

Aaron Spettl · Thomas Werz ·
Carl E. Krill III · Volker Schmidt

Parametric representation of 3D grain ensembles in polycrystalline microstructures

Received: date / Accepted: date

Abstract As a straightforward generalization of the well-known Voronoi construction, Laguerre tessellations have long found application in the modelling, analysis and simulation of polycrystalline microstructures. The application of Laguerre tessellations to real (as opposed to computed) microstructures—such as those obtained by modern 3D characterization techniques like x-ray microtomography or focused-ion-beam serial sectioning—is hindered by the mathematical difficulty of determining the correct seed location and weighting factor for each of the grains in the measured volume. In this paper, we propose an alternative to the Laguerre approach, representing grain ensembles with convex cells parametrized by orthogonal regression with respect to 3D image data. Applying our algorithm to artificial microstructures and to microtomographic data sets of an Al-5 wt% Cu alloy, we demonstrate that the new approach represents statistical features of the underlying data—like distributions of grain sizes and coordination numbers—as well as or better than a recently introduced approximation method based on the Laguerre tessellation; furthermore, our method reproduces the local arrangement of grains (*i.e.*, grain shapes and connectivities) much more accurately. The additional computational cost associated with orthogonal regression is marginal.

Keywords polycrystalline microstructure · Laguerre tessellation · orthogonal regression · convex cells

A. Spettl · V. Schmidt
Institute of Stochastics, Ulm University, Helmholtzstr. 18, 89069 Ulm, Germany
Tel.: +49-731-50-23555
Fax: +49-731-50-23649
E-mail: aaron.spettl@uni-ulm.de

T. Werz · C. E. Krill III
Institute of Micro and Nanomaterials, Ulm University, Albert-Einstein-Allee 47, 89081 Ulm, Germany

1 Introduction

Inspired by experience gained from working with metals, mankind has long employed various processing steps—such as mechanical deformation or heat treatment—to improve the properties of crystalline solids. It wasn't until the past 100 years, however, that materials scientists were able to place this empirical knowledge on a firm scientific footing, attributing processing-induced property changes to the generation and annihilation of lattice defects (grain boundaries, dislocations, point defects, *etc.*) constituting the microstructure of crystalline materials. And only during the past decade has a determination of the true three-dimensional arrangement of various microstructural elements—such as the network of grain boundaries spanning a polycrystalline specimen—become routinely feasible, thanks to powerful characterization techniques like x-ray microtomography [19, 1], focused-ion-beam-based serial sectioning [7, 25] and diffraction-based x-ray microscopy [21, 17].

The increasing utilization of these new methods to map out polycrystalline microstructures in 3D has spurred newfound interest in developing mathematical models for space-filling grain ensembles—*i.e.*, tessellations of space—with emphasis placed on matching not only the statistically averaged grain morphology in real microstructures, but also their grain connectivities and local environments. A popular starting point for such modelling efforts is the Laguerre tessellation (see [15] or Appendix A of the present paper), which is a generalization of the Voronoi construction, with weighting factors assigned to “seed points” offering rather flexible control of the resulting cell sizes. As with the conventional Voronoi diagram, the Laguerre tessellation consists of a set of non-overlapping convex cells that fill space.

Because Laguerre tessellations are simple to define and construct mathematically, they are attractive for the modelling of grain ensembles and their dynamics. For example, static polycrystalline structures have been represented by Laguerre tessellations based on close-packed spheres with randomly assigned radii [3], and the coarsening of grains has been modelled and simulated using 2D and 3D Laguerre tessellations [24, 22, 23, 28]. Furthermore, Laguerre tessellations have been applied to the modelling of other materials, like open and closed-cell foams [11, 10]. Interestingly, it can be proven that every normal tessellation in 3D that consists entirely of convex cells is a Laguerre tessellation [9]. A *normal* tessellation is one in which adjacent cells are juxtaposed face-to-face, sharing not only a face, but also its edges and vertices; furthermore, each planar face is contained in exactly two cells, each edge borders exactly three cells, and each vertex is shared by exactly four cells. These conditions are typically satisfied by real polycrystalline microstructures, as well, despite the fact that most grains are slightly non-convex thanks to the presence of curved boundaries. If we approximate the latter by planar surfaces bounded by lines connecting the corresponding grain vertices, then we can map any such real microstructure to a convex tessellation and, therefore, to a Laguerre tessellation.

Although it is straightforward to compare the statistical properties of cell sizes, shapes and environments in a given Laguerre tessellation to the same properties evaluated for the grains of a real polycrystalline specimen, it is no easy task to determine the Laguerre tessellation that corresponds most closely to a particular

grain mapping. This would entail finding the set of seed points and weighting factors that, when inserted into the Laguerre construction, leads to a network of grain boundaries best matching the one measured experimentally. But in a Laguerre tessellation, there is no guarantee that the seed point for a given cell actually lies within that cell, as the seed point location depends strongly on the positions and weighting factors of all nearby seed points, which are also unknown. To solve this optimization problem, one could, for example, devise a strategy for identifying and iteratively varying the seed point locations and weighting factors until a predetermined convergence criterion is fulfilled. We refer to the output of such an approach as a *Laguerre approximation* to the experimental microstructure. A prescription for constructing a Laguerre approximation was formulated recently by Lyckegaard *et al.* [12], who proposed starting from the centroids and volume-equivalent radii of individual grains and then applying an additional optimization step. The resulting representation faithfully reproduces the statistical properties of the experimental microstructure, but, according to the authors of [12], their Laguerre approximation is not always able to capture local grain configurations precisely. While this method aims primarily at reconstructing the full morphology when only grain centres and grain volumes are known, it is also applicable to voxelated data.

