Available online at www.sciencedirect.com
ScienceDirect
Journal of Molecular Graphics and Modelling 25 (2007) 595–604
Journal of Molecular Graphics and Modelling
www.elsevier.com/locate/JMGM

Quaternions in molecular modeling

Charles F.F. Karney *

Sarnoff Corporation, Princeton, NJ 08543-5300, USA

Received 18 July 2005; received in revised form 4 April 2006; accepted 4 April 2006
Available online 18 April 2006

Abstract

Quaternions are an important tool to describe the orientation of a molecule. This paper considers the use of quaternions in matching two conformations of a molecule, in interpolating rotations, in performing statistics on orientational data, in the random sampling of rotations, and in establishing grids in orientation space. These examples show that many of the rotational problems that arise in molecular modeling may be handled simply and efficiently using quaternions.
© 2006 Elsevier Inc. All rights reserved.

Keywords: Quaternions; Rotations; Molecular simulation

1. Introduction

Quaternions were introduced in the mid-nineteenth century by Hamilton [1,2] as an extension of complex numbers and as a tool for manipulating 3-dimensional vectors. Indeed Maxwell used them to introduce vectors in his exposition of electromagnetic theory ([3], §10–11). However, unlike complex numbers which occupy a central role in the development of algebra, quaternions found no similar place in mathematics and, with the introduction of modern vector notation by Gibbs [4], quaternions fell out of favor by the end of the nineteenth century. Nevertheless, quaternions excel as a way of representing rotations of objects in 3-dimensional space. They are economical to work with (both in terms of storage and computation); but more importantly they offer a clean conceptual framework which allow several problems involving rotations to be easily solved.

Basic quaternion algebra is well covered in Hamilton's papers [1,2], which are both accessible and readable. These papers may be supplemented with a wealth of on-line resources [5,6]. Many authors over the past 20 years have "rediscovered" the application of quaternions to rotations and it is with some trepidation that this author inflicts another paper on the subject on the scientific community. However, within the molecular modeling community, quaternions are quite narrowly applied. This paper therefore briefly reviews quaternion algebra and then describes their applications to a broad range of rotational problems in molecular modeling. Much of this material has appeared before—but often scattered about in journals for fields unrelated to molecular modeling. I have, therefore, endeavored to organize the material, to generalize it, and to present it with a consistent notation, with the hope this affords a deeper appreciation of the power of quaternions in describing rotations and encourages their wider adoption in molecular modeling.

The outline of this paper is as follows. After introducing quaternions and their use in describing rotations, we tackle various applications. First we review the quaternion method for computing the least-squares fit of two conformations of the same molecule. We also see how to include molecular inversions and discuss why the least-squares fit is a poor choice to describe the orientation of a flexible molecule. We next show how to interpolate smoothly between two orientations and that this corresponds to rotating the molecule at constant angular velocity. In order to carry out statistics on orientational data, we give a robust definition of the mean orientation showing how to transform the deviations from the mean to 3-dimensional space so that familiar statistical tools may be employed. In Monte Carlo applications, we need to be able to select a random orientation uniformly; we show that this is trivially accomplished in quaternion space and we also consider the problem of making random incremental rotations. Finally, it is frequently useful to impose a grid on orientation

* Tel.: +1 609 734 2312; fax: +1 609 734 2662.
E-mail address: ckarney@sarnoff.com.

1093-3263/$ – see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmgm.2006.04.002

space and we illustrate how this may be done with applications to quadrature and searching.

2. Quaternions

The original notation for quaternions [1] paralleled the convention for complex numbers

$$\mathbf{q} = q_0\mathbf{u} + q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k},$$

which obey the conventional algebraic rule for addition and multiplication by scalars (real numbers) and which obey an associative non-commutative rule for multiplication where $\mathbf{u}$ is the identity element and

$$\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = -\mathbf{u}, \quad \mathbf{ij} = -\mathbf{ji} = \mathbf{k}, \quad \mathbf{jk} = -\mathbf{kj} = \mathbf{i},$$
$$\mathbf{ki} = -\mathbf{ik} = \mathbf{j}.$$

It is frequently useful to regard quaternions as an ordered set of 4 real quantities which we write as

$$\mathbf{q} = [q_0, q_1, q_2, q_3], \tag{1}$$

or as a combination of a scalar and a vector

$$\mathbf{q} = [q_0, \mathbf{q}], \tag{2}$$

where $\mathbf{q} = [q_1, q_2, q_3]$. A "scalar" quaternion has zero vector part and we shall write this as $[q_0, \mathbf{0}] = q_0\mathbf{u} = q_0$. A "pure" quaternion has zero scalar part $[0, \mathbf{q}]$.

In the scalar–vector representation, multiplication becomes

$$\mathbf{pq} = [p_0q_0 - \mathbf{p}\cdot\mathbf{q}, \quad p_0\mathbf{q} + q_0\mathbf{p} + \mathbf{p}\times\mathbf{q}],$$

where "·" and "×" are the vector dot and cross products. The conjugate of a quaternion is given by

$$\bar{\mathbf{q}} = [q_0, -\mathbf{q}];$$

the squared norm of a quaternion is

$$|\mathbf{q}|^2 = \mathbf{q}\bar{\mathbf{q}} = q_0^2 + q_1^2 + q_2^2 + q_3^2,$$

and its inverse is

$$\mathbf{q}^{-1} = \bar{\mathbf{q}}/|\mathbf{q}|^2.$$

Quaternions with $|\mathbf{q}| = 1$ are called unit quaternions, for which we have $\mathbf{q}^{-1} = \bar{\mathbf{q}}$.

The quaternion $\mathbf{q}$ can also be represented as a 2×2 complex matrix,

$$\begin{pmatrix} q_0 + iq_1 & q_2 + iq_3 \\ -q_2 + iq_3 & q_0 - iq_1 \end{pmatrix},$$

or as a 4×4 real matrix,

$$\mathbf{Q}(\mathbf{q}) = \begin{pmatrix} q_0 & q_1 & q_2 & q_3 \\ -q_1 & q_0 & -q_3 & q_2 \\ -q_2 & q_3 & q_0 & -q_1 \\ -q_3 & -q_2 & q_1 & q_0 \end{pmatrix}; \tag{3}$$

in these forms, quaternion multiplication becomes matrix multiplication.

The notation we adopt here is to use light-face italics for scalar quantities, bold roman for 3-dimensional vectors and 3×3 matrices, bold sans serif for quaternions and 4×4 matrices. Quaternion multiplication is indicated by $\mathbf{pq}$, while "·" is used to indicate matrix-vector and vector–vector (including quaternion-quaternion) contractions and in this context $\mathbf{q}$ and $\mathbf{v}$ are treated as *column* vectors. Thus, we may write $|\mathbf{q}|^2 = \mathbf{q}^\mathsf{T}\cdot\mathbf{q}$. We also find that $\mathbf{pq} = \mathbf{p}^\mathsf{T}\cdot\mathbf{Q}(\mathbf{q})$, with $\mathbf{Q}$ given by Eq. (3). Consistent with Eqs. (1) and (2), we shall number quaternion indices starting at 0 and vector indices from 1.

3. Rotations

The chief application of quaternions to molecular modeling lies in their use to represent rotations. Consider a unit quaternion

$$\mathbf{q} = [\cos(\theta/2), \mathbf{v}\sin(\theta/2)], \tag{4}$$

where $|\mathbf{v}| = 1$, and define an operator $R_\mathbf{q}$ on 3-dimensional vectors by

$$[0, R_\mathbf{q}(\mathbf{x})] = \mathbf{q}\,[0, \mathbf{x}]\,\bar{\mathbf{q}}. \tag{5}$$

