



Nearly uniform sampling of crystal orientations

Romain Quey, Aurélien Villani, Claire Maurice

► To cite this version:

Romain Quey, Aurélien Villani, Claire Maurice. Nearly uniform sampling of crystal orientations. *Journal of Applied Crystallography*, 2018, 51 (4), pp.1162 - 1173. 10.1107/S1600576718009019 . hal-01850591

HAL Id: hal-01850591

<https://hal.science/hal-01850591>

Submitted on 27 Jul 2018

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.



Nearly uniform sampling of crystal orientations

Romain Quey, Aurélien Villani and Claire Maurice

J. Appl. Cryst. (2018). **51**, 1162–1173



IUCr Journals
CRYSTALLOGRAPHY JOURNALS ONLINE

Copyright © International Union of Crystallography

Author(s) of this paper may load this reprint on their own web site or institutional repository provided that this cover page is retained. Republication of this article or its storage in electronic databases other than as specified above is not permitted without prior permission in writing from the IUCr.

For further information see <http://journals.iucr.org/services/authorrights.html>



Nearly uniform sampling of crystal orientations

Romain Quey,* Aurélien Villani and Claire Maurice

Mines Saint-Etienne, Univ. Lyon, CNRS, UMR 5307 LGF, Centre SMS, F-42023 Saint-Etienne, France. *Correspondence e-mail: romain.quey@mines-stetienne.fr

Received 5 March 2018

Accepted 20 June 2018

Edited by G. Kosterz, ETH Zurich, Switzerland


The authors dedicate this article to the memory of Julian Driver.

Keywords: orientation sampling; crystal symmetry; Thomson problem; quaternions; forward modelling.

A method is presented for generating nearly uniform distributions of three-dimensional orientations in the presence of symmetry. The method is based on the Thomson problem, which consists in finding the configuration of minimal energy of N electrons located on a unit sphere – a configuration of high spatial uniformity. Orientations are represented as unit quaternions, which lie on a unit hypersphere in four-dimensional space. Expressions of the electrostatic potential energy and Coulomb's forces are derived by working in the tangent space of orientation space. Using the forces, orientations are evolved in a conventional gradient-descent optimization until equilibrium. The method is highly versatile as it can generate uniform distributions for any number of orientations and any symmetry, and even allows one to prescribe some orientations. For large numbers of orientations, the forces can be computed using only the close neighbourhoods of orientations. Even uniform distributions of as many as 10^6 orientations, such as those required for dictionary-based indexing of diffraction patterns, can be generated in reasonable computation times. The presented algorithms are implemented and distributed in the free (open-source) software package *Neper*.

1. Introduction

Distributions of three-dimensional orientations are routinely used in studies of the physical and mechanical properties of crystalline materials. Examples of experimental studies involve forward simulations of electron backscatter diffraction (EBSD) (Roşca *et al.*, 2014), electron transmission diffraction (Rauch & Véron, 2014) and high-energy X-ray diffraction (Suter *et al.*, 2006). In these studies, simulated diffraction patterns are compared with experimental ones to identify the experimental orientations (or phases), in an approach referred to as 'dictionary-based indexing' (Rauch & Véron, 2014). Correct indexing requires orientation distributions (for simulating the diffraction patterns) that pave orientation space with a step size of no more than $1\text{--}2^\circ$, which can only be attained by using as many as $10^6\text{--}10^7$ orientations (in the absence of crystal symmetry). For computational efficiency, it is important to minimize this number by distributing orientations as uniformly as possible. Examples of numerical studies in which orientation distributions are used concern simulations of the deformation of polycrystals. Random orientations are usually assigned to the grains, but uniformly distributed orientations could be used to improve orientation space coverage when analysing the orientation dependency of a mechanical or physical property (Wong & Dawson, 2010; Quey *et al.*, 2012, 2015; Wielewski *et al.*, 2017), or to reduce the variability of the polycrystal response in homogenization simulations (Nygårds, 2003). Such numerical studies typically use polycrystals with $10^2\text{--}10^4$ grains, so that orientation sets are significantly smaller than for dictionary-based indexing.



© 2018 International Union of Crystallography

The general problem of generating uniform distributions of (three-dimensional) orientations, or ‘uniform sampling of orientation space’, has been studied for a long time (Helming, 1997; Yang & Chen, 2006; Karney, 2007; Yershova *et al.*, 2010; Yan & Chirikjian, 2012) but only rarely in the context of crystalline materials, which involve symmetries. In the past few years, two methods have been proposed to generate orientation distributions for dictionary-based indexing (Roşca *et al.*, 2014; Larsen & Schmidt, 2017). Both methods parameterize orientations as unit quaternions to take advantage of the proportionality between the intrinsic distance separating unit quaternions and the angle made by their corresponding orientations (Morawiec, 2004). The problem is then to generate a uniform distribution of points on the unit-quaternion hypersphere. Roşca *et al.* (2014) combine two homochoric projections to transform a regular grid of points in the cube into a regular grid of points on the unit-quaternion hypersphere. Larsen & Schmidt (2017) generate a Voronoi tessellation (and its dual, the Delaunay triangulation) by considering the unit quaternions as seed points and combine several optimizations to maximize the minimum distance between any orientation pair. The method of Roşca *et al.* (2014) is direct and therefore very computationally efficient, and it provides orientation grids that are iso-latitudinal (as needed for the spherical harmonics expansions used in texture-related computations). However, it applies only to specific numbers of orientations (n^3), generates orientations that may be relatively close to each other (as equal volumes about orientations do not guarantee equal distances between them) and accounts for crystal symmetry only *a posteriori*, by rejecting orientations located outside the fundamental region, which leads to undesirable boundary effects. The method of Larsen & Schmidt (2017) provides orientation distributions which are effectively more uniform (although not iso-latitudinal), but the method may not behave well for relatively small numbers of orientations owing to an ‘intended use of Euclidean geometries’ (*i.e.* the use of the extrinsic Euclidean distance in four-dimensional space instead of the intrinsic distance on the unit-quaternion hypersphere; Larsen & Schmidt, 2017). Finally, it should be noted that both methods (Roşca *et al.*, 2014; Larsen & Schmidt, 2017) involve relatively complex mathematical developments.

In this article, we describe a relatively simple and versatile method that provides improved uniform sampling of crystal orientations. In §2, we formulate the method and describe how to solve it. Simplifications are made for the case of large numbers of orientations. In §3, we apply the method to several practical cases and compare the results with those of previous methods. In §4, we close the article and present our conclusions. All algorithms described in this article are implemented and distributed in the free (open-source) software package *Neper* (Quey, 2018).

2. Method

The principle of the method is to adapt the Thomson problem (Thomson, 1904) to orientations represented as unit quater-

nions. The original Thomson problem consists of finding the configuration of minimal energy for a set of electrons located on a unit sphere (Thomson, 1904). It is well known that this configuration is characterized by a uniform spatial distribution of electrons, for which all electrons occupy about the same surface area of the sphere and exhibit similar distances to their first neighbours. Although formulated back in 1904, the Thomson problem remains open, and several methods have been proposed to solve it, such as gradient-descent optimization (Erber & Hockney, 1991), global optimization (Altschuler *et al.*, 1994), simulated annealing algorithms (Xiang *et al.*, 1997), Monte Carlo algorithms (Wales *et al.*, 2009; Bondarenko *et al.*, 2015) and genetic algorithms (Morris *et al.*, 1996; Kanimozhi *et al.*, 2016). A major difficulty is that the number of local minima grows exponentially with the number of electrons (Stillinger & Weber, 1984; Erber & Hockney, 1991). However, these minima are very close to the global minimum (Stillinger & Weber, 1984; Erber & Hockney, 1991), so that gradient-descent optimization should provide nearly optimal configurations. Another difficulty is that the problem becomes computationally intensive for large numbers of electrons (say, more than 10 000), as all electrons interact with each other. This is not dealt with in most studies (such as Erber & Hockney, 1991; Altschuler *et al.*, 1994; Xiang *et al.*, 1997; Wales *et al.*, 2009; Bondarenko *et al.*, 2015; Morris *et al.*, 1996; Kanimozhi *et al.*, 2016), which rather focus on finding, and analysing, nearly optimal configurations for relatively small numbers of electrons (less than 1000), and it will therefore need a specific treatment.

When the Thomson problem is transposed to orientations, several differences arise in its formulation. First, while electrons can interact along chords of the unit sphere (in three-dimensional space), orientations (represented as unit quaternions) are defined only on the unit-quaternion hypersphere (the 3-sphere in four-dimensional space) and, as such, should interact along its geodesics. Second, in contrast to an electron being a standalone object, an orientation is represented by two (opposite) quaternions or even more (in the presence of crystal symmetry). In this section, we start by recalling the definition and some useful properties of quaternions (§2.1). We then present a general formulation for interactions between orientations in the general case and in the presence of crystal symmetry (§2.2), and treat the case of large numbers of orientations (§2.3). Finally, we provide a simple and efficient method to solve the problem (§2.4).