Lacking a satisfactory solution to this complex optimization problem, we propose an alternative route to the extraction of a set of convex polyhedra that closely matches the grain mapping of any given real polycrystal. Our approach offers most of the advantages of the Laguerre tessellation without the burden of having to find seed points and weighting factors for the individual cells. In our method, grains are described parametrically, with the individual faces obtained by orthogonal regression of planes; subsequently, the cells are constructed by an appropriate combination of the planar faces. Just like the Laguerre tessellation, our parametric representation of the grains affords significant data reduction, while facilitating the estimation of structural characteristics for individual grains. The proposed algorithm is explained in Section 2. We test the algorithm by applying it to artificial 3D image data and to a real microtomography data set obtained from an Al-5 wt% Cu alloy, comparing the algorithm's performance to the Laguerre approximation devised by Lyckegaard *et al.* [12] (Section 3). We observe that both our approach and the Laguerre approximation successfully reproduce statistical features of the real ensemble of grains, such as the distributions of grain sizes and coordination numbers, but our algorithm delivers a significantly closer match to local grain properties, such as the grain size, shape and neighbourhoods, at the cost of a slightly more complex parametrization.

2 Representation of grains by parametric cells

In this section, we describe a new algorithm for extracting convex cells directly from 3D image data. It is assumed that the voxels belonging to each grain are labelled with a unique number—that is, the processed image I is given by $\{I(x, y, z) \in \{1, \dots, N\} : (x, y, z) \in W\}$, where $N \geq 1$ denotes the number of grains and $W \subset \mathbb{N}^3$ is the (convex) grid of voxel coordinates; $\mathbb{N} = \{0, 1, \dots\}$. Furthermore, we assume that all voxels belonging to a given grain with label $i \in \{1, \dots, N\}$ —*i.e.*, $R_i = \{(x, y, z) \in W : I(x, y, z) = i\}$ —form a connected region.

A convex cell can be described by an intersection of half-spaces, in which the boundary of each half-space is aligned along a single planar face of the cell. By looking at the individual faces of a given grain—that is, the voxels bordering a grain and one of its neighbours—we can fit a plane to those voxels using orthogonal regression and then define a closed cell from the set of all planes delimiting the given grain. A 2D illustration of this idea is shown in Figure 1(a).




Fig. 1 Schematic illustration of the cell extraction algorithm in 2D: (a) orthogonal regression is used to fit lines to pixels at each boundary between two grains; taken together, the set of lines surrounding the grain defines the extracted cell. (b) Implementation of this idea in 2D using half-planes: orthogonal regression is used to fit lines to pixels (shown here only for two grain boundaries), the corresponding half-planes are computed (grey), and the intersection of all half-spaces (dark grey) is taken to define the cell. The filled circle denotes the grain’s centroid, which must be located within the half-planes.

A prerequisite for this approach to succeed is that the grains under consideration be at least nearly convex. For convex grains having planar faces, the algorithm works perfectly, in the sense that the precision is limited only by the voxel resolution. Non-planar faces between grains are approximated by planar faces; the presence of non-convex grains can introduce artefacts, as discussed in Section 3.3.

We now describe the extraction of planes for individual grain faces and the construction of cells from sets of planar faces.

2.1 Extraction of a plane for an individual grain face

In a typical polycrystalline microstructure, the voxels in simultaneous contact with two grains i and j are located approximately along a plane in 3D. We wish to detect this plane using orthogonal regression. The plane $P_{\mathbf{n},d} \subset \mathbb{R}^3$ parametrized with a (unit-length) normal vector $\mathbf{n} = (n_1, n_2, n_3) \in \mathbb{R}^3$ and signed distance $d \in \mathbb{R}$ from the origin is given by

$$P_{\mathbf{n},d} = \{(x, y, z) \in \mathbb{R}^3 : n_1x + n_2y + n_3z + d = 0\}.$$

The set of voxels adjacent to two grains with labels i and j is

$$N_{i,j} = \{(x, y, z) \in W : \mathcal{N}_{26}(x, y, z) \cap R_i \neq \emptyset \text{ and } \mathcal{N}_{26}(x, y, z) \cap R_j \neq \emptyset\},$$

where $\mathcal{N}_{26}(x, y, z)$ denotes the set of 26 voxels that share a face, edge or corner with the voxel at (x, y, z) (*i.e.*, the voxel neighbourhood defined by a distance less than or equal to $\sqrt{3}$ from (x, y, z)). Note that $N_{i,j} = N_{j,i}$, which ensures that the detected planes do not depend on the order of processing.

The idea of orthogonal regression is to minimize the (squared) distance of points (in our case the voxel coordinates in $N_{i,j}$) from the plane that is to be detected. Formally, we want to find the (global) minimum of the function $f : \mathbb{R}^3 \times \mathbb{R} \mapsto [0, \infty)$ denoted by

$$f(\mathbf{n}, d) = \sum_{(x,y,z) \in N_{i,j}} (n_1x + n_2y + n_3z + d)^2$$

for a normal vector $\mathbf{n} = (n_1, n_2, n_3) \in \mathbb{R}^3$ with length $|\mathbf{n}| = 1$ and signed distance $d \in \mathbb{R}$ from the origin. Orthogonal regression is therefore a least-squares problem, which can be solved by singular value decomposition [4]. In essence, for a given set of boundary voxels, it can be shown that the centroid $\mathbf{c} \in \mathbb{R}^3$ of the voxels lies on the plane to be extracted, and a normal vector $\mathbf{n} \in \mathbb{R}^3$ to the plane is given by the right-singular vector corresponding to the smallest singular value of a matrix containing all voxel coordinates shifted by $-\mathbf{c}$ [4].

