Mathematics 308—Fall 1996

Motion in three dimensions

In this section we begin to think about how to draw objects in three dimensions. There are no built-in routines for this purpose in PostScript, and we shall have to start from scratch. There are a number of somewhat technical, but nonetheless interesting, things to be dealt with which do not occur in two dimensions.

The overall scheme we shall adopt is the following: we are given a three-dimensional object of some kind, such as a cube or a sphere, possibly with something drawn on its surface. We shall imagine this object in what could be called an ideal position, where it is initially placed. In order to draw it, one way of proceeding might be to allow ourselves to move around to see it from various positions, and then imagine it as being drawn on a surface between us and the object, the way a painter classically renders things. But instead of this, for various reasons, I shall do something complementary: we shall imagine that we are situated at a fixed point in space—in fact at the origin—and that it is the object which moves instead of us. As for the problem of rendering the $3D$ object in $2D$, we shall do this in two different ways, **orthogonal projection** and **perspective**.

We must therefore attack two main problems: (1) how to describe the possible ways in which an object can move in three dimensions, and (2) how to render a three-dimensional object in two dimensions, which is in the end the dimension of the surface we draw on. If we were to try to draw only linear objects the second problem would be very simple, but in fact we shall eventually see how to use the routine `mkpath` to draw $3D$ parametrized paths, which leads to a technique which is at once more complicated and more rewarding.

In this part of the notes we attack the first of these problems. I begin by recalling some vector algebra.

1. Dot and cross products

In any number of dimensions we define the dot product of two vectors

$$u = (x_1, x_2 \ldots, x_n), \quad v = (y_1, y_2, \ldots, y_n)$$

to be

$$u \bullet v = x_1 y_1 + x_2 y_2 + \cdots x_n y_n \ .$$

The relation between dot products and geometry is expressed by the cosine rule for triangles, which asserts that if θ is the angle between u and v then

$$\cos \theta = \frac{u \bullet v}{\|u\| \, \|v\|} \ .$$

In particular u and v are perpendicular when $u \bullet v = 0$.

In $3D$ there is another kind of product. If

$$u = (x_1, x_2, x_3), \quad v = (y_1, y_2, y_3)$$

then their **cross product** $u \times v$ is the vector

$$(x_2 y_3 - y_2 x_3, x_3 y_1 - x_1 y_3, x_1 y_2 - x_2 y_1) \ .$$

This formula can be remembered if we write the vectors u and v in a 2×3 matrix

$$\begin{bmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{bmatrix}$$

and then for each column of this matrix calculate the determinant of the 2×2 matrix we get by crossing out in turn each of the columns. The only tricky part is that *with the middle coefficient we must reverse sign.* Thus

$$u \times v = \left(\begin{vmatrix} x_2 & x_3 \\ y_2 & y_3 \end{vmatrix}, - \begin{vmatrix} x_1 & x_3 \\ y_1 & y_3 \end{vmatrix}, \begin{vmatrix} x_1 & x_2 \\ y_1 & y_2 \end{vmatrix}, \right) .$$

Here

$$\begin{vmatrix} a & b \\ & c & d \end{vmatrix} = ad - bc .$$

The geometrical significance of the cross product is contained in these rules:

- *The length of $w = u \times v$ is the area of the parallelogram spanned in space by u and v.*

- *It lies in the line perpendicular to the plane containing u and v and its direction is determined by the right hand rule.*
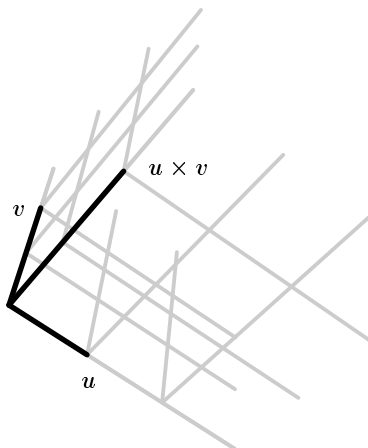
The cross product $u \times v$ will vanish only when u and v are multiples of one another.

In these notes, the main use of dot products and cross products will be in calculating various projections and related things.

