor (inclusion radius distribution, for example), yet are still measured to have different properties, one can readily hypothesize alternative statistical descriptors that may give rise to the discrepancy (interparticle spacing distribution, for example). Moreover, one would not be burdened with keeping track of the absolute position of every inclusion particle during this comparison, as required in a deterministic approach, and statistical descriptions are readily amenable to sampling of smaller subsets while simultaneously keeping track of the uncertainties that might arise from such sampling.

The development of a robust reconstruction algorithm that can use three or less 2D inputs to rigorously produce a fully quantifiable microstructure is desirable, but faces two issues common to inverse problems - that exact analytical solutions either may not exist or may exist in local optima that fail to match the desired criteria, or that multiple solutions are possible without indication as to the superiority of one solution over the others. Although general circumventions of these issues often do not exist, it is possible to mitigate them in varying degrees by a selection of dynamic, ideally self-correcting, adjustments to the solution space. Such adjustments can include n-point spatial correlations¹⁸, which are useful in determining the importance of specific spatial structures to the effective microstructure properties; minimization of objective error functions, often employed

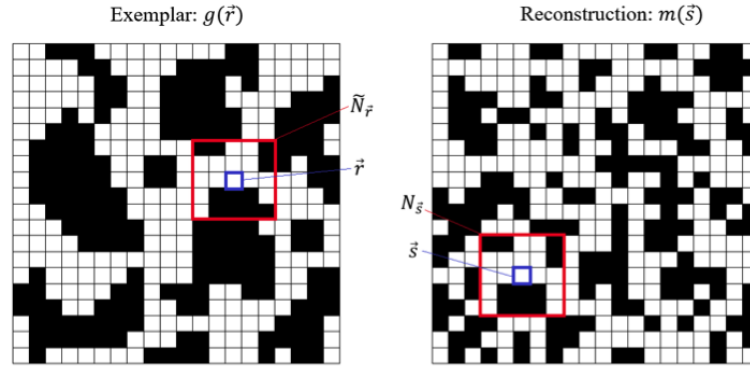


Figure 2: The neighborhood vectors $\tilde{N}_{\vec{r}}$ for each point in exemplar space $g(\vec{r})$ are matched with their nearest neighbor $N_{\vec{s}}$ in the reconstruction $m(\vec{s})$. Although each neighborhood pair is unlikely to perfectly match, as shown here with an error of 1, FLANN is able to quickly find the best pairs that minimize the local error, and thus, by extension the global error.

in Monte-Carlo methods⁵; reweighting schemes that uniformly sample features within the reconstruction to fairly represent all points within the microstructure¹⁵; and randomization conditions within the algorithm to nudge the convergence out of a potential local minimum¹⁹.

2 Methods

A robust, computationally efficient algorithm written in Python, the Hierarchical Algorithm for the Reconstruction of Exemplars (HARE), was created to reconstruct volumetric microstructures using a single 2D exemplar for isotropic materials or two to three orthogonal 2D exemplars for isotropic or anisotropic materials. This work builds upon the hybrid MATLAB/C++-based algorithm of Turner and Kalidindi¹² with a code written entirely in Python to provide extensibility and harness the power of modern open-source scientific computing libraries, toolkits and machine learning algorithms, as well as automated DREAM.3D pipelines.

The study of the mathematical and computational principles necessary for synthetic microstructure generation is an active area of research²⁰. A number of approaches have been developed, largely in the context of research focusing on graphics and image processing algorithms. In particular, Kopf et al.¹⁵ developed an algorithm for three-dimensional texture synthesis based on the input of exemplars. Although this method was developed primarily for applications in computer graphics, it has been readily assimilated into the materials research community to address the longstanding issue of synthetic microstructure generation, to which it is particularly well suited; when a material is characterized, this is often done via optical or scanning electron microscopy of a polished flat surface of the material. The resulting dataset is a 2D matrix of color values, to which the method of Kopf et al.¹⁵ is applicable. The method becomes simpler if a threshold is performed on the initial color value input dataset, so that the exemplars are segmented into two phases such that each pixel of the 2D cross-section can take only one of two discrete values (0 or 1). In the discussion that follows, it will be assumed that this has already been done.

The generation of a statistically-equivalent 2D slice from an input exemplar is described first, after which the method will be extended to three dimensions. Consider the $K \times K$ exemplar shown in Figure 2, where K is the side-length in units of pixels. As described above, each pixel is associated with a phase ID, which in the case of this binary composite is either 0 (white) or 1 (black). Let the vector $\vec{r} = (r_1, r_2)$ denote the position of each pixel and let the scalar-valued function $g(\vec{r})$ denote the phase ID of the pixel located at position \vec{r} ; thus, $g(\vec{r})$ is a full description of the exemplar.