2.2 Determination of grains sharing a common face

We first consider the issue of deciding which pairs of grains R_i and R_j share a common face, since this determines whether a plane between the grains should be extracted or not. Note that the rule $N_{i,j} \neq \emptyset$ does not suffice for this purpose, because there may be voxels touching two grains that share only a common vertex or edge. Fitting a plane to these voxels has no physical justification and would likely generate odd and unpredictable cell shapes. To avoid this situation, we apply a threshold criterion, with the value of the threshold determined from an inspection of the data in question. For the sample data sets discussed in Section 3, we interpret the square root of the number of voxels in $N_{i,j}$ as an “effective face diameter,” under the provisional assumption that the voxels in $N_{i,j}$ arise primarily from a shared grain face. We then calculate the ratio between the effective face diameter and the mean volume-equivalent diameter of the two grains being evaluated. If this ratio is larger than the threshold value—typically in the range of 0.2 to 0.25—then we accept the voxels in $N_{i,j}$ as belonging to a grain face, and we apply the algorithm of Section 2.1 to extract the corresponding best-fit plane.

2.3 Cell representation by the intersection of half-spaces

As explained above, we can fit a plane to each face of a grain R_i , and each of these planes restricts the volume of the corresponding extracted cell; consequently, the cell is a convex polytope. Convex polytopes can be defined by their vertices (the polytope itself is then given by the convex hull of the vertices, sometimes called the V-representation) or by an intersection of half-spaces (H-representation). The latter is directly applicable to our approach, because the appropriate half-spaces are directly related to the detected planes.

A half-space $H_{\mathbf{n},d}$ in \mathbb{R}^3 is given by

$$H_{\mathbf{n},d} = \{(x, y, z) \in \mathbb{R}^3 : n_1x + n_2y + n_3z + d \leq 0\}, \quad (\star)$$

where $\mathbf{n} = (n_1, n_2, n_3) \in \mathbb{R}^3$ is a (unit-length) normal vector, and $d \in \mathbb{R}$ is the signed distance from the origin. It is straightforward to obtain these half-spaces from the previously detected planes $P_{\mathbf{n},d}$, which are the boundaries of the half-spaces. Of course the proper side for the half-space must also be determined; here, a simple criterion is to require the half-space in (\star) to include the centroid of R_i —*i.e.*, the grain centroid fulfils the inequality given in the definition of $H_{\mathbf{n},d}$.

Finally, following the steps previously described, we can extract planes for all faces of a given grain, determine the corresponding half-spaces and obtain a closed cell by taking their intersection. A schematic illustration is shown in Figure 1(b). For m half-spaces of a single grain with $m \geq 3$ and $\{(\mathbf{n}^{(k)}, d^{(k)}) \in \mathbb{R}^3 \times \mathbb{R}, k = 1, \dots, m\}$, their intersection and thus the convex cell C is given by

$$C = \{(x, y, z) \in \mathbb{R}^3 : n_1^{(k)}x + n_2^{(k)}y + n_3^{(k)}z + d^{(k)} \leq 0, k = 1, \dots, m\}.$$

3 Evaluation

We evaluate the quality of the proposed cell extraction algorithm by computing structural characteristics. First, using the standard Laguerre construction (see Appendix A), we generate an artificial data set consisting of a space-filling ensemble of convex cells (taking on the role of the grains in a polycrystalline microstructure), to which we apply our algorithm in order to test the accuracy of the resulting cell extraction. Second, we apply our algorithm to the grains of a real polycrystalline material mapped experimentally by x-ray microtomography.

In both cases, we compare our approach to the Laguerre approximation introduced by Lyckegaard *et al.* [12]. In the latter, seed points are placed initially at the centroids of individual grains, and a weighting factor equal to the volume-equivalent radius of the associated grain is assigned to each seed point. An additional optimization step is carried out to improve the result by shifting the location of seed points depending on adjacent cells. From the analysis carried out in [12], Lyckegaard *et al.* conclude that this cell-extraction algorithm reproduces statistical features like the distribution of cell sizes and the distribution of the number of faces quite well, but discrepancies can appear in more-sophisticated local characteristics, such as the sizes and shapes of cell neighbourhoods.

Throughout Section 3, when estimating structural characteristics we employ the Miles–Lantuéjoul correction [14, 8] to avoid bias arising from grains touching the boundary of the grid, and, whenever possible, we prefer to calculate density functions rather than histograms, as the shape of the latter is highly sensitive to the chosen bin widths and positions. In this work, we determine density functions by (non-parametric) kernel density estimation [20].

3.1 Artificial data set with convex cells

Here, we describe the construction of an artificial polycrystalline microstructure as a realization of a random Laguerre tessellation. We then apply our cell-extraction

algorithm and compare the resulting cells to those of the original known tessellation.

3.1.1 Generation of the artificial data set

A random Laguerre tessellation is generated by considering a random marked point process [6], the points of which are taken to be seed points for Laguerre cells. Numerical marks are assigned randomly to the points, constituting weighting factors that influence the resulting sizes of the Laguerre cells. More precisely, we employ a Poisson-Laguerre tessellation [9], which is based on the well-known (homogeneous) Poisson point process. In particular, the point pattern is generated as follows. For a given intensity λ of the homogeneous Poisson process, the number of points N in a cuboid $W \subset \mathbb{R}^3$ is Poisson-distributed with expectation value given by the product of λ and the volume of W . For a realization n of the number of points, the point pattern itself is denoted by the set of coordinates $\{(x_i, y_i, z_i), i = 1, \dots, n\}$; these points are realized by a uniform distribution on W , and they do not influence each other.