Multiplying out the quaternion product, we find

$$R_\mathbf{q}(\mathbf{x}) = \mathbf{R}(\mathbf{q})\cdot\mathbf{x},$$

where $\mathbf{R}(\mathbf{q})$ is the tensor

$$\mathbf{R}(\mathbf{q}) = (q_0^2 - |\mathbf{q}|^2)\mathbf{I} + 2\mathbf{qq} + 2q_0\mathbf{I}\times\mathbf{q} \tag{6}$$

$$= \mathbf{vv} + \cos\theta(\mathbf{I} - \mathbf{vv}) + \sin\theta\,\mathbf{I}\times\mathbf{v}, \tag{7}$$

where $\mathbf{aa}$ is the parallel projector [$(\mathbf{aa})\cdot\mathbf{b} = (\mathbf{a}\cdot\mathbf{b})\mathbf{a}$] and $\mathbf{I}\times\mathbf{a}$ is the cross operator [$(\mathbf{I}\times\mathbf{a})\cdot\mathbf{b} = \mathbf{a}\times\mathbf{b}$] ([4], §113). Eq. (7) is the conventional tensor representation for a right-handed rotation of θ about an axis $\mathbf{v}$ through the origin ([4], §126). Eq. (6) may be written in component form as

$$\mathbf{R}(\mathbf{q}) = \begin{pmatrix} 1 - 2q_2^2 - 2q_3^2 & 2q_1q_2 - 2q_0q_3 & 2q_1q_3 + 2q_0q_2 \\ 2q_2q_1 + 2q_0q_3 & 1 - 2q_3^2 - 2q_1^2 & 2q_2q_3 - 2q_0q_1 \\ 2q_3q_1 - 2q_0q_2 & 2q_3q_2 + 2q_0q_1 & 1 - 2q_1^2 - 2q_2^2 \end{pmatrix}. \tag{8}$$

The definition, Eq. (5), gives $R_\mathbf{p}(R_\mathbf{q}(\mathbf{x})) = R_\mathbf{pq}(\mathbf{x})$, so that $\mathbf{pq}$ corresponds to composing rotations (with the rotation by $\mathbf{q}$ performed first). We also find that $R_\mathbf{q} = R_{-\mathbf{q}}$; i.e., $\mathbf{q}$ and $-\mathbf{q}$ give the same rotation—changing the sign of $\mathbf{q}$ is equivalent to increasing θ by 2π in Eq. (4). Unit quaternions satisfy $q_0^2 + q_1^2 + q_2^2 + q_3^2 = 1$ and the quaternion representation of rotations are as points on a hypersphere $\mathbb{S}^3$ with opposite points identified. For future reference, we note that the (3-dimensional) area of $\mathbb{S}^3$ is $2\pi^2$.

Because $\mathbf{q}$ and $-\mathbf{q}$ give the same rotation, some care needs to be taken when comparing two orientations represented by $\mathbf{q}_a$ and $\mathbf{q}_b$. The rotation, $\mathbf{q} = \mathbf{q}_b\bar{\mathbf{q}}_a$, moves from $\mathbf{q}_a$ to $\mathbf{q}_b$. When inverting Eq. (4) to determine the rotation angle θ between the two orientations, we should, if necessary, change the sign of $\mathbf{q}$

to ensure that $q_0 \geq 0$, so that $\theta \in [0, \pi]$. A simple metric for closeness is given by $\cos(\theta/2) = |\mathbf{q}_a^{\mathsf{T}} \cdot \mathbf{q}_b|$.

Describing rotations with quaternions has a number of benefits. They offer a compact representation of rotations. Compared to Euler angles, they are free of singularities. Rotations may be composed more efficiently using quaternions than by matrix multiplication. Also in contrast to rotation matrices, it is easy to maintain a quaternion's unit normalization (merely divide it by $|\mathbf{q}|$). However the chief benefit is that the representation of a rotation as point on $\mathbb{S}^3$ allows us to derive many important results concerning rotations in a simple coordinate-free way.

There is one application where the matrix representation of rotations is more efficient that the quaternion representation. If we wish to apply the same rotation to many points, then we should form the rotation matrix using Eq. (8) and transform the points by matrix multiplication.

The conventional representation for rotations that is most closely allied to quaternions is the axis–angle representation, where the rotation is given by a vector $\mathbf{s} = \theta\mathbf{v}$ which denotes a rotation of $\theta = |\mathbf{s}|$ about an axis $\mathbf{v} = \mathbf{s}/|\mathbf{s}|$. It is useful to have an analytic relation between the quaternion and axis–angle representations and this is provided by the quaternion exponential [1],

$$\exp([0, \mathbf{s}/2]) = \mathbf{q}, \tag{9}$$

where $\mathbf{q}$ is given by Eq. (4), This definition of the exponential follows from its series expansion. Similarly the inverse operation is given by the quaternion logarithm

$$\ln \mathbf{q} = [0, \mathbf{s}/2 + 2\pi n\mathbf{v}], \tag{10}$$

where n is an integer.

It is useful here to make a distinction between "orientation" and "rotation". We imagine that our molecule has some arbitrary but definite reference state. We apply a rotation and a translation (jointly referred to as a "displacement") to this reference state and so bring the molecule to a new orientation and position (jointly referred to as a "configuration").

4. Least-squares fit

Given two conformations of the same molecule, it is often useful to be able to determine how close the conformations are. In order to do this, we can rigidly move one conformation so that it nearly coincides with the other and then determine the difference in the positions of the corresponding atoms. Thus, if we are given two sets of atomic positions $\{\mathbf{x}_k\}$ and $\{\mathbf{y}_k\}$ with $k \in [1, N]$ together with a set of atomic "weights" $\{w_k\}$, we wish to determine the (rigid) displacement T which minimizes

$$E = \frac{1}{W}\sum_k w_k |\mathbf{y}_k - T(\mathbf{x}_k)|^2, \tag{11}$$

where $W = \sum_k w_k$. Here w_k is merely a statistical weight of an atom—it is not necessarily related to the atomic mass. The two sets of atomic positions are ordered which presumes that we can identify corresponding atoms. (This is not necessarily a simple matter, if, for example, we are dealing with a molecule with several identical branches.) The displacement $T = (\mathbf{q}, \mathbf{d})$ may be expressed as a rotation about an axis through the origin followed by a translation, i.e., $T(\mathbf{x}) = R_{\mathbf{q}}(\mathbf{x}) + \mathbf{d}$.

This problem has been considered by many authors and a review of various approaches is given by Flower [7]. Using quaternions to describe the rotation leads to an elegant and robust solution. An early use of quaternions in this context is to solve the problem formulated by Wahba [8,9], the determination of the attitude of a spacecraft given the directions of several objects relative to the craft. The resulting "$\mathbf{q}$-method" is described by Keat ([10], §A.3) and by Lerner ([11], §12.2.3) who both credit the invention of the method to Paul B. Davenport (1968). The generalization to matching points (as opposed to directions) was considered by Faugeras and Hebert [12] who independently found the same method for determining the orientation. Their method was subsequently rediscovered by Horn [13], by Diamond [14], and by Kearsley [15]. The derivation of Faugeras and Hebert is one of the clearest, and we briefly summarize it here including the straightforward generalization of including arbitrary weights w_k.

If we demand that the variation of E with respect to $\mathbf{d}$ vanish, we find that

$$\mathbf{d} = \langle \mathbf{y} \rangle - R_{\mathbf{q}}(\langle \mathbf{x} \rangle), \tag{12}$$

where $\langle \ldots \rangle$ denotes the sample average,

$$\langle X \rangle = \frac{1}{W}\sum_k w_k X_k. \tag{13}$$

Eq. (11) may now be written as

$$E = \frac{1}{W}\sum_k w_k |\mathbf{y}'_k - R_{\mathbf{q}}(\mathbf{x}'_k)|^2, \tag{14}$$

where $\mathbf{x}'_k = \mathbf{x}_k - \langle \mathbf{x} \rangle$ and $\mathbf{y}'_k = \mathbf{y}_k - \langle \mathbf{y} \rangle$. Using Eq. (5), Eq. (14), becomes

$$E = \frac{1}{W}\sum_k w_k \left| [0, \mathbf{y}'_k] - \mathbf{q}\,[0, \mathbf{x}'_k]\,\bar{\mathbf{q}} \right|^2. \tag{15}$$

Because, the norm of a quaternion is unchanged on multiplying it by a unit quaternion, we may right-multiply the kernel of Eq. (15) by $\mathbf{q}$ to give

$$E = \frac{1}{W}\sum_k w_k \left| [0, \mathbf{y}'_k]\,\mathbf{q} - \mathbf{q}\,[0, \mathbf{x}'_k] \right|^2. \tag{16}$$