We wish to produce an $L \times L$ -pixel microstructure that is statistically equivalent to the exemplar described by $g(\vec{r})$. Analogously to $g(\vec{r})$ and in keeping with standard notation in the literature, let $m(\vec{s})$ describe the phase ID values over all pixels in the reconstruction. To generate this statistically-equivalent reconstruction, Kopf et al.¹⁵ introduced the concept of local-neighborhood similarity. Each pixel in $m(\vec{s})$ is centered on an $n \times n$ -pixel neighborhood $N_{\vec{s}}$, as shown in Figure 2 with $n=5$. Starting with a randomized $m(\vec{s})$, we would like to iteratively make all neighborhoods $N_{\vec{s}}$ resemble the neighborhoods of $g(\vec{r})$.

To achieve this convergence between the reconstruction and the exemplar, each pixel of $m(\vec{s})$ must find the neighborhood $\tilde{N}_{\vec{r}}$ of $g(\vec{r})$ that best matches its own local neighborhood $N_{\vec{s}}$. In Figure 2, notice that $\tilde{N}_{\vec{r}}$ differs from $N_{\vec{s}}$ only at the bottom right pixel. This confirms that, as required, $\tilde{N}_{\vec{r}}$ is the closest-matching neighborhood




Figure 3: The reconstruction is driven by the nearest neighbor search, matching neighborhoods in the exemplar to those in the reconstruction. The red and green stars and associated boxes indicate sample neighborhoods and their closest match in the reconstruction; the yellow star describes the subsequent voting process. The red and green stars both vote for the yellow star to turn white, and only pixels in the yellow box in the reconstruction (which are chosen by the nearest neighbor matching) can cast votes for the yellow star.

of $N_{\vec{s}}$ in the exemplar. If all of the phase ID values of the pixels in the neighborhood $\tilde{N}_{\vec{r}}$ equal those in $N_{\vec{s}}$, the reconstruction is in agreement with the exemplar. Otherwise, the two models are in disagreement. This discrepancy is characterized by the error function $E(\vec{s})$ and is defined for each pixel in the reconstruction $m(\vec{s})$.

$$E(\vec{s}) = \sqrt{\sum_{\vec{u}} (N_{\vec{s}}(\vec{u}) - \tilde{N}_{\vec{r}}(\vec{u}))^2} \quad (1)$$

A new vector \vec{u} is introduced here, ranging from (1,1) to (n,n), that denotes the position of pixels within each neighborhood. For example, in Figure 2, $N_{\vec{s}}(\vec{u} = (5, 5)) = 0$ and $\tilde{N}_{\vec{r}}(\vec{u} = (5, 5)) = 1$, where the origin of \vec{u} is chosen as the top left corner of each neighborhood; this is the only discrepancy between the two neighborhoods, as mentioned previously. The sum over \vec{u} allows Eq. 1 to capture the total discrepancy between the neighborhoods. $E(\vec{s})$ is minimized when $N_{\vec{s}} = \tilde{N}_{\vec{r}}$, or when the $n \times n$ neighborhood $N_{\vec{s}}$ of the pixel at position \vec{s} in the reconstruction is identical, in terms of phase ID values, to the best-matching $n \times n$ neighborhood $\tilde{N}_{\vec{r}}$ in the exemplar. $E(\vec{s})$ is maximized when the neighborhoods are opposites, i.e., the black and white representation of the neighborhoods are inverted. The global error functional is shown in Eq. 2, and captures the total discrepancy between $m(\vec{s})$ and $g(\vec{r})$ by summing over all local neighborhood discrepancies.

$$E_{global} = \sum_{\vec{s}} E(\vec{s}) = \sum_{\vec{s}} \sqrt{\sum_{\vec{u}} (N_{\vec{s}}(\vec{u}) - \tilde{N}_{\vec{r}}(\vec{u}))^2} \quad (2)$$

Equation 2 is described as a functional, rather than a function, because it is a function of $m(\vec{s})$, which is itself a function describing the microstructure. The target reconstruction seeks to minimize this global error. This is done in an iterative fashion, alternating between a search phase and an optimization phase¹⁵. In the search phase, the closest matching neighborhoods $\tilde{N}_{\vec{r}}$ are found for all \vec{s} . This is a nearest neighbors search, and previous researchers in solid texture synthesis have generally found the approximate nearest neighbor (ANN) algorithms implemented with the Fast Library for Approximate Nearest Neighbors (FLANN C++ library) to provide an acceptable compromise between speed and accuracy^{21,22}. The details of the internal working of the ANN algorithms are not described here, but their purpose is to organize the total set of possible neighborhoods $\tilde{N}_{\vec{r}}$ of the exemplar into a k-dimensional (or k-d) tree in such a way that the approximate nearest neighborhood of any $N_{\vec{s}}$ can be found with significantly less computational expenditure than if the exact nearest neighborhood had been sought²¹.

The 3D microstructure is generated by stochastically minimizing a reconstruction loss defined as the sum of element-wise differences between each pixel's local neighborhood and its corresponding nearest-neighbor patch in the parallel exemplar, summed over the three orthogonal exemplars. Thus, in a given plane, each stochastic update searches through the fixed space of all size N^2 patches in the corresponding exemplar. For a fixed neighborhood size, this search space increases quadratically with the exemplar size, while for a fixed exemplar size, the search space decreases quadratically with neighborhood size. However, this benefit comes at the expense of accuracy in the reconstructed microstructure over length scales larger than the neighborhood size.




