Representing Rotations

Gabriel Taubin [*]

Representing Rotations: A three-dimensional rotation can be represented as an orthogonal 3×3 matrix Q (with the transpose equal to inverse $QQ^t = I$ and unit determinant $|Q| = 1$). The result of applying a rotation to a three-dimensional vector p is obtained by multiplying the matrix by the vector Qp. A more geometrically intuitive way of describing a three-dimensional rotation is as a turn of angle θ around a unit-length three-dimensional vector u, with the positive direction of rotation specified by the right hand rule. Let us denote such rotation $Q(\theta, u)$. The matrix representation of this rotation can be computed using Rodrigues' formula as follows

$$Q(\theta,u) = I + s\,U + (1-c)\,U^2 \quad \text{where} \quad \begin{cases} s &= \sin(\theta) \\ c &= \cos(\theta) \end{cases} \quad (1)$$

and U is the *skew-symmetric* ($U^t = -U$) matrix

$$U = \begin{pmatrix} 0 & -u_z & u_y \\ u_z & 0 & -u_x \\ -u_y & u_x & 0 \end{pmatrix} .$$

corresponding to the vector product by the vector u:

$$U\,v = u \times v = \begin{pmatrix} u_y\,v_z - u_z\,v_y \\ u_z\,v_x - u_x\,v_z \\ u_x\,v_y - u_y\,v_x \end{pmatrix} \quad \forall\, v \in \mathbb{R}^3 .$$

For each pair of unit length linearly independent vectors n_i and n_j, the matrix Q_{ij} corresponding to the rotation that minimizes the turning angle amongst all the rotations that transform the vector n_i into n_j can be computed using Rodrigues' formula without explicit determination of the angle of rotation, with

$$\begin{cases} c &= n_i^t n_j \\ s &= \sqrt{1-c^2} \\ u &= n_i \times n_j / s \ \in \mathbb{R}^3 \end{cases} \quad (2)$$

This can be extended with continuity to the case $n_i = n_j$ by defining Q_{ij} as the identity matrix. It cannot be extended to the case $n_i = -n_j$ because of lack of uniqueness: for every unit-length vector u orthogonal to n_i, the rotation $Q(\pi, u)$ transforms n_i into n_j, and no rotation of angle less than π transforms n_i into n_j.

The Exponential Map: Rodrigues' formula es an efficient algorithm to evaluate the *exponential map*

$$\begin{array}{lcl} \exp : \mathbb{R}^3 & \to & SO(3) \\ \quad v & \mapsto & e^V = I + \sum_{n=1}^{\infty} \frac{1}{n!}\,V^n \ , \end{array}$$

where a three-dimensional vector v is represented as the product of a non-negative *magnitude* θ and a unit-length *phase vector* u (the value of u is irrelevant if $\theta = 0$), and V is the skew-symmetric matrix defined by v ($V = \theta U$). The identity $Q(\theta, u) = e^{\theta U}$ is easy to prove using the series expansion of $\cos(\theta)$ and $\sin(\theta)$ and the identity $U^3 + U = 0$, which can be verified by direct expansion.

The exponential map with $\mathbb{R}^3$ as domain is surjective, i.e., every rotation has a representation as an exponential of a skew symmetric matrix. To obtain a local parameterization we need to restrict the domain to a smaller open subset of $\mathbb{R}^3$. To determine such subset, note that the map $\theta \mapsto e^{\theta U}$ with fixed u is 2π-periodic, and for $\pi \le \theta \le 2\pi$ we have

$$e^{\theta U} = e^{(2\pi - \theta)(-U)} ,$$

i.e., the rotation of angle θ around u is equal to the rotation of angle $2\pi - \theta$ around $-u$. When the domain is restricted to the open ball of radius π

$$\Omega_\pi = \{v : \|v\| < \pi\} = \{\theta u : 0 \le \theta < \pi \ , \ \|u\| = 1\},$$

the exponential map becomes $1-1$, but rotations of angle π cannot be represented. This is acceptable within the framework of mesh processing because surface normal rotations of angle π do not occur in practice. The image of Ω_π through the exponential map, i.e., the set of rotations of angle less than π, is an open neighborhood of the identity in the group of three-dimensional rotations $SO(3)$ (which is a three-dimensional Lie group).