We need to minimize Eq. (16) subject to the constraint $|\mathbf{q}| = 1$. Because the kernel is linear in $\mathbf{q}$, it can be written as

$$[0, \mathbf{y}'_k]\,\mathbf{q} - \mathbf{q}\,[0, \mathbf{x}'_k] = \mathbf{A}_k \cdot \mathbf{q}, \tag{17}$$

where $\mathbf{A}_k$ is a 4×4 skew matrix

$$\mathbf{A}_k = \mathbf{A}(\mathbf{y}'_k + \mathbf{x}'_k, \mathbf{y}'_k - \mathbf{x}'_k),$$

with

$$\mathbf{A}(\mathbf{a},\mathbf{b}) = \begin{pmatrix} 0 & -\mathbf{b}^{\mathsf{T}} \\ \mathbf{b} & \mathbf{I}\times\mathbf{a} \end{pmatrix} = \begin{pmatrix} 0 & -b_1 & -b_2 & -b_3 \\ b_1 & 0 & -a_3 & a_2 \\ b_2 & a_3 & 0 & -a_1 \\ b_3 & -a_2 & a_1 & 0 \end{pmatrix}.$$

Substituting this into Eq. (16), we obtain

$$E = \frac{1}{W}\sum_k w_k\, \mathbf{q}^{\mathsf{T}}\cdot\mathbf{A}_k^{\mathsf{T}}\cdot\mathbf{A}_k\cdot\mathbf{q} = \mathbf{q}^{\mathsf{T}}\cdot\mathbf{B}\cdot\mathbf{q}, \tag{18}$$

where $\mathbf{B} = \langle \mathbf{A}_k^{\mathsf{T}}\cdot\mathbf{A}_k\rangle$ is a 4×4 symmetric matrix which has real eigenvalues, $0 \le \lambda_0 \le \lambda_1 \le \lambda_2 \le \lambda_3$. Setting $\mathbf{q}$ to the eigenvector corresponding to λ_0 gives the minimum value for $E = \lambda_0$.

In summary, the best fit is achieved by subtracting the mean positions from the original sets of points to give $\{\mathbf{x}'_k\}$ and $\{\mathbf{y}'_k\}$, forming the matrices $\mathbf{A}_k$ and $\mathbf{B}$, and determining the minimum eigenvalue λ_0 of $\mathbf{B}$. The optimal rotation is given by setting $\mathbf{q}$ to the corresponding eigenvector of $\mathbf{B}$ and the optimal translation is found from Eq. (12). The mean squared error for this fit is λ_0.

This procedure has two attractive features. The rotation obtained is a proper rotation (without an inversion); this is usually the desired result. Secondly, degenerate molecules are treated satisfactorily. For example if one or both of the sets $\{\mathbf{x}_k\}$ and $\{\mathbf{y}_k\}$ is collinear, then the best fit is no longer unique. The result will be that there will be multiple minimum eigenvalues of $\mathbf{B}$ with distinct eigenvectors. The general solution is obtained by setting $\mathbf{q}$ to a linear combination of these eigenvectors. The method does require finding the eigenvalues and eigenvectors of a 4×4 matrix. However there are many numerical libraries [16–18] which solve such problems and the results are accurate to round-off for small symmetric matrices such as $\mathbf{B}$. A fast method of determining just the required eigenvector and associated eigenvalue in order to determine the attitude of a spacecraft is given in ([19], §III). However, in applications to molecular modeling, it is probably preferable merely to invoke a library eigenvector routine.

Horn [13] considered including a scaling in the transformation T in Eq. (11). This is quite easily accommodated. However there seems little need to include such an effect in molecular modeling.

Diamond [20] considers the case where inversions are allowed. This is easily achieved by substituting $-\mathbf{x}'_k$ for $\mathbf{x}'_k$ in Eq. (14). Eq. (18) then involves a matrix $\mathbf{B}'$ where

$$\mathbf{B}' = 2\langle|\mathbf{x}'|^2 + |\mathbf{y}'|^2\rangle\mathbf{I} - \mathbf{B}. \tag{19}$$

Consequently the rotation giving the best inverted fit is the eigenvector with the *greatest* eigenvalue of $\mathbf{B}$, λ_3. Because the sum of the eigenvalues of $\mathbf{B}$ is its trace, $4\langle|\mathbf{x}'|^2 + |\mathbf{y}'|^2\rangle$, we can express the mean squared error for the inverted fit as $\frac{1}{2}(\lambda_0 + \lambda_1 + \lambda_2 - \lambda_3)$. Thus, once the eigenvalues of $\mathbf{B}$ have been computed we immediately determine whether the inverted fit will be better than the proper fit.

Coutsias et al. [21] provide an interesting extension of this method. Suppose the atomic positions $\{\mathbf{x}_k\}$ represent a model of a molecule which depends on a set of parameters $\{\alpha_i\}$, for example, the torsion angles of a protein backbone. By considering the gradient of E in parameter space $\partial E/\partial\alpha_i$, they provide a method for determining the parameter values which result in the best fit to a given crystal structure.

One other interesting consequence of the result for the best fit is that the rotation is not a continuous function of the configurations of the molecules. Let us suppose that $\{\mathbf{x}_k\}$ gives the position of the atoms in a molecule in some predefined configuration and suppose that $\{\mathbf{y}_k\}$ gives the atom positions during the course of a dynamical simulation of the molecule. If the forces acting on the atoms are finite then $\mathbf{y}_k$ is a C^1 function (twice differentiable). During the course of the deformation of the molecule, $\mathbf{B}$ and its eigenvalues change. In the typical case, the two smallest eigenvalues exchange roles and $\mathbf{q}$ switches from one direction in $\mathbb{R}^4$ to an orthogonal direction. This results in the orientation of the best fit changing discontinuously by 180°.

In modeling a flexible molecule, it is frequently useful to separate the external degrees of freedom, namely position and orientation, from the internal degrees of freedom. This allows, for example, translational and rotational symmetry to the system to be enforced and correlations between the motions of atoms within a molecule to be studied. This begs the question of how best to define the position and orientation of a molecule. Taking the position to be the center of mass is often the obvious choice. The position (so defined) evolves according to Newton's second law driven by the total force on the molecule. It is not possible to keep track of the orientation in an analogous fashion by integrating the total angular momentum, because flexible bodies can change their orientation with zero angular momentum—witness the ability of a cat always to land on its feet. A possible definition of the orientation is the best fit orientation to a reference conformation; i.e., we define $\mathbf{o}_R(A)$ as the best fit orientation, expressed as a quaternion, of the molecule in conformation A relative to a reference conformation R. Here again this choice has the attractive feature that the whole molecule is included in the definition. There are two problems with this prescription. Firstly, the difference in orientations between two conformations A and B depends, in general, on the choice of reference conformation, namely

$$\mathbf{o}_R(B)\overline{\mathbf{o}_R(A)} \neq \mathbf{o}_S(B)\overline{\mathbf{o}_S(A)}.$$

(This is easily demonstrated for simple triatomic molecules.) Thus this definition of orientation entails a degree of "arbitrariness" absent in our definition of position. A second more serious defect arises from the discussion in the previous paragraph. Recovering the actual configuration of the molecule from the orientation defined in this way is numerically unstable (by a flip of 180°!) whenever the lowest eigenvalues cross. This would also lead to large and discontinuous apparent internal motions of the molecule with small changes in the atoms' true positions. A better choice would therefore be to make the fit to some rigid (or nearly rigid) subcomponent of the molecule [22]. Although this still yields an arbitrary definition of orientation (depending on the choice of reference subcomponent), the resulting orientation varies continuously under continuous deformations of the molecule. An extensive discussion of

how to separate the orientation from the internal motions of a flexible molecule is given by Littlejohn and Reinsch [23].

5. Interpolating rotations

The power of the quaternion representation of rotations is evident when we consider the problem of interpolating between two orientations of a molecule. (This application might arise in the animation of a molecular simulation.) Suppose we wish to interpolate between $\mathbf{q}_a$ and $\mathbf{q}_b$. Because these quaternions and their interpolants lie on the unit sphere $\mathbb{S}^3$, the shortest path will be a great circle whose parametric equation is given by [24]

$$\mathbf{q}(\phi) = \frac{\mathbf{q}_a \sin(\theta - \phi) + \mathbf{q}_b \sin(\phi)}{\sin(\theta)}, \tag{20}$$

where $\cos\theta = \mathbf{q}_a^\mathsf{T} \cdot \mathbf{q}_b$. In the computer animation community this "spherical linear interpolation" operation is denoted by $\mathrm{Slerp}(\mathbf{q}_a, \mathbf{q}_b; u) = \mathbf{q}(u\theta)$ [24]. As ϕ is increased from 0 to 2π, $\mathbf{q}(\phi)$ becomes successively $\mathbf{q}_a$, $\mathbf{q}_b$, $-\mathbf{q}_a$, $-\mathbf{q}_b$, and finally returns to $\mathbf{q}_a$. During this operation the corresponding 3-dimensional rotation has increased by 4π. If $\mathbf{q}_a^\mathsf{T} \cdot \mathbf{q}_b \geq 0$, then $0 \leq \phi \leq \theta$ takes $\mathbf{q}(\phi)$ smoothly from $\mathbf{q}_a$ to $\mathbf{q}_b$. If, on the other hand, $\mathbf{q}_a^\mathsf{T} \cdot \mathbf{q}_b < 0$, then a shorter path is found with $0 \geq \phi \geq \theta - \pi$ which takes $\mathbf{q}(\phi)$ smoothly from $\mathbf{q}_a$ to $-\mathbf{q}_b$.