We choose the cuboid for simulation to have the size $W = [0, 1000]^3$ and λ to be 5×10^{-6} , which means that the expected number of points in W is 5000. The random marks assigned to the points are independent and distributed according to a gamma distribution with shape parameter $\alpha = 50$ and rate parameter $\beta = 0.75$ (expectation value α/β and variance α/β^2) [5]. This combination of parameter values was found to generate cells similar to the experimental data presented in Section 3.2. A 2D section of a realization of this random Laguerre tessellation is shown in Figure 2.




Fig. 2 Cross-section through a 3D artificial data set generated by Poisson-Laguerre tessellation.

3.1.2 Extraction of cells and evaluation

Given that the Poisson-Laguerre tessellation described above has been discretized to a cubic grid, it is straightforward to apply the cell-extraction algorithm proposed

in Section 2. For the sake of comparison, we also compute the Laguerre approximation of Lyckegeard *et al.* [12] for the same artificial data set. Figure 3 shows the original tessellation overlaid with cells extracted by the two methods. To quantify the degree of agreement between the extracted and original cells, we compare them at the voxel level. In light of the near-perfect match between the boundary voxels extracted by our approach and those in the original data set, it is clear that all structural characteristics, like cell sizes, cell faces and cell neighbourhoods, are also identical.




Fig. 3 Cross-section through the original Poisson-Laguerre tessellation, with boundary voxels shaded grey. Superimposed in black are the boundaries of cells extracted by (a) orthogonal regression and (b) Laguerre approximation [12].

In order to quantify the number of voxels correctly assigned, we examine the set of voxels of an extracted cell \hat{C} , which should be nearly identical to that of the corresponding original cell C . We define the fraction of correct voxels as the number of voxels in the intersection $\hat{C} \cap C$ divided by the number of voxels in C . The fraction of correct voxels is a number lying between 0 and 1, with values near unity implying a nearly perfect match to the original voxels. Note that an extracted cell that happens to be larger than the original cell can also yield a high value for this fraction, even though the extracted cell is not a perfect fit, but in this case the adjacent cells will automatically manifest much lower values for the fraction of correct voxels.

Figure 4 shows the estimated density function of the fraction of correct voxels. For cells constructed by orthogonal regression, the mean fraction value is about 99%, which constitutes a very good overall fit. This is not surprising, because our algorithm is constructed to match the individual faces as closely as possible with planes, and the artificial data set consists entirely of convex cells having planar faces. The Laguerre approximation [12], on the other hand, is generated by a simpler parametrization. Because the original data set is itself a Laguerre tessellation, it is clear that the cells extracted by the algorithm in [12] could have conceivably matched perfectly throughout the sample, but the fact that the mean fraction of correct voxels was only 84% for the Laguerre approximation indicates

that determination of the proper seed point locations and weighting factors from cell boundaries alone is a non-trivial endeavour.

The voxel-based comparisons presented in Figures 3 and 4 indicate that for (in our sense) optimal data, the extraction of cells by orthogonal regression works nearly perfectly. Despite the artificial data set having been generated by the Laguerre construction, the Laguerre approximation [12] does not describe the data nearly as well. For the experimental data set considered in the next section, we extend our analysis beyond the voxel level to include characteristics like the sizes and shapes of grains and their neighbourhoods.




Fig. 4 Estimated density function plotted against the fraction of voxels correctly assigned in the artificial data set for cells extracted by orthogonal regression and by Laguerre approximation [12].

3.2 Experimental data set of a polycrystalline alloy

After a short description of the method by which 3D microstructural data were obtained for a polycrystalline alloy, we examine the cells determined by orthogonal regression and present various measures for quantifying the extent to which the extracted cells properly represent the real microstructure.

3.2.1 Data description

Cylindrical samples of an Al-5 wt% Cu alloy (4 mm diameter, 4 mm height) were characterized using a SkyScan 1172 laboratory tomograph at an isotropic resolution of $2 \mu\text{m}$ (voxel side length). Absorption-contrast microtomography yields a three-dimensional mapping of the local absorption of x-rays, variations in which can arise, for instance, from concentration gradients present at the boundaries between different phases. When Al-Cu alloys with low copper content are heated above the solidus temperature, phase separation occurs into solid, Al-rich grains surrounded by a liquid phase having a significantly higher concentration of Cu atoms [13,16]. Upon subsequent cooling to room temperature, the liquid phase crystallizes into an eutectic mixture of two solid phases, but the local Cu concentration in the formerly liquid regions remains enhanced, effectively marking the boundaries of the Al-rich grains in the tomographic reconstruction (Figure 5(a)). Because coverage of the grains by the eutectic mixture was incomplete, we applied

a watershed image-processing algorithm to a Euclidean distance transform of the binary images to fill in gaps in the network of grain boundaries (Figure 5(b)) [2, 18]. Additional details concerning the measurement and data segmentation procedures can be found in [27].




Fig. 5 (a) Cross-section through a tomographic reconstruction of an Al-5 wt% Cu alloy measured after cooling from above the solidus down to room temperature; lighter regions indicate the presence of the phase Al_2Cu , which forms during crystallization of the liquid phase surrounding the solid, Al-rich grains. (b) Same cross-section following binarization and application of the watershed transformation.

3.2.2 Extraction of cells and evaluation

As mentioned above, the watershed transformation was used to reconstruct the network of grain boundaries, thereby uniquely identifying connected sets of voxels as belonging to the interior regions of individual grains in the sample. This is precisely the starting point considered in Section 2 for the extraction of cells from a 3D mapping of a polycrystalline microstructure. Application of this method to the Al-5 wt% Cu data set yields the cell system illustrated in Figure 6(a) in black, which is superimposed on the boundary network of the original tomographic data plotted in grey.