2.1. Orientations as unit quaternions

Quaternions are four-dimensional vectors that can be used to describe three-dimensional rotations (Hamilton, 1844). In fact, there is an homomorphism between the multiplicative group of unit quaternions, $SU(2)$, and the group of rotations in \mathbb{R}^3 , $SO(3)$ (Morawiec, 2004). The space of quaternions, \mathbb{H} , is described by the real axis and three imaginary axes along the vectors **i**, **j** and **k**, so that a quaternion, **q**, can be written as $\mathbf{q} = q_0 + q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k}$ (q_0 being the real component and q_{1-3} being the imaginary components). The base vectors **i**, **j**

and \mathbf{k} follow the Hamilton rules for multiplication: $\mathbf{ijk} = \mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = -1$, $\mathbf{ij} = \mathbf{k}$, $\mathbf{jk} = \mathbf{i}$, $\mathbf{ki} = \mathbf{j}$, $\mathbf{ji} = -\mathbf{k}$, $\mathbf{kj} = -\mathbf{i}$ and $\mathbf{ik} = -\mathbf{j}$. A unit quaternion, \mathbf{q} , can be defined from a rotation of angle θ around an axis \mathbf{r} ($\|\mathbf{r}\| = 1$, where ' $\|\cdot\|$ ' denotes the Euclidean norm) as

$$\mathbf{q} = \cos(\theta/2) + \mathbf{r} \sin(\theta/2). \quad (1)$$

Unit quaternions lie on the unit 3-sphere centred at the origin of \mathbb{H} , denoted \mathbb{S}^3 and simply referred to as the 'unit-quaternion hypersphere' throughout the article. For future reference, we note that the (three-dimensional) surface area of \mathbb{S}^3 is $2\pi^2$. Moreover, opposite quaternions, \mathbf{q} and $-\mathbf{q}$, which are antipodal points on \mathbb{S}^3 , correspond to rotations of (θ, \mathbf{r}) and $(\theta + 2\pi, \mathbf{r})$, respectively, and so represent the same orientation. It is therefore sufficient to consider only positive quaternions, *i.e.* quaternions of non-negative real parts, to represent orientations (although it is not required in this work). These quaternions occupy the positive hemisphere of \mathbb{S}^3 , \mathbb{S}^3_+ . Two rotations, \mathbf{q}_1 and \mathbf{q}_2 , both expressed in the reference coordinate system, can be combined into a single rotation by quaternion multiplication as $\mathbf{q}_2\mathbf{q}_1$. Quaternion multiplication is associative, distributive but not commutative. The inverse rotation is given by the conjugate of the original quaternion, which corresponds to inverting the rotation axis [see equation (1)]. Finally, the misorientation between two orientations, \mathbf{q}_1 and \mathbf{q}_2 , expressed in the reference coordinate system, is simply given by $\mathbf{q}_2\mathbf{q}_1^{-1}$.

2.2. Orientation interactions

We formulate the problem of the interaction of orientations represented as electrically charged quaternions. The aim is to determine the energy of the orientation set and the forces at orientations that derive from it. For simplicity, we assume values of Coulomb's constant (k) and the elementary charge (e) equal to 1. We consider a set of two orientations before incorporating crystal symmetry. We then generalize the formulation to a set of N orientations.

2.2.1. Case of a set of two orientations. Let us consider a set of two orientations, \mathbf{q}_i and \mathbf{q}_j , and determine its energy, E , and the resulting forces at these orientations, \mathbf{f}_i and \mathbf{f}_j ($\mathbf{f}_j = -\mathbf{f}_i$), of norm f . As pointed out before, \mathbf{q}_i and \mathbf{q}_j are defined on \mathbb{S}^3 . They therefore interact on \mathbb{S}^3 (not in its embedding space, \mathbb{H}), and so the distances and directions between them are to be measured along geodesics (not chords) of \mathbb{S}^3 ; *i.e.* the intrinsic distance is to be used. This is consistent with the fact that the difference between two orientations, measured by the misorientation angle, is proportional to the intrinsic distance between their quaternions. Working in \mathbb{S}^3 is important to measure distances between orientations properly, especially when the number of orientations is small. It also follows that the forces at orientations are tangential to \mathbb{S}^3 , as can be easily pictured in the three-dimensional case. In other words, the force at an orientation belongs to the tangent space at that orientation. It should therefore be possible to compute it in that space.

The tangent space at an orientation is obtained by logarithm mapping (Morawiec, 2004), so that the position of \mathbf{q}_j in the tangent space at \mathbf{q}_i , \mathbf{d}_{ij} , is given by

$$\mathbf{d}_{ij} = \ln(\mathbf{q}_j\mathbf{q}_i^{-1}) = (\theta_{ij}/2)\mathbf{r}_{ij}, \quad (2)$$

where $\mathbf{q}_j\mathbf{q}_i^{-1}$ corresponds to the misorientation between \mathbf{q}_i and \mathbf{q}_j , and θ_{ij} and \mathbf{r}_{ij} are the corresponding rotation angle and axis (see §2.1). It should be noted that, even if \mathbf{q}_i and \mathbf{q}_j are taken as positive quaternions, $\mathbf{q}_j\mathbf{q}_i^{-1}$ is of arbitrary sign, with the real part defined in $[-1, 1]$, and so $\theta_{ij} \in [0, 2\pi]$. \mathbf{d}_{ij} is a pure imaginary quaternion, and the space to which it belongs is the three-dimensional ball of radius π . The geodesics passing through \mathbf{q}_i on \mathbb{S}^3 correspond to straight lines passing through the origin in tangent space, and distances are conserved along these lines (Morawiec, 2004). The distance between \mathbf{q}_i and \mathbf{q}_j (either on \mathbb{S}^3 or in tangent space), d_{ij} , is

$$d_{ij} = \|\mathbf{d}_{ij}\| = \theta_{ij}/2, \quad (3)$$

and the (signed) direction between \mathbf{q}_i and \mathbf{q}_j is equal to \mathbf{r}_{ij} . Since $\theta_{ij} \in [0, 2\pi]$, we have $d_{ij} \in [0, \pi]$. As the orientation represented by \mathbf{q}_j can also be represented by $-\mathbf{q}_j$, \mathbf{q}_i interacts with both \mathbf{q}_j and $-\mathbf{q}_j$. We may compute the position of $-\mathbf{q}_j$ in the tangent space at \mathbf{q}_i , \mathbf{d}'_{ij} , as

$$\mathbf{d}'_{ij} = \ln(-\mathbf{q}_j\mathbf{q}_i^{-1}) = (\theta'_{ij}/2)\mathbf{r}'_{ij}. \quad (4)$$

Noting that \mathbf{d}_{ij} and \mathbf{d}'_{ij} are mapping opposite quaternions, and using equation (1), it is easy to show that $\theta'_{ij} = 2\pi - \theta_{ij}$ and $\mathbf{r}'_{ij} = -\mathbf{r}_{ij}$. Therefore, \mathbf{d}'_{ij} can be written as

$$\mathbf{d}'_{ij} = -\frac{2\pi - \theta_{ij}}{2}\mathbf{r}_{ij}, \quad (5)$$

and it follows that [also using equation (3)]

$$d'_{ij} = \|\mathbf{d}'_{ij}\| = \frac{2\pi - \theta_{ij}}{2} = \pi - d_{ij}. \quad (6)$$

As $d_{ij} \in [0, \pi]$, we also have $d'_{ij} \in [0, \pi]$. It can finally be noted that \mathbf{d}_{ij} and \mathbf{d}'_{ij} are parallel but point in opposite directions, and that if one of \mathbf{d}_{ij} or \mathbf{d}'_{ij} belongs to the three-dimensional ball of radius $\pi/2$, the other belongs to the three-dimensional spherical shell of inner radius $\pi/2$ and outer radius π , as illustrated in Fig. 1(a). The same calculations could be carried out by considering $-\mathbf{q}_i$ instead of \mathbf{q}_i , which would provide the same distances (d_{ij} and $\pi - d_{ij}$). The total energy of the set of two orientations (E) results from the interactions between the four pairs $(\mathbf{q}_i, \mathbf{q}_j)$, $(\mathbf{q}_i, -\mathbf{q}_j)$, $(-\mathbf{q}_i, \mathbf{q}_j)$ and $(-\mathbf{q}_i, -\mathbf{q}_j)$, and is given by

$$E = 2\left(\frac{1}{d_{ij}} + \frac{1}{\pi - d_{ij}}\right). \quad (7)$$

The force at \mathbf{q}_i , \mathbf{f}_i , derives from the energy that \mathbf{q}_i forms with \mathbf{q}_j and $-\mathbf{q}_j$, which is equal to $E/2$. Moreover, as the force is repulsive, \mathbf{f}_i is opposite to \mathbf{r}_{ij} . So, the force at \mathbf{q}_i can be written in scalar and vector forms as

$$f = -\frac{d(E/2)}{dd_{ij}}, \quad \mathbf{f}_i = -f\mathbf{r}_{ij}. \quad (8)$$