Eq. (20) is derived using simple geometrical arguments applied to $\mathbb{S}^3$ and the same result is obtained for the great-circle interpolation for $\mathbb{S}^n$. For $\mathbb{S}^3$, the result can also be expressed as

$$\mathbf{q}(\phi) = (\mathbf{q}_b \bar{\mathbf{q}}_a)^{\phi/\theta} \mathbf{q}_a.$$

This relation has the interpretation: rotate to $\mathbf{q}_a$ and then rotate a fraction ϕ/θ to the path from $\mathbf{q}_a$ to $\mathbf{q}_b$. The operation $\mathbf{q}^u$ is defined by [1]

$$\mathbf{q}^u = \exp(u \ln \mathbf{q}).$$

In fact this interpolation scheme results in the molecule undergoing rotation at constant angular velocity. In order to show this, consider a body rotating at ω about a unit axis $\mathbf{v}$. The evolution of the orientation $\mathbf{q}$ satisfies the differential equation

$$\dot{\mathbf{q}} = [0, (\omega/2)\mathbf{v}]\, \mathbf{q}. \tag{21}$$

This is easily solved (e.g., by using finite differences and passing to the limit $\delta t \to 0$) to give

$$\begin{aligned}\mathbf{q}(t) &= \exp([0, (\omega t/2)\mathbf{v}])\mathbf{q}(0) \\ &= [\cos(\omega t/2), \mathbf{v}\sin(\omega t/2)]\, \mathbf{q}(0),\end{aligned}$$

which agrees with Eq. (20) with the substitutions $\phi = \omega t/2$, $\mathbf{q}_a = \mathbf{q}(0)$ and $\mathbf{q}_b = [0, \mathbf{v}]\, \mathbf{q}(0)$.

If we wish to interpolate between two configurations of a rigid molecule, we are free to specify a point, $\mathbf{x}_0$, in the reference molecule which will move with constant velocity. If the initial and final configurations are given by $T_a = (\mathbf{q}_a, \mathbf{d}_a)$ and $T_b = (\mathbf{q}_b, \mathbf{d}_b)$, with $\mathbf{q}_a^\mathsf{T} \cdot \mathbf{q}_b \geq 0$, then the required interpolation is achieved by increasing u from 0 to 1 with the orientation given by $\mathbf{q}(u\theta)$ and the translation given by

$$(\mathbf{d}_a + R_{\mathbf{q}_a}(\mathbf{x}_0))(1-u) + (\mathbf{d}_b + R_{\mathbf{q}_b}(\mathbf{x}_0))u - R_{\mathbf{q}(u\theta)}(\mathbf{x}_0).$$

6. Mean orientation

The mean of directional quantities has frequently presented difficulties [25]. Let us assume we have N samples of some directional quantity with weights w_k for $k \in [1, N]$ and $\sum_k w_k = W$. In the case where the samples are angles (e.g., the dihedral angles of a molecular bond) or directions (e.g., the orientations of a diatomic molecule), there is a well established procedure ([25], §2.2.1, §9.2.1): express the directions as unit vectors in $\mathbb{R}^2$ or $\mathbb{R}^3$, $\mathbf{n}_k$, and determine $\langle \mathbf{n} \rangle$ where we take the sample average according to Eq. (13). Now the mean direction is given by $\langle\langle \mathbf{n} \rangle\rangle = \langle \mathbf{n} \rangle / |\langle \mathbf{n} \rangle|$, while $1 - |\langle \mathbf{n} \rangle|$, a quantity lying in $[0, 1]$, is the "circular variance" ([25], §2.3.1) or "spherical variance" ([25], §9.2.1). Here $\langle \ldots \rangle$ is defined as a simple weighted arithmetical average, Eq. (13), while $\langle\langle \ldots \rangle\rangle$ denotes the physically relevant mean of a quantity.

This procedure cannot be directly applied to unit quaternions used to represent rotations because of the indistinguishability of $\pm\mathbf{q}$. Instead, we view $\{\mathbf{q}_k\}$ as *axes* ([25], §1.1, §9.1) in $\mathbb{R}^4$, and define $\langle\langle \mathbf{q} \rangle\rangle$ as the unit quaternion about which the weighted moment of inertia of $\{\mathbf{q}_k\}$ is minimum ([26], §3). Thus we wish to minimize

$$\begin{aligned}L &= \frac{1}{W} \sum_k w_k\, \mathbf{q}_k^\mathsf{T} \cdot (\mathbf{I} - \langle\langle \mathbf{q} \rangle\rangle \langle\langle \mathbf{q} \rangle\rangle^\mathsf{T}) \cdot \mathbf{q}_k \\ &= \frac{1}{W} \sum_k w_k\, \langle\langle \mathbf{q} \rangle\rangle^\mathsf{T} \cdot (\mathbf{I} - \mathbf{q}_k \mathbf{q}_k^\mathsf{T}) \cdot \langle\langle \mathbf{q} \rangle\rangle \\ &= \langle\langle \mathbf{q} \rangle\rangle^\mathsf{T} \cdot (\mathbf{I} - \langle \mathbf{q}\mathbf{q}^\mathsf{T} \rangle) \cdot \langle\langle \mathbf{q} \rangle\rangle.\end{aligned}$$

The minimum value of L is given by the minimum eigenvalue of $\mathbf{I} - \langle \mathbf{q}\mathbf{q}^\mathsf{T} \rangle$ and $\langle\langle \mathbf{q} \rangle\rangle$ is corresponding eigenvector. The resulting L, which is a quantity lying in $[0, \frac{3}{4}]$, then provides a measure of the variance of the rotations. This definition of the mean has a number of desirable properties: it is invariant when the signs of the $\mathbf{q}_k$ are changed; it is independent of the order of the samples; and it transforms properly if the samples are transformed.

This prescription can also be applied to determine the mean direction of objects whose symmetry makes $\mathbf{n}$ and $-\mathbf{n}$ indistinguishable (for example, the orientation of the diatomic molecule N_2).

Suppose we wish to determine the mean configuration of a rigid molecule, i.e., the mean of $\{T_k = (\mathbf{q}_k, \mathbf{d}_k)\}$. We are free to choose a point $\mathbf{x}_0$ in the reference molecule whose position in the mean configuration coincides with its mean position. (Compare this with the discussion of interpolating configurations in the previous section.) A suitable definition for the mean configuration is then

$$\langle\langle T \rangle\rangle = (\langle\langle \mathbf{q} \rangle\rangle, \langle \mathbf{d} \rangle + \langle R_{\mathbf{q}}(\mathbf{x}_0) \rangle - R_{\langle\langle \mathbf{q} \rangle\rangle}(\mathbf{x}_0)). \tag{22}$$

Frequently, we need more precise information about the distribution of configurations than its variance. We might need

to know how much the rotation about different axes are correlated or whether rotational and translational motions are coupled. It is also desirable to be able to fit model distributions to a set of samples. For these purposes, it is useful to be able to map rotations onto $\mathbb{R}^3$ so that standard statistical tools can be employed. We require that the mapping be measure preserving (constant Jacobian) to simplify the use of the transformed rotations.

We have already introduced the axis–angle representation of rotations. We may make the restriction $|\mathbf{s}| \le \pi$ and so map the hemisphere $q_0 \ge 0$ of $\mathbb{S}^3$ onto a ball of radius π in $\mathbb{R}^3$. Unfortunately, the mapping, Eq. (9), does not have constant Jacobian. We can correct this by defining [27] a new "turn" vector $\mathbf{u}$ with the properties

$$\mathbf{u} \parallel \mathbf{s}, \tag{23a}$$

$$|\mathbf{u}| = \left(\frac{|\mathbf{s}| - \sin|\mathbf{s}|}{\pi}\right)^{1/3}. \tag{23b}$$

This is an extension of the Lambert azimuthal equal-area projection providing a measure-preserving mapping of the hemisphere $q_0 \ge 0$ of $\mathbb{S}^3$ onto the unit ball $\mathbb{B}^3$. Eq. (23) is well behaved at the boundary, $|\mathbf{u}| = 1$; however on this boundary antipodal points are identified. The inverse mapping has an infinite derivate at $|\mathbf{u}| = \sqrt[3]{2n}$ for integer $n \ne 0$ which corresponds to shells in $\mathbf{u}$ space which map to the origin. This inverse of Eq. (23) is easily implemented via Newton's method supplemented by Taylor series at the origin and at $\sqrt[3]{2n}$.

