# Averages of Rotations and Orientations in 3-space

*Leandra Vicci*
*Microelectronic Systems Laboratory*
*Department of Computer Science*
*University of North Carolina at Chapel Hill*
*10 August 2001*

## 1 Introduction

Finding the mean (average) of $n$ arbitrary 3D rotations or orientations is conceptually challenging both in its algebra and its geometric interpretation. To my knowledge, no solutions have been found for the general case, although a number of approximate solutions have been proposed [3, 4, 5] to accurately and efficiently handle those limiting cases typical of engineering interest. This report proposes an exact, closed form solution which is extended from the 2D complex number domain where the averaging methods are widely accepted.

Representing rotations as rotation (unit) quaternions [1] is convenient in that a rotation composed of a sequence of successive rotations can be simply represented as the (non-commutative) product of the successive rotation quaternions. But this very property is problematical for averaging.

A mean average is conventionally taken as a sum of averagands divided by their number. If applied to rotation quaternions however, the result is not in general a rotation quaternion; moreover, there is no sensible geometric interpretation of such a procedure. The multiplicative analogue of this procedure, taking the $n^{th}$ root of the product of rotation quaternions *does* return a rotation quaternion. However the result is dependent on the order of multiplication, while a meaningful average should be order independent.

This report explores the extension of rotations by unit complex numbers in the complex plane to rotations by unit quaternions in 3-space. Unit complex numbers $u = e^{i\phi}$ form a proper subspace of the rotation quaternions. Multiplication of any complex number $z$ by $u$ will yield a number $z'$ with identical magnitude, rotated by an angle $\phi$ about the origin in the complex plane. Analogous to quaternion rotations, a complex rotation may be composed of a product (commutative in this case) of successive rotations. In averaging, the commutative property of complex multiplication provides a unique result, the principal $n^{th}$ root of which does indeed geometrically represent the average rotation. Mapped into logarithmic space, the averaging procedure becomes the conventional sum of averagands divided by their number. We'll show how this logarithmic mapping can be extended from complex to quaternion space, where the summation remains commutative, thus satisfying the order independence requirement for averaging.

Distinct from rotations are orientations which pose an additional set of problems, which we shall discuss in the second half of this report.

We shall use notational representations of, e.g., $r$ for real or complex numbers, $\boldsymbol{v}$ for vectors, $\hat{\boldsymbol{v}}$ for unit vectors and $Q$ for quaternions. This notation notwithstanding, it should be understood that all these quantities should be regarded as equivalent to their quaternion form.

## 2 Review of rotations in the complex plane

Let us introduce the rotation of a general complex number $z = x + i\,y$ by a unit magnitude complex number, which without loss of generality can be represented as $u = \cos\phi + i\,\sin\phi$, and which we shall call a *rotation complex.*

Consider the complex product $z' = uz = zu = \big(x\,\cos\phi - y\,\sin\phi\big) + i\big(x\,\sin\phi + y\,\cos\phi\big)$ which represents a counterclockwise (CCW) rotation of $z$ by an angle $\phi$ about the origin (figure 1).

Since $e^{i\phi} = \cos\phi + i\,\sin\phi = u$, we may write $\log u = i\phi$, which we shall call a *rotation imaginary.*

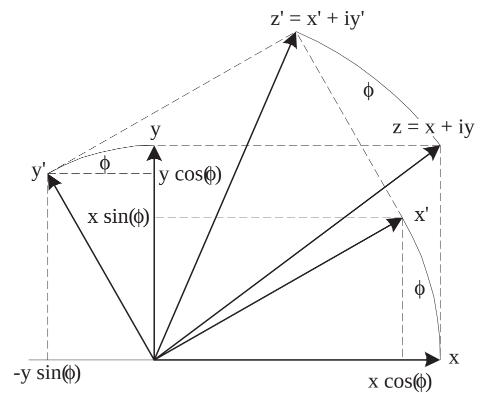


Figure 1: A complex rotation

A succession of rotations $(u_1, u_2, \cdots)$ may be composed to produce the result $z' = (\cdots((z\,u_1)u_2)\cdots u_N) = z\prod_{n=1}^{N}(u_n)$, representing a rotation of $z$ to the state after the last of the successive rotations. In complex logarithmic space, the corresponding result is $\log z' = \log z + i\sum_{n=1}^{N}(\phi_n)$.

One could imagine calculating a mean average of $N$ rotations by applying them successively as just described, then appropriately decimating the result by $N$. For rotation complexes, $u_{avg} = \left(\prod_{n=1}^{N} u_n\right)^{1/N}$; for rotation imaginaries, $\phi_{avg} = \left(\sum_{n=1}^{N} \phi_n\right)/n$.

A fine point in geometric interpretation should be appreciated here. Composing a final rotation from a succession of individual rotations implies the rotations occur in a particular order which may be significant to the result. A set of rotations to be averaged, however, does not involve any notion of succession, so the result should be independent of the order in which the rotations are considered.

Algebraicly, this means whatever operation is used to aggregate the individual rotations for averaging *must* be commutative. In the complex number domain, both addition and multiplication are commutative so averaging either rotational complexes or rotational imaginaries as described above could be valid.

On the other hand, a composite rotation *may* depend on the inherent order in which the rotations occur. For 2D rotations this is not the case, consistent with commutativity of complex multiplication. But looking ahead to 3D rotations, the order in which rotations are composed is definitely significant, suggesting non-commutative rotation operators.

