



HAL
open science

**Optimal polyhedral description of 3D polycrystals:
method and application to statistical and synchrotron
X-ray diffraction data**

Romain Quey, Loïc Renversade

► **To cite this version:**

Romain Quey, Loïc Renversade. Optimal polyhedral description of 3D polycrystals: method and application to statistical and synchrotron X-ray diffraction data. *Computer Methods in Applied Mechanics and Engineering*, In press, 330, pp.308-333. 10.1016/j.cma.2017.10.029 . hal-01626440

HAL Id: hal-01626440

<https://hal.science/hal-01626440>

Submitted on 30 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal polyhedral description of 3D polycrystals: method and application to statistical and synchrotron X-ray diffraction data

Romain Quey¹ and Loïc Renversade¹

¹*Univ Lyon, IMT Mines Saint-Etienne, Centre SMS, CNRS, LGF, F – 42023 Saint-Etienne FRANCE*

Abstract

A methodology is presented for optimal polyhedral description of 3D polycrystals from experimental properties. This is achieved by determining, by optimization, appropriate attributes of the seeds of Laguerre tessellations. The resulting tessellations are optimal in the sense that no further improvements are possible using convex geometries. The optimization of Laguerre tessellation combines a new, computationally-efficient algorithm for updating tessellations between iterations to a generic optimization algorithm. The method is applied to different types of experimental data, either statistical, such as grain size distributions, or grain-based, as provided by synchrotron X-ray diffraction experiments. It is then shown how the tessellations can be meshed for finite-element simulations. The new method opens the way to more systematic and quantitative analyses of microstructural effects on properties. The presented algorithms are implemented and distributed in the free (open-source) software package Neper.

Keywords: Microstructure, Laguerre tessellation, optimization, micromechanics

1 Introduction

The influence of the microstructure on the mechanical and physical properties of polycrystalline materials has been known for a long time. The average grain size, for example, is responsible for the so-called “Hall-Petch strengthening” [1, 2] and is important for fatigue resistance [3, 4]. The grain size distribution has similar effects, albeit generally less pronounced [5, 6]. Of course, these phenomena are directly related to the stress and strain heterogeneities developing at the microstructure scale [7]. In the past 15 years, large efforts have been conducted to analyse experimentally the local deformation of polycrystalline materials in 3D using synchrotron X-ray diffraction. These methods provide both the grain structure of polycrystalline materials [8, 9] and the local stresses and lattice rotations [10–13] that develop as they are deformed. Numerical methods are available to simulate these phenomena [14–22], e.g. the widely-used finite element method [14–20]. Simulations can then be compared to experiments [19, 20] or used alone to analyse systematically the influence of microstructural attributes, on condition that appropriate polycrystal models are available.

As the (morphological) properties of polycrystalline microstructures result from material’s chemistry and elaboration, they can be simulated using physical models, e.g. the phase field model [23, 24]. Geometrical models, in comparison, are conceptually simpler and computationally cheaper. A polycrystal can for example be represented by a packing of ellipsoids combined with simple growth [25]. These approaches generally provide microstructures of high fidelity but do not offer full control on them. Moreover, the microstructures are in general complex, as they can include general (non-convex) grain shapes which can only be described on grids of voxels, which is not adapted to standard, conformal meshing [26]. Such is also the case for polycrystal images provided by synchrotron X-ray diffraction [8, 9], which motivated the development of dedicated meshing procedures [27, 28]. In contrast, geometrical models such as Voronoi or Laguerre tessellations can be described in a compact, scalar fashion, using sets of points, lines, surfaces and volumes. A major advantage is that they can be meshed with relatively standard approaches [26], while their cells being necessarily convex is often a minor limitation for representing polycrystals [29]. Voronoi tessellations, in particular, have been widely used in numerical simulations in the past few years [7, 19, 20, 30–35] but only partially reproduce elementary polycrystal properties such as their grain size distribution [26].

Laguerre tessellations are a generalization of Voronoi tessellations and are becoming increasingly popular for modelling materials such as polycrystals [36, 37] or foams [38]. Similarly to Voronoi tessellations, Laguerre tessellations are generated from a set of seeds [39], but the seeds are more delicate to set, so that different strategies were developed for different applications. For reproducing grain size distributions, seeds are often determined from a dense packing of spheres [36, 37]. For foams, seeds can be determined from the geometry and arrangement of their walls [38]. For grain centroids and volumes provided by synchrotron X-ray diffraction experiments, seeds can also be determined using simple schemes [40]. These strategies are both application-specific and suboptimal, for example due to constraints on the seed attributes resulting from the correlation between spheres centres and radii in dense sphere packing [36, 37].

In this work, a general methodology is developed for generating polycrystal models from any type of experimental properties, using Laguerre tessellations. The approach is to consider the desired (experimental) grain properties as an input of the problem, from which the attributes of the seeds are determined. As will be explained in the following, the method provides, irrespective of the application, optimal results, i.e. the best possible tessellations matching the prescribed properties and made of convex cells. It should be stressed that under investigation in this article will be so-called “single-scale” polycrystals, such as those typically encountered in single-phase (and some multi-phase) materials, in contrast to polycrystals characterized by grain subdivision (often arising during elaboration), such as pearlite, bainite or martensite in steels, lamellar Ti-6Al-4V, etc. The article is organised as follows. We first describe the generic method for tessellation generation in Section 2. We then show how it can be applied to different types of experimental data, either statistical (Section 3) or grain-based such as those provided by synchrotron X-ray diffraction (Sections 4 and 5). We then discuss the results and show how the tessellations can be used in finite element simulations (Section 6), before closing the article with conclusions (Section 7). Throughout the article, we use a consistent terminology, where “experimental”, “polycrystal” and “grain” refer to input data which we aim to replicate, while “numerical”, “tessellation” and “cell” refer to output data supplied by the method. All algorithms described in this article are implemented and distributed in the free (open-source) software package Neper [41].

2 Method

The method is based on optimization of Laguerre tessellations. In Section 2.1, we recall the principle of Laguerre tessellations and introduce the general optimization framework in which we use them. In Section 2.2, we describe a method to update Laguerre tessellations during optimization.

2.1 Principle

2.1.1 Laguerre tessellation

Mathematically, a Laguerre tessellation is a partition of an nD space in a collection of nD entities that fill the space with no overlaps and no gaps. These entities are polyhedra and are formally defined as zones of influence of a particular set of weighted points, or seeds. Each seed, S_i , is defined by its coordinates, \mathbf{x}_i , and a non-negative weight, w_i . Being given a spatial domain $D \in \mathcal{R}^n$ and a set of seeds $\{S_i(\mathbf{x}_i, w_i), i = 1, \dots, N\}$, every seed S_i is associated a Laguerre cell C_i as follows,

$$C_i = \{P(\mathbf{x}) \in D \mid d_L(P, S_i) < d_L(P, S_j) \quad \forall j \neq i\} \quad \text{with} \quad d_L(P, S_\bullet) = d_E(P, S_\bullet)^2 - w_\bullet \quad (1)$$

d_L is referred to as “power distance” and d_E is the Euclidean distance. Under the effect of distinct seed weights, boundaries between cells become non-equidistant between seeds, and cells may not contain their associated seeds or be empty. The larger the weight of a seed (relative to the weights of other seeds), the bigger its associated cell. The special case of Voronoi tessellation is obtained when all weights are equal. Figure 1 provides an example of a 3D Laguerre tessellation generated from seeds of random positions and weights.

Laguerre tessellations are normal tessellations, i.e. are comprised of convex cells intersecting along planar faces, straight edges, and vertices [39]. From a topological point of view, Laguerre tessellations are equivalent to polycrystals, which are comprised of grains intersecting along grain boundaries, triple lines and quadruple points. An important property of Laguerre tessellations is that, in contrast to Voronoi tessellations, they can describe any normal tessellations, whatever the morphologies or arrangement of the cells [39]. More specifically, Laguerre tessellations can include contrasts in cell size that cannot be attained using Voronoi tessellations, and Laguerre cells can exhibit geometries which are not possible with Voronoi cells (due to the fact that the faces of Voronoi tessellations are necessarily equidistant between

seeds). As a consequence, Laguerre tessellations can be considered as both general and optimal to represent tessellations composed of convex cells. Of course, this may be verified only in the absence of constraints on (or across) seed attributes.

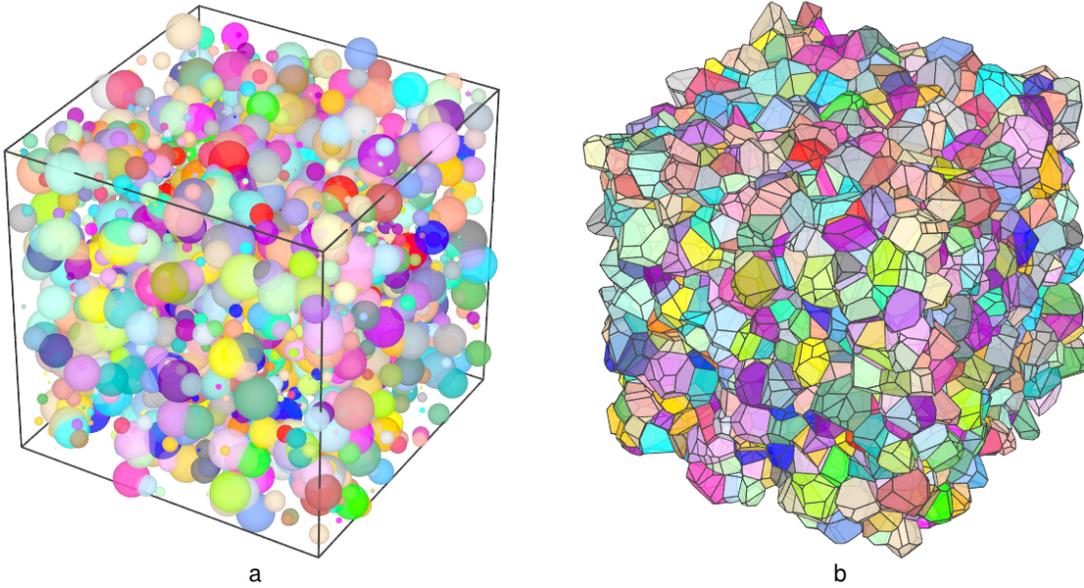


Figure 1: Example of a 3D Laguerre tessellation. (a) 3,000 weighted seeds with random and uncorrelated positions and weights. Seed coordinates are randomly distributed according to a uniform distribution in $[0, 1]$, while seed weights elevated to the power of 4 (w_i^4) are randomly distributed according to a uniform distribution in $[0, 0.25]$. Each seed is shown as a sphere of centre and radius equal to the seed position and square root of weight, respectively (note that spheres can overlap). (b) Associated Laguerre tessellation, which contains 2,895 non-empty cells. Periodicity conditions are applied. Note that large weights tend to generate big cells, but with noticeable variations between cells (depending on the local seed attributes).

2.1.2 Optimization problem

The process of determining seed attributes resulting in specified cell properties can be expressed as an optimization problem as illustrated on Figure 2; that is, minimizing an objective function over a set of variables while accounting for potential constraints. Here, the optimization problem must be defined in such a way that the generality of Laguerre tessellations is retained.

The variables of the optimization problem are the coordinates and the (square roots of the) weights of the seeds. For a 3D tessellation comprising N cells, this leads to $4N$ variables. The coordinates of the seeds are allowed to take any value; that is, seeds can be located inside or out of the domain.¹ A first constraint, which results from the definition of a Laguerre tessellation (Equation 1), is that all seed weights must be non-negative, $w_i \geq 0$ ($i = 1, \dots, N$). A second constraint must be considered to avoid empty cells. It cannot be expressed from the optimization variables (which describe the seeds) but can be introduced in the objective function under the form of a penalty term:

$$\mathcal{O}_p = \mathcal{O} + pN_e \quad (2)$$

where \mathcal{O}_p is the penalized objective function, \mathcal{O} is the nominal objective function (defined in Sections 3, 4 and 5), N_e is the number of empty cells and p is an arbitrarily high penalization factor (taken equal to 10^3). This simple penalization is sufficient when the initial tessellation contains no (or only a few) empty cells, as is the case when a proper initial solution is chosen (see Sections 3, 4 and 5). It should be noticed that neither of the two constraints affect the ability of Laguerre tessellations to represent any normal tessellation: while the first one is intrinsic to Laguerre tessellations, the

¹In the case of a periodic tessellation, to a seed located out of the domain (in the optimization variables) is associated an image seed located inside the domain (in the tessellation computation), shifted by the periodicity distance. This is necessary to properly compute the tessellation.

second one is equivalent to the fact that only N seeds are available to generate N (non-empty) cells, which is sufficient since any additional seed (of empty associated cell) would not affect other cells and therefore is unnecessary to the Laguerre tessellation overall. Finally, the objective function quantifies the difference between desired properties and current properties. Its expression in different applications will be provided in Sections 3, 4 and 5.

The nature of the optimization problem dictates the choice of the resolution algorithm. First, for typical tessellations, comprising a thousand cells or more, the number of optimization variables is large ($4N$). Second, the objective function generally is a non-linear function of unknown gradient. The resolution of the optimization problem therefore necessitates a large-scale, non-linear and gradient-free optimization algorithm. Finally, the large dependency of the cell properties of Laguerre tessellations on the seed attributes suggests that a local optimization may be sufficient (given a proper initial solution), and this will be demonstrated empirically in Sections 3, 4 and 5. Two existing algorithms were found to fulfill all previous requirements: Subplex [42] and Praxis [43], as available in the NLOpt free (open-source) library [44]. Subplex is derived from the well-known Nelder-Mead simplex algorithm [45] and decomposes high-dimensional problems into a series of low-dimensional problems that can be handled by the simplex algorithm, while Praxis is derived from the direction-set method [46], which assumes a local quadratic form of the objective function and minimizes it along repeatedly-updated directions of search space. An interesting property of both algorithms is that they modify only a few variables between successive iterations (typically, 1 to 3), independently on the total number of variables (and so on N). As a result, only a few cells are modified between successive iterations (typically, 10 to 50), which makes it possible to update the tessellation, as described in Section 2.2. In practice, Praxis consistently offers better performances than Subplex and will be used in this article.

Optimization can be terminated based on different criteria, such as an absolute or relative error on the objective function, an absolute or relative error on the seed attributes, a maximal value of the objective function, a maximal number of iterations, or even a maximal computation time [44]. In this article, an absolute error criterion on the objective function is used:

$$\Delta \mathcal{O}_{\min}(i) = \mathcal{O}_{\min}(i - 40N) - \mathcal{O}_{\min}(i) < \varepsilon \quad (3)$$

where \mathcal{O}_{\min} is the minimal value of \mathcal{O} , and the absolute error, $\Delta \mathcal{O}_{\min}$, is measured over $40N$ iterations (10 times the number of variables). In this article, we use $\varepsilon = 10^{-6}$.

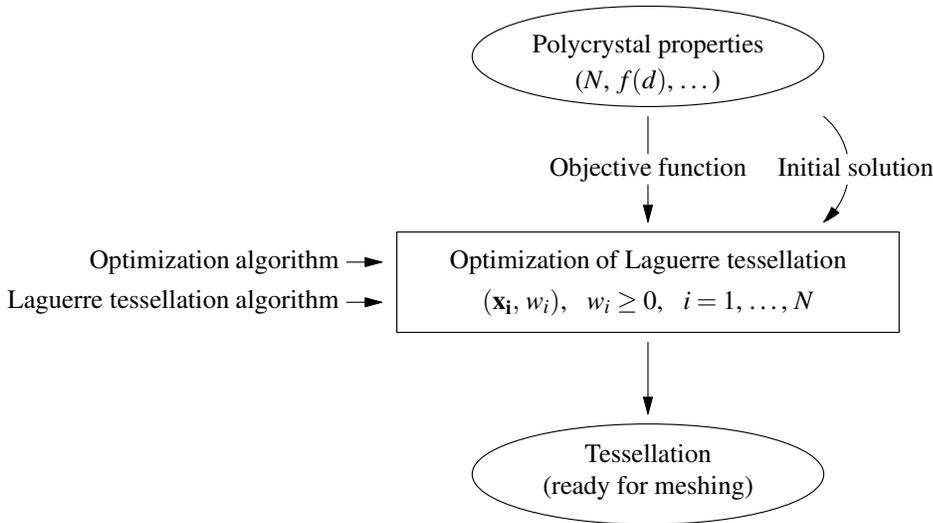


Figure 2: Principle of tessellation generation from polycrystal properties. The optimization of Laguerre tessellation and its associated algorithms are described in Section 2. The input data, the objective function, the initial solution and the output data are described, for different applications, in Sections 3, 4 and 5.