To assess the accuracy of cell extraction, we again consider the fraction of correct voxels, as defined in Section 3.1; the estimated density function is plotted in Figure 7. As with the artificial data set, we see that orthogonal regression applied to data from a real sample yields very high values for the mean fraction of correct voxels (0.96), with only 5% of all cells having a value smaller than 0.9. The Laguerre approximation [12] achieves a mean fraction of about 0.87, but 5% of the cells manifest a value below 0.64. This quantitative discrepancy is reflected in obvious qualitative differences between the overlays shown in Figure 6.

3.2.3 Comparison of further structural characteristics

In this section, we consider important structural characteristics of polycrystalline microstructures with respect to grain sizes, shapes and neighbourhoods, compar-




Fig. 6 Cross-section through a tomographic reconstruction of Al-5 wt% Cu following image processing as described in the text; voxels at grain boundaries are shaded grey. Superimposed in black are the boundaries of cells extracted by (a) orthogonal regression and (b) Laguerre approximation [12].




Fig. 7 Estimated density function plotted against the fraction of voxels correctly assigned in the Al-5 wt% Cu data set for cells extracted by orthogonal regression and by Laguerre approximation [12].

ing the corresponding values for cells extracted by orthogonal regression and by Laguerre approximation of the same experimental Al-5 wt% Cu data set.

A natural structural characteristic is the distribution of grain sizes. Although the grains themselves are not perfectly spherical, it is common to equate the size of a given grain to the diameter of a sphere having the same volume. The estimated density functions for grain sizes computed in this manner are plotted in Figure 8(a). We see there that all of the distributions take on a similar shape, but some discrepancies are apparent between the experimental data set and the Laguerre approximation, particularly with respect to smaller grains, which tend to be represented by Laguerre cells that are somewhat too large. The near-linearity of a scatter plot of the diameter of each extracted cell against its corresponding tomographic grain diameter attests to the high accuracy our cell extraction algorithm (Figure 8(b)). The few outliers can be attributed to the presence of non-convex cells in the real microstructure (see Section 3.3).

Another interesting structural characteristic is the sphericity of grains [26]. *Sphericity* is defined to be the ratio of the surface area of the volume-equivalent

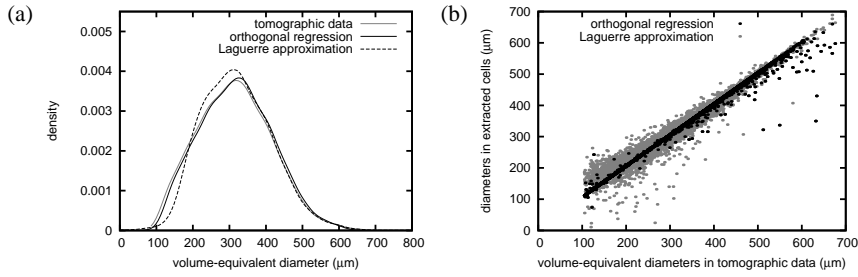


Fig. 8 (a) Estimated density functions for the volume-equivalent grain diameter: distributions calculated directly from tomographic data as well as from cells extracted by orthogonal regression or Laguerre approximation [12]. (b) Scatter plot of the volume-equivalent diameter of extracted cells vs. the corresponding grain diameter measured by tomography.

sphere to the surface area of the corresponding grain or cell. Small values for this ratio imply that the shape differs significantly from that of a sphere—*e.g.* the true shape could be elongated in a particular direction. As shown in Figure 9, neither orthogonal regression nor the Laguerre approximation [12] does a satisfactory job of representing the sphericity of the experimental grains: the mean sphericity is approximately equal in all three cases, but the distribution of sphericities is wider for both types of extracted cells than for the original grains. Detailed investigation reveals that discrepancies are particularly noticeable for small grains, which tend to be rather spherical in the experimental data set. Consequently, the approximation of planar faces (which is enforced by the cell extraction algorithms) can be rather poor for these grains (particularly in light of the fact that small grains have few faces), and the sphericity of the cells extracted for these grains never approaches unity. Therefore, the differences in sphericity that are clearly visible in Figure 9 can be attributed primarily to the constraint of convexity on each cell in the extracted cell network (which, in turn, entails that all cell boundaries be planar).

A further important characteristic of space-filling ensembles of grains is the coordination number, which denotes the number of neighbour grains sharing a common face with a given grain. Figure 10(a) compares the relative frequency of coordination numbers determined directly from experimental data to the same quantity evaluated for microstructural representations produced by the two cell-extraction algorithms. For the tomographic data set, the identification of shared grain faces was performed using the algorithm described in Section 2.2.

When assessing grain adjacency, it is of interest to determine whether the topology of the system is represented properly. Figure 10(b) presents histograms for the number of incorrectly assigned neighbours, which is determined by counting all differences in the adjacency list—*i.e.*, by summing up the number of missing neighbours and the number of additional neighbours. As already discussed in [12], the cells extracted by Laguerre approximation are not perfect with respect to the local arrangement of cells. Cell extraction by orthogonal regression manifests similar deficiencies in the proper representation of this particular structural characteristic, although they are somewhat less severe than in the case of Laguerre approximation.




Fig. 9 Grain sphericity: (a) tomographic data acquired from Al-5 wt% Cu; (b) cells extracted by orthogonal regression; (c) cells extracted by Laguerre approximation [12]. In each row, the density function for sphericity is plotted on the left, and the dependence of sphericity on the grain/cell size is illustrated graphically on the right.