The reader is encouraged to try composing two simple 3D rotations, each by a quarter turn or $\pi/2$ radians about $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{z}}$ respectively. If the first rotation is about $\hat{\boldsymbol{x}}$ and the second is about $\hat{\boldsymbol{z}}$, the coordinate frame is transformed as $[x, y, z] \mapsto [y, z, x]$, but if the order is reversed, the transformation is $[x, y, z] \mapsto [z, -x, -y]$.

## 3 Quaternions as 3D rotation operators

Quaternions are hypercomplex numbers (which comprise the reals, complexes, quaternions, and Cayley numbers), and thus share many of the properties of the complexes. In particular, they work well for representing 3D rotations, and we shall see how they may be used to obtain rotational averages.

### 3.1 Definition and basic properties of quaternions

The quaternion $Q = [w, x, y, z]$ represents a 4-variate polynomial $1 \cdot w + i \cdot x + j \cdot y + k \cdot z$, where the symbolic coefficients $i, j, k$ have the following properties of both vector cross products and complex multiplication,

$$\begin{gathered} ij = k,\ ji = -k, \\ jk = i,\ kj = -i, \\ ki = j,\ ik = -j, \\ i^2 = -1,\ j^2 = -1,\ k^2 = -1. \end{gathered}$$

Accordingly, $i, j, k$ represent three mutually orthogonal unit vectors $\hat{\boldsymbol{i}}, \hat{\boldsymbol{j}}, \hat{\boldsymbol{k}}$ respectively, and the quaternion $[0, x, y, z]$ represents a 3-vector $(\hat{\boldsymbol{i}}x + \hat{\boldsymbol{j}}y + \hat{\boldsymbol{k}}z)$. Moreover, $i, j, k$ each also represents a complex quantity $\sqrt{-1}$, and the quaternions $[w, x, 0, 0], [w, 0, y, 0], [w, 0, 0, z]$ represent three orthogonal spaces of complex numbers. Other definitions are,

Conjugate: $\overline{Q} = [w, -x, -y, -z]$,

Norm: $|Q| = \sqrt{Q\overline{Q}} = \sqrt{\overline{Q}Q} = \sqrt{w^2 + x^2 + y^2 + z^2}$.

The following quaternion properties are proved in [1],

Addition: $Q + P = P + Q, \qquad Q + (P + R) = (Q + P) + R,$

Distributive: $P(Q + R) = PQ + PR, \qquad (P + Q)R = PR + QR,$

Multiplication: $zQ = Qz, \qquad QP \neq PQ, \qquad \overline{QP} = \overline{P}\,\overline{Q}, \qquad Q(PR) = (QP)R,$

Unit quaternion: $|U| = 1, \qquad U\overline{U} = 1 \quad \Rightarrow \quad \overline{U} = U^{-1}.$

### 3.2 Rotations by purely complex quaternions

We define purely complex quaternions to be $Q_x = [w, x, 0, 0]$. The algebra of $Q_x$ is isomorphic with the algebra of complex numbers. Thus we may rely on the properties of complex numbers for all operations on $Q \in Q_x$.

Let $U = [\cos\phi, \sin\phi, 0, 0]$. Clearly $U \in Q_x$ and $|U| = 1$, i.e., it is a purely complex unit quaternion. Let $Q = [w, x, y, z]$ be an arbitrary quaternion which we may write as $Q = Q_\parallel + Q_\perp$, where $Q_\parallel = [w, x, 0, 0] \in Q_x$ and $Q_\perp = [0, 0, y, z] \perp Q_x$. More precisely, their vector components are parallel and perpendicular, respectively.

Then $UQ = U(Q_\parallel + Q_\perp) = UQ_\parallel + UQ_\perp$. Since $[U, Q_\parallel] \in Q_x$, the first term is simply a rotation of $Q_\parallel$ by an angle $\phi$ in the complex $wx$ plane. Expanding the second term

$UQ_\perp = [0, 0, y\cos\phi - z\sin\phi, y\sin\phi + z\cos\phi]$, we find a formally similar rotation of $Q_\perp$ by an angle $\phi$ in the $yz$ plane.

Similarly, $Q\overline{U} = (Q_\parallel + Q_\perp)\overline{U} = Q_\parallel\overline{U} + Q_\perp\overline{U}$. In this case, $Q_\parallel\overline{U} \mapsto qu^* = u^*q$ in the corresponding complex algebra, where the complex conjugate $u^*$ represents a rotation of $q$ by an angle of $-\phi$. Thus, $Q_\parallel\overline{U}$ represents a rotation of $Q_\parallel$ by an angle of $-\phi$ in the $wx$ plane. However expanding $Q_\perp\overline{U} = [0, 0, y\cos\phi - z\sin\phi, y\sin\phi + z\cos\phi]$, we find it represents a rotation of $Q_\perp$ by angle $+\phi$ in the $yz$ plane.

We may compose these products $(UQ)\overline{U} = UQ\overline{U}$ to define a rotation operator, $\mathcal{R}(Q, U) = UQ\overline{U}$, which rotates any quaternion $Q$ by an angle $2\phi$ about $U$. In particular, the $y$ and $z$ components of $Q$ are rotated about $\hat{\boldsymbol{x}}$ (in the $yz$ plane), while the $w$ and $x$ components are left invariant. The vector $V = [0, x, y, z]$ is just a subspace of $Q$ and is rotated in the same manner about $U$.