Replacing E by its expression in equation (7) and applying the derivative, we finally obtain

$$f = \frac{1}{d_{ij}^2} - \frac{1}{(\pi - d_{ij})^2}, \quad \mathbf{f}_i = -\left[\frac{1}{d_{ij}^2} - \frac{1}{(\pi - d_{ij})^2} \right] \mathbf{r}_{ij}. \quad (9)$$

Fig. 1(b) shows the evolutions of E and f as a function of d_{ij} . As expected, a minimal value of E and a zero value of f are obtained for $d_{ij} = \pi/2$ ($\theta_{ij} = \pi$), i.e. when \mathbf{q}_i and \mathbf{q}_j represent opposite orientations [there actually is a family of quaternions \mathbf{q}_j representing orientations opposite to \mathbf{q}_i , which corresponds to the sphere of radius $d_{ij} = \pi/2$ in Fig. 1(a) ($\theta_{ij} = \pi$, arbitrary \mathbf{r}_{ij})].

Finally, it should be noted that equations (7) and (9) differ from the corresponding concepts known from electrostatics in several ways. First, antipodal symmetry leads to an equilibrium configuration at a finite distance, where forces vanish. Second, and strictly speaking, the force equilibrium between the two orientations, $\mathbf{f}_j = -\mathbf{f}_i$, can only be written under the concept of parallel transport. Indeed, \mathbf{f}_i and \mathbf{f}_j have the same magnitude (f), but they belong to different tangent spaces and so have different directions in \mathbb{H} . Only when \mathbf{f}_j is transported from \mathbf{q}_j to \mathbf{q}_i (along the geodesics between \mathbf{q}_j and \mathbf{q}_i) does the standard expression of force equilibrium hold true. This will, however, not be used in the following since the forces at orientations will always be computed in their respective tangent spaces.

2.2.2. Crystal symmetry. In the presence of crystal symmetry, an orientation, represented by quaternion \mathbf{q} , can be described by several equivalent quaternions, \mathbf{q}^k , $k = 1, \dots, n_c$, given by

$$\mathbf{q}^k = \mathbf{u}^k \mathbf{q}, \quad (10)$$

where \mathbf{u}^k is a symmetry operator. For cubic crystal symmetry, $n_c = 24$ and \mathbf{u}^k can represent the identity, the three rotations of $\pi/2$ about each of the three $\langle 100 \rangle$ crystal directions, the rotation of π about each of the six $\langle 110 \rangle$ crystal directions and the two rotations of $2\pi/3$ about each of the four $\langle 111 \rangle$ crystal directions. To compute E and \mathbf{f}_i , all equivalent orientations must be taken into account. Let us consider the interaction

between \mathbf{q}_i and \mathbf{q}_j^k . The misorientation axis and angle between \mathbf{q}_i and \mathbf{q}_j^k are denoted as \mathbf{r}_{ij}^k and θ_{ij}^k , respectively. \mathbf{q}_j can be replaced by \mathbf{q}_j^k in equations (2) and (3) to compute the distance between \mathbf{q}_i and \mathbf{q}_j^k , d_{ij}^k . The total energy of the pairs formed by \mathbf{q}_i and all \mathbf{q}_j^k s (and their negatives) can simply be written as the sum of the energies of all orientation pairs [which are of the form of equation (7)]. It should then be noted that, owing to symmetry considerations, the same result would be obtained by considering \mathbf{q}_i^k ($k = 1, \dots, n_c$) instead of \mathbf{q}_i . The total energy, therefore, is

$$E = 2n_c \sum_{k=1}^{n_c} \left(\frac{1}{d_{ij}^k} + \frac{1}{\pi - d_{ij}^k} \right). \quad (11)$$

The net force at an orientation, \mathbf{f}_i , can be written as the sum of the forces between the orientation (\mathbf{q}_i) and all other orientations (\mathbf{q}_j^k) [which are of the form of equation (9)]:

$$\mathbf{f}_i = - \sum_{k=1}^{n_c} \left[\frac{1}{d_{ij}^k} - \frac{1}{(\pi - d_{ij}^k)^2} \right] \mathbf{r}_{ij}^k. \quad (12)$$

2.2.3. General case. Let us now consider a set of N orientations, \mathbf{q}_i , $i = 1, \dots, N$, with crystal symmetry. The total energy of the orientation set (E) can be written as the sum of the energies of all orientation pairs [see equation (11)]:

$$E = 2n_c \sum_{i=1}^N \sum_{j=i+1}^N \sum_{k=1}^{n_c} \left(\frac{1}{d_{ij}^k} + \frac{1}{\pi - d_{ij}^k} \right). \quad (13)$$

The net force at an orientation (\mathbf{f}_i) can be written as the sum of the forces between the orientation and all other orientations [see equation (12)]:

$$\mathbf{f}_i = - \sum_{j=1, j \neq i}^N \sum_{k=1}^{n_c} \left[\frac{1}{d_{ij}^k} - \frac{1}{(\pi - d_{ij}^k)^2} \right] \mathbf{r}_{ij}^k. \quad (14)$$

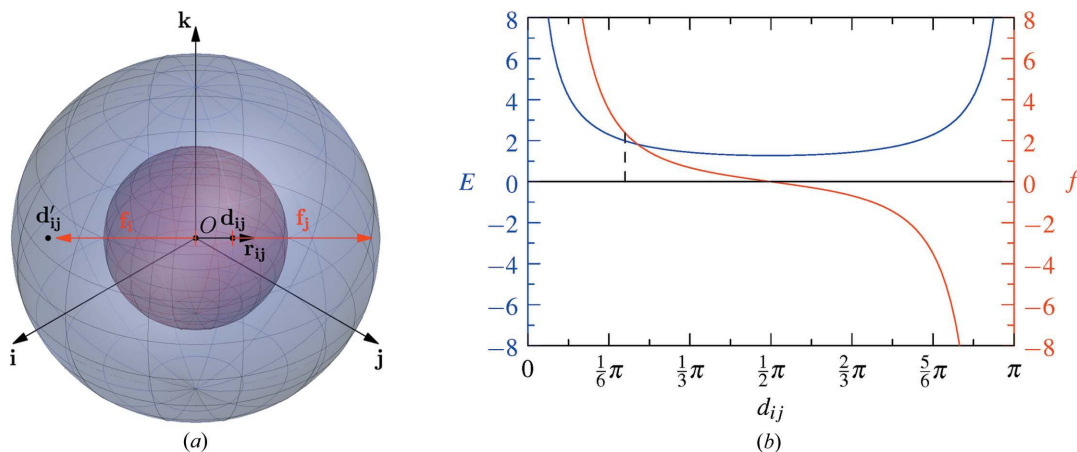


Figure 1 Interaction between two orientations, \mathbf{q}_i and \mathbf{q}_j . (a) Interaction shown in the tangent space at \mathbf{q}_i (which is at the origin, O), where the inner sphere has a radius of $\pi/2$ and the outer sphere has a radius of π . As \mathbf{d}_{ij} and \mathbf{d}_{ij}' are distant by π , if one of them belongs to the inner ball (\mathbf{d}_{ij} here), the other belongs to the outer spherical shell (\mathbf{d}_{ij}' here). \mathbf{f}_i and \mathbf{f}_j tend to move the origin and the closest of \mathbf{d}_{ij} and \mathbf{d}_{ij}' (\mathbf{d}_{ij} here) away from each other until $d_{ij} = d_{ij}' = \pi/2$ (and so $\theta_{ij} = \pi$), i.e. when \mathbf{q}_i and \mathbf{q}_j represent opposite orientations. (b) Evolution of the energy (E) and the force (f) as a function of the distance (d_{ij}), where the dashed line indicates the configuration illustrated in (a).

As expected from the original Thomson problem (Thomson, 1904) and as will be seen in §3, a minimal value of the energy (E) and zero values of the forces (\mathbf{f}_i) are obtained when orientations are uniformly distributed on \mathbb{S}^3 .