On the set of rotations of angle less than π the exponential map has a well defined inverse, the logarithm. The logarithm of a rotation Q of angle less than π can be computed following these steps

$$\begin{cases} c &= (1 - \text{trace}\,(Q))/2 \\ V &= (Q - Q^t)/2 \\ s &= \|v\| \\ u &= v/s \\ \theta &= \text{angle}_{[0,\pi)}(s,c) \ . \end{cases} \quad (3)$$

In our application, where we only need to evaluate the logarithm for rotations defined by vector products of unit-length vectors, instead of the first four steps of equation 3, we perform the three steps of equation 2. In either case, determining the angle θ from $s = \cos(\theta)$ and $c = \cos(\theta)$ is a rather expensive computation, which we would like not to perform very often, if ever.

Averaging Rotations: Let $Q_1, \ldots, Q_N$ be rotation matrices of angle less than π, and let $v_1, \ldots, v_N \in \Omega_\pi$ with $v_j = \theta_j u_j$ and $Q_j = e^{v_j}$ for $j = 1, \ldots, N$. We can define the scaled weighted average of these rotations, with scale factor λ and weights $w_1, \ldots, w_N$ as

$$Q = e^v \quad \text{where} \quad v = \lambda \sum_{j=1}^{N} w_j\,v_j \in \mathbb{R}^3 .$$

This is well defined as long as v stays within Ω_π. Since Ω_π is symmetric with respect to the origin, and convex, a sufficient condition for v to remain within Ω_π is that

$$|\lambda| \sum_{j=1}^{N} |w_j| \le 1 .$$

This parameterization is robust and works well in our applications. The problem with this approach is that we cannot avoid evaluating inverse trigonometric functions to get the angles θ_j, and then forward trigonometric functions to obtain the final result after scaling and averaging. To avoid computing forward and inverse trigonometric functions we need to consider other parameterizations.

[*]Brown University, Division of Engineering,Box D, Providence, RI 02912

Quaternions: Since in equation 1 we have $s^2 + c^2 = 1$, let us consider the following mapping as a candidate to replace the exponential

$$\begin{cases} \Omega_1 & \rightarrow & SO(3) \\ su & \mapsto & I + s\,U + (1-c)\,U^2 \end{cases}$$

where Ω_1 is the open unit ball in $\mathbb{R}^3$, $0 \leq s < 1$, $c = \sqrt{1-s^2}$, and $|u| = 1$. Rather than the angle of rotation itself, we use the sine of the angle as magnitude for the parameter vector $v = su$. This parameterization is also $1-1$, and neither inverse nor forward trigonometric functions need to be computed to do the averaging and scaling in the domain of this parameterization, because in this case $R^{-1}(Q_{ij}) = n_i \times n_j$. Only one square root is needed to evaluate $R(v)$ after the scaling and weighted averaging:

$$\begin{cases} \Omega_1 & \rightarrow & SO(3) \\ v & \mapsto & I + V + \frac{1-\sqrt{1-\|v\|^2}}{\|v\|^2}\,V^2\ , \end{cases}$$

where $V = sU$ is the skew-symmetric matrix corresponding to the vector $v = su$.

The problem is that, since c is non-negative here, the image of Ω_1 through this parameterization is the set of rotations of angle less than $\pi/2$, and so, some of the rotations to be averaged may not have a corresponding preimage in Ω_1. To solve this problem we regard the magnitude $s = \|v\|$ of the vectors in Ω_1 not as $\sin(\theta)$, but as $\sin(\theta/2)$. Since

$$\begin{cases} 1-\cos(\theta) & = & 2\sin(\theta/2)^2 & = & 2s^2 \\ \sin(\theta) & = & 2\sin(\theta/2)\cos(\theta/2) & = & 2sc \end{cases}$$

the following parameterization

$$\begin{cases} \Omega_1 & \rightarrow & SO(3) \\ su & \mapsto & I + 2sc\,U + 2s^2\,U^2 \end{cases} \tag{4}$$

covers the same open set of rotations as the exponential map. Note that since in this case $0 \leq \sin(\theta/2)\ , \cos(\theta/2) \leq 1$, we can safely compute c as $\sqrt{1-s^2}$, and by combining terms we can evaluate this parameterization also with a single square root:

$$\begin{cases} \Omega_1 & \rightarrow & SO(3) \\ v & \mapsto & I + 2\sqrt{1-\|v\|^2}\,V + 2\,V^2\ , \end{cases} \tag{5}$$

where $V = sU$ is again the skew-symmetric matrix corresponding to the vector $v = su$. But in this case one additional square root to compute the inverse parameterization $R^{-1}(Q_{ij}) = n_i \times n_j \,/ \sqrt{2 + 2\,n_i^t n_j}$ is needed for each neighbor