Note that from the symmetry in definition of quaternions, everything we showed here for $U \in Q_x$ applies equally for $U \in Q_y = [w, 0, y, 0]$ and $U \in Q_z = [w, 0, 0, z]$; that is, the classes $U_x, U_y, U_z$, are all purely complex unit quaternions. Accordingly, we can write,

$$\begin{aligned}
\mathcal{R}(Q, U_x) &= U_x Q \overline{U}_x, \quad U_x = [\cos\phi, \sin\phi, 0, 0];\\
\mathcal{R}(Q, U_y) &= U_y Q \overline{U}_y, \quad U_y = [\cos\phi, 0, \sin\phi, 0];\\
\mathcal{R}(Q, U_z) &= U_z Q \overline{U}_z, \quad U_z = [\cos\phi, 0, 0, \sin\phi];
\end{aligned}$$

where $\mathcal{R}$ rotates $Q$ by an angle $2\phi$ about $U_k$, $k = x, y, z$.

### 3.3 Rotations by any unit quaternion

Let us now extend these results to rotations by *any* unit quaternion, which can be represented as $U = [cos(\phi), \hat{\boldsymbol{u}}\, sin(\phi)] \equiv [cos(\phi), x\, sin(\phi), y\, sin(\phi), z\, sin(\phi)]$, where $\hat{\boldsymbol{u}}$ is an arbitrary unit vector, and $x^2 + y^2 + z^2 = 1$.

Consider two rotations, $U_x \parallel \hat{\boldsymbol{x}}$ and $U_z \parallel \hat{\boldsymbol{z}}$, which we apply sequentially to $U$. Let us pick $U_x$ such that $U \mapsto U'$ is rotated into the $xy$ plane, i.e., $u'_z = 0$; and pick $U_z$ such that $U' \mapsto U''$ is rotated into the $xz$ plane, i.e., $u''_y = 0$. Now $U'' \parallel \hat{\boldsymbol{x}}$ and is therefore a pure complex quaternion, as are $U_x$ and $U_y$ by hypothesis. Formally we write, $U'' = U_z(U_x U \overline{U}_x)\overline{U}_z = U_z U_x U \overline{U}_x \overline{U}_z$. Multiplying both sides on the left by $\overline{U}_x\overline{U}_z$ and on the right by $U_z U_x$ we obtain $\overline{U}_x\overline{U}_z U'' U_z U_x = \overline{U}_x\overline{U}_z(U_z U_x U \overline{U}_x \overline{U}_z)U_z U_x = U$, showing that $U$ can be composed from three purely complex unit quaternions.

We may therefore write the geometric operation

$$\mathcal{R}(Q, U), \quad U = [\cos\phi, \hat{\boldsymbol{u}}\sin\phi],$$

which rotates $Q$ about $U$ by an angle $2\phi$, where $\mathcal{R}(Q, U) = \mathcal{R}(\mathcal{R}(\mathcal{R}(Q, U''), \overline{U}_z), \overline{U}_x)$.

## 4 Averaging 3D rotations with quaternions

Notice that the composition of $\mathcal{R}(Q, U)$ in the last section is order dependent both in its notation and in its algebra. As pointed out in the 2D case, there is no notion of order in averaging, so trying to average, for example, $U_x$, $U_z$, and $U''$ multiplicatively would be nonsensical. It is easy to prove that multiplicative averaging of $U_1 \dots U_n$ works if $\hat{\boldsymbol{u}}_1 = \hat{\boldsymbol{u}}_2 = \dots \hat{\boldsymbol{u}}_n$, but that is not very general; in fact it is the 2D case in the plane normal to $\hat{\boldsymbol{u}}_1$.

Instead, let us attempt to extend the 2D averaging of rotations in logarithmic space. But first we must have definitions for quaternion logarithms $log_e(Q)$ and exponentials $e^Q$, and understand their properties. We shall approach this task using power series expansions for sines, cosines, and the exponential function.

In preparation, we establish the properties of integer powers of unit vector quaternions. By simple component wise multiplication of two vector quaternions $P = [0, \boldsymbol{p}]$ and $Q = [0, \boldsymbol{q}]$ it is easily shown that $PQ = \boldsymbol{pq} = -\boldsymbol{p} \cdot \boldsymbol{q} + \boldsymbol{p} \times \boldsymbol{q}$, where the product $\boldsymbol{pq}$ indicates quaternion multiplication of the vectors while $\boldsymbol{p} \cdot \boldsymbol{q}$ and $\boldsymbol{p} \times \boldsymbol{q}$ represent the conventional vector dot and cross products respectively.

Let us represent an arbitrary unit vector quaternion in the vector notation $\hat{\boldsymbol{u}}$, while recalling that the product $\hat{\boldsymbol{u}}\hat{\boldsymbol{u}}$ represents quaternion multiplication. Thus $\hat{\boldsymbol{u}}^2 = \hat{\boldsymbol{u}}\hat{\boldsymbol{u}} = -\hat{\boldsymbol{u}} \cdot \hat{\boldsymbol{u}} = -1$. Clearly then, $\hat{\boldsymbol{u}}^3 = -\hat{\boldsymbol{u}}$, and $\hat{\boldsymbol{u}}^4 = 1$. By definition, $\hat{\boldsymbol{u}}^0 = 1$, and in general, $\hat{\boldsymbol{u}}^n = \hat{\boldsymbol{u}}^{n-4}$, $n > 3$.