2.2 Laguerre tessellation update

Typical tessellation optimization involves a large number of iterations, of the order of 10^5 – 10^6 for $N = 1,000$. At each iteration, the tessellation and its associated objective function must be computed. As only a few seeds are modified

between successive iterations and for computational efficiency, the tessellation and its associated objective function are updated from their values at the previous iteration. In this section, we provide a method to update a Laguerre tessellation; i.e., given one or several modified seeds, we determine, using the neighbour relations between cells, which cells are modified (these cells are then computed in the standard way, see Section 2.1.1). For the presentation of the method, we assume only one modified seed and tackle all theoretically possible situations. The case of several modified seeds can be treated by superposition. The update of the objective function will be discussed for different applications in Sections 3, 4 and 5.

2.2.1 General case: modified-seed cell and first-neighbour cells

The general (and simplest) case of tessellation update corresponds to a relatively small change of seed attributes, as illustrated on Figure 3. When a seed (S_i) is modified during an iteration, j , its associated cell (C_i) is modified. The neighbours of C_i in the initial configuration, \mathcal{N}_i^{j-1} , are also modified, as are the neighbours of C_i in the final configuration, \mathcal{N}_i^j . Therefore, *the cells to update are the ones associated with the modified seed and the initial and final neighbouring cells*,

$$\mathcal{C}_i^j = C_i \cup \mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j \quad (4)$$

Again, for a relatively small change of S_i , \mathcal{N}_i^{j-1} and \mathcal{N}_i^j intersect by a large amount, but they may be different, as can be seen on Figure 3.

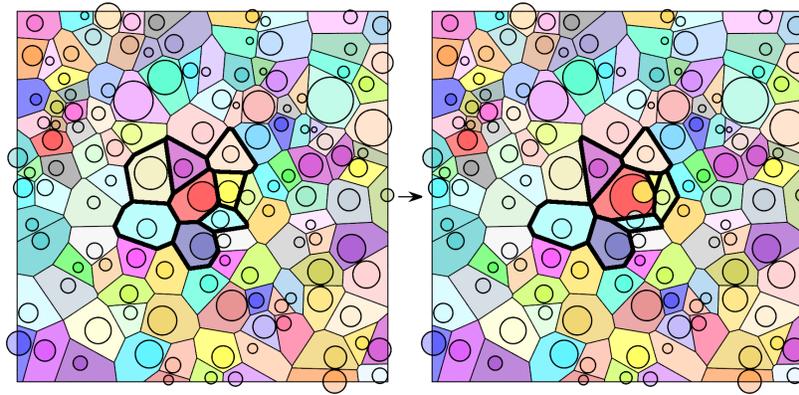


Figure 3: General case of tessellation update when one seed is modified. Seeds are represented by spheres of centres and radii equal to the seed coordinates and square roots of weights, respectively. The position and weight of the red, central seed (S_i , of cell C_i) are modified during the iteration. The left image shows the initial configuration, with C_i 's initial neighbours (\mathcal{N}_i^{j-1}) highlighted. The right image shows the final configuration, with C_i 's final neighbours (\mathcal{N}_i^j) highlighted. Note that $\mathcal{N}_i^{j-1} \neq \mathcal{N}_i^j$. Also note that, in the final configuration, the yellow seed covered by the red, central one is not located in its associated cell (as is possible in Laguerre tessellations).

2.2.2 Special case: higher-order-neighbour cells

In some other cases, and especially when the change of seed attributes is large, more delicate situations can arise, as illustrated on Figure 4. When a tessellation update results in (or starts from) empty cells, the general algorithm may not detect all modified cells. (It should be recalled that empty cells are only avoided *a posteriori*, by penalization, and so can be present at some iterations.)

Figure 4a illustrates a situation for which empty cells are present in the final configuration. By applying the general algorithm, the seed's cell and its initial and final neighbours (\mathcal{N}_i^{j-1} and \mathcal{N}_i^j) are updated; however, cells that should have become empty still are present, leading to an anomalous tessellation exhibiting cell overlaps. These cells are neither initial nor final neighbours of C_i and therefore cannot be detected from C_i . They can actually be detected from C_i 's neighbours, $\mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j$; namely, if one of them lost a neighbour once updated, this neighbour necessarily changed and is to be

updated. Figure 4b illustrates the inverse situation, where empty cells are present in the initial configuration, which leads to an anomalous tessellation exhibiting cells gaps. These cells also can be detected from C_i 's neighbours, $\mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j$; namely, if one of them gained a neighbour once updated, this neighbour necessarily changed and is to be updated.

Both situations can be treated as follows: *during tessellation update, if an updated cell loses or gains a neighbour, this neighbour is to be updated.* This is not limited to cells updated by the general algorithm ($\mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j$), but actually applies to all updated cells. This means that the resulting set of updated cells is not limited to first or second neighbours of C_i but instead can include any higher-order neighbours (and therefore any cell). On Figure 4, this would happen in the case of a larger weight change. It should be noted that cells updated during this process switch either from non-empty to empty or from empty to non-empty. Let $\Delta(\bullet)$ denote the set of “gained” and “lost” neighbours of cell(s) \bullet . The set of modified cells then becomes

$$\mathcal{C}_i^j = C_i \cup \left(\mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j \right) \cup \Delta \left(\mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j \right) \cup \Delta \left(\Delta \left(\mathcal{N}_i^{j-1} \cup \mathcal{N}_i^j \right) \right) \cup \dots \quad (5)$$

Finally, an extreme situation can occur where C_i occupies the full domain in the initial configuration and is empty in the final configuration (or the opposite). In such a situation, both previous algorithms would fail to detect cells to update (as C_i has no neighbours in the initial and final configurations), but it is straightforward to see that all cells need to be updated.

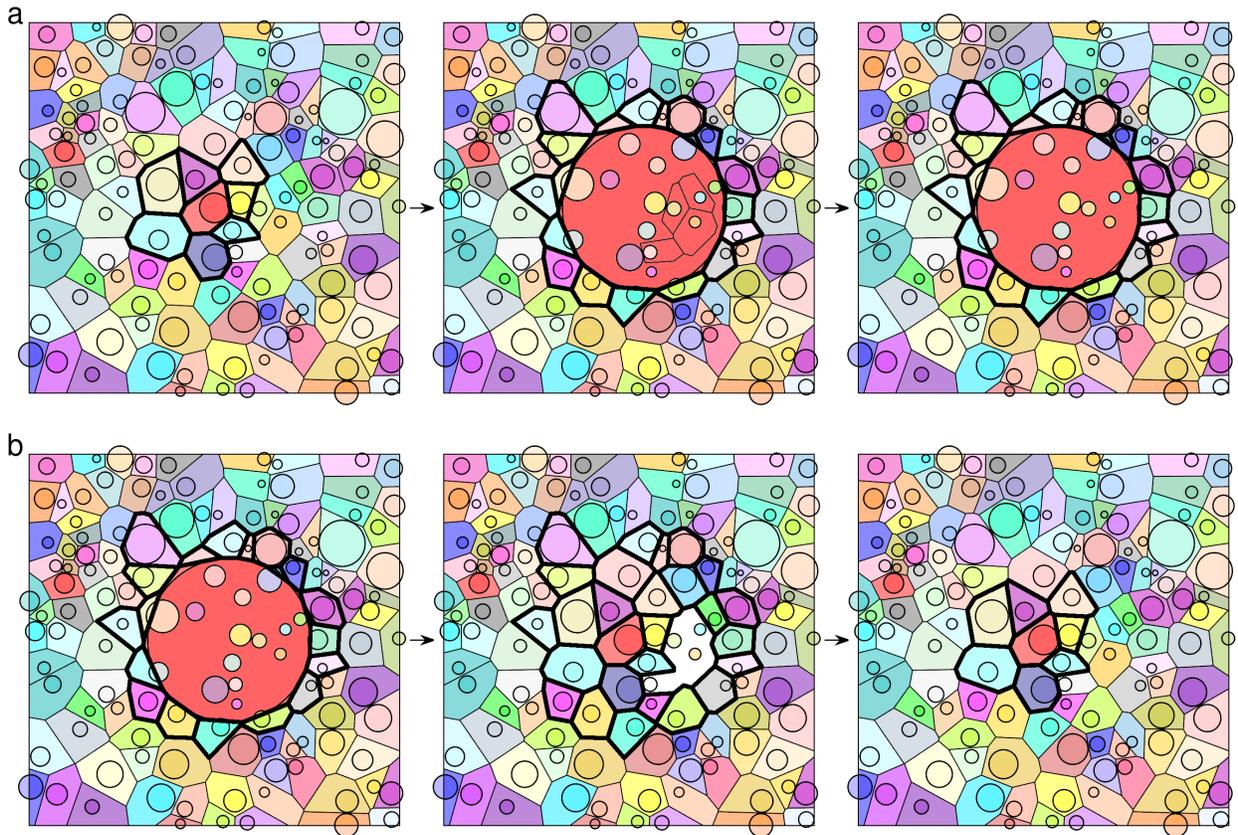


Figure 4: Special cases of tessellation update involving empty cells. The weight of the red, central seed (S_i , of cell C_i) is modified to a large extent during optimization. The left images show initial configurations, the middle images show intermediate configurations as generated by the general algorithm, and the right images show final configurations. (a) Case with empty cells in the final configuration and for which the general algorithm generates cell overlaps (resulting from a large seed weight increase). (b) Case with empty cells in the initial configuration and for which the general algorithm generates cell gaps (resulting from a large seed weight decrease).

2.3 Validation

The efficiency of the tessellation update algorithm, and more specifically the detection of modified cells, can be analysed quantitatively. To do so, a periodic 1,000-cell Voronoi tessellation is subjected to random seed perturbations to mimic changes occurring during optimization. At each iteration, one seed attribute (chosen randomly) is altered by a specific amount, Δx_i , and the modified cells are detected as described in Section 2.2. To keep a known initial state (and unlike what actually occurs during optimization), the seed change is then reversed before the next iteration. Figure 5 provides results on the number of modified cells during random perturbations, for magnitudes Δx_i ranging from $10^{-1} \langle d \rangle$ down to $10^{-3} \langle d \rangle$, where $\langle d \rangle$ denotes the average cell diameter (see Section 3), as is typically encountered at the beginning and end of optimization, respectively. First, it appears that the number of modified cells, $|\mathcal{C}_i^j|$, vary significantly, from 9 to 39 cells for $\Delta x_i = 10^{-1} \langle d \rangle$, with an average of 23. This is, as expected, much lower than the total number of cells (1,000). Moreover, the lower Δx_i , the lower $|\mathcal{C}_i^j|$; at $\Delta x_i = 10^{-3} \langle d \rangle$, the average is 16.65, which is close to $1 + \langle n_f \rangle$, where $\langle n_f \rangle = 15.2$ is the average number of faces of a cell. This confirms that for low perturbations, modified cells are only the cell of the modified seed and its first neighbours.

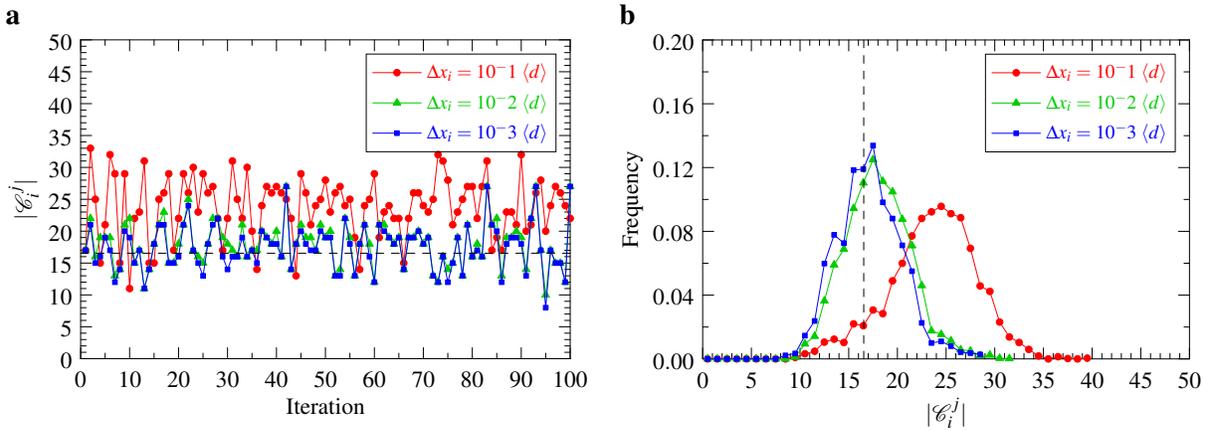


Figure 5: Number of modified cells ($|\mathcal{C}_i^j|$) during tessellation update in the case of random perturbation of seed attributes of different magnitudes (Δx_i). (a) Number of modified cells as a function of iteration number (limited to 100) and (b) distribution of the number of modified cells during the whole random perturbation process (100,000 iterations). The dashed lines correspond to a value of $1 + \langle n_f \rangle$.

3 Application to statistical grain properties

Polycrystalline morphologies are commonly described in terms of their average grain sizes and grain size distributions. Grain shapes are more rarely considered but may also be important as, in principle, polycrystals with the same grain size distributions can exhibit appreciably different grain shapes. Experimental distributions of 3D materials can either be measured directly [47–49] or determined from 2D measurements and stereology [50]. In this section, given a domain (D) and a number of cells (N), tessellations are generated from prescribed grain size and shape distributions.

3.1 Metrics

The size of a grain can be defined as the diameter of its sphere of equivalent volume, d . For each grain of a polycrystal, it is then possible to determine a relative value, $d/\langle d \rangle$, where $\langle d \rangle$ denotes the average diameter over all grains. In this section and unless mentioned otherwise, $d/\langle d \rangle$ will simply be referred to as “diameter” (or “size”). A possible shape metric is the sphericity, s [29], which is defined as the ratio between the surface area of the sphere of equivalent volume and the surface area of the grain, and takes a maximal value of 1 for a spherical grain.

Experimental grain size ($d/\langle d \rangle$) and sphericity (s) distributions of different materials are provided on Figure 6 [29, 47–49]. For a grain-growth material, i.e. a material that underwent grain boundary migration by capillarity effects, $d/\langle d \rangle$

generally follows a lognormal distribution of average 1 (by definition) and standard deviation 0.35, and $1 - s$ follows a lognormal distribution of average 0.145 and standard deviation 0.030 [29]. In comparison, Voronoi tessellations provide less variable sizes [26], of standard deviation 0.15, and lower sphericities, of average 0.810 and standard deviation 0.044.