This mapping was introduced [27] to allow distributions of orientations to be fit using a mixture of Gaussians [28]. Given a set of sample orientations $\{\mathbf{q}_k\}$, we compute the mean orientation, $\langle\langle\mathbf{q}\rangle\rangle$. The deviations of the samples from the mean are then given by the rotations $\{\mathbf{q}_k\langle\langle\bar{\mathbf{q}}\rangle\rangle\}$ and these are mapped to a set of turns $\{\mathbf{u}_k\}$. Because these are points in $\mathbb{R}^3$, we may fit them with a 3-dimensional Gaussian with zero mean and with covariance matrix $\langle\mathbf{u}^\mathsf{T}\mathbf{u}\rangle$.

This procedure can be extended to fits of molecular configurations. In this case, the deviations from the mean configuration, Eq. (22), is mapped into a point in $\mathbb{R}^6$; the resulting Gaussian fit will capture the correlation between the translational and rotational degrees of freedom.

In closing this section, we mention an alternative way of fitting quaternion orientational data with analytic functions, namely in terms of spherical harmonics. The normal (3-dimensional) spherical harmonics can be generalized to 4 (and higher) dimensions [29,30] and the orthogonality relation allows the coefficients of the harmonics to be computed simply. The $\pm\mathbf{q}$ symmetry merely results in the odd harmonics dropping out. However in typical molecular interactions, the relative orientation of the molecules is tightly constrained which means that a large number of spherical harmonics will be needed to represent the orientational distribution. For such applications, a representation in terms of localized functions, such as Gaussians, is preferable.

7. Random orientation

In Monte Carlo simulations [31], it is sometimes necessary to select a molecule with a random and uniform position and orientation, for example, when attempting to insert a molecule into a simulation box during a grand canonical simulation [32]. Choosing a random position is straightforward. However, we need to be careful to select the random orientation uniformly or else detailed balance will be violated (when balancing insertions and deletions). One possibility is to choose a random turn $\mathbf{u}$ in $\mathbb{B}^3$ and to convert this to a quaternion. However, it is much simpler to sample directly in quaternion space.

Let us first establish the requirement for "uniform" sampling of orientations. Composing 3-dimensional rotations is carried out by the multiplication of unit quaternions; but we know that $\mathbf{pq} = \mathbf{p}^\mathsf{T} \cdot \mathbf{Q}(\mathbf{q})$, where $\mathbf{Q}(\mathbf{q})$, given in Eq. (3), is orthonormal if $\mathbf{q}$ is a unit quaternion. Thus 3-dimensional rotations map into a rigid rotation of $\mathbb{S}^3$; a uniform density on $\mathbb{S}^3$ is invariant to such rotations. It follows that the task of sampling a random orientation reduces to picking a random unit quaternion uniformly on $\mathbb{S}^3$.

Marsaglia [33] provides one prescription: select x_1 and y_1 uniformly in $(-1, 1)$ until $s_1 = x_1^2 + y_1^2 < 1$; similarly, select x_2 and y_2 uniformly in $(-1, 1)$ until $s_2 = x_2^2 + y_2^2 < 1$; then

$$\mathbf{q} = [x_1, y_1, x_2\sqrt{(1 - s_1)/s_2}, y_2\sqrt{(1 - s_1)/s_2}]$$

is uniformly distributed on $\mathbb{S}^3$.

A more transparent and symmetric method (which generalizes to sampling points on $\mathbb{S}^n$ ([34], §7.1)) is to pick 4 normal deviates g_i for $i \in [0, 4)$ and to set

$$\mathbf{p} = [g_0, g_1, g_2, g_3], \quad \mathbf{q} = \mathbf{p}/|\mathbf{p}|.$$

Although this method is less efficient than Marsaglia's, the overall impact in the context of a molecular simulation is probably tiny. Both of these methods return points uniformly over the whole of $\mathbb{S}^3$ rather that over just one hemisphere. In most applications, this is of no consequence.

Other representations of rotation yield more complex rules for obtaining random orientations. For example, with Euler angles, we would sample uniformly the first and third angles and the cosine of the second angle. If the orientation is given in axis–angle space, $\mathbf{s}$, then the axis, $\mathbf{s}/|\mathbf{s}|$, should be chosen uniformly on $\mathbb{S}^2$, and the rotation angle, $|\mathbf{s}|$, should be sampled from $[0, \pi]$ with probability $(2/\pi)\sin^2(|\mathbf{s}|/2)$. Of course, this simplifies when $\mathbf{s}$ is transformed to $\mathbf{u}$ space, Eq. (23), leading to a uniform distribution in $\mathbb{B}^3$.

A related problem is selection of random rotational moves for use in a Monte Carlo simulation [31]. This method requires that detailed balance be satisfied, which, in the absence of torque bias, means that the probability of selecting the new orientation is symmetric under interchange of old and new orientations. Because we are typically interested in small changes in orientation, it is most convenient to select the rotation in axis–angle space as $\exp([0, \mathbf{s}])$ and to set the new

orientation

$$\mathbf{q}' = \exp([0, \mathbf{s}])\mathbf{q},$$

where $\mathbf{s}$ is selected from an even distribution, $p(\mathbf{s}) = p(-\mathbf{s})$. (This result follows because the Jacobian factor is even in $\mathbf{s}$.) Usually, we wish the choice of rotation axis to be isotropic, and in that case we have $p(\mathbf{s}) = p(|\mathbf{s}|)$. Thus we might select $\mathbf{s}$ uniformly in a sphere of radius Δ. Rao et al. [35] select $|\mathbf{s}|$ uniformly in $[0, \Delta]$ (which results in a distribution which is singular at the origin in $\mathbf{s}$ space). An attractive choice of distribution is a 3-dimensional Gaussian

$$p(\mathbf{s}) = \frac{\exp(-\frac{1}{2}|\mathbf{s}|^2/\Delta^2)}{(2\pi)^{3/2}\Delta^3}.$$

Not only is this simple to sample from, but it allows torque bias to be included in a simple manner. Torque bias is implemented [35] by multiplying the a priori probability of selecting a move by $\exp(\lambda\beta\mathbf{t}^{\mathsf{T}} \cdot \mathbf{s})$, where β is the inverse temperature, λ a constant (usually taken to be $\frac{1}{2}$), and $\mathbf{t}$ is the torque on the molecule. If the "starting" distribution is a Gaussian then the torque bias factor merely shifts the Gaussian to give

$$p(\mathbf{s}) = \frac{\exp(-\frac{1}{2}|\mathbf{s} - \lambda\beta\Delta^2\mathbf{t}|^2/\Delta^2)}{(2\pi)^{3/2}\Delta^3}.$$

This offers two simplifications over the original procedure [35]: (a) it is trivial to sample from a shifted Gaussian and (b) the acceptance probability, which involves the ratio of the forward and reverse a priori probabilities, is also easy to compute and, in particular, it does not require the evaluation of a normalization factor for the distribution. Similar considerations obviously apply to the application of force bias for translational moves, as has been discussed by Rossky et al. [36]. Indeed, in the case of moving molecules, we would naturally perform a combined translational and orientational move applying both force and torque bias simultaneously. There are often strong gradients in the forces in molecular simulations and a direct application of force bias in this case can lead to poor sampling because certain transitions are effectively disallowed. In such cases, it is prudent to limit the effect of the bias by limiting the shift in the Gaussian, if necessary, to ensure that there a finite probability (at least 5–10%, say) of the sampled move being in the opposite direction to the force. This ensures that the molecule can effectively explore configuration space because small steps are always permitted and it provides a simpler "safety" mechanism than the distance scaling of λ proposed by Mezei [37].