2.3. Case of a large number of orientations

As the energy and the forces [provided by equations (13) and (14)] depend on all orientation pairs, the time taken to compute them can be relatively long [$O(N^2)$] when the number of orientations (N) is large. However, it should be noted first that only the forces (not the energy) must be computed during the resolution (see §2.4), and second that the forces should be mostly driven by short-distance interactions, as the (many) long-distance interactions yield low forces that tend to cancel each other out. It should therefore be sufficient to compute the forces in close neighbourhoods of the orientations. In contrast to what was done in §2.2, it is here useful to write the expression of the forces by representing all quaternions (positive and negative) equivalently, so that equation (14) becomes

$$\mathbf{f}_i = - \sum_{\mathbf{q} \in \{\mathbf{q}_j^k, -\mathbf{q}_j^k\}, j \neq i} \frac{1}{d(\mathbf{q}_i, \mathbf{q})^2} \mathbf{r}(\mathbf{q}_i, \mathbf{q}), \quad (15)$$

where $d(\mathbf{q}_i, \mathbf{q})$ and $\mathbf{r}(\mathbf{q}_i, \mathbf{q})$ denote the distance and the misorientation axis between \mathbf{q}_i and \mathbf{q} , respectively (and correspond, respectively, to d_{ij} and \mathbf{r}_{ij} for $\mathbf{q} \in \{\mathbf{q}_j^k\}$, and to d'_{ij} and \mathbf{r}'_{ij} for $\mathbf{q} \in \{-\mathbf{q}_j^k\}$). An approximate value of \mathbf{f}_i can then be obtained by considering only the quaternions located in the neighbourhood of \mathbf{q}_i :

$$\mathbf{f}_i \simeq - \sum_{\substack{\mathbf{q} \in \{\mathbf{q}_j^k, -\mathbf{q}_j^k\}, j \neq i \\ d(\mathbf{q}_i, \mathbf{q}) \leq \delta d_r}} \frac{1}{d(\mathbf{q}_i, \mathbf{q})^2} \mathbf{r}(\mathbf{q}_i, \mathbf{q}), \quad (16)$$

where the neighbourhood corresponds to a hyperspherical cap (on \mathbb{S}^3) of radius δd_r and centre \mathbf{q}_i . δ is a multiplicative factor (to be determined in §3.3), and d_r is the so-called ‘average orientation radius’ and corresponds to the radius of a hyperspherical cap of (three-dimensional) surface area S , defined as

$$S = \pi^2 / N^*. \quad (17)$$

Here, $N^* = N n_c$. S is related to the angular radius of the corresponding hyperspherical cap, θ_r , by (Morawiec, 2004)

$$S = \pi(\theta_r - \sin \theta_r). \quad (18)$$

Knowing N and n_c , θ_r can be computed using equations (17) and (18), and d_r is given by $\theta_r/2$. Knowing δd_r , neighbours are rapidly found using a k -dimensional tree technique [in $O(N \log N)$] (Blanco & Rai, 2014).

2.4. Resolution

The energy of an orientation distribution and the resulting forces being known [equations (13) and (14)], the energy can be minimized using a conventional gradient-descent method, which iteratively evolves the orientations along (and proportionally to) their respective forces.

2.4.1. General formulation. During an iteration (l), each orientation evolves from its initial position, $\mathbf{q}_i^{(l-1)}$, to its final position, $\mathbf{q}_i^{(l)}$. The rotation during the iteration, $\Delta \mathbf{q}_i^{(l-1)}$, is computed from the force at the beginning of the iteration, $\mathbf{f}_i^{(l-1)}$, as

$$\Delta \mathbf{q}_i^{(l-1)} = \alpha^{(l-1)} \mathbf{f}_i^{(l-1)}, \quad (19)$$

where $\alpha^{(l-1)}$ is the step size of the iteration (l) and controls the degree of evolution of the orientations. $\alpha^{(l-1)}$ is the same for all orientations and is computed as described in §2.4.2. [Should some orientations be constrained to their nominal values (as will be done in Fig. 10, §3.4), $\Delta \mathbf{q}_i^{(l-1)} = \mathbf{0}$ would be used.] The orientation evolution during the iteration (*i.e.* the calculation of $\mathbf{q}_i^{(l)}$ from $\mathbf{q}_i^{(l-1)}$ and $\Delta \mathbf{q}_i^{(l-1)}$) can be expressed as a simple additive composition, although it needs more discussion. It should be recalled that while \mathbf{q}_i is a unit quaternion, *i.e.* a point on \mathbb{S}^3 , \mathbf{f}_i and therefore $\Delta \mathbf{q}_i$ belongs to the tangent space at \mathbf{q}_i . This space is the space of pure imaginary quaternions and is centred at the origin of the space of unit quaternions, (1, 0, 0, 0). $\Delta \mathbf{q}_i$ is therefore tangential to \mathbb{S}^3 at (1, 0, 0, 0), not at \mathbf{q}_i . An additive composition requires $\Delta \mathbf{q}_i$ to be rotated to \mathbf{q}_i (through parallel transport), which leads to $\Delta \mathbf{q}_i^*$:

$$\Delta \mathbf{q}_i^* = \Delta \mathbf{q}_i \mathbf{q}_i. \quad (20)$$

The orientation evolution during an iteration can then be written as

$$\mathbf{q}_i^{(l)} = \mathbf{q}_i^{(l-1)} \widehat{+} \Delta \mathbf{q}_i^{*(l-1)}, \quad (21)$$

i.e. the additive composition is subjected to a normalization (represented by ‘ $\widehat{+}$ ’). Normalization is necessary since, during the iteration, and as $\Delta \mathbf{q}_i^*$ is tangential to \mathbb{S}^3 , \mathbf{q}_i moves along \mathbb{S}^3 but also deviates from it to second order.

2.4.2. Step size and termination criterion. The step size ($\alpha^{(l-1)}$) used in equation (19) must be set appropriately throughout optimization, so that the fewest possible iterations are needed. This could be done by determining, at each iteration and through a line search, the value of α that most decreases the energy. However, this would require many (expensive) computations of E (one for each value of α tested). Instead, we adopt some assumptions to compute values of $\alpha^{(l-1)}$ from information known *a priori* or from the previous iteration.

The value of α at the first iteration, $\alpha^{(0)}$, is defined so that it provides a maximal decrease of E , *i.e.* a minimal value of $E^{(1)}$, for a typical random orientation distribution. A minimal value of $E^{(1)}$ is obtained for specific amplitudes of the rotations, $\|\Delta \mathbf{q}_i^{(0)}\|$, or, equivalently [as all rotations evolve according to equation (21)], for a specific value of their average, $\langle \|\Delta \mathbf{q}_i^{(0)}\| \rangle$. The average amplitude of the rotations ($\langle \|\Delta \mathbf{q}_i^{(0)}\| \rangle$) can be expressed as a fraction, β , of the average orientation radius (d_r). This can be formulated as

$$\langle \|\Delta \mathbf{q}_i^{(0)}\| \rangle = \beta d_r \text{ with } \langle \|\Delta \mathbf{q}_i^{(0)}\| \rangle \text{ such that } E^{(1)} \text{ is minimum.} \quad (22)$$

By combining equations (19) and (22), $\alpha^{(0)}$ can be written as

$$\alpha^{(0)} = \frac{\beta d_r}{\langle \|\mathbf{f}_i^{(0)}\| \rangle}, \quad (23)$$

where $\langle \|\mathbf{f}_i^{(0)}\| \rangle$ corresponds to the average of the forces at orientations in the initial configuration. Note that $\mathbf{f}_i^{(0)}$ is of the form $1/d_{ij}^2$ [neglecting the term in $1/(\pi - d_{ij})^2$ in equation (9), as in §2.3] and therefore depends mostly on interactions within the close neighbourhood of $\mathbf{q}_i^{(0)}$. The distances to the orientations within the neighbourhood (d_{ij}) are proportional to d_r , and so $\langle \|\mathbf{f}_i^{(0)}\| \rangle$ is inversely proportional to $1/d_r^2$:

$$\langle \|\mathbf{f}_i^{(0)}\| \rangle = \gamma/d_r^2. \quad (24)$$

By combining equations (23) and (24), $\alpha^{(0)}$ can finally be written as

$$\alpha^{(0)} = (\beta/\gamma)d_r^3. \quad (25)$$

It is then interesting to determine the values of β and γ independently, which can be done using, as values of $\mathbf{q}_i^{(0)}$, a random distribution of 10^5 orientations (considered sufficiently large to give representative results). The value of β , as defined in equation (22), can be determined by a line search (see Fig. 2), which provides $\beta = 0.8$. In other words, at the first iteration, a maximal decrease of the energy (E) is obtained when orientations are rotated on average by 80% of the average orientation radius (d_r). The value of γ , as defined in equation (24), can be computed directly, which provides $\gamma = 3.4$.