Note that this is the parameterization associated with the representation of rotations by quaternions. Quaternions are particularly popular in Computer Graphics because only four parameters are needed to represent a rotation, as opposed to nine in matrix form, and composition of rotations corresponds to the product of quaternions. A quaternion is a pair (a, b), where $a \in \mathbb{R}$ is a scalar, and $b \in \mathbb{R}^3$ is a vector. The product of two quaternions (a_1, b_1) and (a_2, b_2) is given by this formula

$$(a_1, b_1) \cdot (a_2, b_2) = (a_1 a_2 - b_1^t b_2, a_1 b_2 + a_2 b_1 + b_1 \times b_2)\ .$$

If $Q = Q(\theta, u)$ is the rotation of angle $0 \leq \theta < \pi$ around a unit-length vector u, and $p \in \mathbb{R}^3$, it is well known that

$$(0, Qp) = (c, su) \cdot (0, p) \cdot (c, -su) \tag{6}$$

where $c = \cos(\theta/2)$ and $s = \sin(\theta/2)$. To prove this identity it is sufficient to expand the two quaternion products to verify that Q is equal to the right hand side expression of equation 4.

Note that in terms of computational cost, the expression $(I + 2sc\,U + 2s^2\,U^2)\,p$ can be evaluated with 21 multiplications and 12 additions, while evaluating Qp expanding the quaternion products in equation 6 requires 24 multiplications and 17 additions, taking into account the computational savings associated with the zeros in the middle factor and result. As a comparison, a three-dimensional matrix vector multiplication requires 27 multiplications and 18 additions, without counting the operations needed to build the matrix from the axis and angle.

Cayley Rational Parameterization: If we also want to minimize the number of square roots, we can use the following less known parameterization of the set of rotations of angle less than π, due to Cayley

$$\begin{cases} \mathbb{R}^3 & \rightarrow & SO(3) \\ v & \mapsto & (I-V)(I+V)^{-1} \end{cases} \tag{7}$$

where V is the skew-symmetric matrix corresponding to the three-dimensional vector v. Since $|I - V| = |I + V| = 1 + \|v\|^2$, this function is well defined for any vector $v \in \mathbb{R}^3$. And the inverse of this parameterization is given by the same formula: if Q is a rotation of angle less than π, then we have $|I+Q| \neq 0$ and $V = (I-Q)(I+Q)^{-1}$. To establish the relation with the other parameterizations, we observe that the expression

$$(I+V)^{-1} = I - \frac{1}{1+\|v\|^2}\,V + \frac{1}{1+\|v\|^2}\,V^2 \tag{8}$$

follows from the identity $V^3 + \|v\|^2 V = 0$, and so

$$Q = (I-V)(I+V)^{-1} = I - \frac{2}{1+\|v\|^2}\,V + \frac{2}{1+\|v\|^2}\,V^2\ ,$$

which can be evaluated without square roots. If we now write $v = -\mu u$ with $\|u\| = 1$ in the last equation, we obtain

$$Q = I + \frac{2\mu}{1+\mu^2}\,U + \frac{2\mu^2}{1+\mu^2}\,U^2\ , \tag{9}$$

which es equal to the parameterization of equation 4 with $\mu = \tan(\theta/2)$.

Computing the inverse parameterization does not involve square roots either

$$R^{-1}(Q_{ij}) = \frac{1 - n_i^t n_j}{\|n_i \times n_j\|^2}\,n_i \times n_j$$

because

$$\tan(\theta/2) = \frac{s}{c} = \frac{2s^2}{2sc} = \frac{1-\cos(\theta)}{\sin(\theta)} = \frac{1 - n_i^t n_j}{\|n_i \times n_j\|}\ .$$