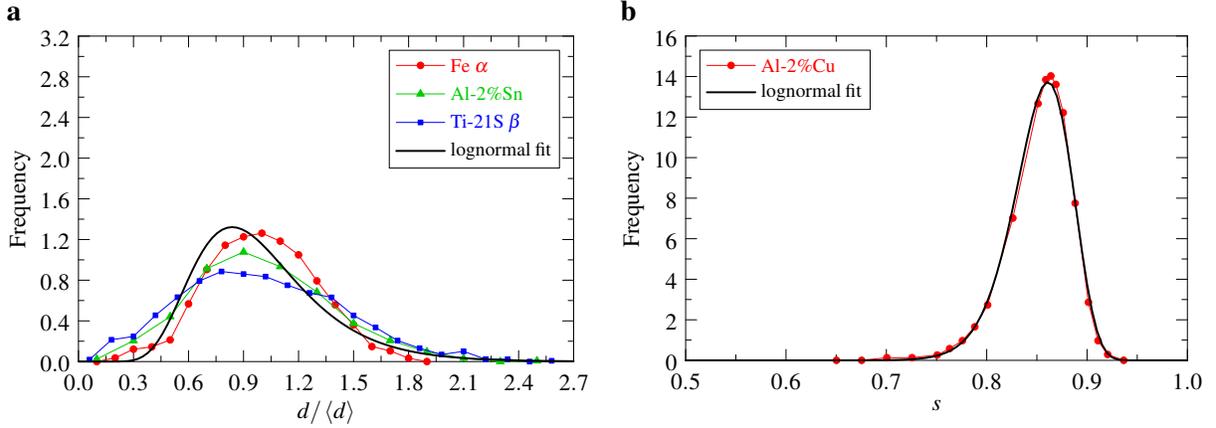


Figure 6: Morphological statistics of the grains of polycrystals: (a) size distribution and (b) sphericity distribution. Fe α data after Ref. [47], Al-2%Sn data after Ref. [48], Ti-21S β data after Ref. [49] and Al-2%Cu data after Ref. [29].

3.2 Objective function

The size and sphericity distributions of the tessellation must be compared to the prescribed ones to define the objective function. This is done independently for each property, and the resulting values are then combined into the objective function.

3.2.1 General expression

Comparing a continuous probability density function (the prescribed distribution), f , to the distribution of a sample of data (the cell values), $\{x_i\}$, is a well-known problem in Statistics that led to so-called “goodness-of-fit” tests [51]. A popular test is the χ^2 test, but more elaborate tests include the Kolmogorov-Smirnov, Kuiper, Cramér-von Mises and Anderson-Darling tests, all of which use the cumulative distribution functions. However, these tests were not designed in the context of optimization, and none of them are efficient in the present context. So, a new test is defined which derives from the Anderson-Darling test. Let F denote the cumulative distribution function corresponding to f and X denote the cumulative distribution function of $\{x_i\}$. F is a continuous function while X is a step function which can be written as

$$X(x) = \frac{1}{N} \sum_{i=1}^N H(x - x_i) \quad \text{with} \quad H(\bullet) = \begin{cases} 0 & \text{if } \bullet < 0 \\ 1 & \text{if } \bullet \geq 0 \end{cases} \quad (6)$$

H is the Heaviside step function. F and X can be compared in terms of a metric, \mathcal{A} , defined as,

$$\mathcal{A} = \int_{-\infty}^{+\infty} \frac{[X(x) - F(x)]^2}{F(x)[1 - F(x)]} dx \quad (7)$$

In principle, integration is made over $]-\infty, +\infty[$, but in practice, it may be restricted to the interval of possible values of x . For general values of F , \mathcal{A} must be computed by numerical integration, i.e. by discretizing the x interval into a finite number of bins. X being a step function, any change in x_i inside individual bins would not change the value of \mathcal{A} , which is problematic for optimization as the gradient of \mathcal{A} with respect to x_i cannot be evaluated properly. An expression of \mathcal{A} including smoothing, \mathcal{A}^s , can be used instead:

$$\mathcal{A}^s = \int_{-\infty}^{+\infty} \frac{[X^s(x) - F^s(x)]^2}{F^s(x)[1 - F^s(x)]} dx \quad \text{with} \quad \begin{cases} F^s = F \circ s \\ X^s = X \circ s \end{cases} \quad (8)$$

where \circ is the convolution product and s is a normal convolution kernel of standard deviation σ_s (small with respect to the one of f , as discussed in Section 3.3.2). When $\mathcal{A}^s = 0$, $F^s = X^s$ and so $F = X$. However, because F is a continuous function and X is a step function, $\mathcal{A}^s > 0$; we will therefore seek for a minimal value rather than a zero value of \mathcal{A}^s , consistent with the termination criterion (Equation 3). Typical F , X , F^s and X^s distributions are provided on Figure 7, from which the definition of \mathcal{A}^s (Equation 8) can be better depicted. The expression of \mathcal{A}^s is such that a (quadratic) difference between F^s and X^s is weighted more highly when occurring at low (near zero) and high (near one) values of F^s (or equivalently, F), i.e. in the two tails of f . This is important since the tails of the distributions are often critical in applications [52].

Equation 8 can be generalized to the case of a multimodal shape of distribution f . Considering that all tails (not only the outer ones) of the distribution should be weighted in similar ways, the expression of \mathcal{A}^s becomes

$$\mathcal{A}^s = \int_{-\infty}^{+\infty} \frac{[X^s(x) - F^s(x)]^2}{\prod_{i=0}^m |F^s(x) - p_i|} dx \quad (9)$$

where $p_0 = 0$, and p_i are the cumulative fractions of the m distribution modes of f ($p_m = 1$). For $m = 1$, Equation 9 is equal to Equation 8.

Finally, let n be the number of properties considered simultaneously (here, the diameter and the sphericity). The n corresponding values of \mathcal{A}^s , \mathcal{A}_i^s , $i = 1, \dots, n$, are combined into the objective function, \mathcal{O} ,

$$\mathcal{O} = \frac{1}{n} \sqrt{\sum_{i=1}^n \mathcal{A}_i^{s2}} \quad (10)$$

For computational efficiency, it is important that \mathcal{A}^s can be updated between successive iterations. Using Eqns. 6 and 9, noting that $H \circ s$ can be expressed using the error function and after elementary mathematical manipulations, X^s can be written as

$$X^s(x) = \frac{1}{2N} \sum_{i=1}^N \left[1 + \operatorname{erf} \left(\frac{x - x_i}{\sqrt{2} \sigma_s} \right) \right] \quad (11)$$

It follows that the value of X^s at iteration j , X^{sj} , can be expressed from the one at iteration $j - 1$, $X^{s(j-1)}$, as

$$X^{sj}(x) = X^{s(j-1)}(x) + \frac{1}{2N} \sum_{i \in \mathcal{C}_i^j} \left[\operatorname{erf} \left(\frac{x - x_i^j}{\sqrt{2} \sigma_s} \right) - \operatorname{erf} \left(\frac{x - x_i^{j-1}}{\sqrt{2} \sigma_s} \right) \right] \quad (12)$$

Equation 12 makes it possible to update the values of X^s , \mathcal{A}^s (using Equation 9) and then \mathcal{O} (using Equation 10) at a cost independent on N . It can be applied directly to sphericity, while the case of the size needs more discussion.

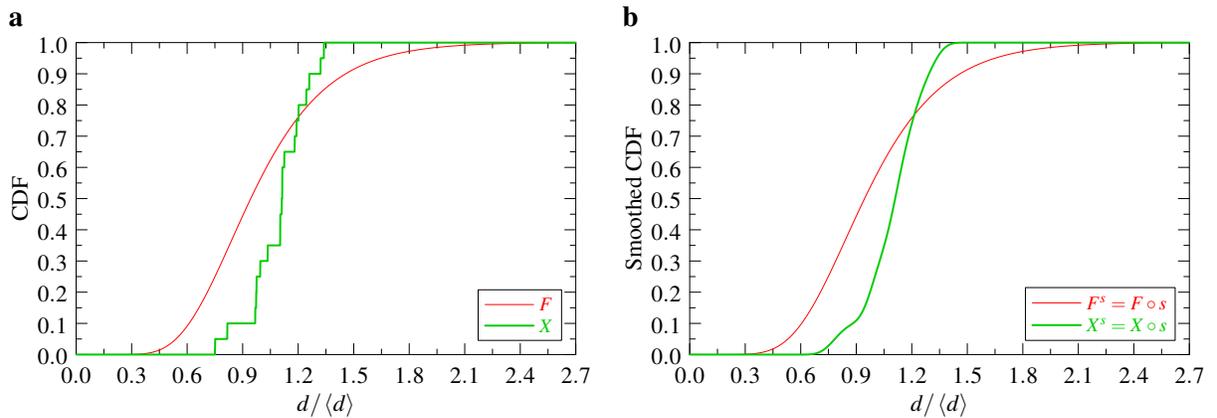


Figure 7: Comparison between prescribed and current $d/\langle d \rangle$ cumulative distributions. (a) Target (F) and current (X) cumulative distributions. (b) Target (F^s) and current (X^s) smoothed cumulative distributions, used for computing the objective function. A tessellation with a limited number of cells (20) is used to highlight the steps of X .

3.2.2 Case of the size distribution

In general, changing the (absolute) diameters (d) of the updated cells changes their average over all cells, $\langle d \rangle$ (contrary to the case of cell volumes). As a consequence, the value of $d/\langle d \rangle$ is changed for all cells, not only updated cells, which prevents updating X^s using Equation 12. This problem can be solved by normalizing, during optimization, the cell diameters (d) not by their current average value, but by their final average value, $\langle d \rangle_p$ (a constant). Given the volume of the domain, V , the number of cells, N , and the prescribed probability density function of the (relative) diameters, f_d , $\langle d \rangle_p$ is given by

$$\langle d \rangle_p = \langle d \rangle_D \left[\int_{-\infty}^{+\infty} f_d(x) x^3 dx \right]^{-\frac{1}{3}} \quad \text{with} \quad \langle d \rangle_D = \left(\frac{6V}{\pi N} \right)^{\frac{1}{3}} \quad (13)$$

$\langle d \rangle_D$ corresponds to the diameter of a cell of average volume. The wider the diameter distribution, the lower $\langle d \rangle_p$. For a Dirac distribution, $\langle d \rangle_p = \langle d \rangle_D$, while for a lognormal distribution of standard deviation 0.35, $\langle d \rangle_p \simeq 0.891 \langle d \rangle_D$.

3.3 Applications

The method is applied to the generation of tessellations from various grain size and sphericity distributions in Section 3.3.1. The influence of the distribution smoothing and discretization parameters is discussed in Section 3.3.2.

3.3.1 Polycrystals of specified grain size and sphericity distributions

As mentioned previously, grain-growth polycrystals are characterized by a $d/\langle d \rangle$ lognormal distribution of average 1 and standard deviation 0.35 and a $1-s$ lognormal distribution of average 0.145 and standard deviation 0.030. Tessellation generation can be started from a random and uniform spatial seed distribution and seed weights equal to $(\langle d \rangle_{target}/2)^2$. This corresponds to a Voronoi tessellation, but considering non-zero weights facilitates optimization since it allows them to decrease. Figure 8 illustrates how a 1,000-cell tessellation evolves during optimization. It can be seen that the tessellation is significantly modified during the first 75,000 iterations, resulting in a strong decrease of \mathcal{O} (Figure 8d), and is then further adjusted until the termination criterion is verified. At the final state, the cell diameter and sphericity distributions closely match the prescribed distributions. Meanwhile, the average number of cell faces is reduced from 15.1 to 13.6 and comes in close agreement with polycrystal values (12–14 [47–49]). This was not monitored by optimization and is only associated to correct size and sphericity distributions, which shows that these metrics are particularly efficient for describing grain-growth polycrystals. Figure 9 shows that the method is applicable to tessellations with large numbers of cells (10,000). About 1,650,000 iterations were needed, in a computation time² of 4.8 hours, while generating a 1,000-cell tessellation requires about 180,000 iterations in a computation time of 12.7 minutes.

The size and sphericity distributions can also be varied (e.g. about their grain-growth values). In Figure 10, the standard deviation of the size distribution is varied from 0 to 0.55 while the average sphericity is varied from 0.9 to 0.7 (standard deviation 0.03). It can be seen that the method successfully generates tessellations from any combination of size and sphericity distributions.

As a last example, a tessellation composed of two cell populations of different size distributions is generated, as illustrated in Figure 11. The two populations are characterized by an average size ratio of 1:2 and size distributions of standard deviations of 5% of their respective averages, and have the same volume fractions. The sphericity distribution has an average of 0.875 and a standard deviation of 0.05. As an initial solution, a (rough) sphere packing is used, where sphere diameters, d_i , are set by sampling the size distribution and sphere positions are set so as to minimize sphere overlapping. Seed positions are set as those of the spheres while seed weights are computed as $w_i = d_i^2/4$, which provides an initial bimodal diameter distribution (see Figure 11b). This is necessary to optimization due to the weighting of all distribution tails in Equation 9. At the final state, the two cell populations of the tessellation match the prescribed distributions and have random spatial arrangements.

²Unless mentioned otherwise, computations are run on a single core of a 2.60 GHz Intel[®] Xeon[®] processor.