The values of α at the following iterations are determined using the Barzilai–Borwein method (Barzilai & Borwein, 1988). In general terms, the Barzilai–Borwein method assumes a quadratic form of the minimized function and computes α from the evolutions, during the previous iteration, of the solution vector and of the gradient of the minimized function. In this work, the minimized function is the energy (E), the solution vector corresponds to the \mathbf{q}_i values and its evolution to the $\Delta\mathbf{q}_i$ values, and the gradient of the minimized function corresponds to the forces (\mathbf{f}_i) and its evolution to the $\Delta\mathbf{f}_i^{(l-1)} = \mathbf{f}_i^{(l)} - \mathbf{f}_i^{(l-1)}$ values. The Barzilai–Borwein method provides two possible values of $\alpha^{(l-1)}$ (which result from two different least-squares problems) (Barzilai & Borwein, 1988):

$$\alpha_1^{(l-1)} = \frac{\sum_{i=1}^N (\Delta\mathbf{q}_i^{(l-2)} \cdot \Delta\mathbf{f}_i^{(l-2)})}{\sum_{i=1}^N (\Delta\mathbf{f}_i^{(l-2)} \cdot \Delta\mathbf{f}_i^{(l-2)})} \quad \text{and} \quad (26)$$

$$\alpha_2^{(l-1)} = \frac{\sum_{i=1}^N (\Delta\mathbf{q}_i^{(l-2)} \cdot \Delta\mathbf{q}_i^{(l-2)})}{\sum_{i=1}^N (\Delta\mathbf{q}_i^{(l-2)} \cdot \Delta\mathbf{f}_i^{(l-2)})}.$$

Which of $\alpha_1^{(l-1)}$ and $\alpha_2^{(l-1)}$ is the most efficient in decreasing E is unknown *a priori*, and a common practice is to alternate between them during optimization. In this work, we use $\alpha_1^{(l-1)}$ on odd iterations and $\alpha_2^{(l-1)}$ on even iterations. Finally, the optimization terminates on the basis of the relative error on the forces (a residual):

$$\frac{\left(\sum_{i=1}^N \|\mathbf{f}_i^{(l)}\|^2\right)^{1/2}}{\left(\sum_{i=1}^N \|\mathbf{f}_i^{(0)}\|^2\right)^{1/2}} < \varepsilon_r. \quad (27)$$

In this article (and unless mentioned otherwise), we use $\varepsilon_r = 10^{-3}$, which provides a good compromise between accuracy and computation time.

3. Results

The method is validated and applied to several practical cases. In §3.1, we describe how the uniformity of an orientation distribution is measured. In §3.2, we analyse the method in detail through the evolution of an orientation distribution during optimization. In §3.3, we discuss the influence of restraining the force calculation to the close neighbourhood of an orientation. In §3.4, we provide several examples of applications, including the case of symmetry and the specification of some orientations, and we compare the results with those of previous methods (Roşca *et al.*, 2014; Larsen & Schmidt, 2017).

3.1. Methodology

The uniformity of orientation distributions is analysed using two metrics. The first metric relates to the orientation distribution function (ODF), $f(\mathbf{q})$, and we will focus on its standard deviation, σ_f . The ODF $[f(\mathbf{q})]$ is computed using a discrete method, in which each orientation (\mathbf{q}_i) [and its symmetry equivalents (\mathbf{q}_i^k)] is represented by a continuous Bingham distribution (which acts as a smoothing kernel) (Morawiec, 2004). The ODF can therefore be written as

$$f(\mathbf{q}) = \frac{1}{Nn_c} \sum_{i=1}^N \sum_{k=1}^{n_c} \left[{}_1F_1\left(\frac{1}{2}, 2, \mathbf{K}\right)^{-1} \exp\left(\mathbf{w}_i^k \mathbf{U} \mathbf{K} \mathbf{U}^t \mathbf{w}_i^k\right) \right] \quad (28)$$

with $\mathbf{w}_i^k = \mathbf{q}_i^k \mathbf{q}_i^{-1}$, where \mathbf{U} is an orthonormal 4×4 matrix, \mathbf{K} is a diagonal matrix describing the shape of the distributions and ${}_1F_1$ is the hypergeometric function with matrix argument normalizing the density of the Bingham distribution. In this work, we consider isotropic Bingham distributions of (unidirectional) standard deviation equal to d_r , as obtained using $\mathbf{U} = \mathbf{I}_4$, $K_{11} = 2/d_r^2$ and $K_{22} = K_{33} = K_{44} = 0$. Considering d_r as the standard deviation provides a level of




Figure 2
Identification of parameter β , entering the definition of $\alpha^{(0)}$ [see equation (25)], using a random distribution of 10^5 orientations.

smoothing that is independent of the number of orientations (N^*).

Different orientation parameterizations can be used to visualize the ODF. In the absence of crystal symmetry, the homochoric parameterization (Roşca *et al.*, 2014), \mathbf{x}_h , is used:

$$\mathbf{x}_h = \mathbf{r} \left[\frac{3}{4}(\theta - \sin \theta) \right]^{1/3}, \quad (29)$$

which maps orientation space onto the three-dimensional ball of radius $(3\pi/4)^{1/3}$ (as will be seen in Fig. 3; §3.2). In the presence of crystal symmetry, it is sufficient to represent the ODF in the fundamental region of orientation space. The Rodrigues parameterization, \mathbf{x}_R , is used:

$$\mathbf{x}_R = \mathbf{r} \tan(\theta/2), \quad (30)$$

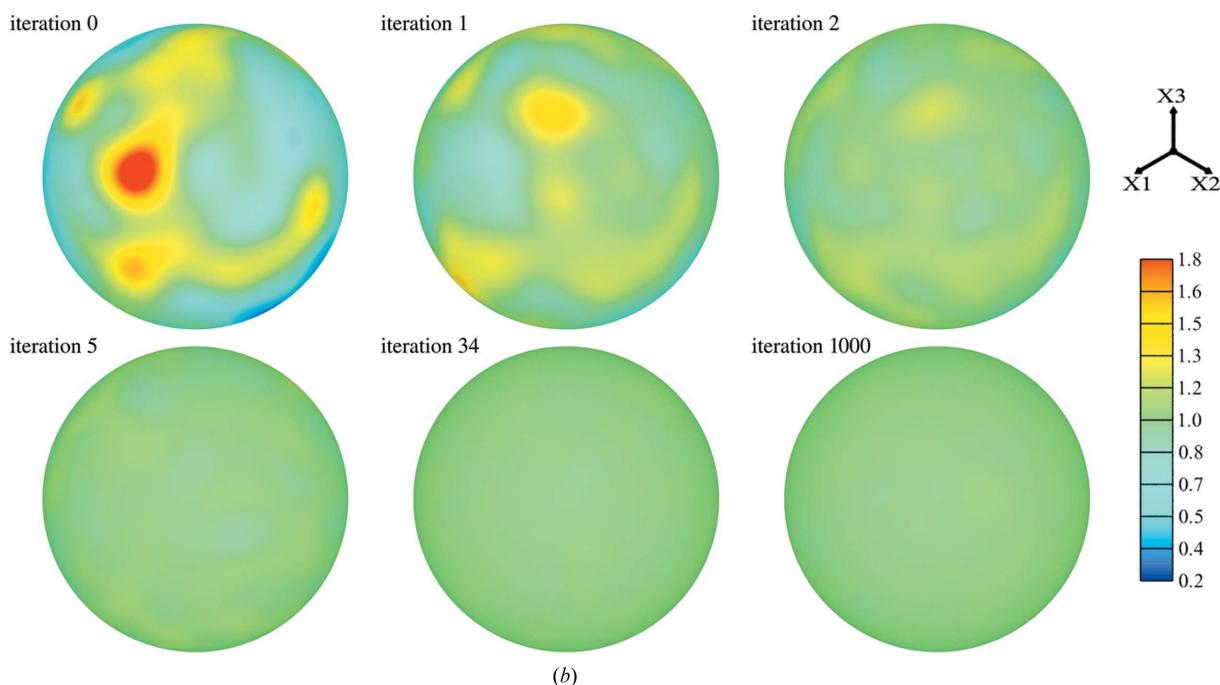



Figure 3 Generation of a uniform distribution of 1000 orientations (no symmetry). (a) Evolution of the energy (E) and the ODF standard deviation (σ_f). The dashed lines represent the asymptotic solutions (1000 iterations). (b) Orientation distribution functions shown on the homochoric projection ball at iterations 0 (initial solution), 1, 2, 5, 34 (final solution) and 1000 (asymptotic solution). Note how the evolution of σ_f is non-monotonic (the same applies to E), which is typical for the Barzilai–Borwein method.

which maps the fundamental region of orientation space into a convex polyhedron whose size and shape depend on the crystal symmetry. For cubic crystal symmetry, the polyhedron is a truncated cube, while for hexagonal crystal symmetry, the polyhedron is a dodecagonal prism (as will be seen in Fig. 5; §3.4). The more uniform an orientation distribution, the closer to 1 its ODF [$f(\mathbf{q})$] and the lower σ_f . σ_f has a value of 0.3 for a typical random orientation distribution and a minimal value of 0 for a perfectly uniform orientation distribution. It should, however, be noted that, as the ODF is computed from a finite number of orientations, σ_f will never reach the minimal value of 0. It instead attains a slightly higher ‘residual’ value that is independent of N^* (because of the chosen dependence of d_r on N^*). The optimal covering in a local area should be the body-centred cubic (b.c.c.) lattice, since this is the best known covering in \mathbb{R}^3 (Larsen & Schmidt, 2017); the corresponding value of σ_f is 0.0021, which can be considered as an absolute minimal value of σ_f .