### 4.1 Exponentials of quaternions

Let us derive an expression for the exponential of a general quaternion $Q = \chi + \hat{\boldsymbol{u}}\phi$, where $\chi$ and $\phi$ are real numbers, and $\hat{\boldsymbol{u}}$ is a unit vector quaternion.

$$
\begin{aligned}
e^Q &= e^{(\chi + \hat{\boldsymbol{u}}\phi)} = e^\chi \cdot e^{\hat{\boldsymbol{u}}\phi} = e^\chi \sum_{\nu=0}^{\infty} \frac{(\hat{\boldsymbol{u}}\phi)^\nu}{\nu!} \\
&= e^\chi \Big( \frac{\phi^0}{0!} + \frac{\hat{\boldsymbol{u}}\phi^1}{1!} - \frac{\phi^2}{2!} - \frac{\hat{\boldsymbol{u}}\phi^3}{3!} + \frac{\phi^4}{4!} + \frac{\hat{\boldsymbol{u}}\phi^5}{5!} - \cdots \Big) \\
&= e^\chi \left( \begin{array}{c} \dfrac{\phi^0}{0!} - \dfrac{\phi^2}{2!} + \dfrac{\phi^4}{4!} - \cdots + \\ \dfrac{\hat{\boldsymbol{u}}\phi^1}{1!} - \dfrac{\hat{\boldsymbol{u}}\phi^3}{3!} + \dfrac{\hat{\boldsymbol{u}}\phi^5}{5!} - \cdots \end{array} \right) = e^\chi \big( \cos\phi + \hat{\boldsymbol{u}} \sin\phi \big)
\end{aligned}
$$

### 4.2 Logarithms of quaternions

In general, we can write $P = |P|\,U$, where $|P|$ is a magnitude and $U$ is a unit quaternion. Moreover, without loss of generality, we can let $|P| = e^\chi$ and $U = cos(\phi) + \hat{\boldsymbol{u}}sin(\phi)$, where $\chi$ and $\phi$ are real numbers, and $\hat{\boldsymbol{u}}$ is a unit vector. Let us define $Q = log(P)$ to be the inverse of the exponential function $e^Q = P$, whence by inspection,

$$
\begin{aligned}
\log P &= \chi + \hat{\boldsymbol{u}}\phi, \text{ where} \\
\chi &= \log |P|, \\
\hat{\boldsymbol{u}} &= U_v / |U_v|, \\
\phi &= \tan^{-1}(|U_v| / U_r), \text{ and} \\
U_r, U_v &= \text{the real and vector parts of } U = P/|P|.
\end{aligned}
$$

Notice that the expression $\hat{\boldsymbol{u}}\phi$ doubly covers the space $\mathbf{R}^3$, so the convention $\phi \geq 0$ may be imposed without loss of generality. This convention is convenient in using $\phi$ to represent the geometric length of a vector pointing in the $\hat{\boldsymbol{u}}$ direction.

### 4.3 Relevant properties of exponentials and logarithms

Consider if $e^{(P+Q)} \stackrel{?}{=} e^P e^Q = e^{\chi_p}\big(\cos\phi_p + \hat{\boldsymbol{u}}_p \sin\phi_p\big) e^{\chi_q}\big(\cos\phi_q + \hat{\boldsymbol{u}}_q \sin\phi_q\big)$. Expanding the right hand side, one term is $\hat{\boldsymbol{u}}_p \sin\phi_p \hat{\boldsymbol{u}}_q \sin\phi_q$ which is commutative iff $\hat{\boldsymbol{u}}_p \| \hat{\boldsymbol{u}}_q$ or $\sin\phi_p = 0$ or $\sin\phi_q = 0$. But $e^{(P+Q)}$ is unconditionally commutative, so equality is only possible if $\hat{\boldsymbol{u}}_p \parallel \hat{\boldsymbol{u}}_q$ or if $P$ or $Q$ or both are scalar. In general, $e^{(P+Q)} \neq e^P e^Q$. Substituting $P = \log R$ and $Q = \log S$ and taking the log of both sides of the inequality proves in general $\log(RS) \neq \log R + \log S$. Thus one of the most useful properties of logarithms is sacrificed in the extension from complex to quaternion domain.

Nevertheless, perfectly valid transformations to and from logarithmic space can be performed on quaternions. More important, the commutativity of addition in logarithmic space has been shown to be decoupled from the corresponding non-commutative multiplication in linear quaternion space. This property enables us to consider logarithmic quaternion space as an excellent candidate domain for averaging general rotations.

### 4.4 Averaging in 3D logarithmic space

A rotation quaternion $U$ is of unit magnitude, i.e., $(e^\chi = 1) \Rightarrow (\chi = 0)$, therefore its logarithm degenerates to $\log U = \hat{\boldsymbol{u}}\phi$. This is simply a vector of length $\phi$ in the $\hat{\boldsymbol{u}}$ direction in a linear 3-space. Let us call this a *rotation vector*. Adopting the convention that $0 \leq \phi \leq \infty$, i.e., representing all instances of $-(\hat{\boldsymbol{u}}\phi)$ as $(-\hat{\boldsymbol{u}})\phi$, gives us a single valued representation of any rotation vector which is unbounded in wrap number.

We now have a formal mapping between rotation quaternions and rotation vectors†. We have also established a correspondence between rotation quaternions and geometric 3-space rotations. This provides us a geometric interpretation for a rotation vector: $\phi\hat{\boldsymbol{u}}$ represents a rotation by angle $2\phi$ about $\hat{\boldsymbol{u}}$.