Finally, some care needs to be taken to treat the possibility of the orientation "wrapping" around. Suppose the sampled $\mathbf{s}$ has $|\mathbf{s}| > \pi$, then the resulting orientation is identical to the wrapped one, $\mathbf{s} - 2\pi\mathbf{s}/|\mathbf{s}|$. To ensure that detailed balance is maintained, the acceptance probability should use the a priori probability for the reverse move $-\mathbf{s}$ (rather than the negative of the wrapped move). A simple expedient for avoiding this problem is simply to reject any move with $|\mathbf{s}| > \pi$.

8. Grids for orientation

In many contexts, it is important to be able to represent the independent variables for a problem on a grid. It is therefore useful to be able to map orientations onto a grid. Possible applications are binning molecular data, implementing cavity bias in orientation [38], performing systematic searching of orientations (where the goal is to provide more regular coverage of orientation space than is achieved by random sampling), and performing integrals over orientation by numerical quadrature [39]. Our goal is to provide a simple rule for covering orientation space with a grid while ensuring that the grid elements are approximately of equal volumes and are not unduly distorted. Here again, representing the orientation as a quaternion provides a reasonable solution.

Recall that unit quaternions lie on a hypersphere $\mathbb{S}^3$. Positions on $\mathbb{S}^3$ can be determined by three angle-like variables. However these are a poor basis for a grid because of singularities in the resulting coordinate system. Instead let imagine surrounding $\mathbb{S}^3$ by a tesseract (the 4-dimensional analogue of the cube) of edge length 2. This consists of eight cells which are $2 \times 2 \times 2$ cubes tangent to $\mathbb{S}^3$. An exemplary cell is given by $\mathbf{p}$ with $p_0 = 1$, $|p_{i\neq 0}| \leq 1$. We need only consider half of the cells of the tesseract because of the identification of $\pm\mathbf{p}$. Thus we choose to consider the four cells for which one of the components of $\mathbf{p}$ is $+1$.

This then forms the basis for a cubical grid for orientation space. This is attractive because cubical grids are simple to index into; they are easy to refine; they have an metric factor which is easy to compute, etc. The overall "wastefulness" of this grids relative to a cubic grid within a domain of $\mathbb{R}^3$ is given by the ratio of the volume of four cells of the tesseract (4×2^3) to the area (really a volume) of a hemisphere of $\mathbb{S}^3$, i.e., $32/\pi^2 \approx 3.24$. This might seem rather profligate. However, if we managed to arrange the grid around the $\mathbb{S}^3$ without any wastage, the grid edge would be reduced by a factor of only $\sqrt[3]{3.24} \approx 1.48$.

Let us divide each of the cells of the tesseract into M^3 grid cubes (of side $2/M$). These cubes can then be projected to $\mathbb{S}^3$ by scaling $\mathbf{p}$ to a unit quaternion. This operation scales the volume of each of the grid cubes by $|\mathbf{p}|^{-4}$—a factor of $|\mathbf{p}|^{-3}$ is due to scaling a volume element linearly by $|\mathbf{p}|^{-1}$ and the last factor of $|\mathbf{p}|^{-1}$ arises from the distortion of the cube during this operation. The maximum scaling occurs at the corners of the tesseract, e.g., $\mathbf{p} = [1, 1, 1, 1]$, where $|\mathbf{p}| = 2$, so that range of volumes for the grid elements is 16 with the maximum distortion being a factor of 2. Mapping between an arbitrary orientation $\mathbf{q}$ and a point in the grid is then achieved as follows. We identify the component q_l of $\mathbf{q}$ which is largest in absolute value and set $\mathbf{p} = \mathbf{q}/q_l$, giving $p_l = 1$ and $p_{i\neq l} \in [-1, 1]$. The grid then consists of $4 \times M \times M \times M$ elements. The resolution of the grid, given by the maximum change in orientations between neighboring grid cells, is approximately $4/M$. (We need to multiply the grid cube edge by 2 to obtain the equivalent rotation angle, because, from Eq. (4), we have $\mathbf{q} = [1, \mathbf{v}\theta/2]$ for θ small.)

When the application is quadrature, it is natural to evaluate the function and to compute the metric factor $|\mathbf{p}|^{-4}$ at the

centers of the grid cubes. For binning, we assign the samples to the grid cube in the obvious way and again use the grid center to compute the metric factor to obtain a sample density.

The cubical grid defined above is suitable for quadrature and searching where the cost of function evaluations is small. Sometimes, however, the cost of function evaluations is so high that it is desirable to find an "optimal" set of grid points. For integrations over $\mathbb{S}^2$, this is a well-studied problem [40] and various integration grid have been given that ensure accuracy to high order [41]. For $\mathbb{S}^3$, various spherical t-designs are known [42,43]. A t-design is a set of points on the sphere such that the average of a polynomial of degree t over the sphere is given by averaging the values of polynomial at those points. Unfortunately t-designs for $\mathbb{S}^3$ are only known for t up to 11 with the 11-design corresponding to 60 orientations.

In order to provide a denser coverage of the sphere we propose the following strategy: consider N sample orientations, corresponding to $2N$ points on $\mathbb{S}^3$. Define a "covering radius", α, as the maximum rotation needed to align an arbitrary orientation with one of the sample orientations. The "coverage", c, is defined by the ratio of the area of $2N$ spherical caps of rotational extent α to the total area of $\mathbb{S}^3$, i.e.,

$$c = \frac{N(\alpha - \sin\alpha)}{\pi} \tag{24}$$

[compare with Eq. (23b)]. For a given N, the optimal configuration of sample orientations is obtained by minimizing α—this gives the "thinnest" coverage, c. Finally, we weight each sample point according to the fraction of orientational space which is closest to it (i.e., in proportion to the volumes of the Voronoi cells); and we set a secondary goal of minimizing the variation in the weights. We expect the resulting sample points and weights to provide robust and accurate estimates of orientational integrals—particularly of experimentally or numerically determined functions which are bounded but which may not have bounded derivatives. The sample points are also suitable for searching orientation space optimally. Finding such optimal sets of points is difficult in practice. So, here, we propose some sets based on the regular 4-dimensional polytopes [44,45], with the results summarized in Table 1.

The 24-orientation set is obtained by placing two 24-cells (or icositetrachora) in their mutually dual configurations to give the set

$$8 \text{ permutations of } \quad [\pm 1, 0, 0, 0], \tag{25a}$$

$$16 \text{ permutations of } \quad \left[\pm\frac{1}{2}, \pm\frac{1}{2}, \pm\frac{1}{2}, \pm\frac{1}{2}\right], \tag{25b}$$

$$24 \text{ permutations of } \quad \left[\pm\frac{1}{\sqrt{2}}, \pm\frac{1}{\sqrt{2}}, 0, 0\right]. \tag{25c}$$

(Each orientation is counted twice here because of the identification of $\pm\mathbf{q}$.) The corresponding Voronoi tessellation is a truncated-cubic tetracontaoctachoron (or 48-cell) which consists of 48 regular truncated cubes [48]. The set of orientations, Eq. (25), is the direct symmetry group for the cube.

Table 1
Coverings of orientation space. N is the number of orientations; α is the covering radius (expressed as a rotation), and c is the coverage, Eq. (24)

N	α	c	Notes
24	62.80° [a]	1.579	Vertices of 2 24-cells, a 7-design[b]
60	44.48° [c]	1.445	Vertices of 600-cell, an 11-design[b]
360	27.78°	2.152	Vertices and cells of 600-cell[d]
50	69.66°	4.426	ZCW3_50[e]
538	32.53°	5.142	ZCW3_538[e]
6044	18.10°	10.051	ZCW3_6044[e]
∞	$5.499/\sqrt[3]{N}$	8.821[f]	Cubic lattice in tesseract[g]
∞	$4.472/\sqrt[3]{N}$	4.745[h]	Body-centered cubic lattice in tesseract[g]
∞	$4.092/\sqrt[3]{N}$	3.635[i]	Cubic lattice in 48-cell[j]
∞	$3.328/\sqrt[3]{N}$	1.956[k]	Body-centered cubic lattice in 48-cell[j]
648	20.83°	1.641	c48u27[l] ($\delta = 0.33582$)
7416	10.07°	2.133	c48u309[l] ($\delta = 0.15846$)
70728	4.71°	2.078	c48u2947[l] ($\delta = 0.07359$)
∞	$3.022/\sqrt[3]{N}$	1.464[m]	Uniform body-centered cubic lattice[n]

[a] $\cos^{-1}\left(\frac{1}{4}(2\sqrt{2}-1)\right)$.
[b] See Ref. [43].
[c] $\cos^{-1}\left(\frac{1}{8}(3\sqrt{5}-1)\right)$.
[d] 60 vertices with weight 1.32870 and 300 cell centers with weight 0.93426.
[e] Euler angles for orientations taken from Ref. [39].
[f] $16\sqrt{3}/\pi$.
[g] Ratio of maximum to minimum weights is 16.
[h] $20\sqrt{5}/(3\pi)$.
[i] $224\sqrt{3}/[(17+12\sqrt{2})\pi]$.
[j] Ignoring boundary effects, ratio of maximum to minimum weights is $64/(17+12\sqrt{2}) = 1.884$.
[k] $280\sqrt{5}/[3(17+12\sqrt{2})\pi]$.
[l] Body-centered cubic lattice in a 48-cell with lattice spacing δ; see Ref. [46].
[m] $5\sqrt{5}\pi/24$.
[n] Conjectured thinnest covering for $N \to \infty$, based on optimal covering of $\mathbb{R}^3$ [47].