The second metric is the function corresponding to the misorientation between an arbitrary orientation and the nearest orientation of the orientation distribution, $\theta_m(\mathbf{q})$, and we will focus on properties such as its average or maximum. Unlike $f(\mathbf{q})$, which is a fundamental and widely used metric, $\theta_m(\mathbf{q})$ is relevant specifically in the context of dictionary-based indexing. Although the two metrics are related, they are not strictly equivalent, as $f(\mathbf{q})$ considers that orientations are represented by a kernel (and therefore represent a small region) while $\theta_m(\mathbf{q})$ depends only on the closest orientation. As a result, the two metrics may not evolve in the same way, as will be seen in the following.

3.2. Orientation evolution during optimization

The evolution of a distribution of 1000 orientations during optimization is illustrated in Fig. 3 (no symmetry). In Fig. 3(a), it can be seen that the energy (E) shows a rapid decrease at the beginning of optimization and then saturates. The energy decreases only from 7.75×10^5 to 7.69×10^5 but, at the same time, the ODF standard deviation (σ_f) decreases significantly, from 0.29 to 0.007. This reveals that the local arrangement of the orientations or, more specifically, its degree of uniformity (which is related to σ_f) represents only a small fraction of the energy (all the more that N is large). Still, minimizing the energy appears to be very efficient in decreasing σ_f . Fig. 3(b) provides the ODF at successive iterations, and it can be seen how the ODF, which initially shows large variations, becomes more and more uniform during optimization to eventually reach values very close to 1. In Fig. 3, the ODF is shown only on the surface of the homochoric region, but this should not be considered a limitation since surface values are representative of interior values; slice views will be provided for comparison with other methods (see Fig. 7 in §3.4).

3.3. Influence of the local calculation of the forces

The influence of restraining the force calculation to a close neighbourhood (as described in §2.3) on the resulting ODF standard deviation (σ_f) can be analysed. The aim is to determine how small the neighbourhood can be (minimal value of δ) while still providing a value of σ_f similar to the nominal value. The procedure is as follows: distributions of 10^5 orientations (no symmetry) are generated for increasing values of δ , and their respective ODF standard deviations (σ_f) are computed. The evolution of σ_f as a function of δ is plotted in Fig. 4 for two values of the residual, $\varepsilon_r = 10^{-3}$ (the default) and $\varepsilon_r = 10^{-4}$. For both values, and as expected, σ_f tends to decrease as δ increases. For $\varepsilon_r = 10^{-3}$, σ_f becomes close to the nominal value (and lower than 0.01) for $\delta \simeq 20$. For larger




Figure 4

Influence of the size of the neighbourhood (δ) on the ODF standard deviation (σ_f). The dashed lines indicate the asymptotic values. A distribution of $N = 10^5$ orientations was used, with no symmetry. For $\varepsilon_r = 10^{-3}$, the fluctuations of σ_f for $\delta > 20$ are due to the relatively large variations of the residual near the termination criterion, which are in turn due to the relatively large value of ε_r (10^{-3}) and to the Barzilai–Borwein method.

Table 1

Computation times for generating uniform distributions for several numbers of orientations (N) and crystal symmetries.

For $N \geq 10^5$, forces are computed in neighbourhoods ($\delta = 20$). Computation times are averaged over 50 simulations (starting from different random orientation distributions). Simulations are run on ten cores (20 threads) of an Intel Xeon CPU E5-2660 v3 processor. Note that the computation times are similar for all symmetries for $N = 10^6$, which is a consequence of the force computation in close neighbourhoods.

| N | Triclinic | Cubic | Hexagonal |
|--------|------------|------------|------------|
| 10^3 | 3 s | 20 s | 13 s |
| 10^4 | 4 min | 36 min | 22 min |
| 10^5 | 17 min | 45 min | 31 min |
| 10^6 | 3 h 30 min | 3 h 30 min | 3 h 30 min |

values of δ , fluctuations of σ_f appear, which are caused by the variations of the residual as it approaches ε_r (at the end of optimization). As can be seen in Fig. 4, $\varepsilon_r = 10^{-4}$ leads to lower values of σ_f (and no apparent fluctuations), albeit at the cost of longer computations and larger values of δ to approach the nominal value. A value of δ of 20 corresponds to 7500 orientations in the neighbourhood, and so restraining the force calculation to close neighbourhoods (with $\delta = 20$) reduces the computation time for $2N^* > 7500$.

3.4. Practical applications

The method is first applied to the generation of uniform distributions of 1000 orientations with cubic and hexagonal crystal symmetries. The ODFs corresponding to the initial, random orientation distributions and to the final, uniform orientation distributions are shown in Fig. 5. Similarly to the case of triclinic symmetry (Fig. 3b), the ODFs corresponding to the uniform orientation distributions exhibit significantly lower standard deviations (σ_f). During the simulation, σ_f actually decreases from 0.27 to 0.007 for cubic symmetry and from 0.32 to 0.007 for hexagonal crystal symmetry.

The method can be applied to arbitrary numbers of orientations (N). Typical computation times are reported in Table 1 for several values of N and several crystal symmetries. Fig. 6 provides the values of σ_f for N ranging from 10^2 to 10^6 in the case of cubic crystal symmetry, for the present method and for previous methods (Roşca *et al.*, 2014; Larsen & Schmidt, 2017). It appears that the present method leads to a value of σ_f of about 0.007 independently of N , which is significantly lower than those obtained by previous methods (Roşca *et al.*, 2014; Larsen & Schmidt, 2017). More specifically, the method of Roşca *et al.* (2014) provides relatively low values of σ_f only for high values of N , but these values are still 10–20 times higher than for other methods, while the method of Larsen & Schmidt (2017) provides a value of 0.024 independently of N . Fig. 6 also shows that, unlike the present method, previous methods apply only to specific values of N [those available online for the method of Larsen & Schmidt (2017)]. A closer comparison can be made between the different methods by inspecting their ODFs, as illustrated in Fig. 7 for $N = 9218$ [$N = 9077$ for Roşca *et al.* (2014)]. It can be seen that, unlike other methods, the method of Roşca *et al.* (2014) leads to strong




Figure 5
Comparison of a random distribution (initial solution) and a uniform distribution (final solution) of 1000 orientations in the presence of crystal symmetry. ODFs are shown in the fundamental regions of Rodrigues space. (a) Cubic crystal symmetry and (b) hexagonal crystal symmetry.

variations of the ODF near the surface of the fundamental region. This results from the fact that crystal symmetry is only taken into account *a posteriori*, by discarding orientations located outside the fundamental region (Roşca *et al.*, 2014). As for the method of Larsen & Schmidt (2017), it is confirmed that the ODF shows stronger variations compared to the present method.