Averaging in this space is appealing in that it is a straightforward procedure of linear operations just as it is for scalars and 3D position vectors. Moreover any case which degenerates to rotations in a unique 2D plane becomes isomorphic with well understood averaging methods in the complex domain.

Therefore we have demonstrated an exact, formally consistent method for calculating an average $U_a$ of $N$ rotation quaternions $U_n$:

$$\begin{aligned} \boldsymbol{v}_n &= \log U_n, \\ \boldsymbol{v}_a &= \frac{1}{N}\sum_{n=1}^{N} \boldsymbol{v}_n, \\ U_a &= e^{\boldsymbol{v}_a}. \end{aligned}$$

† Note: the exponential mapping is actually many to one which complicates matters, but which may be addressed with suitable bookkeeping, and will not be discussed here. The same problem exists in the 2D complex domain.

## 5 Averaging 3-space orientations

Let us take an orientation to be a quantification of angular pose in space. This differs from a rotation, which implies an action taking an initial orientation to a rotated one. At first glance, it seems their geometrical interpretations should possess identical formal properties in whatever algebra is used to represent them. But some simple examples should serve to illustrate a fundamental distinction.

### 5.1 A distinction between rotation and orientation in 2D

To represent a 2D plane in polar form, we are obliged to pick a branch line at some particular angle, say $\phi_b = \pi$ radians. Let us now choose two orientations sensibly displaced CW and CCW respectively from $\phi_b$ by some angle $\delta < \pi/2$. As quantified in this space, these orientations are, $O_1 = (\pi - \delta)$, and $O_2 = (-\pi + \delta)$. Also consider the rotations $R_1 = (\pi - \delta)$ and $R_2 = (-\pi + \delta)$ representing actions taking $\phi = 0 \to O_1$ and $0 \to O_2$ respectively.

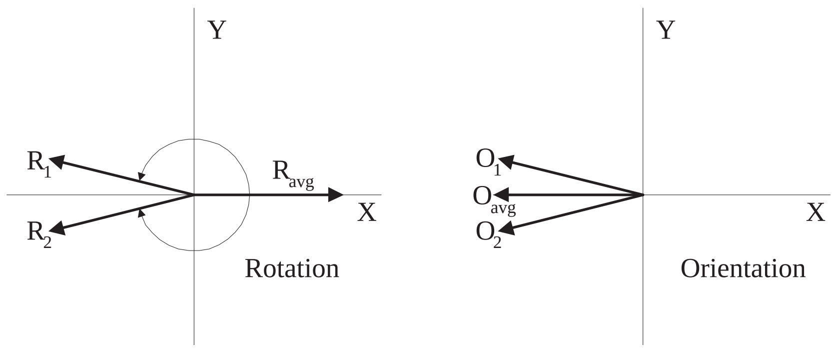


Fig 2: Average orientation differs from average rotation

According to the procedure for averaging rotations, their average is $(R_1 + R_2)/2 = 0$, which makes intuitive sense in terms of the actions represented. But this is not what we had in mind for the average of the orientations. An average of $\pi$ is the sensibly intuitive one we would prefer.

Another example is three orientations of $-2\pi/3,\ 0,\ 2\pi/3$, respectively. The corresponding rotations average to 0, but the three way symmetry in this case gives three equally valid candidate average orientations.

### 5.2 Statement of the problem

Even in the 2D case, averaging orientations incurs complications which arise from existence of branch lines in the geometric representation of orientations. It makes sense to consider rotations with unbounded wrap number, e.g., $\phi$ is distinct from $24\pi + \phi$ which is 12 entire turns more than $\phi$. But both corresponding orientations are sensibly measurable as equal to $\phi$, normally chosen to be in the principal branch. In the above example, this branch is $-\pi < \phi \leq \pi$.

Thus there are two distinct parts of this problem. The first is to find an acceptable definition of what we mean by average orientations in 2D, and to find a procedure whereby we can calculate them. The other is to find a suitable algebraic extension of the 2D definition to the 3D case.

Naturally, we would like to test our hypothetical definitions and extensions against our geometric intuition of what we mean by an average orientation, as we have done for rotations above. Unfortunately, things can get intuitively murky even in 2D, and hopelessly so in 3D except in certain limiting cases.

Therefore it makes better sense to rely on the formal properties of whatever algebraic system we use to calculate orientation averages. We shall attempt to develop a suitable algebraic method having a plausible geometrical interpretation, then to intuitively test the method in certain visualizable limiting cases.

### 5.3 Averaging 2D orientations

In spite of the first example given in section 5.1, it is not unusual for the calculated average rotation to also equal what we expect for the average orientation. Thus we are motivated to find a procedure which degenerates to one formally equivalent to the procedure for averaging rotations in such cases.

Figure 3 shows three distributions of orientations, 1 through 7. The angles in figure 3a have been selected so the average orientation is precisely equal to orientation 4. The distributions in 3b and 3c are successively expanded about orientation 4, like unfolding a paper fan, such that the ratios of angles between orientations are preserved. Thus, the average orientation in all three cases is expected to be precisely equal to orientation 4.