For both cell-extraction routines, erroneous neighbour assignments can be traced to the constraint of cell convexity. When the orthogonal regression of planes is applied to the boundaries of a non-convex grain, it is possible for a neighbour to be lost when the voxels of the corresponding shared face do not generate a half-space that restricts the volume of the resulting cell. Approaches based on Laguerre tessellations do not suffer from this particular problem, but there—as we have seen above—it is much more difficult to obtain a good overall fit to the individual grain faces, which may, in turn, lead to incorrectly assigned neighbours. In the following section, we discuss in greater detail the cell-extraction artefacts that can arise from non-convex grains.

3.3 Influence of non-convexity

The proposed algorithm aims to obtain an exact representation of individual grains. For an ensemble of convex grains with planar faces, our cell-extraction method

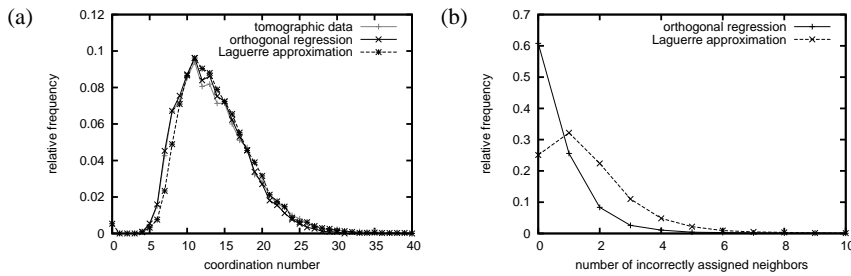


Fig. 10 (a) Distribution of the coordination number of grains in tomographic image data as well as of cells extracted by orthogonal regression and Laguerre approximation [12]. (b) Distribution of the number of incorrectly assigned neighbours: comparison of cells extracted by orthogonal regression to Laguerre approximation, evaluated in both cases with respect to grain neighbour assignments determined from experimental data.

works perfectly. However, the orthogonal regression approach does not explicitly enforce constraints like the face-to-face adjacency of cells or the matching of edges and vertices; consequently, it is highly likely that the conditions for normality of a tessellation will be violated to a certain extent when our algorithm is applied to a microstructure containing non-convex grains. We now discuss two types of artefacts that can arise from non-convexity.

First, our algorithm works by fitting a plane to the boundary between two grains. This ensures that we obtain a good fit with respect to the position of the boundary itself, but the approach places no constraints on the locations of the edges and vertices delineating the shared face. Consider, for example, the 2D illustration in Figure 11(a). Grain number 1 shares a curved boundary with grain 3 below, which is clearly non-convex. By fitting lines to this boundary and to the boundary between grains 1 and 2, we obtain the vertex labelled V . The same procedure applied to the boundaries of grain 2, however, yields a different vertex position V' for the same line separating cells 1 and 2. Such non-matching of vertices in 2D—or edges and vertices in 3D—will generally result in small gaps—*i.e.* regions not covered by any extracted cell—or even in small overlaps between neighbouring cells. Examination of our data found that approximately 97% of all voxels were included within an extracted cell, but this number also encompasses artefacts of the second type, which we discuss next.

A second type of artefact occurs when a detected half-space cuts off too much volume from a cell. This can be a particularly serious problem when the segmentation of an experimental data set fails to detect a boundary between two grains, falsely grouping the voxels together as a single grain. The union of the two grains is usually highly non-convex, frequently taking on a “dog-bone” shape. Then, because each planar face is detected individually and used to construct the resulting cell by taking the intersection of the corresponding half-spaces, some half-spaces may not restrict the volume of the cell at all, or, even worse, they may strongly affect the overall cell shape at the wrong location. This generally results in extracted cells that are significantly smaller than the corresponding grains, as illustrated in Figure 11(b). (Such an obvious loss of volume between the original grain and the extracted cell could be a useful tool for automatically detecting underlying segmentation errors.)

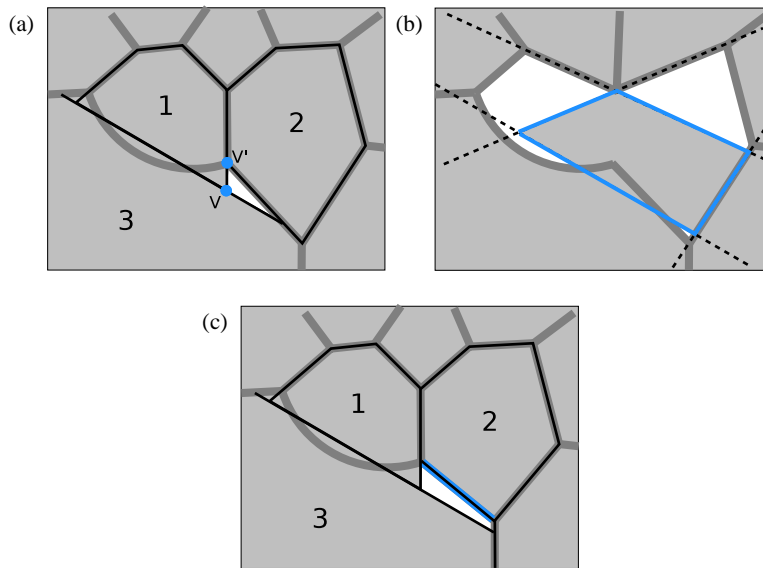


Fig. 11 Schematic illustration of artefacts caused by non-convex grains. (a) Edges and vertices of adjacent cells are not forced to be equal, which may result in small gaps (unshaded white region). (b) Two grains recognized as one in the segmentation cannot be represented properly by a convex cell (detected cell bordered in solid blue lines, restricting planes dashed, lost volume in white); (c) A planar face can be “cut off” from one grain when another face more strongly restricts the volume of the extracted cell (the face marked in blue has no effect on the spatial extent of cell 3); since the former face may still be relevant for an adjacent cell (cell 2 in the illustrated example), a gap can form between the extracted cells.