The different methods can also be compared in terms of the distribution of the misorientations to the nearest orientation $[\theta_m(\mathbf{q})]$, as illustrated in Fig. 8. A distribution of 262 144 orientations [259 385 orientations for the method of Roşca *et al.* (2014)] with cubic crystal symmetry is considered, and the frequency distribution of $\theta_m(\mathbf{q})$ was computed from the values at 10^8 arbitrary orientations. It can be seen that the present method and the method of Larsen & Schmidt (2017) generate similar distributions, with an average value of $\sim 0.63^\circ$, which is lower than that of the method of Roşca *et al.* (2014), 0.67° . Larger differences arise for the maximal values: 1.84° for the method of Roşca *et al.* (2014), 0.99° for the method of Larsen & Schmidt (2017) and 1.17° for the present method. As a matter of fact, 0.2% of the values of $\theta_m(\mathbf{q})$ exceed 0.99° for the present method, while 4.5% do so for the method of Roşca *et al.* (2014). The fact that the method of Larsen & Schmidt (2017) produces a lower maximal value of $\theta_m(\mathbf{q})$ (but a higher value of σ_f) can be attributed to its being specifically designed for that purpose. Actually, the Thomson problem solved in this work and the packing problem solved by Larsen & Schmidt (2017) are two specializations of the generalized Thomson problem of minimizing a Riesz energy of the form $\sum_{i \neq j} 1/d_{ij}^s$

(Dragnev *et al.*, 2002). For the Thomson problem, $s = 1$ and the energy depends on the interaction between all pairs, while for the packing problem, $s = \infty$ and only the orientation pair with the smallest distance matters (Dragnev *et al.*, 2002). This explains why the present method provides better results on σ_f (which depends on interactions within a neighbourhood) while the method of Larsen & Schmidt (2017) provides better results on the maximum value of $\theta_m(\mathbf{q})$ (which is local).

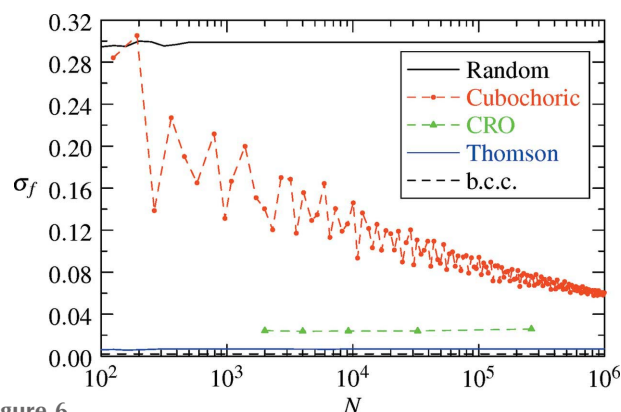


Figure 6
ODF standard deviation (σ_f) for uniform orientation distributions in the case of cubic crystal symmetry, as obtained by different methods: cubochoric after Roşca *et al.* (2014), CRO after Larsen & Schmidt (2017) and Thomson (this work). Values for a random distribution and for the b.c.c. lattice are provided for reference. The cubochoric and CRO data are shown by the symbols while the dashed lines are added for visualization.

The improvement provided by a uniform orientation distribution (compared to a random orientation distribution) can finally be analysed from a different perspective: instead of determining how σ_f can be decreased (for a fixed N), we can determine how N can be decreased (for a fixed σ_f), as illustrated in Fig. 9. This corresponds to looking at how the number of orientations can be decreased while still getting the same ‘coverage’ of orientation space (measured by σ_f). This can be done by computing the ODFs slightly differently than previously, using the same smoothing kernel for all values of N . We can then analyse how σ_f changes as a function of N , for random and uniform orientation distributions. In Fig. 9, it can

be seen that, in the case of cubic crystal symmetry, a uniform distribution of 130 orientations has the same value of σ_f as a random distribution of 1000 orientations.

A last application concerns the generation of a uniform orientation distribution including prescribed orientations. This is desirable in some applications, for example to analyse the behaviour of specific (experimental) orientations (Quey *et al.*, 2012, 2015). Fig. 10 shows a distribution of 1000 orientations among which 82 orientations were positioned along the $X_i \parallel \langle 100 \rangle$ and $X_i \parallel \langle 110 \rangle$ orientation fibres. It can be seen that, despite the constraints imposed by the prescribed orientations, the final orientation distribution is as uniform as in the

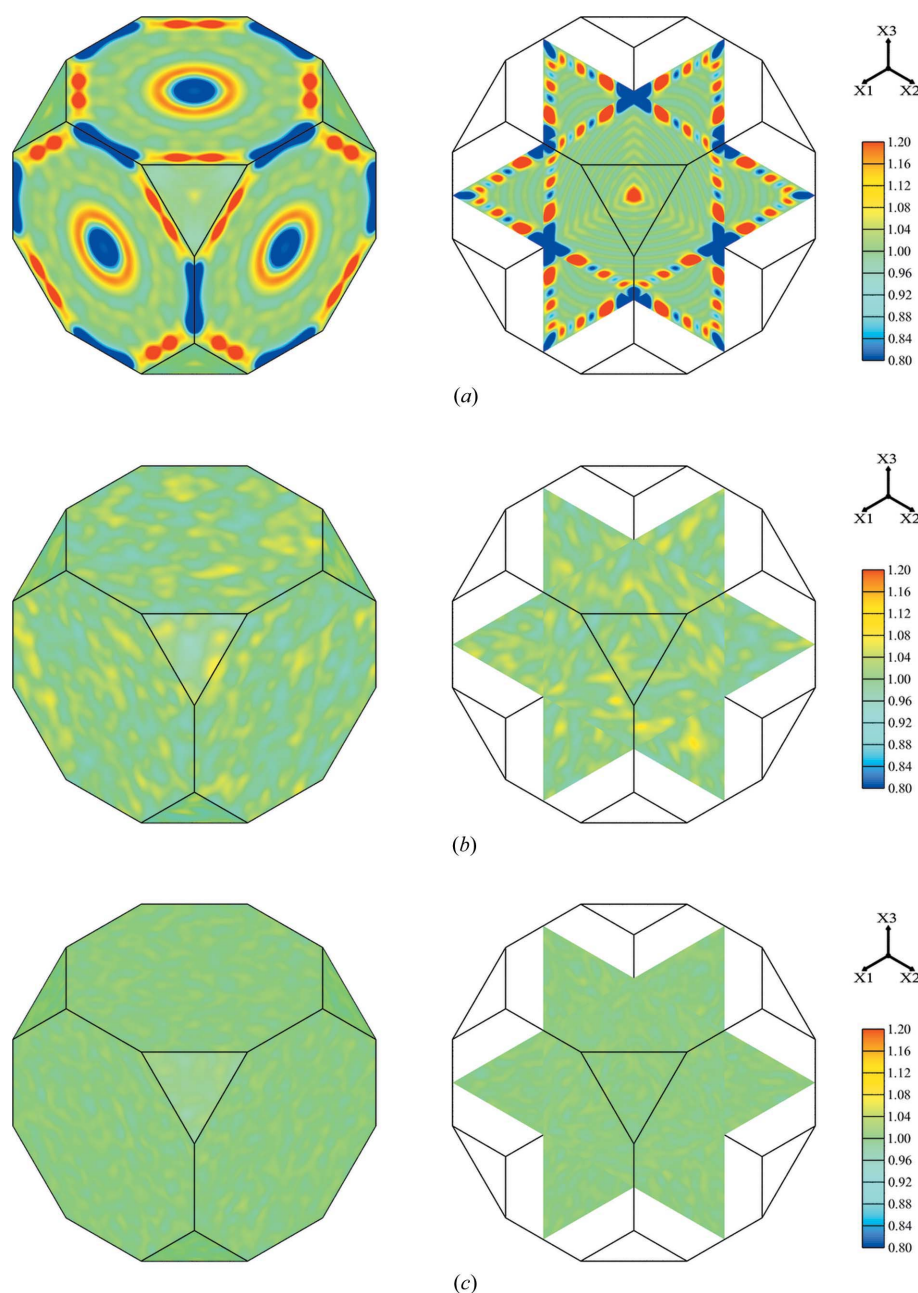


Figure 7

ODFs provided by different methods, for ~9200 orientations and cubic crystal symmetry: (a) Roşca *et al.* (2014) (9077 orientations), (b) Larsen & Schmidt (2017) (9218 orientations) and (c) this work (9218 orientations). ODFs are shown as surface and slice views. Note that the colour scale has been reduced to [0.8, 1.2] compared to Figs. 3 and 5.

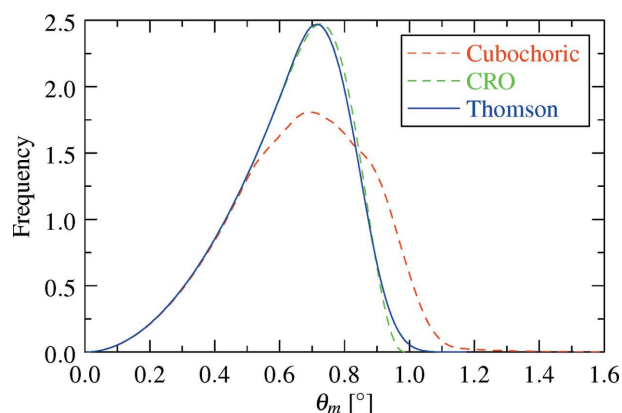


Figure 8
Frequency distributions of the misorientations to the nearest orientation $[\theta_m(\mathbf{q})]$ in the case of cubic crystal symmetry, as obtained by different methods: cubochoric after Rošca *et al.* (2014), CRO after Larsen & Schmidt (2017) and Thomson (this work). $N = 259\,385$ for cubochoric (after Rošca *et al.*, 2014), and $N = 262\,144$ for CRO (after Larsen & Schmidt, 2017) and Thomson (this work). The average values are 0.66, 0.63 and 0.63° for the three methods, respectively, and the maximal values are 1.84 (beyond the limits shown), 0.99 and 1.17° for the three methods, respectively. The cubochoric distribution has been scaled to correct for its slightly smaller number of orientations.