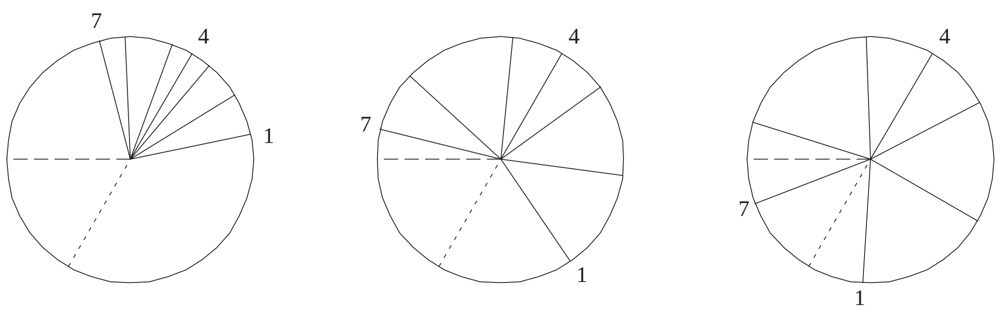


Fig 3a Fig 3b Fig 3c

Now let us look at the averages of the corresponding rotations. In figures 3a and 3b, they are fortuitously equal to rotation 4, but in 3c the average is closer to rotation 3. What's different about 3c?

We have shown the branch line for rotations as a dashed line in figure 3, and have also shown a dotted line directed precisely opposite from the average orientation. Our geometric intuition leads us to interpret orientations CCW from the average up to the dotted line to be positive displacements from the average, and CW ones to be negative. In this sense, the dotted line is the natural branch line for the orientation average.

The disparity between orientation and rotation averages occurs whenever the two sectors delineated by the orientation and rotation branch lines both contain at least one of the averagands. If for example in figure 3c the rotation branch were picked to lie anywhere between orientations 1 and 7, the average rotation would be exactly equal to rotation 4.

Therefore, choosing the correct branch line, and averaging rotations mapped onto the principle branch will consistently yield an average rotation equal to the expected average orientation, as desired.

But if we were given the distribution shown in figure 3c *de novo*, it might be quite difficult to judge by eye what we would expect as the average orientation, so an intuitive guess for the orientation branch line is problematical. Indeed, in the three way symmetric example mentioned in section 5.1, there are three equivalently correct choices.

Thus a mechanism is needed to select the right branch, or branches if more than one exist. Moreover, some measure is needed to quantify the relative quality of the multiple possible candidates.

### 5.3.1 The mean and standard deviation of 2D orientations

Let us pick an arbitrary orientation $\phi$ and define an oppositely directed branch line $\phi_b = (\phi + \pi)$. We wish to consider an ensemble of orientations $(\phi_i,\ i = 1 \ldots n)$ mapped onto the principle branch of $\phi_b$. Let's look at the linear and quadratic functions $L(\phi) = \sum_{i=1}^{n}(\phi - \phi_i)/n$ and $Q(\phi) = \sum_{i=1}^{n}(\phi - \phi_i)^2/n$ respectively, for $(0 \leq \phi < 2\pi)$ radians.

Figure 4 shows graphs of these functions corresponding to the three distributions shown in figure 3. A number of enlightening features are immediately apparent.

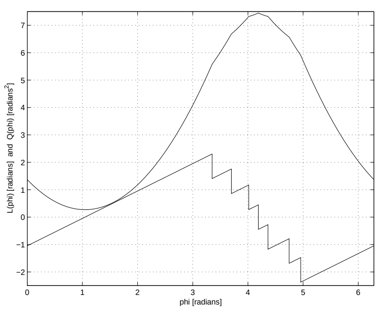


Fig 4a

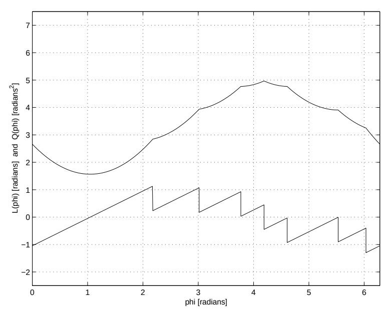


Fig 4b

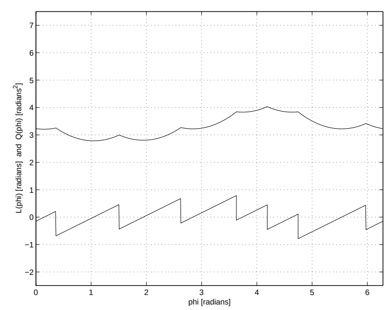


Fig 4c

First, note that $L(\phi)$ is a piecewise linear function, with a constant slope of 1. There are $n$ discontinuities, all of size $-2\pi/n$, occurring at locations $\phi_b = \phi_i$, so that $L(0) = L(2\pi)$. Excluding discontinuities, the zeros $\mu_k$ of $L(\phi)$ by definition satisfy a requirement for being a mean average. Clearly there is just one such solution, at 1.05 radians, in figure 4a and less obviously just one, also at 1.05 radians, in figure 4b. In figure 4c however, there are seven such solutions, one of which is indeed at 1.05 radians. This ambiguity is quite consistent with the difficulty mentioned above of judging which branch line is "correct."

Turning now to the function $Q(\phi)$, it is apparent (and can easily be shown formally) that it has a local minimum occurring at each location $\mu_k$. At each of these locations, $\sigma_k^2 = Q(\mu_k)$ provides a measure of the mean square deviation of the $\phi_i$ from $\mu_k$. It is also thereby a measure of how "good" the $k^{th}$ average is: intuitively, the average $k$ with the smallest $\sigma_k^2$ is the "best" one.