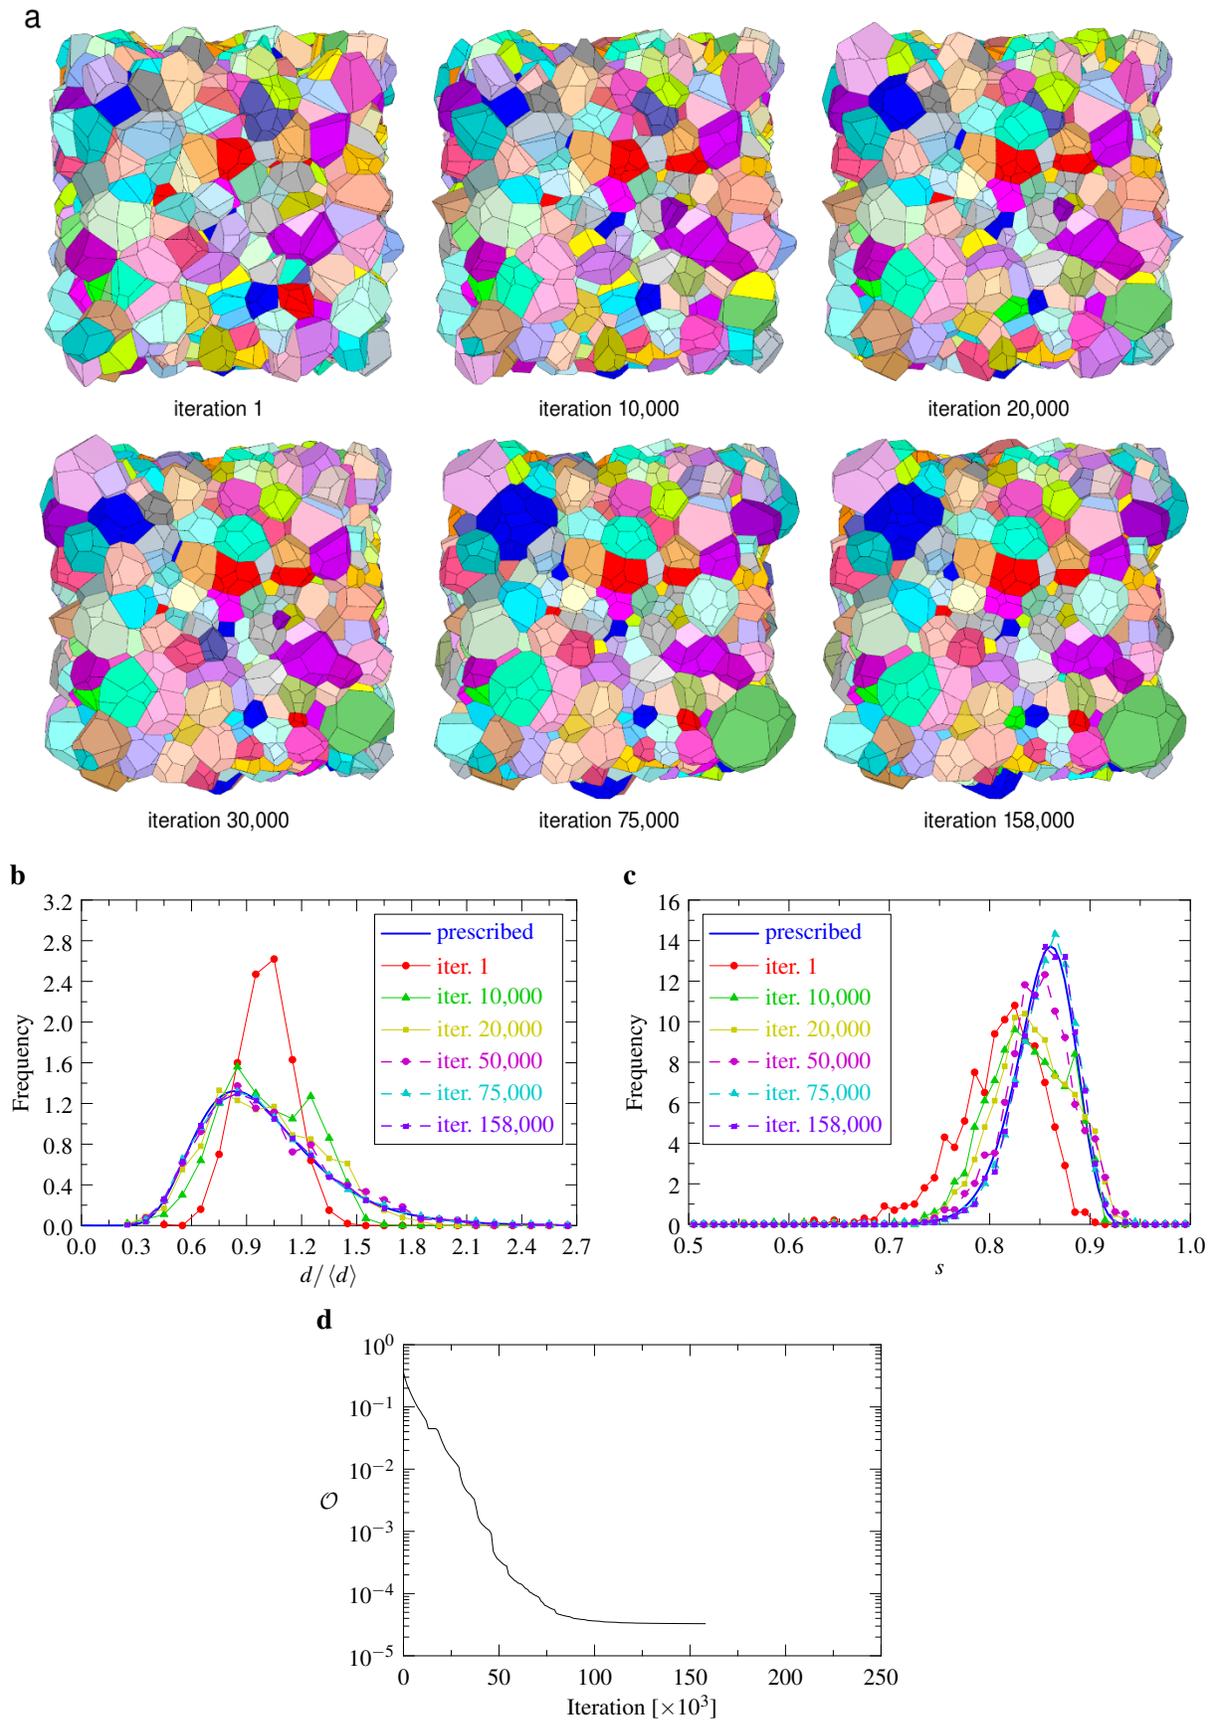


Figure 8: Tessellation generation from grain-growth properties. (a) Tessellations at several optimization stages up to the final iteration (158,000), (b) evolution of the size distribution, (c) evolution of the sphericity distribution and (d) evolution of the objective function.

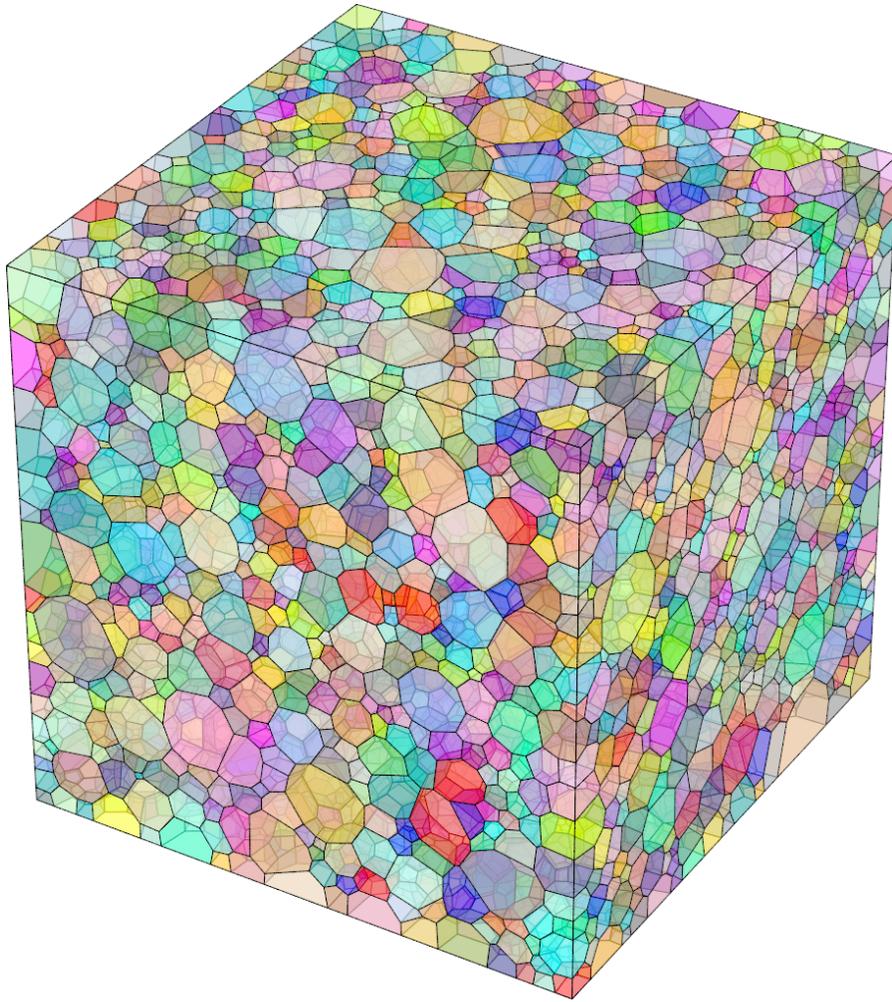


Figure 9: Example of a 10,000-cell tessellation with grain-growth properties.

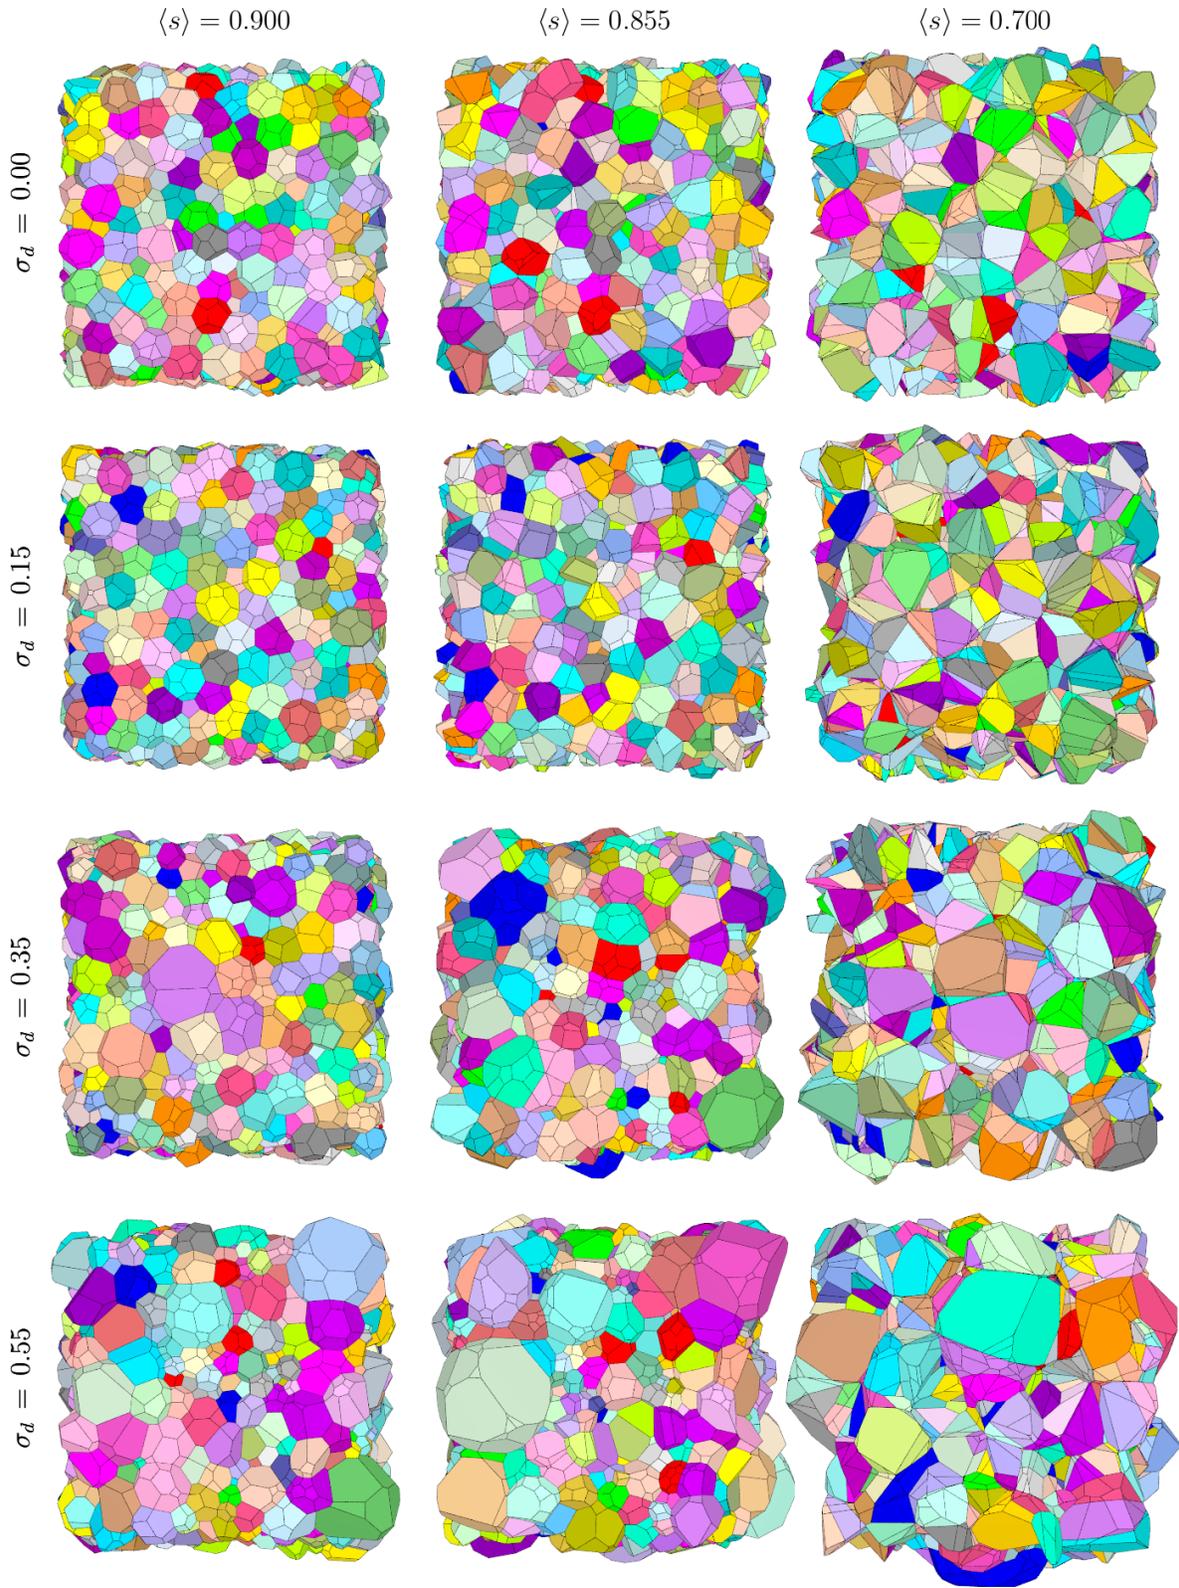


Figure 10: 3D periodic tessellations with different size (top to bottom) and sphericity (left to right) distributions. σ_d stands for the standard deviation of the size distribution (average 1), and $\langle s \rangle$ stands for the average of the sphericity distribution (standard deviation 0.03). The image at $(\sigma_d = 0.35, \langle s \rangle = 0.855)$ represents grain-growth properties (see Figure 8a).

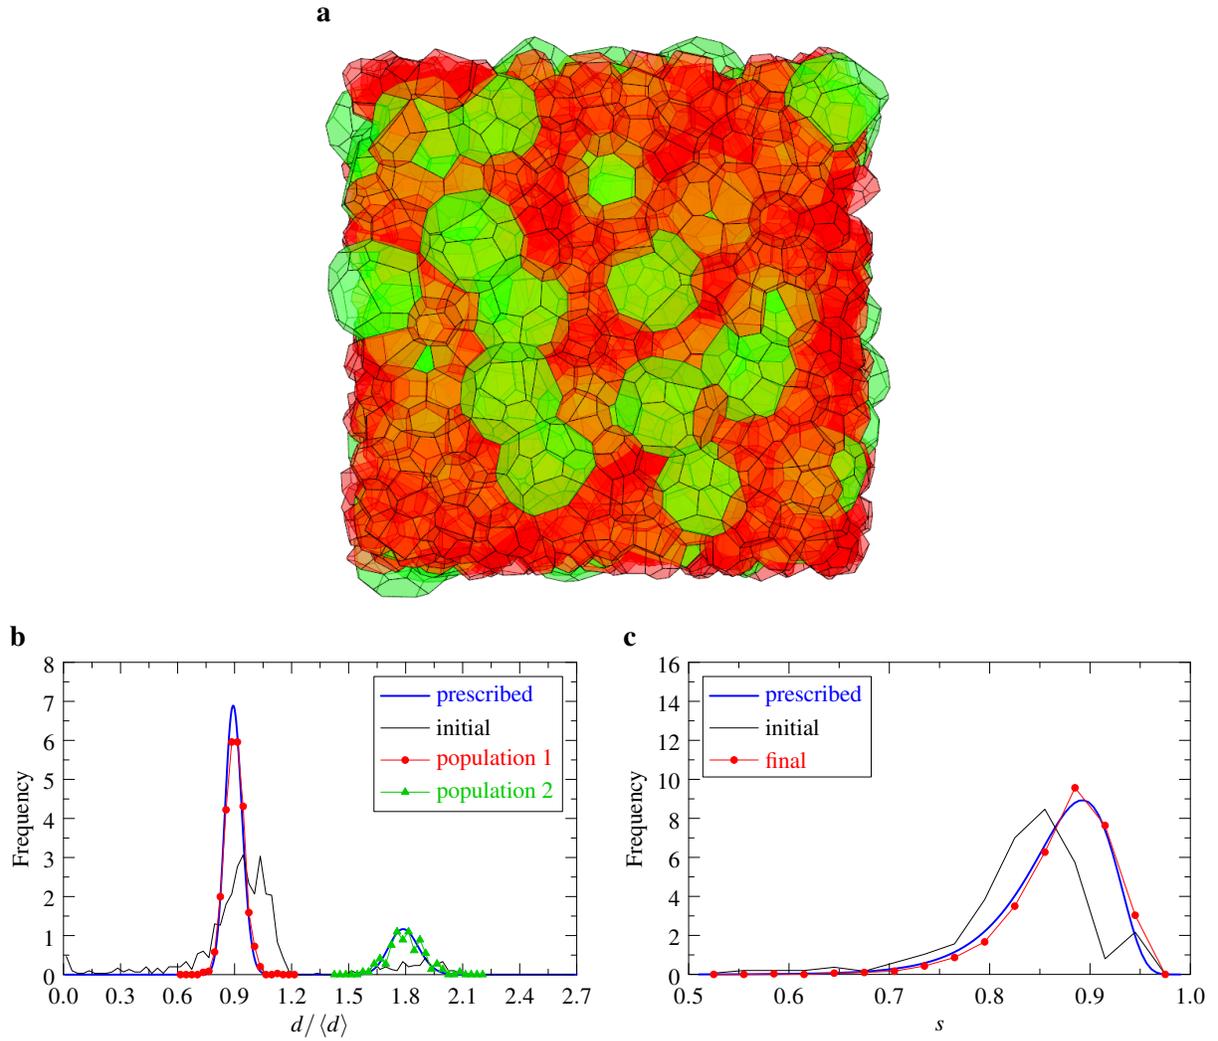


Figure 11: Example of a tessellation containing two cell populations with an average size ratio of 1:2 and equal volume fractions. (a) Tessellation with the two populations uniformly coloured, (b) size distributions and (c) sphericity distributions.