The misclassification of two grains as a single grain is clearly a fault of segmentation and not of cell extraction. Nevertheless, the artefact of a half-space having no effect or even the wrong effect on the cell shape can occur under less-extreme circumstances, as well. The enforced approximation of a curved boundary by a planar face can generate a half-space that “cuts off” other fitted planar faces in such a manner that a “cut off” plane disappears as a face for one cell, but the same plane forms the boundary of an adjacent cell. Such a case is illustrated in Figure 11(c). Here, the plane fitted to the boundary between grains 2 and 3 is only relevant for cell 2, since, for cell 3, this plane is cut off by the half-space generated from the (non-planar) boundary between grains 1 and 3. This phenomenon can account for instances in which cells extracted by our algorithm have different neighbours than the corresponding grains in the experimental data set.

In spite of such artefacts, our strategy for extracting (parametric) cells performs quite well in general for experimental data corresponding to typical polycrystalline microstructures. As expected, the high accuracy noted in voxel-based comparisons carries over to statistical characteristics of the cell ensembles, as well. The Laguerre approximation [12] is likewise able to represent global statistical features faithfully, but it is less suited to the proper description of more sophisticated structural characteristics of space-filling grain ensembles. It should be noted that Laguerre tessellations have the clear advantage of, by definition, always generating a tessellation—*i.e.*, it is impossible for there to be gaps between

cells or for edges and vertices not to coincide. Unfortunately, it is these same constraints that make it so difficult to determine the “best” seed points and weighting factors for accurately approximating a measured polycrystalline microstructure by the Laguerre construction.

4 Conclusions

We have presented a new algorithm for extracting parametric cells from 3D image data that is based on the orthogonal regression of individual grain faces. Applied to an artificial data set consisting of convex grains, the cell-extraction algorithm generates a nearly perfect match, as quantified by a voxel-based comparison. For experimental 3D image data, the new algorithm also works quite well, performing significantly better in most respects than a recently proposed Laguerre approximation algorithm [12]. The latter conclusion is based on evaluations of statistical features like distributions of grain size and the number of neighbours (coordination number), but also on local characteristics like the number of incorrectly assigned neighbours per grain. The higher accuracy in representation comes at the cost of a slightly more complex parametrization. It should be noted that voxelated (image) data are needed for our approach, while [12] requires only grain centres and volumes, which may be easier to obtain. A further disadvantage is the possibility for artefacts like small gaps between adjacent cells in the case of clearly non-convex grains. The latter problem was negligible for the real microstructure studied in this work.

Acknowledgements The authors would like to thank D. Molodov of the Institute of Physical Metallurgy and Metal Physics, RWTH Aachen, for sample preparation; the Institute of Orthopaedic Research and Biomechanics, Ulm University, for x-ray microtomography beamtime; and especially Uwe Wolfram for assistance with the tomography measurements and numerous discussions. Furthermore, the authors are grateful to the Deutsche Forschungsgemeinschaft for funding through NSF/DFG Materials World Network Project KR 1658/4-1.

A Laguerre tessellations

The Laguerre tessellation [15]—a weighted version of the well-known Voronoi diagram—has long been a popular choice for modelling polycrystalline grain structures (see *e.g.* [3,24,22,23,28]). In a Voronoi diagram, each “seed point” (also called a “spring point”) creates exactly one cell, as every point in space is assigned to its nearest seed point. In 3D, the result is a partition of space into convex polyhedra. When constructing a Laguerre tessellation, every seed point is given an additional weight r^2 , which permits finer control over the cell sizes. The weight r^2 can be interpreted as the squared radius of a sphere centred on the seed point—see Figure 12 for an illustration in 2D.

Formally, a 3D Laguerre tessellation is defined as follows. Given a (locally finite) set $S = \{(\mathbf{x}_i, r_i), i \in I\} \subset \mathbb{R}^3 \times \mathbb{R}^+$ of seed points \mathbf{x}_i with radii r_i , the Laguerre cell of (\mathbf{x}_i, r_i) with respect to S is given by

$$C((\mathbf{x}_i, r_i), S) = \left\{ \mathbf{y} \in \mathbb{R}^3 : |\mathbf{y} - \mathbf{x}_i|^2 - r_i^2 \leq |\mathbf{y} - \mathbf{x}_j|^2 - r_j^2, (\mathbf{x}_j, r_j) \in S \right\}.$$

Then, the Laguerre tessellation is the set of all Laguerre cells $\{C((\mathbf{x}_i, r_i), S), i \in I\}$. Note that it is possible for a seed point to create no cell at all, provided an adjacent seed point has a sufficiently large weight. For the same reason it is also possible for a seed point not to be contained within




Fig. 12 Illustration of Laguerre tessellation in 2D: seed points with weights pictured as circles, along with the resulting Laguerre cell boundaries.

the cell generated by that seed point. These properties complicate the determination of the seed point locations and corresponding weights when only the resulting cells are known.