Figure 4: The multithreading within the reconstruction must be accompanied by a patching scheme to properly coordinate the actions of each section with its neighbor(s) to produce microstructures that contain physically relevant data. However, the patching can be circumvented to produce non-cubic volumes such as plates that also carry statistical similarities to the exemplar, in effect producing a range of reconstructions at once (only two such plates are shown here, but many horizontal section boundaries can be seen on the bottom-left cube).

In their review of the methods described above, Turner et al.¹² clarified the following point: each time the solid texture synthesis algorithm goes through an iteration, the values of each pixel in the reconstruction are updated in order to decrease the global error determined from Eq. 2. This update can be described as a weighted average of the “votes” each pixel receives from all neighborhoods to which it belongs²⁰. Consider pixel A and three of its neighbors B, C, and D in the reconstruction $m(\vec{s})$, and focus on the update to the phase ID value of pixel A. Say that from the perspective of pixel B, its local error would decrease if pixel A were switched to the other phase, but from the perspectives of pixel C and pixel D, pixel A should remain the same to minimize their local errors. If the weights of the “votes” of B, C, and D are equal, pixel A will remain unchanged in the next iteration of the reconstruction in pursuit of the minimization of global error, as calculated from Eq. 2. This mechanism is expressed as¹⁵.

$$m(\vec{s}) = \frac{\sum_{j=1}^{n^2} \omega_{\vec{s}, \vec{s}_j} \tilde{N}_{\vec{s}_j}(\vec{s})}{\sum_{j=1}^{n^2} \omega_{\vec{s}, \vec{s}_j}} \quad (3)$$

Note that the notation in Eq. 3 has been modified slightly to accommodate the relatively large amount of information referenced without using too many subscripts. Consider the pixel at \vec{s} in the reconstruction. It participates in n^2 neighborhoods, because it is within the neighborhood of each of the $n^2 - 1$ pixels that surround it in addition to its own. The index $j = 1, 2, \dots, n^2$ is the counter over those pixels, \vec{s}_j , each of which “votes”, with the weight $\omega_{\vec{s}, \vec{s}_j}$, as to what the phase ID $m(\vec{s})$ should be set to (either 0 or 1) during the next iteration such that its own local error is decreased. These votes are equal to $\tilde{N}_{\vec{s}_j}(\vec{s})$, which denotes the phase ID of the pixel within $\tilde{N}_{\vec{s}_j}$ that corresponds to the location of \vec{s} . This is best understood through an example.

Suppose that $m(\vec{s})$, as illustrated in Figure 3, is the initial randomized “guess” for the reconstruction of the exemplar. Let \vec{s}_y , \vec{s}_r , and \vec{s}_g denote the pixels marked by the yellow, red, and greens stars in $m(\vec{s})$. We would like to know what the phase ID value of \vec{s}_y (i.e., $m(\vec{s}_y)$) should be updated to in the next iteration in order to decrease the global error functional $E(m(\vec{s}))$. To do this, we will consider the individual votes of \vec{s}_r and \vec{s}_g , although in reality all pixels in the yellow neighborhood will contribute if the neighborhood size is chosen to be 3×3 . The best-matching neighborhood of \vec{s}_r is outlined with red in the exemplar $g(\vec{r})$ and is denoted by $\tilde{N}_{\vec{s}_r}$, and similarly for the green star with \vec{s}_g and $\tilde{N}_{\vec{s}_g}$. As mentioned previously, these best-matching neighborhoods were determined via the ANN search algorithm using the FLANN library²². Notice that the local error of \vec{s}_r would

be zero if the yellow star-marked pixel were white, or $m(\vec{s}_y) = 0$. Similarly, the pixel at \vec{s}_g votes for $m(\vec{s}_y)$ to be 0. The weighted average of such votes is calculated over all pixels in the yellow-outlined neighborhood, and the final result is rounded to 0 or 1. The result will be the value of $m(\vec{s}_y)$ in the next update of the reconstruction. Note that this rounding operation is not always performed in solid texture synthesis algorithms, but doing so mitigates the possibility of obtaining a blurry reconstruction¹⁶. The reweighting prescription given in Eq. 4 is implemented after each iteration of the reconstruction

$$\omega'_{\vec{s}, \vec{s}_j} = \frac{\omega_{\vec{s}, \vec{s}_j}}{1 + \max[0, H(\tilde{N}_{\vec{s}_j}(\vec{s})) - \theta]} \quad (4)$$

where H denotes the position histogram and $\theta = \frac{1}{K^2}$, which represents the (normalized) probability of each pixel of the exemplar being sampled in the reconstruction if all such pixels are sampled uniformly¹⁶. This reweighting permits a dynamic adjustment of the relative importance of each pixel to the reconstruction between each iteration, selectively choosing undersampled pixels wherever possible.