The vertices of the 600-cell (or hexacosichoron) [45] are given by Eqs. (25a) and (25b) together with

$$96 \text{ even permutations of } \quad \left[\pm\frac{\sqrt{5}+1}{4}, \pm\frac{\sqrt{5}-1}{4}, \pm\frac{1}{2}, 0\right].$$

In this case, the Voronoi tessellation is the dual of the 600-cell, namely the 120-cell (or hecatonicosachoron). Because the Voronoi cells are dodecahedra which are nearly spherical, the resulting 60 orientations gives a particularly thin covering of orientation space. A good covering is also provided by adding the centers of the tetrahedral cells of the 600-cell.

For comparison, we list in Table 1 the data for some of the ZCW3 orientation sets used by Edén and Levitt [39]. These are obtained by taking sets of points appropriate for integrating of a periodic unit cube [49–51] and mapping this set to the space of three Euler angles. There are two potential problems with this approach: (1) even though the metric of orientation space is treated properly, the mapping from Euler angles to orientation space is not distance-preserving and we expect this to degrade the properties of a mesh and (2) because one of the Euler angles is not periodic, functions in orientation space do not obey the constraints assumed in constructing the sets of sample points. (More complete data for the ZCW3 sets is available in [46].)

Finally, Table 1 provides various strategies for constructing an arbitrarily fine grid. We start with gridding the tesseract on

which we easily impose a cubical grid (see above). However, the optimal sphere covering of $\mathbb{R}^3$ is body-centered cubic [47], and such a grid results in a thinner covering. Still better coverings can be found by starting with the 48-cell which has a typical cell (a truncated cube),

$$p_0 = 1, \quad |p_{i\neq 0}| \leq \sqrt{2} - 1, \quad |p_1| + |p_2| + |p_3| \leq 1. \tag{26}$$

The other cells are obtained by multiplying $\mathbf{p}$ by the members of Eq. (25). A cubic or body-centered cubic grid can easily be placed within each cell. For example, a body-centered lattice can be obtained with

$$p_0 = 1, \quad \mathbf{p} = [k, l, m]\delta/2,$$

subject to the constraint Eq. (26), where k, l, and m are either all even integers or all odd integers. Table 1 gives three examples of such grids. The disadvantage of grids in 48-cells is that care must be taken to treat the faces of the cell correctly. The triangular faces of the truncated cubes slice cut through the grid cells at an angle and the octagonal faces fit together with a 45° twist. It is therefore necessary to resort to numerical methods to determine the volume of the Voronoi cells near the faces. The resulting data for the weights and examples of other body-centered cubic grids in the 48-cell with $\alpha \geq 0.65°$ are given in [46].

One special searching problem is determining the volume of the smallest rectangular box (whose edges are parallel to the coordinate axes) into which a given molecule fits. This problem arises in the study of a single protein bathed in a solvent. In order to eliminate boundary effects, it is possible to construct a periodic system and, for efficiency, we wish the volume of the periodic cell to be minimum. We can solve this problem by systematically sampling over all orientations using our grid. However, because of the symmetries of a cube, Eq. (25), there are 24 equivalent orientations which minimize the volume and we can restrict the search to $1/24$ of orientation space by searching only in Eq. (26). We should point out that for the purposes of mimicking a single solute molecule in a solvent with a periodic system, the "best" computational box is not given by fitting a single image of the solute into a box but rather by the more challenging problem of optimally fitting the solute molecules into its neighboring images [52].

The emphasis in the section is on covering all orientation space with a grid. In many molecular modeling applications, the orientation may be quite restricted, e.g., when considering the orientation of a ligand in a protein binding pocket, and we may elect to restrict the integration (or search) to a set of orientations which differ from the mean rotation by at most Θ. If we express the deviation from the mean as a turn vector, Eq. (23), integrations may be carried out in (3-dimensional) turn space with the range of integration restricted to the ball $|\mathbf{u}| \leq \sqrt[3]{(\Theta - \sin\Theta)/\pi}$. Because the mapping to turn space is volume preserving, the integrals are exact. In addition, provided that $\Theta \lesssim \pi/2$, the mapping to turn space entails little distortion ($\lesssim 12\%$) and standard numerical methods for integrating in a ball $\mathbb{B}^3$ can be used.

9. Discussion

Quaternions are an ideal "fiducial" representation [53] of orientation in a molecular simulation. They provide an economical format for program input and output and as the internal representation of orientation. There is little redundancy in the representation—there is just the normalization constraint on its four elements and this is easily tested and corrected. At a given numerical precision, quaternions cover orientation space uniformly. Most operations involving orientation can be carried out directly and efficiently with quaternions and they can be converted to other representations as needed. The basic operation of composing rotations is most cheaply performed with quaternions. On the other hand, if we need to rotate a large molecule it is quicker to convert the quaternion to a rotation matrix, Eq. (8), and to perform matrix-vector multiplication than to apply Eq. (5) directly.

In comparison, other representation suffer serious drawbacks. Rotations cannot be easily composed when expressed as Euler angles. Picking a random orientation is more awkward when rotation matrices are used. In neither of these representations is it easy to interpolate between two orientations or to compute the mean orientation.

Although quaternions may be unfamiliar to some readers, we only needed to use quaternion algebra in the rule for composing rotations and in deriving the least-squares fit. In carrying out the other tasks, we just used the fact that rotations are represented by opposite points on $\mathbb{S}^3$ and this provides a "natural" metric for rotations. In working with $\mathbb{S}^3$, we are able to carry over geometrical concepts from $\mathbb{S}^2$ or use straightforward extensions from Euclidean space, $\mathbb{R}^3$, to $\mathbb{R}^4$.

A curious and non-obvious property of rotations which is evident from their representation on $\mathbb{S}^3$, with $\pm\mathbf{q}$ identified, is that rotations do not form a simply connected group. Thus, if we rotate an object by 360° it returns to its original orientation but with the sign of $\mathbf{q}$ changed. This means that we cannot continuously deform the path that the object took to reduce it to a point. However, we *can* do this if we rotate the object by 720°. This property of rotations is an immediate consequence of their representation as a pair of points $\pm\mathbf{q}$ on $\mathbb{S}^3$ and good visual illustrations of this property are provided by the Dirac belt trick [54] and the Phillipine wine dance [55].

Acknowledgments

This work was supported by the U.S. Army Medical Research and Materiel Command under Contract No. DAMD17-03-C-0082. The views, opinions, and findings contained in this report are those of the author and should not be construed as an official Department of the Army position, policy, or decision. No animal testing was conducted and no recombinant DNA was used.

References

[1] W.R. Hamilton, On a new species of imaginary quantities connected with a theory of quaternions, Proc. R. Irish Acad. 2 (1844) 424–434. http://www.maths.tcd.ie/pub/HistMath/People/Hamilton/Quatern1.