3.3.2 Influence of distribution smoothing and discretization

In this section, the influences on tessellation quality of the smoothing parameter (σ_s) and the bin width used to discretize the statistical distributions are analysed in the case of tessellation generation from grain-growth properties.

Concerning the smoothing parameter, it should be recalled that smoothing is introduced as a way to make the objective function (\mathcal{O}) differentiable, as needed for optimization. In theory, equal smoothed distributions ($F^s = X^s$) provide equal distributions ($F = X$) whatever the smoothing kernel s (and so its standard deviation σ_s), and so a minimal value of \mathcal{O} should provide the same final distributions independently on σ_s . However, optimization is terminated using an error criterion on \mathcal{O} , which may lead to slightly different distributions. The procedure for determining the influence of σ_s on the quality of the tessellation is to generate tessellations using different values of σ_s and to compare them to the prescribed distributions without smoothing (i.e. using Equation 7 or equivalently $\sigma_s = 0$ in Equation 8); the corresponding value of \mathcal{O} is denoted \mathcal{O}^y . Using \mathcal{O}^y , F and X are therefore compared directly and independently on σ_s , which enables to measure its effects. A constant ratio of 10 is used between the σ_s values of the size and sphericity distributions (which is close to the ratio between their standard deviations). The bin width used to discretize the distributions is set to 10^{-3} for the size and 10^{-4} for the sphericity; its influence will be analysed afterwards. Results are provided on Figure 12a. First, smoothing clearly facilitates optimization, as non-zero σ_s values consistently result in a faster decrease of \mathcal{O}^y . The smaller the value of σ_s (but still non-zero), the smaller the final value of \mathcal{O}^y . It should be noted, however, that all σ_s values lower than 0.05 provide similar \mathcal{O}^y values until about 75,000 iterations, i.e. when the tessellation distributions become visually close to the prescribed ones on Figure 8; the following iterations only serve for fine adjustments (all the more than σ_s is small). Therefore, σ_s values of 0.05 for the size and 0.005 for the sphericity appear sufficient to obtain a good-quality tessellation while limiting the number of iterations compared to smaller values. At the opposite, higher values of σ_s fail to provide low enough values of \mathcal{O}^y , even at large iteration numbers.

The influence of the bin width, Δx , can be analysed by following the same procedure: the bin width is varied and the resulting values of \mathcal{O}^y are compared. $\sigma_s = 0.05$ and 0.005 are used for the size and sphericity distributions, respectively. Similarly, a constant ratio of 10 is used between the bin widths of the size and sphericity distributions. Results are provided on Figure 12b. It appears that the evolution of \mathcal{O}^y is almost independent on Δx up to the value of σ_s . Bin widths of 0.01 for the size and 0.001 for the sphericity are chosen.

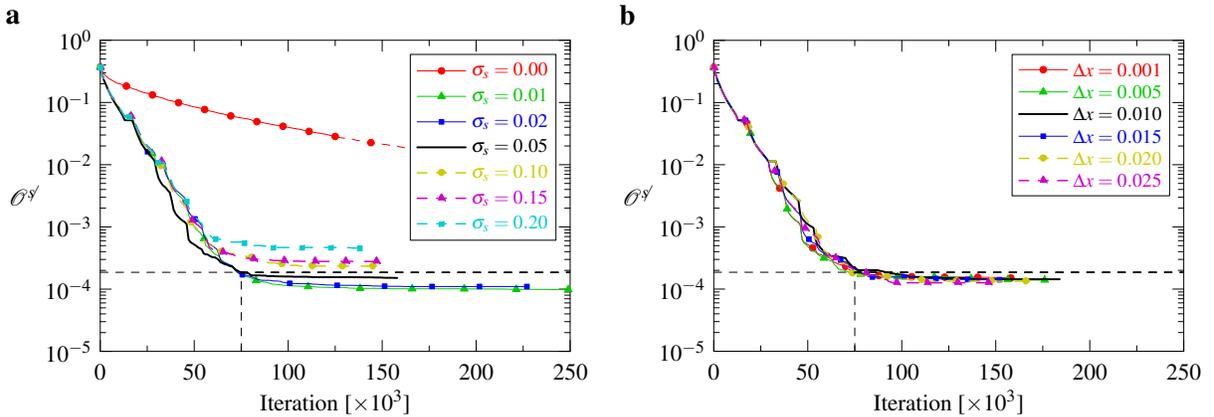


Figure 12: Influence of distribution smoothing and discretization on tessellation quality. (a) Evolution of \mathcal{O}^y during optimization for different values of σ_s and (b) evolution of \mathcal{O}^y during optimization as a function of the discretization bin width (Δx). Values of σ_s and Δx apply to the size distribution, while $\sigma_s/10$ and $\Delta x/10$ are used for the sphericity distribution. On (a), the $\sigma_s = 0$ curve is truncated for clarity. The black curves correspond to the chosen values of σ_s and Δx ($\sigma_s = 0.05$ and $\Delta x = 0.010$). The dashed lines represent the apparent convergence state on Figure 8.

4 Application to grain centroids and volumes

The centroids, μ_i , and volumes, V_i , of the individual grains of a 3D polycrystal can be obtained by 3D X-ray diffraction microscopy (3DXRD) [10], so as their lattice orientations and elastic strains. Also known is the spatial extent of the polycrystal (usually from conventional tomography), from which the domain (D) is defined, and (of course) the number of grains (N). In this section, tessellations are generated from grain centroids and volumes.

4.1 Metrics

While the grain centroids can be used directly, the grain volumes (V_i) are used to compute grain diameters (d_i),

$$d_i = \left(\frac{6}{\pi} V_i \right)^{\frac{1}{3}} \quad (14)$$

which are of the same dimension as centroids' coordinates and can be combined with them to define the objective function.

4.2 Objective function

An efficient metric to quantify the difference between the grain and cell centroid / diameter pairs, (μ_i, d_i) and (μ_i^*, d_i^*) , respectively, can be obtained from Figure 13. Spheres are constructed from the centroid / diameter pairs and are compared using two distances between them, δ_a and δ_b , measured along the line passing through their centres. The difference between the two spheres (and therefore between a grain and a cell), δ_i , can be defined as

$$\delta_i = \sqrt{\delta_a^2 + \delta_b^2} \quad (15)$$

which, after elementary geometrical considerations, can be written as

$$\delta_i = \sqrt{d_E(\mu_i, \mu_i^*)^2 + \frac{1}{4}(d_i - d_i^*)^2} \quad (16)$$

The objective function can then be computed as

$$\mathcal{O} = \frac{1}{N \langle d \rangle} \sqrt{\sum_{i=1}^N \delta_i^2} \quad (17)$$

Factors $1/N$ and $1/\langle d \rangle$ are introduced for normalization, where N is the total number of grains and $\langle d \rangle$ is the average grain diameter. \mathcal{O} takes a minimal value of 0 when all cell centroid / diameter pairs $((\mu_i^*, d_i^*))$ are equal to grain ones $((\mu_i, d_i))$.

The objective function can be updated between successive iterations ($j-1$ and j) using

$$\mathcal{O}^j = \sqrt{\mathcal{O}^{j-1^2} + \frac{1}{N^2 \langle d \rangle^2} \sum_{i \in \mathcal{C}_i^j} (\delta_i^{j^2} - \delta_i^{j-1^2})} \quad (18)$$

4.3 Applications

The method is applied to tessellation generation from the centroids and volumes of a 131-grain polycrystal, as illustrated on Figure 14. A full image of this polycrystal was actually obtained experimentally (see Section 5), but the image is here cast into grain centroids / volumes to obtain data free of the experimental uncertainties encountered in 3DXRD. As initial solution, a straightforward possibility is to position the seeds on the grain centroids, while weights are obtained from the grain diameters,

$$\mathbf{x}_i^* = \mu_i, \quad w_i^* = \left(\frac{d_i}{2} \right)^2 \quad (19)$$

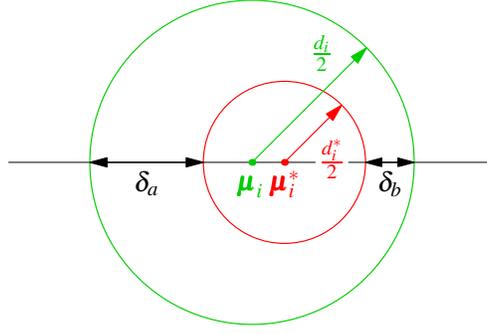


Figure 13: Comparison between grain and cell centroid / diameter pairs, (μ_i, d_i) and (μ_i^*, d_i^*) , respectively, from which spheres are constructed. The line passing through the centres is considered to define δ_a and δ_b , from which the difference between the spheres (δ_i) is computed.

Tessellation generation involved about 1,030,000 iterations in a computation time of 4.1 hours. Figure 14 compares the centroids and volumes of the grains and cells, at the initial and final states. The correspondence is significantly improved by optimization, as the objective function is decreased from 0.0107 to 0.000942. The agreement can also be analysed for individual cells (Figure 14d,e). The average error on the centroid positions (relative to the average diameter) is reduced from 10.1% to 0.7%, on the volumes (relative to the average volume) from 14.4% to 3.2% and on the diameters (relative to the average diameter) from 10.8% to 1.3%. Final values are significantly lower than the typical experimental uncertainty of a few percents [53]. It should also be noticed from Figure 14b,c that optimization leads to a significant change of morphology. The change of morphology can be quantified from the rate of intersection, I^V , between the cells of the initial and final tessellations, C_i^{ini} and C_i^{fin} , respectively. I^V can be written as

$$I^V = \frac{1}{V} \sum_{i=1}^N C_i^{\text{ini}} \cap C_i^{\text{fin}} \quad (20)$$

and is equal to 84.4%. In other words, the initial and final tessellations differ by 15.6%.

5 Application to raster polycrystal images

3D images of polycrystals can be obtained by synchrotron X-ray techniques such as “diffraction contrast tomography” (DCT) [8] and “high-energy diffraction microscopy” (HEDM) [9]. Other, destructive techniques combine optical or electron microscopy to serial sectioning [49, 54, 55]. Of course, 2D-only characterizations are also often used [56, 57]. All of these techniques provide raster images, i.e. images composed of voxels. In this section, the domain (D) of the tessellations is identified to the one of the polycrystals and tessellations are generated from the polycrystal images.

5.1 Metrics

Depending on the experimental technique, the raster polycrystal image can be defined, at the voxel scale, in terms of grain label (e.g. in DCT) or lattice orientation (e.g. in HEDM). Here, grain labels are used, which can be obtained from lattice orientations using standard techniques. Let \mathcal{M} be the set of all voxels, v_k , $k = 1, \dots, M$, of the polycrystal. Let I be the grain label function. A grain, G_i ($i = 1, \dots, N$), is defined as

$$G_i = \{v_k \in \mathcal{M} \mid I(v_k) = i\} \quad (21)$$

G_i defines a grain completely in terms of its morphology and location in the domain.

5.2 Objective function

The grains and cells are compared directly, using the sets of voxels of the grains (G_i) and the polyhedral shapes of the cells (C_i). From a computational point of view, it is important for this to be done efficiently in spite of the very large number of voxels of raster polycrystal images (typically 100–1000³). This is made possible by an appropriate definition and a rapid evaluation of the objective function.

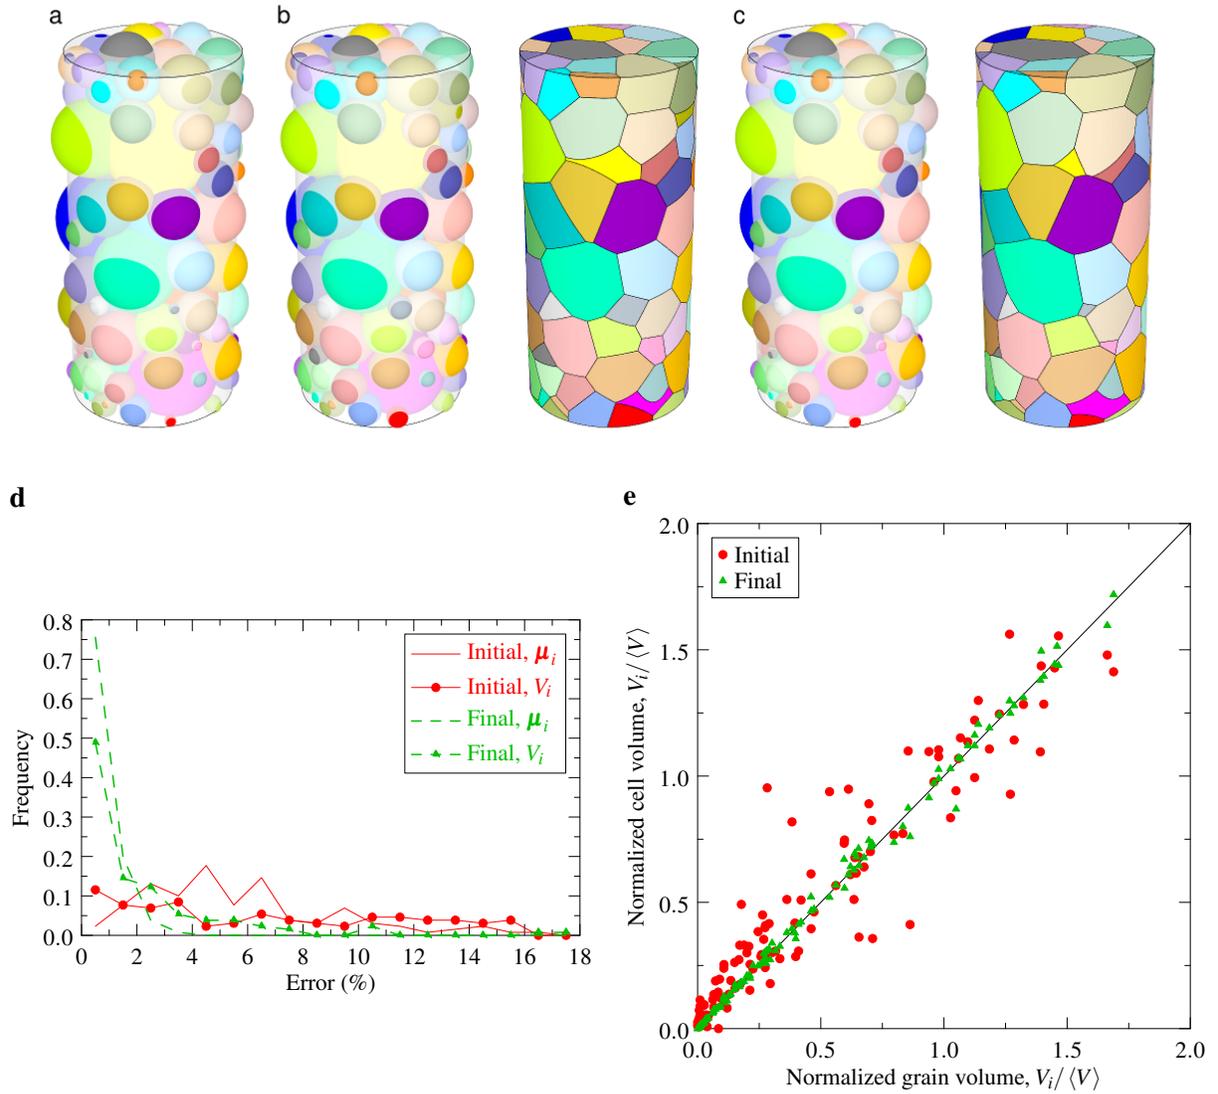


Figure 14: Tessellation generation from grain centroids and volumes. (a) Grain centroids and volumes (grains are represented as spheres), (b) initial cell centroids and volumes, and corresponding tessellation, (c) final cell centroids and volumes, and corresponding tessellation, (d) difference between the grain and cell centroids and volumes (computed as $d_E(\mu_i, \mu_i^*) / \langle d \rangle$ and $\Delta V_i / \langle V \rangle$), and (e) direct comparison between the grain and cell volumes. On (e), the scale is limited to 2 for clarity.