Histogram reweighting provides a mathematical bridge between microstructure reconstruction and statistical mechanics by unbiasing a statistical distribution away from its initial state. Consider a Monte Carlo simulation that, through random initialization, begins near a metastable state that is not representative of the equilibrium properties of the system. Depending on the height of nearby energy barriers, the simulation may never escape its local minimum. To avoid this situation, one can artificially increase the energy of a particular configuration based on the number of times that configuration has been visited along the simulation's trajectory, thus biasing the trajectory away from them. Similarly, over-sampled patches in the exemplar can be down-weighted based on how frequently they are used in the reconstruction, thus ensuring representative sampling of all of the input data. This technique is commonly used when training neural networks as well, particularly in the reverse phase of training restricted Boltzmann machines.

The convergence is further accelerated by a hierarchical upsampling across several resolution levels, implemented after a set number of iterations. Exemplars can be successively downsampled to half the previous resolution to produce a coarse structure, and the lowest resolution is used as the first input exemplar. Initial use of this coarsest exemplar allows the reconstruction to capture large-scale features that might otherwise be missed with only one level of resolution, and circumvents the initial guess issue in inverse problems. Although it is always better to start with an input as close to the desired solution as possible, this can be a difficult requirement to fulfill due to the vast solution space available to a 2D-to-3D generalization; a randomized structure as an initial guess is sufficient for most purposes, and is where the reconstruction begins. The use of multiple levels then permits the initial guess of subsequent levels to use the output from the previous level, allowing the new guess to more closely match the input exemplars while simultaneously increasing the resolution between each level to capture more and more features. When downsampling the exemplars, every set of four pixels becomes a single pixel. The downsampled phase ID value (0 = black and 1 = white) is determined using a rounded average. When the average is calculated to be 0.5, the downsampled pixel is chosen randomly.

To allow for inevitable increases in reconstruction size, the algorithm was wrapped in an asynchronous multiprocessing scheme which divides the reconstruction into multiple sections, enabling a concurrent evaluation of all sections in each iteration while piping the result outside the reconstruction to be plotted on screen. This implementation dramatically sped up the reconstruction: in a reconstruction containing three levels of 40 iterations each, neighborhood sizes in each level of 5, 11, and 19, respectively, and a full reconstruction size of 256^3 , the parallelization scheme reduced the runtime from about 37 hours to just 8.5 hours when each level is divided into six sections, a reduction of almost 80%.

The multiprocessing scheme is accompanied by a compensatory patching tool, serving as a communicator between separate processors to coordinate the actions of each section. A padded buffer at the edges of each section is matched with a similar buffer at the edge of the adjacent section(s) to coordinate the weighted phase ID votes and subsequent updated phase values. Without this patching, distinct divisions between sections are readily visible, rendering the reconstruction unusable for any physical application. However, this nominal defect was adapted to extend the algorithm into non-cubic volumes, namely the reconstruction of plate-like structures. Without the patching tool, the reconstruction can be considered as independently executing as many times as the number of sections determined by the user, with the caveat of stacking each section with the others to produce a cubic volume. The patching scheme is illustrated in Figure 4. A plate reconstruction, then, can be extracted by cropping and tiling the two thin plate exemplars with themselves to match the dimensions of the largest exemplar and then running the algorithm as normal. In this case, however, we exploit the segregation that occurs without the patching by setting the number of sections equal to the number of tiles of the thin exemplars needed to match the large exemplar's dimensions, rounded up to the nearest integer. Each section is then sequestered from the others and saved to its own automatically generated directory; in this manner, we actually achieve multiple reconstructions of the plate for the same computational cost as one.




Figure 5: Multiphase materials can be reconstructed by isolating non-ordered pairs of values in a grayscale image and reconstructing each phase pair separately. Each reconstruction is inserted into a master array, where each "pixel" contains the value from each reconstruction at that same point. This can be extended to arbitrary numbers of phases, although many phases in one image raises the difficulties of characterizing very low volume fraction materials. Here, we combine two RGB images obtained by EDS (a)²³ into a single image (b), then convert to grayscale and segment (c). The segmented image is then split according to each combination of $\binom{p}{2}$, where p is the number of phases, from which a binary reconstruction can proceed as normal (d). Here, three binary phase reconstructions are illustrated.

The generalization to multiphase materials would provide a useful avenue for the exploration of microstructures with known chemical or phase data. We have developed a prototype methodology that characterizes spatial correlations between phase pairs within a multiphase material by reconstructing each pair individually. To effect maximum information preservation, the final reconstruction for each pair is superimposed into a four-dimensional master array, from which a number of three-dimensional nano- or microstructures can be extracted. Results from energy-dispersive X-ray spectroscopy (EDS) images of Si and Al in a simulated nuclear waste feed²⁴ are shown in Figure 5. Although superimposed phase-pair reconstructions are only a rudimentary attempt to generate multiphasic microstructures, this implementation reduced the complexities of ternary (and greater) reconstructions despite increasing computational time. In its current implementation, the runtime t increases as

$$t = 0.5p^2 - 0.5p \quad (5)$$

where p is the number of phases, reflecting the significant increase in combinatorial pairs as the number of phases increases; however, we believe the trade-off between the greatly expanded informational scope versus the time investment to be worthwhile. When multiprocessing across multiple nodes is fully enabled (rather than partitioning a single node), the nodal count will likely negatively offset the p^2 term in Eq. 5. The reconstructions were executed on a 192-core Dell PowerEdge distributed large-memory system with 32 cores per node and 2.25 TB total memory and a LINPACK rating of 3.78 TFlops. In future versions of HARE, adapting this combinatorial approach to the exemplars would permit the reconstruction of RGB images as opposed to grayscale, significantly expanding the capabilities of the algorithm.