[2] W.R. Hamilton, On quaternions, Proc. R. Irish Acad. 3 (1847) 1–16. http://www.maths.tcd.ie/pub/HistMath/People/Hamilton/Quatern2.
[3] J.C. Maxwell, A Treatise on Electricity and Magnetism, vol. 1, third ed., Clarendon Press, Oxford, 1892.
[4] J.W. Gibbs, E.B. Wilson, Vector Analysis, Yale University Press, 1901.
[5] Quaternion, Wikipedia, The Free Encyclopedia, http://en.wikipedia.org/wiki/Quaternion.
[6] E.W. Weinstein, Quaternion, From MathWorld, http://mathworld.wolfram.com/Quaternion.html.
[7] D.R. Flower, Rotational superposition: A review of methods, J. Mol. Graph. Mod. 17 (1999) 238–244.
[8] G. Wahba, Problem 65–1: A least squares estimate of spacecraft attitude, SIAM Rev. 7 (1965) 409.
[9] G. Wahba, J.L. Farrell, J.C. Stuelpnagel, R.H. Wessner, J.R. Velman, J.E. Brock, R. desJardins, Solution of problem 65–1: A least squares estimate of spacecraft attitude, SIAM Review 8 (1966) 384–386.
[10] J.E. Keat, Analysis of least-squares attitude determination routine DOAOP, Technical Report CSC/TM–77/6034, Computer Sciences Corp., February 1977, http://charles.karney.info/papers/keat77.pdf.
[11] G.M. Lerner, Three-axis attitude determination, in: J.R. Wertz (Ed.), Spacecraft Attitude Determination and Control, vol. 73, Astrophysics and Space Science Library, Reidel, Dordrecht, 1978, pp. 420–428.
[12] O.D. Faugeras, M. Hebert, A 3-D recognition and positioning algorithm using geometrical matching between primitive surfaces, in: Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Held in Karlsruhe, W. Germany, August 8–12, 1983, pp. 996–1002.
[13] B.K.P. Horn, Closed-form solution of absolute orientation using unit quaternions, J. Opt. Soc. Am. A 4 (1987) 629–642.
[14] R. Diamond, A note on the rotational superposition problem, Acta Cryst. A 44 (1988) 211–216.
[15] S.K. Kearsley, On the orthogonal transformation used for structural comparisons, Acta Cryst. A 45 (1989) 208–210.
[16] E. Anderson, Z. Bai, C. Bischof, et al., LAPACK Users' Guide, third ed., SIAM, 1999, http://www.netlib.org/lapack/.
[17] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, F. Rossi, GNU Scientific Library Reference Manual, second ed., Network Theory Ltd., 2002 http://www.gnu.org/software/gsl/.
[18] R. Pozo, Template Numerical Toolkit for linear algebra: High performance programming with C++ and the Standard Template Library, Int. J. Supercomput. Appl. High Perform. Comput. 11 (1997) 251–263. http://math.nist.gov/tnt/.
[19] M.D. Shuster, S.D. Oh, Three-axis attitude determination from vector observations, J. Guid. Control 4 (1981) 70–77.
[20] R. Diamond, Chirality in rotational superposition, Acta Cryst. A 46 (1990) 423.
[21] E.A. Coutsias, C. Seok, K.A. Dill, Using quaternions to calculate RMSD, J. Comput. Chem. 25 (2004) 1849–1857.
[22] P.H. Hünenberger, A.E. Mark, W.F. van Gunsteren, Fluctuation and cross-correlation analysis of protein motions observed in nanosecond molecular dynamics simulations, J. Mol. Biol. 252 (1995) 492–503.
[23] R.G. Littlejohn, M. Reinsch, Gauge fields in the separation of rotations and internal motions in the n-body problem, Rev. Mod. Phys. 69 (1997) 213–275.
[24] K. Shoemake, Animating rotation with quaternion curves, Computer Graphics 19 (1985) 245–254, SIGGRAPH '85, San Fransisco, Jul 22–26, 1985.
[25] K.V. Mardia, P.E. Jupp, Directional Statistics, J. Wiley & Sons, 1999.
[26] D. Rancourt, L.-P. Rivest, J. Asselin, Using orientation statistics to investigate variations in human kinematics, Appl. Statist. 49 (2000) 81–94.
[27] C.F.F. Karney, J.E. Ferrara, S. Brunner, Method for computing protein binding affinity, J. Comput. Chem. 26 (2005) 243–251. http://arxiv.org/abs/cond-mat/0401348.
[28] A.P. Dempster, N.M. Laird, D.B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, J. R. Statist. Soc. B 39 (1977) 1–38.
[29] C.Müller, Spherical Harmonics, Lecture Notes in Mathematics, vol. 17, Springer-Verlag, 1966.
[30] C.Müller, Analysis of Spherical Symmetries in Euclidean Spaces, Applied Mathematical Sciences, vol. 129, Springer-Verlag, 1998.
[31] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, Equation of state calculations by fast computing machines, J. Chem. Phys. 21 (1953) 1087–1092.
[32] D.J. Adams, Grand canonical ensemble Monte Carlo for a Lennard-Jones fluid, Mol. Phys. 29 (1975) 307–311.
[33] G. Marsaglia, Choosing a point from the surface of a sphere, Ann. Math. Stat. 43 (1972) 645–646.
[34] M.E. Muller, Some continuous Monte Carlo methods for the Dirichlet problem, Ann. Math. Stat. 27 (1956) 569–589.
[35] M. Rao, C. Pangali, B.J. Berne, On the force bias Monte Carlo simulation of water: Methodology, optimization and comparison with molecular dynamics, Mol. Phys. 37 (1979) 1773–1789.
[36] P.J. Rossky, J.D. Doll, H.L. Friedman, Brownian dynamics as smart Monte Carlo simulation, J. Chem. Phys. 69 (1978) 4628–4633.
[37] M. Mezei, Distance-scaled force biased Monte Carlo simulation for solutions containing a strongly interacting solute, Mol. Simulation 5 (1991) 405–408.
[38] M. Mezei, A cavity-biased (T,V, μ) Monte Carlo method for the computer simulation of fluids, Mol. Phys. 40 (1980) 901–906.
[39] M. Edén, M.H. Levitt, Computation of orientational averages in solid-state NMR by Gaussian spherical quadrature, J. Magn. Reson. 132 (1998) 220–239. http://www.mhl.soton.ac.uk/research/software/Orientations.
[40] A.D. McLaren, Optimal numerical integration on a sphere, Math. Comp. 17 (1963) 361–383.
[41] V. Lebedev, D. Laikov, A quadrature formula for the sphere of the 131st algebraic order of accuracy, Dokl. Math. 59 (1999) 477–481.
[42] P. Delsarte, J.M. Goethals, J.J. Seidel, Spherical codes and designs, Geom. Dedicata 6 (1977) 363–388.
[43] N.J.A. Sloane, R.H. Hardin, P. Cara, Spherical designs in four dimensions, in: Proceedings of the Information Theory Workshop, Paris, March 31 to April 4, 2003, IEEE Press, pp. 253–257, http://www.research.att.com/~njas/doc/paris2.pdf.
[44] H.S.M. Coxeter, Regular Polytopes, second ed., Macmillan, 1963.
[45] P.D. Bourke, Regular polytopes (Platonic solids) in 4D (2003), http://astronomy.swin.edu.au/~pbourke/polyhedra/platonic4d/.
[46] C.F.F. Karney, Nearly optimal coverings of orientation space (2006) http://charles.karney.info/orientation/.
[47] R.P. Bambah, On lattice coverings by spheres, Proc. Nat. Inst. Sci. India 20 (1954) 25–52.
[48] G. Olshevsky, Convex uniform polychora based on the icositetrachoron (24-cell) (2004) http://members.aol.com/Polycell/section3.html.
[49] S.K. Zaremba, Good lattice points, discrepancy, and numerical integration, Ann. Mat. Pura. Appl. (ser. 4) 73 (1966) 175–180.
[50] H. Conroy, Molecular Schrödinger equation. VIII. A new method for the evaluation of multidimensional integrals, J. Chem. Phys. 47 (1967) 5307–5318.
[51] V.B. Cheng, H.H. Suzukawa, M. Wolfsberg, Investigations of a nonrandom numerical method for multidimensional integration, J. Chem. Phys. 59 (1973) 3992–3999.
[52] H. Bekker, J.P. van den Berg, T.A. Wassenaar, A method to obtain a near-minimal-volume molecular simulation of a macromolecule, using periodic boundary conditions and rotational constraints, J. Comput. Chem. 25 (2004) 1037–1046.
[53] W. Kahan, J.D. Darcy, How Java's floating-point hurts everyone everywhere, ACM Workshop on Java for High-Performance Network Computing, March 1998, Stanford University, http://www.cs.berkeley.edu/~wkahan/JAVAhurt.pdf.
[54] G. Egan, Dirac's belt trick, Java Applet (2000) http://gregegan.customer.netspace.net.au/APPLETS/21/21.html.
[55] G.K. Francis, L.H. Kauffman, D.J. Sandin, Air on the Dirac strings, Film produced at Electronic Visualization Laboratory, University of Illinois, Chigaco, 1993 http://www.evl.uic.edu/hypercomplex/html/dirac.html.