5.2.1 General expression

A grain and a cell can be compared using the distance between their boundaries, as illustrated on Figure 15. First, for each grain (G_i), the set of boundary voxels (i.e. the voxels that are adjacent to at least another grain), G_i^b , is considered and can be formally defined as

$$G_i^b = \{v_k \in G_i \mid \exists l \in [1, n_n] \mid I(N_l(v_k)) \neq i\} \quad (22)$$

where $N_l(v_k)$, $l = 1, \dots, n_n$ are v_k 's face neighbours ($n_n = 6$ in 3D and 4 in 2D). The difference between a grain (G_i) and its associated cell (C_i), δ_i , can then be defined as

$$\delta_i = \sqrt{\sum_{v_k \in G_i^b} d_E(v_k, C_i)^2} \quad (23)$$

The objective function can then be computed as

$$\mathcal{O} = \frac{2}{n_v \langle d \rangle} \sqrt{\sum_{i=1}^N \delta_i^2} \quad \text{with} \quad n_v = \sum_{i=1}^N |G_i^b| \quad (24)$$

where $|G_i^b|$ is the number of elements of G_i^b . Factors $2/n_v$ and $1/\langle d \rangle$ are introduced for normalization, where $n_v/2$ is the total number of boundary voxels divided by 2 (due to the double covering of grain boundaries, as discussed in the following), and $\langle d \rangle$ is the average diameter. The distance between a voxel and a cell is defined as positive when the voxel centre is located out of the cell and zero otherwise, and is computed using the GJK algorithm [58,59]. It follows from this definition that $\delta_i = 0$ when a cell contains all of its associated grain boundary voxels, irrespectively of the actual correspondence between the grain and cell boundaries (or grain and cell); δ_i therefore is not a true measure of the difference between a grain and a cell. As can be seen on Figure 15, this is however counterbalanced by the fact that grain boundaries are described by two covers of voxels, one for each of the adjacent grains, and that when one cover of a grain boundary is located inside the associated cell (resulting in zero distances), the other cover is not (resulting in non-zero distances). Both covers of a grain boundary acquire zero distances only when the grain and cell boundaries coincide. This definition of the distance function both ensures that the distance at a particular location of a grain boundary is taken into account only once and makes it possible to use the standard GJK algorithm [58,59]. \mathcal{O} takes the minimal value of 0 when each cell contains all of its associated grain boundary voxels, i.e. when the polycrystal and tessellation coincide.

The objective function can be updated between successive iterations ($j-1$ and j) using

$$\mathcal{O}^j = \sqrt{\mathcal{O}^{j-1^2} + \frac{4}{n_v^2 \langle d \rangle^2} \sum_{i \in \mathcal{C}_i^j} (\delta_i^{j^2} - \delta_i^{j-1^2})} \quad (25)$$

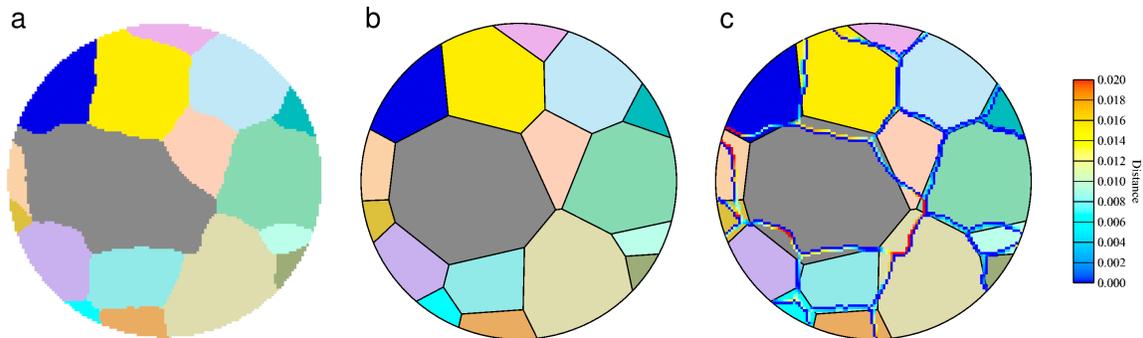


Figure 15: Comparison between a polycrystal and a tessellation. (a) Polycrystal, (b) tessellation and (c) boundary voxels of each grain and distance to their associated cells. Note the double cover of grain boundaries. The polycrystal corresponds to the top surface of the one shown in Figure 17.

5.2.2 Control points

As pointed out previously, typical polycrystal images provided by synchrotron X-ray diffraction techniques contain a large number of voxels. As can be seen on Figure 15, the resulting number of boundary voxels of a grain also is relatively large (of the order of 10^4 in 3D), so that δ_i is expensive to compute (Equation 23). It is however possible to use fewer voxels, or (in more general terms) “control points”, to describe grain boundaries.

The control points are obtained by mapping the boundary voxels on a grid, the resolution of which, r , is defined as the ratio between the average diameter ($\langle d \rangle$) and the grid voxel size. For each grain (G_i), and for each voxel of the grid, the boundary voxels located in the grid voxel are merged into a single control point, p_k , defined as their barycenter, and the number of merged voxels is denoted as w_k . Results are shown on Figure 16 and consist of control points that are evenly-spaced along grain boundaries while the spatial resolution in the direction perpendicular to the grain boundary is retained. This would not be the case when using methods such as random sampling or preliminary scaling of the polycrystal image.

The difference between a grain and its associated cell (δ_i , Equation 23) can then be approximated as

$$\delta_i \simeq \sqrt{\sum_{p_k \in G_i^{b'}} w_k d_E(p_k, C_i)^2} \quad (26)$$

and is used for computing the objective function (Equation 24). An appropriate value of the grid resolution (r) will be determined in the case of tessellation generation from real polycrystals in Section 5.3.2.

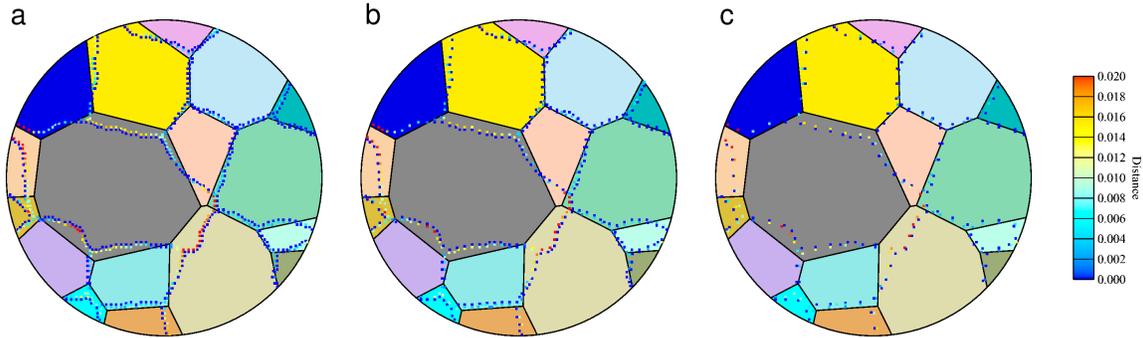


Figure 16: Control points used to compare a polycrystal and a tessellation. The images correspond to different control-point resolutions: (a) $r = 15$, (b) $r = 10$ and (c) $r = 5$. To be compared to the initial resolution, $r = 30$ (Figure 15c).

5.3 Applications

The method is applied to full-resolution polycrystal images in Section 5.3.1. Control points are then used to reduce computational cost, and their influence on the resulting tessellations is discussed in Section 5.3.2.

5.3.1 DCT polycrystals

The method is applied to a 131-grain Al polycrystal and to a 379-grain NaCl polycrystal, both of which were characterized by DCT on beamline ID11 of ESRF [8]. The two polycrystals contain an average of 210,285 and 32,317 voxels per grain, respectively, which corresponds to 11,545 and 3,549 boundary voxels per grain, respectively. As initial solution, for both polycrystals, the coordinates and weights of the seeds are set from the centroids and volumes of the grains, as in Section 4.

The polycrystals and tessellations are shown on Figures 17 and 18, while quantitative results are reported in Table 1. Iteration numbers and computation time will be discussed in Section 5.3.2. Optimization significantly improves the tessellation, as testified by the decrease of the objective function, but also by the evolution of other, more explicit metrics. First, the average distance between the polycrystal and tessellation boundary networks relative to the average grain size, $\langle d_E^b \rangle / \langle d \rangle$, can be used. It is closely related to the distances between the boundary voxels and their associated cells appearing in the objective function (Equation 26); however, by definition and in first approximation, half of these distances are zero (on half of the two covers of a grain boundary). To compute $\langle d_E^b \rangle$ properly, these zero distance values are

discarded. The values of $\langle d_E^b \rangle / \langle d \rangle$ are decreased by optimization down to 4.3% and 6.1% for the two polycrystals, respectively. Second, the intersection rate between the polycrystal and the tessellation, I^V , can be computed,

$$I^V = \frac{1}{V} \sum_{i=1}^N G_i \cap C_i \quad (27)$$

The values reach 91.0% and 87.7% for the two polycrystals, respectively. It should be noticed that the values of both $\langle d_E^b \rangle / \langle d \rangle$ and I^V are relatively close to the DCT resolution [60].

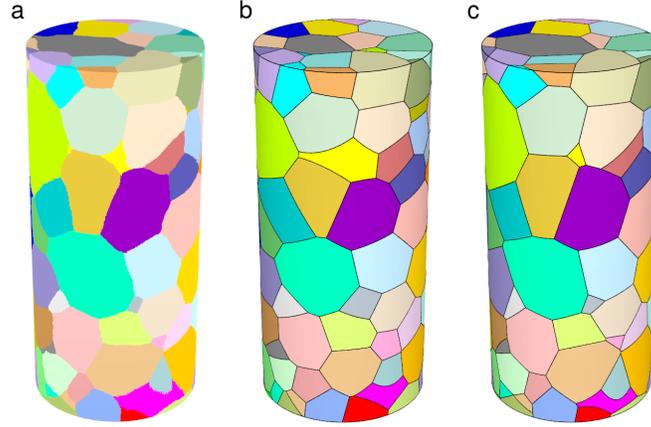


Figure 17: Tessellation generation from the raster image of a 131-grain Al polycrystal. (a) DCT polycrystal image, (b) initial tessellation constructed from the centroids and volumes of the grains, and (c) final tessellation. DCT polycrystal courtesy H. Proudhon (Mines ParisTech, France) and W. Ludwig (ESRF, France).

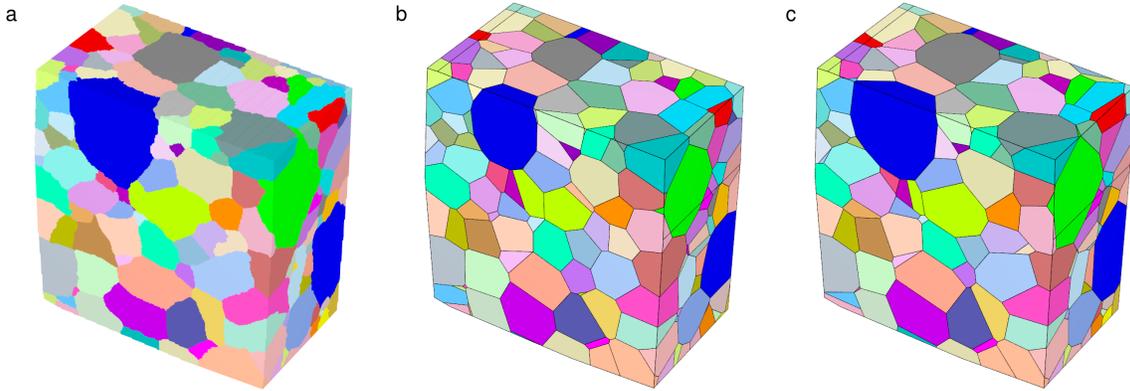


Figure 18: Tessellation generation from the raster image of a 379-grain NaCl polycrystal. (a) DCT polycrystal image, (b) initial tessellation constructed from the centroids and volumes of the grains, and (c) final tessellation. DCT polycrystal courtesy A. Dimanov (LMS, France), M. Bornert (Navier, France) and W. Ludwig (ESRF, France).

Polycrystal	$\mathcal{O} [\times 10^{-4}]$		$\langle d_E^b \rangle / \langle d \rangle$ [%]		I^V [%]	
	Initial	Final	Initial	Final	Initial	Final
131-grain Al	1.45	0.56	11.1	4.4	83.5	91.0
379-grain NaCl	1.97	0.80	11.2	6.1	80.9	87.7

Table 1: Tessellation generation from raster polycrystal images: initial and final values of the objective function (\mathcal{O}), the average relative distance to the grain boundaries ($\langle d_E^b \rangle / \langle d \rangle$) and the intersection rate (I^V).

5.3.2 Influence of the number of control points

In this section, the influence of using control points on the quality of the resulting tessellation is analysed. For each value of the grid resolution (r), a tessellation is generated by optimization as previously and compared to the polycrystal using \mathcal{O} , $\langle d_E^b \rangle / \langle d \rangle$ and I^V . Results are compared to those of the tessellations generated from the full-resolution polycrystal images (Section 5.3.1), which correspond to grid resolutions of 73.8 and 39.5, respectively.

Results are provided in Figure 19. First and as expected, the number of control points strongly depends on r in similar ways for the two polycrystals (Figure 19a). All metrics (\mathcal{O} , $\langle d_E^b \rangle / \langle d \rangle$ and I^V) have values close to the full-resolution ones for r values down to about 5, while lower values decrease tessellation quality, consistently for the two polycrystals. At $r = 5$, the number of control points per grain is reduced to about 100, i.e. 1–3% of the original, while the decrease of accuracy is limited to 0.1–0.5% on \mathcal{O} , $\langle d_E^b \rangle / \langle d \rangle$ and I^V . The number of iterations is nearly independent on r , with about 20,000 and 70,000 iterations for the two polycrystals, respectively. However, using control points decreases computation time from 2.3 hours and 3.1 hours down to 7.7 minutes and 11.6 minutes for the two polycrystals, respectively (each time using 20 calculation cores).