3 Results and Discussion

Figure 1a illustrates the process of reconstructing a synthetic 3D microstructure of a HfAl_3 -Al metal matrix composite from a segmented 2D optical microscopy image³. The reconstruction algorithm is performed for 100 iterations to generate a reconstructed "microstructure" of size 50×50 . The progression from the initial randomized reconstruction at the first and one hundredth iterations of the algorithm is shown in Figure 1b. It

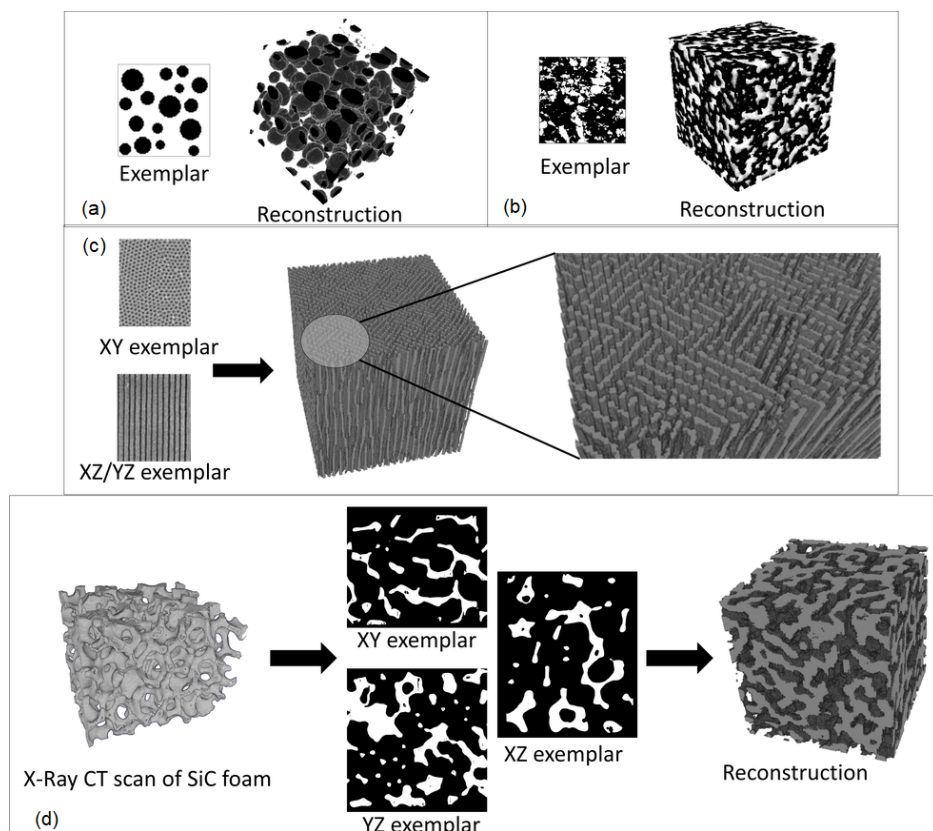


Figure 6: HARE reproduces isotropic and anisotropic materials from computer-generated images as well as digital scans of real materials. Three-dimensional microstructures generated from (a) a single exemplar of a material with spherical inclusions, (b) a single exemplar of a two-phase material, (c) two orthogonal exemplars of a fibrous material, and (d) three orthogonal exemplars of a foam.

can be seen that the general form of the exemplar is visible by the first iteration, suggesting that the rather small size of the exemplar allowed for rapid convergence of the microstructure. The following 99 iterations cleaned up the reconstruction further to produce a somewhat imperfect tiled pattern of the exemplar that nonetheless captures the overall shape of the exemplar.

The impetus for this algorithm centers on the motivation to produce a statistically similar reconstruction of the exemplars. To verify that the reconstructed volumes are indeed representative of the exemplars, we surveyed and analyzed a range of morphological statistics across several exemplar-reconstruction pairs. DREAM.3D software^{25–27}, a tool that is well-suited for analysis of volumetric microstructures, can be used to tally features of the reconstructed microstructure for comparison to the exemplar. For a given reconstruction, DREAM.3D can be used to quantify characteristics, such as the volume fraction of phases, equivalent sphere diameter, distance between features, radius of neighboring features, neighboring features in contact or other parameters of interest. The mathematics of DREAM.3D's workings are available in a series of papers establishing the framework for the recognition and analysis of n-dimensional solids and their associated moment invariants^{28–32}. It should be noted that while DREAM.3D is also adept at microstructure synthesis according to morphological statistics, it currently only generates semi-regular polygonal and ellipsoidal features, which are not space-filling and thus can miss smaller irregular features. HARE is able to construct fully irregular features and performs quite well in tandem with DREAM.3D. DREAM.3D can provide both the two- and three-dimensional statistics of the exemplars and the microstructure, respectively, while a tool such as MATLAB can be subsequently employed to compare the statistical distributions.