References

1. Aagesen, L.K., Fife, J.L., Lauridsen, E.M., Voorhees, P.W.: The evolution of interfacial morphology during coarsening: A comparison between 4D experiments and phase-field simulations. *Scr. Mater.* **64**(5), 394–397 (2011)
2. Brunke, O., Odenbach, S., Beckmann, F.: Quantitative methods for the analysis of synchrotron- μ CT datasets of metallic foams. *Eur. Phys. J.-Appl. Phys.* **29**(1), 73–81 (2005)
3. Fan, Z., Wu, Y., Zhao, X., Lu, Y.: Simulation of polycrystalline structure with Voronoi diagram in Laguerre geometry based on random closed packing of spheres. *Comp. Mater. Sci.* **29**, 301–308 (2004)
4. de Groen, P.P.N.: An introduction to total least squares. *Nieuw Archief voor Wiskunde, 4th Series* **14**, 237–253 (1996)
5. Hogg, R.V., Craig, A.T.: *Introduction to Mathematical Statistics*, 4th edn. Macmillan, New York (1978)
6. Illian, J., Penttinen, A., Stoyan, H., Stoyan, D.: *Statistical Analysis and Modelling of Spatial Point Patterns*. J. Wiley & Sons, Chichester (2008)
7. Konrad, J., Zaeferrer, S., Raabe, D.: Investigation of orientation gradients around a hard Laves particle in a warm-rolled Fe3Al-based alloy using a 3D EBSD-FIB technique. *Acta Mater.* **54**(5), 1369–1380 (2006)
8. Lantuéjoul, C.: On the estimation of mean values in individual analysis of particles. *Microsc. Acta* **5**, 266–273 (1980)
9. Lautensack, C.: *Random Laguerre Tessellations*. Ph.D. thesis, Universität Karlsruhe, Weiler bei Bingen (2007)
10. Lautensack, C.: Fitting three-dimensional Laguerre tessellations to foam structures. *J. Appl. Stat.* **35**(9), 985–995 (2008)
11. Lautensack, C., Sych, T.: 3D image analysis of open foams using random tessellations. *Image Anal. Stereol.* **25**, 87–93 (2006)
12. Lyckegaard, A., Lauridsen, E.M., Ludwig, W., Fonda, R.W., Poulsen, H.F.: On the use of Laguerre tessellations for representations of 3D grain structures. *Adv. Eng. Mater.* **13**(3), 165–170 (2011)
13. Massalski, T.B.: *Binary Alloy Phase Diagrams - AlCu-Phasediagram*, vol. 1, 3rd edn. ASM International (1996)
14. Miles, R.E.: On the elimination of edge effects in planar sampling. In: E.F. Harding, D.G. Kendall (eds.) *Stochastic Geometry*, pp. 228–247. J. Wiley & Sons, New York (1974)
15. Okabe, A., Boots, B., Sugihara, K., Chiu, S.N.: *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd edn. J. Wiley & Sons, Chichester (2000)

-
16. Pompe, O., Rettenmayr, M.: Microstructural changes during quenching. *J. Cryst. Growth* **192**(1-2), 300–306 (1998)
 17. Reischig, P., King, A., Nervo, L., Viganó, N., Guilhem, Y., Palenstijn, W.J., Batenburg, K.J., Preuss, M., Ludwig, W.: Advances in X-ray diffraction contrast tomography: flexibility in the setup geometry and application to multiphase materials. *J. Appl. Crystallogr.* **46**(2), 297–311 (2013)
 18. Roerdink, J.B.T.M., Meijster, A.: The watershed transform: definitions, algorithms, and parallelization strategies. *Fundam. Inform.* **41**, 187–228 (2001)
 19. Salvo, L., Cloetens, P., Maire, E., Zabler, S., Blandin, J.J., Buffière, J.Y., Ludwig, W., Boller, E., Bellet, D., Josserond, C.: X-ray micro-tomography an attractive characterisation technique in materials science. *Nucl. Instrum. Methods Phys. Res. B* **200**, 273–286 (2003)
 20. Silverman, B.W.: *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, London (1986)
 21. Sorensen, H.O., Schmidt, S., Wright, J.P., Vaughan, G.B.M., Techert, S., Garman, E.F., Oddershede, J., Davaasambuu, J., Paithankar, K.S., Gundlach, C., Poulsen, H.F.: Multigrain crystallography. *Z. Kristallogr.* **227**(1), 63–78 (2012)
 22. Telley, H., Liebling, T.M., Mocellin, A.: The Laguerre model of grain growth in two dimensions: I. Cellular structures viewed as dynamical Laguerre tessellations. *Philos. Mag. B* **73**(3), 395–408 (1996)
 23. Telley, H., Liebling, T.M., Mocellin, A.: The Laguerre model of grain growth in two dimensions: II. Examples of coarsening simulations. *Philos. Mag. B* **73**(3), 409–427 (1996)
 24. Telley, H., Liebling, T.M., Mocellin, A., Righetti, F.: Simulating and modelling grain growth as the motion of a weighted Voronoi diagram. *Mater. Sci. Forum* **94–96**, 301–306 (1992)
 25. Uchic, M.D., Holzer, L., Inkson, B.J., Principe, E.L., Munroe, P.: Three-dimensional microstructural characterization using focused ion beam tomography. *MRS Bull.* **32**, 408–416 (2007)
 26. Wadell, H.: Volume, shape and roundness of quartz particles. *J. Geol.* **43**(3), 250–280 (1935)
 27. Werz, T., Baumann, M., Wolfram, U., Krill III, C.E.: Particle tracking during Ostwald ripening using time-resolved laboratory x-ray microtomography (2013). Submitted.
 28. Xue, X., Righetti, F., Telley, H., Liebling, T.M.: The Laguerre model for grain growth in three dimensions. *Philos. Mag. B* **75**(4), 567–585 (1997)