A tabulation of the $\mu_k$ in figure 4c is shown here in order of increasing $\sigma_k^2$. Apparently,

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| $\mu_k$ | 1.05 | 1.95 | 0.15 | 2.84 | 5.54 | 3.74 | 4.64 |
| $\sigma_k^2$ | 2.78 | 2.81 | 3.21 | 3.22 | 3.22 | 3.83 | 3.83 |

the mean at 1.05 is the best, but not by much (and purely by accident). In fact, the spread of all the $\sigma_k^2$ is not large. In contrast to figures 4a and 4b, this case calls into question the entire concept of a unique mean average of orientations. Indeed, we need look no further than the three way symmetric example mentioned in section 5.1 to find a case in which this concept is meaningless.

Therefore the formal definition of a mean average of orientations must allow for a multiple valued result. There are bounds on how many, however. There must be at least one, the proof of which is left to the reader. And since each piecewise linear segment of $L(\phi)$ may contain at most one zero, there must be no more than $n$ of them.
Accordingly, let us define the mean average(s) of $n$ orientations $\phi_i$ in a 2D space as,

$$\mu_k = \frac{1}{n}\sum_{i=1}^{n}\phi_i; \qquad \textbf{(mean)}$$

correspondingly, let us define the variance(s) as

$$\sigma_k^2 = \frac{1}{n}\sum_{i=1}^{n}(\mu_k - \phi_i)^2, \qquad \textbf{(variance)}$$

and the standard deviation(s) as

$$\sigma_k = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\mu_k - \phi_i)^2}, \qquad \textbf{(standard deviation)}$$

where for each $k$, a branch is selected such that for all $i$, $\big(-\pi < (\mu_k - \phi_i) \leq \pi\big)$.

### 5.3.2 Algorithm for averaging 2D orientations

The properties of $L(\phi)$ and $Q(\phi)$ provide a straightforward and efficient way of calculating the mean averages of $n$ orientations $\phi_i$ and their corresponding variances. We start with a list of $n$ angles $\psi_i = (\phi_i + \pi)$ at which the function $L(\psi_i)$ will be discontinuous. From this we calculate the lists $L_i^- = L(\psi_i)$ and $L_i^+ = (L_i^- - 2\pi/n)$, which represent respectively,

$$L_i^- = \lim_{\phi\to\psi_i}^{CCW} L_i(\phi) \qquad \text{and} \qquad L_i^+ = \lim_{\phi\to\psi_i}^{CW} L_i(\phi).$$

The $\psi_i$ delimit $n$ piecewise linear segments $S_i$ of $L(\phi)$. The signs of $L_i^+$ and $L_{(i+1)}^-$ (modulo $n$) differ iff a zero of $L(\phi)$ occurs in $S_i$, a property used to detect the existence of a mean in $S_i$. If one exists, it may be calculated as $\mu_i = (\psi_i - L_i^+)$, or equivalently as $\mu_i = (\psi_{(i+1)} - L_{(i+1)}^-)$. Finally, given the calculated $\mu_i$, the $\sigma_i^2$ are calculated directly from the definition.

### 5.4 Extension to 3D

Let us now address the other part of the problem, that of extending this result to 3 dimensions. We wish to find a method that degenerates to a formal identity with the method for rotations in those cases where the results are the same. We also require that the method be formally identical to the 2D method in that degenerate case.

We begin by observing that in the 2D case, the calculations were carried out in the logarithmic or *rotation imaginary* space, even though the geometric figures were shown in the linear or *rotation complex* space. Accordingly, we shall work in the 3D logarithmic or *rotation vector* space rather than the linear or *rotation quaternion* space.

Following the 2D development, we limit the range of the 3D orientation averagands to the principal branch in rotation quaternion space. In rotation vector space, this maps to the sphere $|\boldsymbol{r}| \leq \pi$, the surface of which represents an angular magnitude of $\pi$ radians of rotation. With the exclusion of one hemispherical surface, which we shall ignore in this development, this space which we shall call *orientation vectors*, represents all possible 3D orientations.

#### 5.4.1 Branches in orientation vector space

While rotation/orientation imaginaries *represent* rotations/orientations in a complex 2-space, they are themselves embedded in a 1D linear space. Rotation/orientation vectors represent rotations/orientations in 3-space, while also themselves being 3-space entities. Thus, an orientation imaginary is graphically represented by a line of length $\phi$ starting from an origin $\phi_0$ while an orientation vector is analogously represented by $\boldsymbol{v}$, a line of length $\phi$ in a 3-space direction $\hat{\boldsymbol{u}}$, starting from an origin $\boldsymbol{v}_0$.

For clarity of illustration, figures 5 and 6 show 2D cross sections $z = 0$ of examples in 3-space in which, again for clarity, all *points* are constrained to lie in the plane $z = 0$.

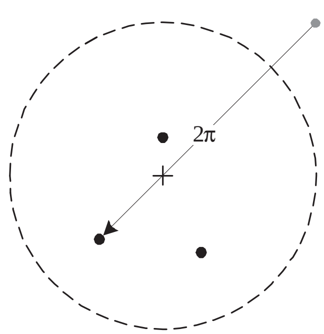


Fig 5a

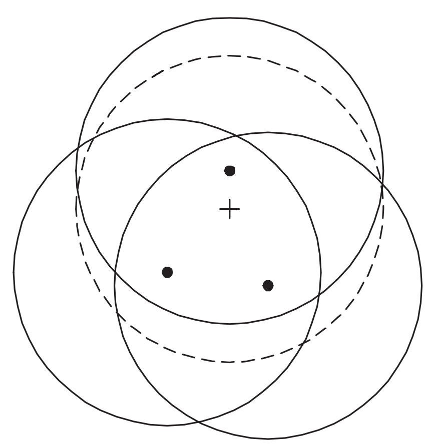

Fig 5b

The averagands are first all mapped into the principal branch about an origin $\phi_0 = 0$. This origin is indicated by a "+", and the branch line by a dashed-line circle (representing actually a spherical branch surface in 3-space). Figure 5a shows shows three averagands as •s, one of which (in gray) must be displaced by a distance $2\pi$ in the direction of the origin to map it into the principal branch.