The reconstruction of materials with different textures is illustrated in Figure 6. The nature of these textures ranges from ideal and isotropic (Figure 6a and b) to the reconstruction of a real anisotropic material from three orthogonal images (such as the SiC foam shown in Figure 6d). The reconstruction of the fibrous material from two orthogonal exemplars shown in Figure 6c results in a 45° arrangement of fibers. For a texture such as this, including a third exemplar would provide the additional information necessary to ensure that the reconstructed microstructure is uniquely defined. The reconstruction of the SiC foam shown in Figure 6d was based upon

three orthogonal slices obtained by X-ray computed tomography. HARE reconstructs a synthetic structure that closely resembles the morphology of the original microstructure.

Repeated tests have demonstrated that use of the histogram reweighting scheme given by Eq. 4 is necessary to generate microstructures that resemble their exemplars; without it, the voting scheme is unable to distinguish between over- and undersampled pixels and adjust the phase ID updates accordingly, the reconstruction overwhelmingly mis-samples exemplar neighborhoods and produces highly unrepresentative structures. We affirm the similar conclusions reached by both Kopf et al.¹⁵ and Turner et al.¹² regarding the necessity of histogram reweighting to correctly match global statistics in conjunction with neighborhood matching to reproduce local statistics. Similarly, the multithreading must be accompanied by a patching scheme to accurately orchestrate the reconstruction at the boundaries between sections; neglecting it renders cubic reconstructions unusable for any material or physical application. However, plate structures can be extracted from a non-patched reconstruction if necessary while retaining the desired statistical equivalence.

The reconstruction of multiphase materials has been developed into a simple, yet prototypic, approach. By isolating pairs of phases, we can avoid a complete overhaul of the voting scheme while preserving the exemplars' structures; an earlier attempt at reconstructing all phases at once showed that the solution quickly converged to either a single phase or eliminated the middle phases after only a few iterations. An added benefit of the new prototype sidesteps the inherently biphasic nature of HARE- instead of the all-or-nothing choice of phases in the cubic and plate reconstructions, wherein the reconstructions for each phase pair are allowed to coexist via superposition. This prototype may be expanded by incorporating more sophisticated methods to preserve low volume fraction phases.

4 Summary and Future Work

The HARE algorithm is a solid foundation for future efforts in multi-scale modeling to support advanced manufacturing and materials discovery. A large part of materials science lies within the development of new materials for specific purposes, and our algorithm provides a convenient and robust base from which to generate statistically representative synthetic microstructures. This work builds upon the hybrid MATLAB/C++-based algorithm of Turner and Kalidindi¹² and offers additional features, such as histogram reweighting, multithreading, and an option for generating plate microstructures. The code was written entirely in Python to provide extensibility and harness the power of open-source scientific computing libraries, toolkits and machine learning algorithms. Python can easily interface with DREAM.3D workflows and run automated pipelines. As mentioned earlier, the "defects" in the reconstruction produced by the algorithm cannot be ignored when analyzing the material across various scales, since the roles of semi-random defects and small-scale heterogeneity have yet to be fully defined in regards to their effect on bulk material properties and consequent performance. These defects are often precisely the features we wish to capture. This algorithm, when coupled with machine learning techniques, can be a powerful tool for materials design³³. For example, synthetically-generated microstructures exhibiting defects could be used to train machine-learning algorithms, such as convolutional neural networks³⁴. Additionally, this algorithm could be incorporated into advanced finite element modeling tools³⁵. With this approach, the incorporation of embedded sensors into additively manufactured parts can be evaluated.

By constructing microstructures of specific texture, the fabrication of materials with chosen macroscale properties emergent from microscale features with performance tailored for a given application can be facilitated. We believe this will be a powerful tool when combined with future multiphase reconstruction capabilities. Currently, the algorithm can only reconstruct two-phase materials. By expanding the multi-phase capabilities the algorithm will become exponentially more useful to researchers who possess chemical and phase data. Further refinements can be incorporated through the inclusion of additional convergence heuristics to narrow the solution space. Two-point spatial correlations would allow FLANN to preferentially select nearest neighbors in the reconstruction that are spatially closest to the associated exemplar neighborhood, the mathematics of which are well understood¹⁸. Similarly, incorporating a simulated annealing sub-algorithm would push solutions out of local minimums toward the global optimum, a technique often used in combinatorial optimizations³⁶. Implementation of these additional heuristics would accommodate low volume percent materials, which currently tend to converge to single-phase reconstructions and are not representative of the materials from which they were taken. As mentioned earlier, the adaptation to accept multivalued exemplars such as those generated by X-Ray Fluorescence Mapping or EDS would greatly expand HARE's ability to reconstruct the chemical or phase composition of a material.