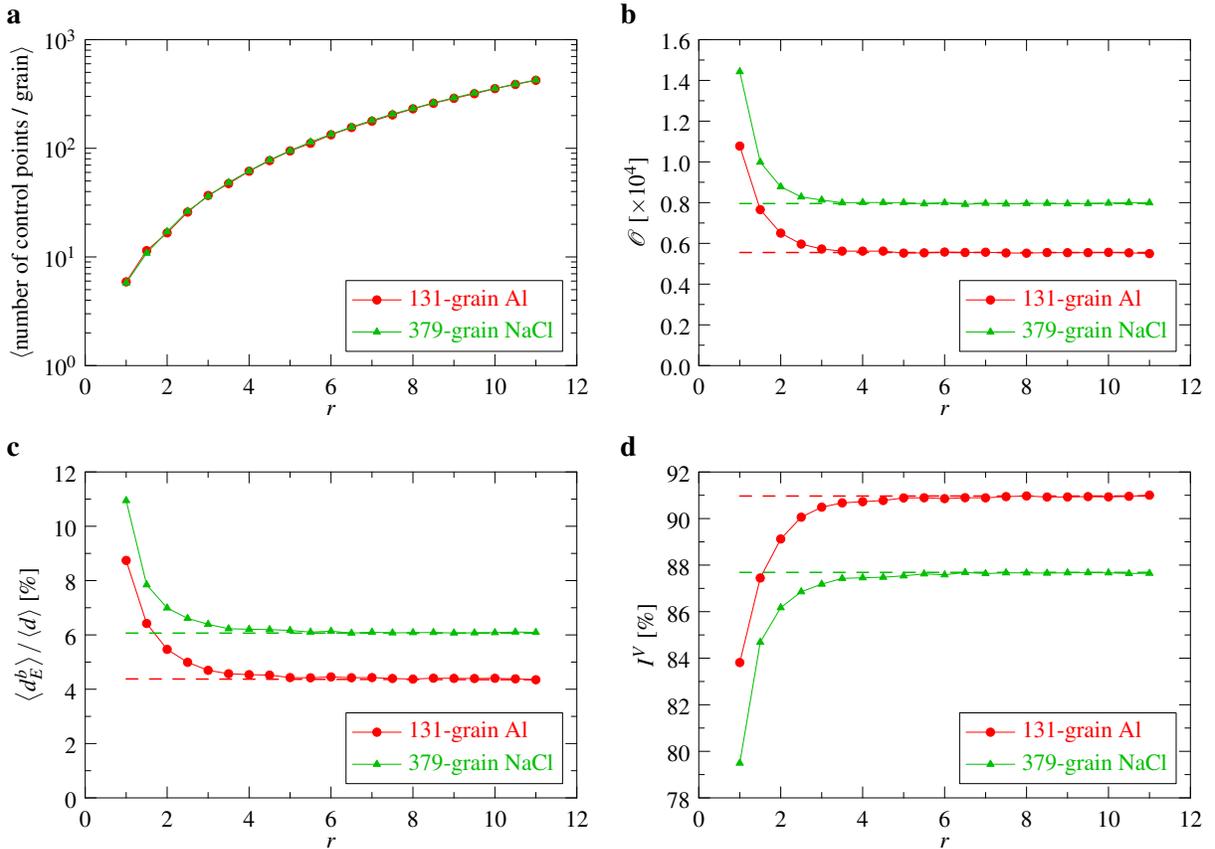


Figure 19: Influence of the grid resolution, r , used to define the control points, on tessellation quality. (a) Evolution of the average number of control points per grain, (b) evolution of the objective function (\mathcal{O}), (c) evolution of the average relative distances between the grain and cell boundaries ($\langle d_E^b \rangle / \langle d \rangle$), and (d) evolution of the intersection rate, I^V . On (b–d), the dashed lines provide the asymptotic values ($r = 73.8$ for 131-grain Al and $r = 39.5$ for 379-grain NaCl).

6 Discussion

6.1 Tessellation generation

The present method for tessellation generation from experimental properties includes several important general characteristics. First, it provides optimal polyhedral representations of polycrystals, which is achieved by using Laguerre tessellations (which can represent any normal tessellation [38]) at their full potential, inside a general optimization framework that set seed attributes without introducing any constraints. This means that the method provides results better than

(or equivalent to) any of its alternatives, as will be discussed below. Laguerre tessellation optimization was made computationally tractable by the development of a new algorithm for Laguerre tessellation update from known seed changes (Section 2), and the use of efficient existing large-scale optimization algorithms [43, 44]. Second, the method is highly-versatile and so can be used for any application: given desired properties, all is needed is to define the objective function as the difference between these properties and the tessellation’s ones. The method is therefore both general in scope and optimal in its achievements, as was shown for different types of experimental inputs.

In the case of grain statistics (Section 3), it was shown that two metrics must be considered simultaneously to properly control the morphology: one for grain size (the diameter) and one for grain shape (the sphericity). These two metrics could be prescribed to replicate experimental grain-growth properties or to construct a “morphology matrix” (Figure 10) where the size and shape statistics are varied independently. This should be useful for systematic analyses of the influence of the microstructure on properties, and recent applications concern thermal resistivity of nanocrystalline materials [61, 62]. Compared to the dense sphere packing approach [36, 37], a first obvious advantage of the present method is that the sphericity is properly taken into account while dense sphere packing necessarily lead to grains of high sphericities. Another advantage is that size is monitored on cells directly in the present approach while only on the spheres in the dense sphere packing approach, which may lead to some inconsistencies.

In the case of 3DXRD data, the method enables to generate a tessellation from individual grain centroids and volumes. The resulting tessellations are useful for interpreting experimental results, for example with respect to grain arrangement, and can be used in finite element simulations. Compared to the direct approach consisting of computing the seed attributes from the grain centroids and volumes (which was used as initial solution), the errors on the centroids and volumes are significantly lower: 0.7% on the centroids, 3.2% on the volumes (0.7% on the diameters) while, for the direct approach: 10.1% on the centroids and 14.4% on the volumes (3.2% on the diameters). The present method also supersedes the one of Lyckegaard and co-authors [40], which provides errors of 6.2% on the centroids and 13.5% on the volumes (9.6% on the diameters). This is an important improvement since such changes can lead to appreciably different grain morphologies and neighbours, as was seen in Section 4.

In the case of polycrystal images provided by DCT and HEDM (Section 5), the main advantage of the present method is the possibility of fully automated meshing. This is at the expense of morphology simplification (as opposed to direct meshing [27, 28]); however, for the materials under investigation, a good agreement was obtained with an intersection rate of about 90%, comparable to experimental uncertainties. This level of agreement would somewhat drop for more complex (non-convex) grain shapes as those generally encountered in more highly alloyed metals [57]. Another advantage of the present method is that the tessellations are the best convex, Laguerre-type approximations of the polycrystal, which can be seen as “cleaned up” descriptors that can be meshed with higher control on element size [26]. This “clean up”, intrinsic to the method, could also be of particular interest in the case of noisier experimental data, as may occur for slicing methods [9, 49], which sometimes include misalignment between the stacked slices. Thanks to its rapidity and robustness, the new method is also an important step toward deeper entanglement of synchrotron X-ray experiments and numerical simulations, as it may enable to generate tessellations and meshes from diffraction data in-line, at the synchrotron beamline, to be used as such or in numerical simulations to guide the rest the experiment.

These are only a few applications and others are well within reach of the method. In all cases, it should be made possible to evaluate the objective function rapidly. As was seen before, optimizing a 1,000-cell tessellation typically involves 10^5 – 10^6 iterations. Tessellation update lasts of the order of 10^{-3} second and the objective function update should not last much longer. This is possible when imposed properties are local variables, such as the size or sphericity, or all other variables in the article. A counterexample is tortuosity [63], which needs to be measured on the full polycrystal even in the case of a local cell change. This difficulty could however be overcome by parallel computation (as was actually used in Section 5).

6.2 Application to the finite element method

In a previous article [26], we presented a method for meshing Voronoi tessellations into high-quality finite elements. We showed that the random character of Voronoi tessellations is problematic for good-quality meshing, which in turn prevented for many years using them in finite element simulations when the number of grains is large or deformation severe. Problems actually arise from the large number of small edges or faces of such tessellations. In Ref. [26], we proposed a correction method, called “regularization”, which replaces small edges and faces by single vertices. Tessellations can then be meshed into high-quality finite elements [26]. It is here of interest to analyse how the edge length distributions of

the new tessellations compare to the ones of Voronoi tessellations and whether regularization can be applied.

Figure 20 provides the edge length distributions of several of the previously-generated tessellations: the 10,000-cell tessellation of Figure 9 and the two tessellations obtained from polycrystal images (Figures 17 and 18). Concerning the grain-growth tessellation (Figure 20a), it appears that the edge length distribution contains a significantly lower number of small edges than for Voronoi tessellations. This can be attributed to the higher cell sphericities, as is apparent on Figure 10. In contrast, tessellations generated from polycrystal images (Figure 20b) exhibit edge length distributions close to the one of Voronoi tessellations. In both cases and as can be seen on Figure 20, regularization can be applied to remove most of the remaining small edges. The resulting tessellations are much more adapted to meshing into high-quality elements. Meshes are shown on Figure 21 and are effectively composed of high-quality elements (as for Voronoi tessellations [26]), which makes them suitable for finite element simulations.

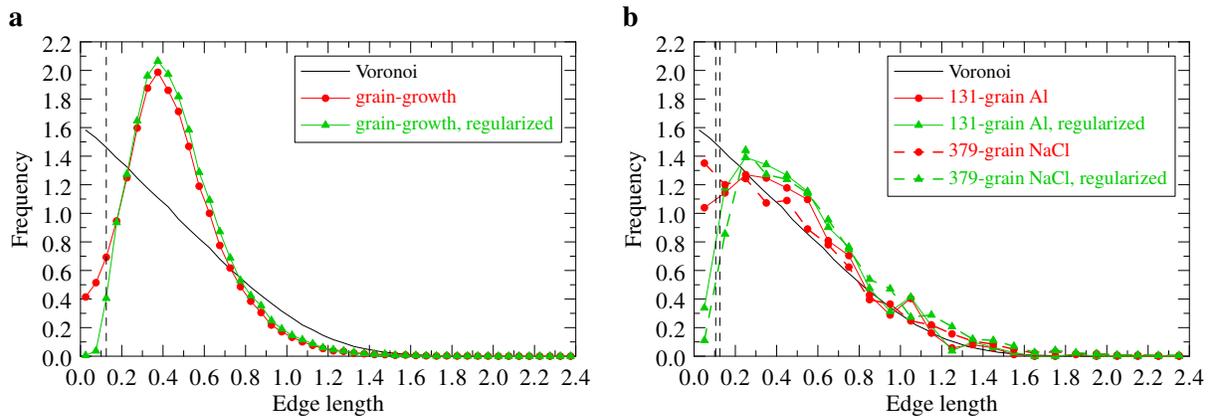


Figure 20: Influence of regularization on the edge lengths of tessellations. (a) 10,000-cell tessellation of Figure 9 and (b) tessellations generated from polycrystal images of Figures 17 and 18. The vertical dashed lines correspond to the edge length thresholds used for regularization (grain-growth: 0.125, 131-grain Al: 0.105, and 379-grain NaCl: 0.125).

7 Conclusion

A new methodology for the generation of optimal polyhedral representations of polycrystals from experimental properties was presented. Laguerre tessellations are used due to their ability to represent any normal tessellation and the attributes of their seeds are set by optimization. This involves a new method for updating tessellations during the optimization process, between successive iterations, which consists of recomputing only the modified cells. The method is general and so can be applied to any type of experimental data. Examples were provided for statistical data and synchrotron X-ray diffraction data. The new method provides a tool to investigate, in a systematic fashion, the influence of the microstructure on physical or mechanical properties of materials. First examples of such studies concern the influence of the grain size distribution on thermal resistivity of nanomaterials [61, 62]. The new method elevates the fidelity of polycrystals used in numerical simulations by incorporating experimentally-derived distributions of the geometric properties to achieve more accurate renditions of the material microstructure. In the absence of experimental data, tessellations generated from grain-growth properties are attractive alternatives to the widely-used Voronoi tessellations.

8 Acknowledgments

This work was partially supported by a public grant overseen by the French National Research Agency (ANR) for project “3DPLASTICITY” (reference ANR-15-CE08-0001). The Al DCT polycrystal was taken from Ref. [64], and the NaCl DCT polycrystal was obtained in the context of the ANR project “MicroNaSel” (reference ANR-10-BLAN-0935) [65, 66]. One of the authors (R.Q.) wishes to thank M. Montanari for discussions about the GJK algorithm. The authors wish to thank P. Dawson and A. Villani for their insightful comments on the article.

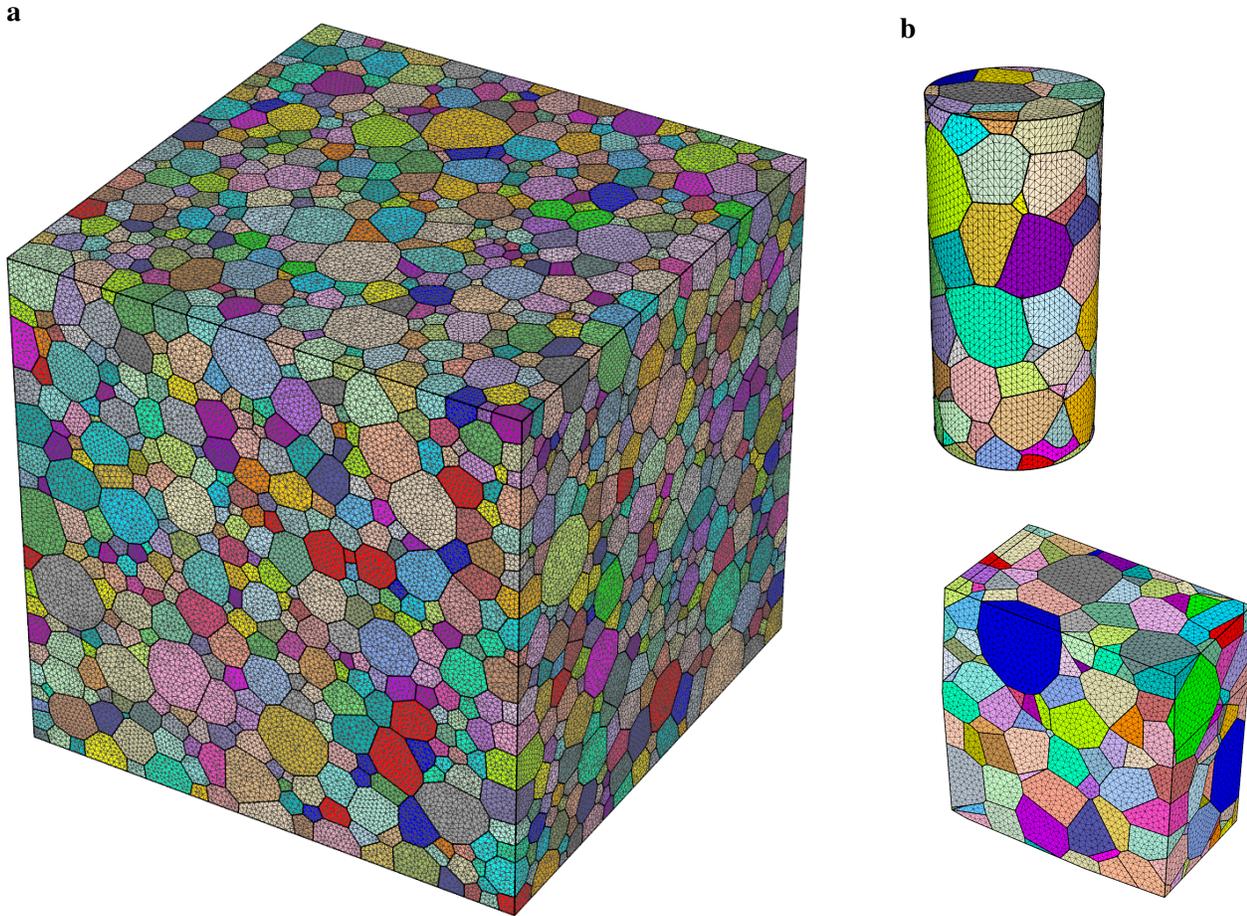


Figure 21: Examples of polycrystal meshes for finite element simulations. (a) 10,000-grain polycrystal corresponding to Figure 9, (b) DCT polycrystals corresponding to Figures 17 and 18.