Figure 9
Comparison of the orientation space coverages (measured by σ_f) obtained for random and uniform orientation distributions, in the case of cubic crystal symmetry. A fixed smoothing kernel is used (which corresponds to the nominal kernel for 1000 orientations). The dashed lines indicate how a uniform orientation distribution of 130 orientations provides the same orientation space coverage as a random orientation distribution of 1000 orientations.

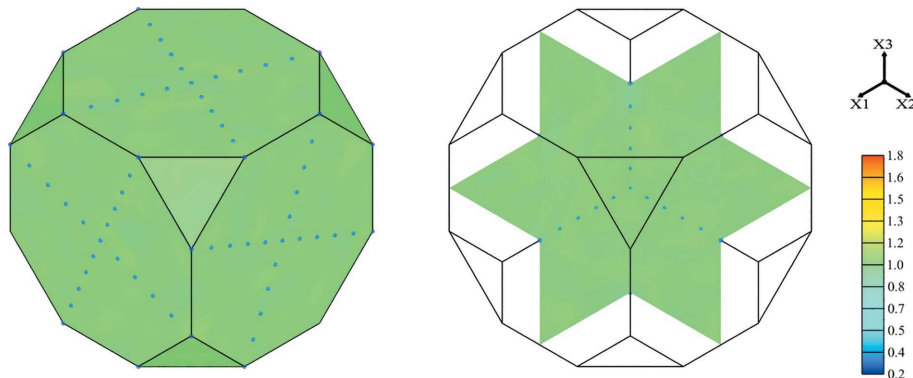


Figure 10
Generation of a uniform orientation distribution in the presence of prescribed orientations. Case of 1000 orientations among which 82 are prescribed, which are located along the $X_i \parallel (100)$ and $X_i \parallel (110)$ fibres. The prescribed orientations are shown by blue dots and the colour field represents the ODF. Cubic crystal symmetry is considered. ODFs are shown as surface and slice views. Note that the ODF is as uniform as in the standard case (Fig. 5a).

standard case (see Fig. 7c), which indicates that the unconstrained orientations have adapted to the prescribed orientations during optimization.

4. Conclusion

By transposing the well known Thomson problem for the equilibrium distribution of electrons on a unit sphere to orientations represented as unit quaternions, we have devised a relatively simple, versatile and efficient method for generating uniform orientation distributions. The key points of the method are the following:

(i) Owing to the fact that three-dimensional orientations represented as unit quaternions are defined on \mathbb{S}^3 (not in its embedding space, \mathbb{H}), they were considered to interact along its geodesics, so that the forces at orientations can be elegantly written in the tangent space at those orientations. This both allows a simple calculation and avoids biases, especially for small numbers of orientations.

(ii) The forces at orientations, being tangential to \mathbb{S}^3 , provide evolution paths and rates (for the orientations) toward states of lower energy, which was used in a conventional gradient-descent optimization. Although this corresponds to local optimization and does not ensure a global minimum, the resulting orientation distributions are highly uniform.

(iii) In the exact formulation, the force at an orientation depends on its interaction with all other orientations, so that a large number of orientations (N) leads to very high computation times [$O(N^2)$]. However, we showed that for more than about 7500 orientations the forces can be computed in close neighbourhoods of the orientations, as forces mostly depend on short-distance interactions. This significantly reduces computation time [$O(N \log N)$], so that large values of N can be used (e.g. 10^6).

(iv) Crystal symmetry can be included naturally, by representing each orientation by all its symmetrically equivalent quaternions. Similarly, orientations that are desired in the distribution can be prescribed.

(v) An efficient resolution and implementation enables one to generate uniform distributions of as many as 10^6

orientations within reasonable computation time. A parallel (MPI) implementation could be used to reduce computation time further.

The uniformity of the orientation distributions was analysed using two metrics: the orientation distribution function $[f(\mathbf{q})]$ and the distribution of the misorientations to the nearest orientation $[\theta_m(\mathbf{q})]$, both of which are of practical importance for crystalline materials. The present method provides highly uniform orientation distributions in terms of the standard deviation of $f(\mathbf{q})$ (σ_f) and the average value of $\theta_m(\mathbf{q})$. Uniform distributions for any number of orientations (from only a few to 10^6 or more) and any crystal symmetry can be generated and used routinely.

Acknowledgements

RQ wishes to thank Adam Morawiec for early discussions on the subject. The authors also thank Adam Morawiec, as well as the two anonymous reviewers, for their thorough comments on the manuscript.

References

- Altschuler, E., Williams, T., Ratner, E., Dowla, F. & Wooten, F. (1994). *Phys. Rev. Lett.* **72**, 2671–2674.
- Barzilai, J. & Borwein, J. (1988). *IMA J. Numer. Anal.* **8**, 141–148.
- Blanco, J. & Rai, P. (2014). *nanoflann: a C++11 Header-Only Library for Nearest Neighbor (NN) Search With KD-trees*, <https://github.com/jlblancoc/nanoflann>.
- Bondarenko, A., Karchevskiy, M. & Kozinkin, L. (2015). *J. Phys. Conf. Ser.* **643**, 012103.
- Dragnev, P., Legg, D. A. & Townsend, D. (2002). *Pac. J. Math.* **207**, 345–358.
- Erber, T. & Hockney, G. M. (1991). *J. Phys. A Math. Gen.* **24**, L1369–L1377.
- Hamilton, W. (1844). *Philos. Mag.* **25**, 10–13.
- Helming, K. (1997). *Textures Microstruct.* **28**, 219–230.
- Kanimozhi, G., Rajathy, R. & Kumar, H. (2016). *Int. J. Electr. Eng. Inf.* **8**, 29–44.
- Karney, C. (2007). *J. Mol. Graphics Modell.* **25**, 595–604.
- Larsen, P. M. & Schmidt, S. (2017). *J. Appl. Cryst.* **50**, 1571–1582.
- Morawiec, A. (2004). *Orientations and Rotations*. Berlin, Heidelberg: Springer.
- Morris, J., Deaven, D. & Ho, K. (1996). *Phys. Rev. B*, **53**, R1740–R1743.
- Nygårds, M. (2003). *Mech. Mater.* **35**, 1049–1057.
- Quey, R. (2018). *Neper: Polycrystal Generation and Meshing*, <http://neper.sourceforge.net>.
- Quey, R., Dawson, P. & Driver, J. (2012). *J. Mech. Phys. Solids*, **60**, 509–524.
- Quey, R., Driver, J. & Dawson, P. (2015). *J. Mech. Phys. Solids*, **84**, 506–527.
- Rauch, E. & Véron, M. (2014). *Mater. Charact.* **98**, 1–9.
- Roşca, D., Morawiec, A. & De Graef, M. (2014). *Modell. Simul. Mater. Sci. Eng.* **22**, 075013.
- Stillinger, F. & Weber, T. (1984). *Science*, **225**, 983–989.
- Suter, R., Hennessy, D., Xiao, C. & Lienert, U. (2006). *Rev. Sci. Instrum.* **77**, 123905.
- Thomson, J. (1904). *London Edinb. Dubl. Philos. Mag. J. Sci.* **7**, 237–265.
- Wales, D., McKay, H. & Altschuler, E. (2009). *Phys. Rev. B*, **79**, 224115.
- Wielewski, E., Boyce, D., Park, J.-S., Miller, M. & Dawson, P. (2017). *Acta Mater.* **126**, 469–480.
- Wong, S. & Dawson, P. (2010). *Acta Mater.* **58**, 1658–1678.
- Xiang, Y., Sun, D., Fan, W. & Gong, X. (1997). *Phys. Lett. A*, **233**, 216–220.
- Yan, Y. & Chirikjian, G. (2012). *2012 IEEE International Conference on Robotics and Automation*, pp. 4254–4259. IEEE.
- Yang, G. & Chen, I.-M. (2006). *IEEE Trans. Rob.* **22**, 869–879.
- Yershova, A., Jain, S., Lavalle, S. & Mitchell, J. (2010). *Int. J. Rob. Res.* **29**, 801–812.