5 Acknowledgments

This work was supported by Albert Kruger of the U.S. Department of Energy’s (DOE) Waste Treatment and Immobilization Plant Project of the Office of River Protection and by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internships Program (SULI). Javier Morales contributed to the development of the Python code. Michael Jackson of Blue Quartz Software provided guidance on DREAM.3D. X-ray CT images of the SiC foam were acquired by Joshua Kane, Austin Matthews and Robert Siefert of the INL. We wish to thank Tom Kielbaso and Michael Velez from UES, Inc. (Dayton, OH) for providing the optical microscopy images of the HfAl₃–Al sample. This work was performed by Battelle Energy Alliance, LLC under DOE Idaho Operations Contract DE-AC07-05ID14517. This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

Nomenclature

\vec{D}_i	Diameter of i^{th} particle in microstructure
\vec{r}	Vector denoting the position of each pixel
\vec{R}_i	Spatial position of i^{th} particle in microstructure
\vec{u}	Neighborhood vector denoting pixel positions
$\tilde{N}_{\vec{r}}$	$n \times n$ pixel neighborhood in the exemplar(s)
$\tilde{N}_{\vec{r}}$	Neighborhood of the pixel at u in exemplar
$E(\vec{s})$	Error function
$g(\vec{r})$	Phase ID value of the pixel at \vec{r} in the exemplar
$H(\tilde{N}_{\vec{s}_j}(\vec{s}))$	Position histogram of phase ID values in neighborhoods in the exemplar
K	Exemplar side dimension (number of pixels)
L	Reconstruction side dimension (number of pixels)
$m(\vec{s})$	Phase ID value of the pixel at \vec{s} in the reconstruction
$N_{\vec{r}}$	$u \times n$ pixel neighborhood in the reconstruction
$N_{\vec{r}}$	Neighborhood of the pixel at u in the reconstruction
t	Computational time
$\omega'_{\vec{s}, \vec{s}_j}$	Weight of a pixel’s vote when updating phase ID values
p	Number of phases
θ	Probability each pixel of the exemplar is sampled in the reconstruction

References

1. Council, N. R. *Integrated Computational Materials Engineering: A Transformational Discipline for Improved Competitiveness and National Security*; The National Academies Press, 2008.
2. Feynman, R. There’s Plenty of Room at the Bottom. *Engineering and Science* **1960**, *23*, 22–36.
3. Guillen, D.; Harris, W. Measurement and Simulation of Thermal Conductivity of Hafnium-Aluminum Thermal Neutron Absorber Material. *Metallurgical and Materials Transactions E* **2016**, *3*, 123–133.
4. Bostanabad, R.; Zhang, Y.; Li, X.; Kearney, T.; Brinson, L.; Apley, D.; Liu, W.; Chen, W. Computational Microstructure Characterization and Reconstruction: Review of the State-of-the-art Techniques. *Progress in Materials Science* **2018**, *95*, 1–41.