References

- [1] E.O. Hall. The Deformation and Ageing of Mild Steel: III Discussion of Results. *Proc. Phys. Soc. London*, 64:747–753, 1951.
- [2] N.J. Petch. The Cleavage Strength of Polycrystals. *J. Iron Steel Inst. London*, 173:25–28, 1953.
- [3] F. Alexandre, S. Deyber, and A. Pineau. Modelling the optimum grain size on the low cycle fatigue life of a Ni based superalloy in the presence of two possible crack initiation sites. *Scr. Mater.*, 50:25–30, 2004.
- [4] T. Hanlon, Y.N. Kwon, and S. Suresh. Grain size effects on the fatigue response of nanocrystalline metals. *Scr. Mater.*, 49:675–680S, 2003.
- [5] P. Lehto, H. Remes, T. Saukkonen, H. Hnninen, and J. Romanoff. Influence of grain size distribution on the Hall-Petch relationship of welded structural steel. *Mater. Sci. Eng. A*, 592:28–39, 2014.
- [6] B. Raesinia, C.W. Sinclair, W.J. Poole, and C.N. Tom. On the impact of grain size distribution on the plastic behaviour of polycrystalline metals. *Model. Simul. Mater. Sci. Eng.*, 16(2):025001, 2008.
- [7] F. Barbe, R. Quey, A. Musienko, and G. Cailletaud. Three-Dimensional characterization of strain localization bands in a high-resolution elastoplastic polycrystal. *Mech. Res. Com.*, 36:762–768, 2009.
- [8] W. Ludwig, A. King, M. Herbig, P. Reischig, J. Marrow, L. Babout, E.M. Lauridsen, H. Proudhon, and Y. Buffière. Characterization of polycrystalline materials by combined use of synchrotron x-ray imaging and diffraction techniques. *J. Microsc.*, 62:22–28, 2010.
- [9] R.M. Suter, D. Hennessy, C. Xiao, and U. Lienert. Forward Modeling Method for Microstructure Reconstruction Using X-ray Diffraction Microscopy: Single Crystal Verification. *Rev. Sci. Instrum.*, 77:123905, 2006.
- [10] H.F. Poulsen, L. Margulies, S. Schmidt, and G. Winther. Lattice rotations of individual bulk grains: Part I: 3D X-ray characterization. *Acta Mater.*, 51:3821–3830, 2003.
- [11] J.V. Bernier, J.V. Park, A.L. Pilchak, M.G. Glavicic, and M.P. Miller. Measuring Stress Distributions in Ti-6Al-4V Using Synchrotron X-Ray Diffraction. *Met. Mat. Trans. A*, 39:3120–3133, 2008.
- [12] J. Oddershede, J. Schmidt, H.F. Poulsen, H.O. Sørensen, J. Wright, and W. Reimers. Determining grain resolved stresses in polycrystalline materials using three-dimensional X-ray diffraction. *J. Appl. Cryst.*, 43:539–549, 2010.
- [13] N.Y. Juul, G. Winther, D. Dale, M.K.A. Koker, P. Shade, and J. Oddershede. Elastic interaction between twins during tensile deformation of austenitic stainless steel. *Scr. Mater.*, 120:1–4, 2016.
- [14] E.B. Marin and P.R. Dawson. On modelling the elasto-viscoplastic response of metals using polycrystal plasticity. *Comput. Methods Appl. Mech. Eng.*, 165:1–21, 1998.
- [15] E.B. Marin and P.R. Dawson. Elastoplastic finite element analyses of metal deformations using polycrystal constitutive models. *Comput. Methods Appl. Mech. Eng.*, 165:23–41, 1998.
- [16] F. Barbe, L. Decker, D. Jeulin, and G. Cailletaud. Intergranular and intragranular behavior of polycrystalline aggregates. Part 1: F.E. model. *Int. J. Plast.*, 17:513–536, 2001.
- [17] F. Barbe, S. Forest, and G. Cailletaud. Intergranular and intragranular behavior of polycrystalline aggregates. Part 2: Results. *Int. J. Plast.*, 17:537–563, 2001.
- [18] A. Musienko and G. Cailletaud. Simulation of inter- and transgranular crack propagation in polycrystalline aggregates due to stress corrosion cracking. *Acta Mater.*, 57:3840–3855, 2009.
- [19] R. Quey, P.R. Dawson, and J.H. Driver. Grain orientation fragmentation in hot-deformed aluminium: Experiment and simulation. *J. Mech. Phys. Solids*, 60:509–524, 2012.

- [20] R. Quey, J.H. Driver, and P.R. Dawson. Intra-grain orientation distributions in hot-deformed aluminium: Orientation dependence and relation to deformation mechanisms. *J. Mech. Phys. Solids*, 84:506–527, 2015.
- [21] I. Benedetti and M.H. Aliabadi. A three-dimensional cohesive-frictional grain-boundary micromechanical model for intergranular degradation and failure in polycrystalline materials. *Comp. Methods Appl. Mech. Eng.*, 265:36–62, 2013.
- [22] I. Benedetti, V. Gulizzi, and V. Mallardo. A grain boundary formulation for crystal plasticity. *Int. J. Plast.*, 83:202–224, 2016.
- [23] C.E. Krill III and L.-Q. Chen. Computer simulation of 3-D grain growth using a phase-field model. *Acta Mater.*, 50:3057–3073, 2002.
- [24] G. Abrivard, E.P. Busso, S. Forest, and B. Appolaire. Phase field modelling of grain boundary motion driven by curvature and stored energy gradients. Part II: Application to recrystallisation. *Philos. Mag.*, 92:3643–3664, 2012.
- [25] L. St-Pierre, E. Héripré, M. Dexet, J. Crépin, G. Bertolino, and N. Bilger. 3D simulations of microstructure and comparison with experimental microstructure coming from O.I.M analysis. *Int. J. Plast.*, 24:1516–1532, 2008.
- [26] R. Quey, P.R. Dawson, and F. Barbe. Large-scale 3D random polycrystals for the finite element method: Generation, meshing and remeshing. *Comput. Methods Appl. Mech. Eng.*, 200:1729–1745, 2011.
- [27] S.-B. Lee, G.S. Rohrer, and A.D. Rollett. Three-dimensional digital approximations of grain boundary networks in polycrystals. *Model. Simul. Mater. Sci. Eng.*, 2:025017, 2014.
- [28] H. Proudhon, J. Li, P. Reischig, N. Guéninchault, S. Forest, and W. Ludwig. Coupling Diffraction Contrast Tomography with the Finite Element Method. *Adv. Eng. Mater.*, 18:903–912, 2016.
- [29] A. Spetttl, T. Werz, C.E. Krill III, and V. Schmidt. Parametric Representation of 3D Grain Ensembles in Polycrystalline Microstructures. *J. Stat. Phys.*, 154:913–928, 2014.
- [30] L.L. Shen, G. Proust, C.K.S. Moy, and G. Ranzi. Three-dimensional crystal plasticity finite element simulation of nanoindentation on aluminium alloy 2024. *Mater. Sci. Eng. A*, 579:41–49, 2013.
- [31] Y. Guilhem, S. Basseville, F. Curtit, J.-M. Stéphan, and G. Cailletaud. Numerical investigations of the free surface effect in three-dimensional polycrystalline aggregates. *Comput. Mater. Sci.*, 70:150–162, 2013.
- [32] D. Gonzalez, I. Simonovski, P.J. Withers, and J. Quinta da Fonseca. Modelling the effect of elastic and plastic anisotropies on stresses at grain boundaries. *Int. J. Plast.*, 61:49–63, 2014.
- [33] D. Ledue, A. Maitre, F. Barbe, and L. Lechevallier. Temperature dependence of the exchange bias properties of ferromagnetic/antiferromagnetic polycrystalline bilayers. *J. Magn. Magn. Mater.*, 372:134–140, 2014.
- [34] E. Ghazvinian, M.S. Diederichs, and R. Quey. 3D random Voronoi grain-based models for simulation of brittle rock damage and fabric-guided micro-fracturing. *J. Rock. Mech. Geotech. Eng.*, 6:506–521, 2014.
- [35] T. Auger, S. Hémerly, M. Bourcier, C. Berdin, M. Martin, and I. Robertson. Crack path in liquid metal embrittlement: experiments with steels and modeling. *Fract. Struct. Integ.*, 35:250–259, 2016.
- [36] T. Suzudo and H. Kaburaki. An evolutionary approach to the numerical construction of polycrystalline structures using Voronoi tessellation. *Phys. Lett. A*, 373:4484–4488, 2009.
- [37] C.R. Morfa, M.M. de Farias, I.P.P. de Morales, E.O.I. de Navarra, and R.R. Valera. Virtual modeling of polycrystalline structures of materials using particle packing algorithms and Laguerre cells. *Comp. Part. Mech.*, 2017.
- [38] C. Lautensack. Fitting three-dimensional Laguerre tessellations to foam structures. *J. Appl. Stat.*, 35:985–995, 2008.
- [39] C. Lautensack and S. Zuyev. Random Laguerre tessellations. *Adv. Appl. Prob.*, 40:630–650, 2008.

- [40] A. Lyckegaard, E.M. Lauridsen, W. Ludwig, R.W. Fonda, and H.F. Poulsen. On the Use of Laguerre Tessellations for Representations of 3D Grain Structures. *Adv. Eng. Mater.*, 13:165–170, 2011.
- [41] R. Quey. Neper: polycrystal generation and meshing (version 3.1). <http://neper.sourceforge.net>.
- [42] T.H. Rowan. *Functional Stability Analysis of Numerical Algorithms*. PhD thesis, 1990.
- [43] R. Brent. *Algorithms for Minimization without Derivatives*. Dover, 2003.
- [44] S.G. Johnson. The NLOpt nonlinear-optimization package. <http://ab-initio.mit.edu/wiki/index.php/NLOpt>.
- [45] J.A. Nelder and R. Mead. A simplex method for function minimization. *Comp. J.*, 7:308–313, 1965.
- [46] M.J.D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Comp. J.*, 7:155–162, 1964.
- [47] C. Zhang, A. Suzuki, T. Ishimaru, and M. Enomoto. Characterization of three-dimensional grain structure in polycrystalline iron by serial sectioning. *Metal. Mater. Trans. A*, 35A:1927–1953, 2004.
- [48] K.M. Döbrich, C. Rau, and C.E. Krill III. Quantitative Characterization of the Three-Dimensional Microstructure of Polycrystalline Al-Sn using X-Ray Microtomography. *Metal. Mater. Trans. A*, 35A:1953–1961, 2004.
- [49] D.J. Rowenhorst, A.C. Lewis, and G. Spanos. Three-dimensional analysis of grain topology and interface curvature in a β -titanium alloy. *Acta Mater.*, 58:5511–5519, 2010.
- [50] E.E. Underwood. *Quantitative stereology*. Addison-wesley, 1970.
- [51] M.A. Stephens. EDF Statistics for Goodness of Fit and Some Comparisons. *J. Am. Stat. Assoc.*, 69:730–737, 1974.
- [52] J.C. Tucker, L.H. Chan, G.S. Rohrer, M.A. Groeber, and A.D. Rollett. Tail Departure of Log-Normal Grain Size Distributions in Synthetic Three-Dimensional Microstructures. *Met. Mater. Trans. A*, 43A:2810–22, 2012.
- [53] L. Renversade, R. Quey, W. Ludwig, D. Menasche, S. Madalli, R.M. Suter, and A. Borbély. Comparison between diffraction contrast tomography and high energy diffraction microscopy on a slightly deformed aluminium alloy. *IUCrJ*, 3:32–42, 2016.
- [54] S.B. Lee, A.D. Rollett, and G.S. Rohrer. Three-Dimensional Microstructure Reconstruction Using FIB-OIM. *Mater. Sci. Forum*, 558–559:915–920, 2007.
- [55] M.P. Echlin, M. Straw, S. Randolph, J. Filevich, and T.M. Pollock. The TriBeam system: Femtosecond laser ablation in situ SEM. *Mater. Charact.*, 100:1–12, 2015.
- [56] Z. Zhao, M. Ramesh, D. Raabe, A.M. Cuitiño, and R. Radovitzky. Investigation of three-dimensional aspects of grain-scale plastic surface deformation of an aluminium oligocrystal. *Int. J. Plast.*, 24:2278–2297, 2008.
- [57] R. Quey, D. Piot, and J.H. Driver. Microtexture tracking in hot-deformed polycrystalline aluminium: Experimental results. *Acta Mater.*, 58:1629–1642, 2010.
- [58] C. Ericson. *Real-Time Collision Detection*. Morgan Kaufmann, 2004.
- [59] M. Montanari, N. Petrinic, and E. Barbieri. Improving the GJK algorithm for faster and more reliable distance queries between convex objects. *ACM Trans. Graph.*, 36:30:1–30:17, 2017.
- [60] M. Syha, A. Trenkle, B. Lödermann, A. Graff, W. Ludwig, D. Weygand, and P. Gumbsch. Validation of three-dimensional diffraction contrast tomography reconstructions by means of electron backscatter diffraction characterization. *J. Appl. Cryst.*, 46:1145–1150, 2013.
- [61] B. Mortazavi, R. Quey, H. Ostadhossein, A. Villani, N. Moulin, A.C.T. Van Duin, and T. Rabczuk. Strong thermal transport along polycrystalline transition metal dichalcogenides revealed by multiscale modelling for MoS₂. *Appl. Mater. Today*, 2017.

- [62] M. Sledzinska, R. Quey, B. Mortazavi, B. Graczykowski, M. Placidi, D. Saleta Reig, D. Navarro Urrios, F. Alzina, L. Colombo, S. Roche, and C.M. Sotomayor Torres. Record Low Thermal Conductivity of Polycrystalline MoS₂ films: Tuning the Thermal Conductivity by Grain Orientation. *ACS Appl. Mater. Interfaces*, doi:10.1021/acsami.7b08811.
- [63] J. Laurencin, R. Quey, G. Delette, H. Suhonen, P. Cloetens, and P. Bleuet. Characterisation of Solid Oxide Fuel Cell Ni₈YSZ substrate by synchrotron X-ray nano-tomography: from 3D reconstruction to microstructure quantification. *J. Power Sources*, 198:182–189, 2012.
- [64] W. Ludwig, A. King, P. Reischig, M. Herbig, E.M. Lauridsen, S. Schmidt, H. Proudhon, S. Forest, P. Cloetens, S. Rolland du Roscoat, J.Y. Buffière, T.J. Marrow, and H.F. Poulsen. New opportunities for 3D materials science of polycrystalline materials at the micrometre lengthscale by combined use of X-ray diffraction and X-ray imaging. *Mater. Sci. Eng. A*, 524:69–76, 2009.
- [65] N. Viganò, A. Tanguy, S. Hallais, A. Dimanov, M. Bornert, K.J. Batenburg, and W. Ludwig. Three-dimensional full-field X-ray orientation microscopy. *Sci. Rep.*, 6, 2016.
- [66] M. Bourcier, M. Bornert, A. Dimanov, E. Héripré, and J.L. Raphanel. Multiscale experimental investigation of crystal plasticity and grain boundary sliding in synthetic halite using digital image correlation. *J. Geophys. Res. B*, 118:511–526, 2013.