Extending the idea of the $\psi_i$ introduced in section 5.3.2, the orientation vector space is partitioned into piecewise linear volumes separated by the branch surfaces about each of the averagands. These surfaces are indicated by the solid line circles of radius $\pi$ about their respective $\bullet$s in figure 5b.

### 5.4.2 Finding the mean averages

The $L(\phi)$ of section 5.3.2 is now extended to $L(\boldsymbol{v})$ which is discontinuous across these surfaces, but continuous within the volumes partitioned thereby. The task now becomes to detect which volumes contain a zero, and for those that do, to locate the zero in the volume. Alternatively, one could look for minima in the extended function $Q(\boldsymbol{v})$.

Allowable mean averages must themselves be orientations, so it is not useful to sample any volume outside the original principal branch. This ensures that no averagand can be further from a mean than the diameter $2\pi$ of the original principal branch. This also ensures that all averagands must be within the principal and second branches of a possible mean, i.e. within a sphere of radius $2\pi$.

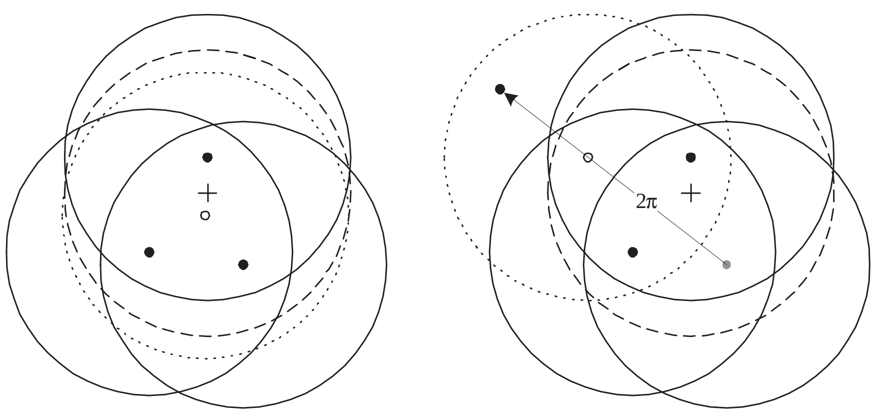


Fig 6a  Fig 6b

Figure 6 shows two sample points for evaluating $L(\boldsymbol{v})$. The point is indicated by a $\circ$ with its branch surface indicated by a dotted line circle. In figure 6a, the sample point cohabits the same volume partition as all three averagands, which is another way of saying all three averagands are within its principal branch. Moreover, it is located near a zero in this partition, as can be seen by casual inspection. In figure 6b, one of the averagands is in a different partition, separated by the branch surface of the sample point, which is another way of saying that this averagand must be remapped to lie in the sample point's principal branch before $L(\boldsymbol{v})$ can be evaluated. This sample point also lies near a zero in its partition, as is apparent by inspection.

### 5.4.3 Algorithm for averaging 3D orientations

*I believe there is a simple method, but have not thought the problem through yet. I encourage the reader to address this problem. Please let me know if you find a useful result!*

## 6 Checks and comparisons

This section will make checks of calculated mean orientations against intuitively tractable limiting cases. *This awaits an algorithm so the cases can be evaluated.*

It also will compare this method with other proposed methods. It will attempt to show some formal equivalence, even if only in the limiting domains for which the other methods are claimed to work. *This awaits my time to get to it!*

## 7 Conclusion

It is entirely possible this work duplicates results reported in some branch of literature of which I am unaware. However it appears to me to have made several new contributions.

First, I have shown how to extend the averaging of 2D rotations to 3D. This was an extension of averaging procedures for rotation complexes to work for rotation quaternions through their respective logarithms.

Second is the insight of how branches must be handled in order to get intuitively acceptable 2D orientation averages, and the realization that the notion of uniqueness is lost in the domain of orientation averages.

Finally, the extension of the 2D orientation average concept to 3D which seems formally plausible. It would be stronger and far more useful if an algorithm were also developed, but that must presently be considered future work.

## 8 References


[1] Vicci, Leandra, *Quaternions and Rotations in 3-Space: The Algebra and its Geometric Interpretation,* TR01-014, Department of Computer Science, UNC, Chapel Hill, NC, 2001.

[2] Weber, Hans, *Averaging Quaternions/Orientations,* http://www.cs.unc.edu/~vicci/quats/weberh/averaging_quaternions.htm, 12 November 1998.

[3] Curtis, W. Dan, Adam L. Janin, Karel Zikan, *A Note on Averaging Rotations,* SIGGRAPH '95 course notes, #9; 0-7803-1363-1/93 ©1993 IEEE.

[4] Horn, Berthold K. P., *Closed-form solution of absolute orientation using unit quaternions,* J. Opt. Soc. Am. A, vol4 no4, April 1987.

[5] Zhuang, Hanqi, Zvi S. Roth, R. Sudhakar, *Practical Fusion Algorithms for Rotation Matrices: A Comparative Study,* Journal of Robotic Systems 9(7), 1992.