5. Garmenstani, H.; Baniassadi, M.; Li, D.; Fathi, M.; Ahzi, S. Semi-inverse Monte Carlo reconstruction of Two-phase Heterogeneous Material using Two-point Functions. *International Journal of Theoretical and Applied Multiscale Mechanics* **2009**, *1*, 134–149.
6. Baniassadi, M.; Mortazavi, B.; Hamedani, H.; H.Garmestani,; S.Ahzi,; M.Fathi-Torbaghan,; D.Ruch,; M.Khaleel, Three-dimensional reconstruction and homogenization of heterogeneous materials using statistical correlation functions and FEM. *Computational Materials Science* **2012**, *51*, 372–379.
7. Garmestani, H.; Baniassadi, M.; Li, D.; Fathi, M.; Ahzi, S. semi-inverse Monte Carlo reconstruction of two phase heterogeneous material using two point functions. *International Journal of Theoretical and Applied Multiscale Mechanics* **2009**, *1*, 134–149.
8. Fullwood, D.; Niezgoda, S.; Adams, B.; Kalidindi, S. Microstructure sensitive design for performance optimization. *Progress in Materials Science* **2010**, *55*, 477–562.
9. Gokhale, A.; Tewari, A.; Garmestani, H. Constraints on microstructural two-point correlation functions. *Scripta Materialia* **2005**, *53*, 989–993.
10. Fullwood, D.; Niezgoda, S.; Kalidindi, S. Microstructure reconstructions from 2-point statistics using phase-recovery algorithms. *Acta Materialia* **2008**, *56*, 942–948.
11. Adams, B.; Garmestani, H.; Saheli, G. Microstructure design of a two phase composite using two-point correlation functions. *Journal of Computer-Aided Materials Design* **2004**, *11*, 103–115.
12. Turner, D.; Kalidindi, S. Statistical construction of 3-D microstructures from 2-D exemplars collected on oblique sections. *Acta Materialia* **2016**, *102*, 136–148.
13. Xu, H.; Dikin, D.; Burkhart, C.; Chen, W. Descriptor-based methodology for statistical characterization and 3D reconstruction of microstructural materials. *Computational Materials Science* **2014**, *85*, 206–216.
14. Yeong, C.; S.Torquato, Reconstructing random media. II. Three-dimensional media from two-dimensional cuts. *Physical Review E* **1998**, *58*, 224–233.
15. Kopf, J.; Fu, C.-W.; Cohen-Or, D.; Deussen, O.; Lischinski, D.; Wong, T.-T. Solid Texture Synthesis from 2D Exemplars. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007)* **2007**, *26*, 2:1–2:9.
16. Chen, J.; Wang, B. High quality solid texture synthesis using position and index histogram matching. *The Visual Computer* **2010**, *26*, 253–262.
17. Sundararaghavan, V. Reconstruction of three-dimensional anisotropic microstructures from two-dimensional micrographs imaged on orthogonal planes. *Integrating Materials and Manufacturing Innovation* **2014**, *3*, 19.
18. Niezgoda, S.; Fullwood, D.; Kalidindi, S. Delineation of the space of 2-point correlations in a composite material system. *Acta Materialia* **2008**, *56*, 5285–5292.
19. Urs, R.; Costa, J. D.; Germain, C. Maximum-likelihood Based Synthesis of Volumetric Textures from a 2D Sample. *IEEE Transactions on Image Processing* **2014**, *23*, 1820–1830.
20. Pietroni, N.; Cignoni, P.; Otaduy, M.; Scopigno, R. A survey on solid texture synthesis. *Journal of LaTeX Class Files* **2007**, *6*, 1–14.
21. Liang, L.; Liu, C.; Xu, Y.; Guo, B.; Shum, H. Real-Time Texture Synthesis by Patch-Based Sampling. *ACM Transactions on Graphics* **2001**, *20*, 127–150.
22. Muja, M. FLANN - Fast Library for Approximate Nearest Neighbors: User Manual. <https://github.com/mariusmuja/flann>.
23. Dixon, D.; Schweiger, M.; Riley, B.; Pokorny, R.; Hirma, P. Temperature Distribution within a Cold Cap during Nuclear Waste Vitrification. *Environmental Science and Technology* **2015**, *49*, 8856–8863.
24. Dixon, D.; Schweiger, M.; Riley, B.; Pokorny, R.; Hirma, P. Temperature Distribution within a Cold Cap during Nuclear Waste Vitrification. *Environmental Science and Technology* **2015**, *49*, 8856–8863.
25. Groeber, M.; Jackson, M. DREAM.3D: A Digital Representation Environment for the Analysis of Microstructure in 3D. *Integrating Materials and Manufacturing Innovation* **2014**, *3*, 5.

26. Rollett, A.; Saylor, D.; Firdy, J.; El-Dasher, B.; Brahme, A.; Lee, S.; Cornwell, C.; Noack, R. Modeling polycrystalline microstructures in 3D. *AIP Conference Proceedings* **2004**, 712, 71–77.
27. Rollett, A.; Lee, S.-B.; Campman, R.; Rohrer, G. Three-dimensional characterization of microstructure by electron back-scatter diffraction. *Annual Review of Materials Research* **2007**, 37, 627–658.
28. Graef, M. *Representation and Reconstruction of Three-dimensional Microstructures in Ni-based Superalloys*; 2010; p 32.
29. MacSleyn, J.; Simmons, J.; Graef, M. On the use of moment invariants for the automated analysis of 3D particle shapes. *Modeling and Simulation in Materials Science and Engineering* **2008**, 16.
30. Mamistvalov, A. n-Dimensional Moment Invariants and Conceptual Mathematical Theory of Recognition n-Dimensional Solids. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **1998**, 20, 819–831.
31. Groeber, M.; Ghosh, S.; Uchic, M.; Dimiduk, D. A framework for automated analysis and simulation of 3D polycrystalline microstructures.: Part 1: Statistical characterization. *Acta Materialia* **2008**, 56, 1257–1273.
32. Simmons, J.; Chuang, P.; Comer, M.; Uchic, M.; Spowart, J.; Graef, M. D. Application and Further Development of Advanced Image Processing Algorithms for Automated Analysis of Serial Section Image Data. *Modelling and Simulations in Materials Science and Engineering* **2009**, 17.
33. Bulgarevich, D. S.; Tsukamoto, S.; Kasuya, T.; Demura, M.; Watanabe, M. Pattern Recognition with Machine Learning on Optical Microscopy Images of Typical Metallurgical Microstructures. *Scientific Reports* **2018**, 8.
34. Kitahara, A.; Holm, E. Microstructure Cluster Analysis with Transfer Learning and Unsupervised Learning. *Integrating Materials and Manufacturing Innovation* **2018**, 7, 148–156.
35. Novascone, S.; Medvedev, P.; Peterson, J. W.; Zhang, Y.; Hales, J. Modeling Porosity Migration in LWR and Fast MOX Fuel using the Finite Element Method. *Journal of Nuclear Materials* **2018**, 508, 226–236.
36. Czyżżak, P.; Jaskiewicz, A. Pareto simulated annealing—a metaheuristic technique for multiple-objective combinatorial optimization. *Journal of Multi-Criteria Decision Analysis* **1998**, 7, 34–47.