(1) The first component of frameworks deals with **projections**. Suppose α to be any vector in space and u some other vector in space. The **projection of** u **along** α is the vector u_0 we get by projecting u perpendicularly onto the line through α.
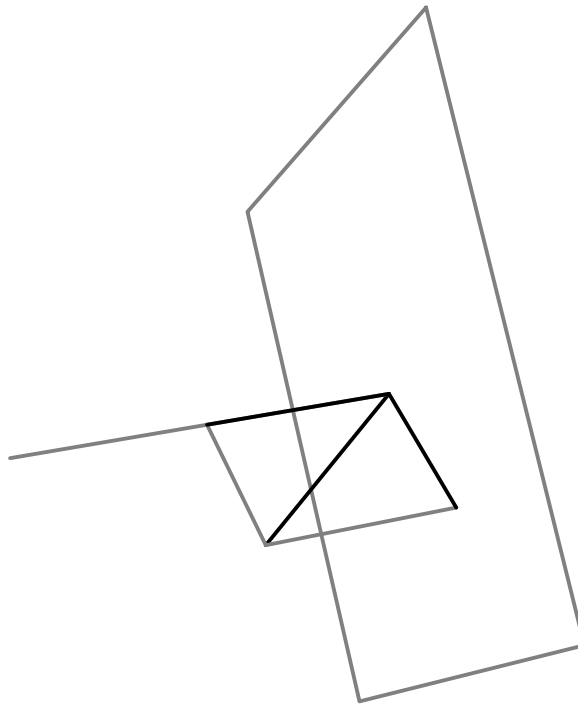
What is this projection? It must be a multiple of α. We can figure out what multiple by using trigonometry. We know three facts: (a) The angle θ between α and u is determined by the formula

$$\cos\theta = \frac{u \bullet \alpha}{\|\alpha\| \, \|u\|} \; .$$

(b) The length of the vector u_0 is $\|u\| \cos\theta$, and this is to be interpreted algebraically in the sense that if u_0 faces in the direction opposite to α this is negative. (c) Its direction is either the same or opposite to α. The vector $\alpha/\|\alpha\|$ is a vector of unit length pointing in the same direction as α. Therefore

$$u_0 = \|u\| \cos\theta \, \frac{\alpha}{\|\alpha\|} = \|u\| \, \frac{u \bullet \alpha}{\|\alpha\| \, \|u\|} \, \frac{\alpha}{\|\alpha\|} = \left(\frac{u \bullet \alpha}{\|\alpha\|^2}\right) \alpha = \left(\frac{u \bullet \alpha}{\alpha \bullet \alpha}\right) \alpha \; .$$

(2) Now let Π be the plane passing through the origin perpendicular to α, and let $u_\perp$ be the projection of u onto Π.
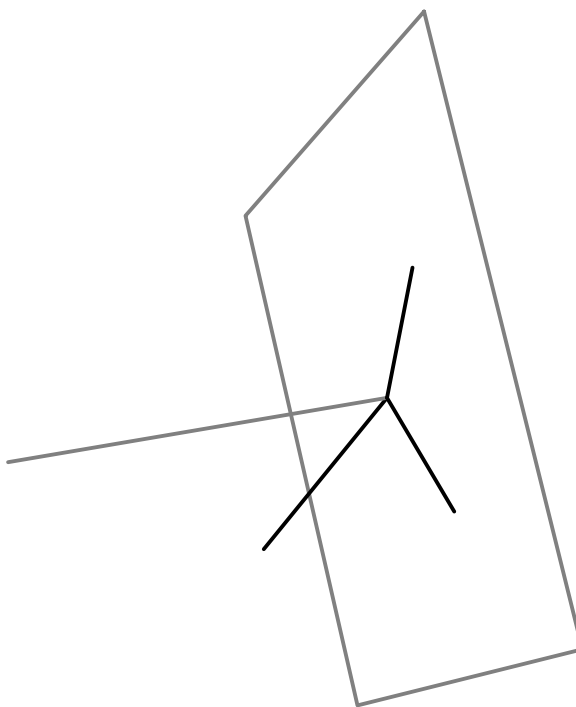
The vector u has the orthogonal decomposition

$$u = u_0 + u_\perp$$

and therefore we can calculate

$$u_\perp = u - u_0 \ .$$

(3) Finally, let u_* be the vector in Π we get by rotating u_* by $90°$ in Π, using the right hand rule to determine what direction of rotation is positive.

How do we calculate u_*? We want it to be perpendicular to both α and $u_\perp$, so it ought to be related to the cross product $\alpha \times u_\perp$. A little thought should convince you that in fact the direction of u_* will be the same as that of $\alpha \times u_\perp$, so that u_* will be a positive multiple of $\alpha \times u_\perp$. We want u_* to have the same length as $u_\perp$. Since α and $u_\perp$ are perpendicular to each other, the length of the cross product is equal to the product of the lengths of α and $u_\perp$, and we must divide by $\|\alpha\|$ to get a vector of length $\|u_\perp\|$. Therefore, all in all

$$u_* = \frac{\alpha}{\|\alpha\|} \times u_\perp \ .$$

Incidentally, in all of this discussion it is only the direction of α that plays a role. It is often useful to normalize α right at the beginning of these calculations, that is to say replace α by $\alpha/\|\alpha\|$.

Exercise 1.1. Write PostScript programs to calculate dot products, cross products, u_0, $u_\perp$, u_*.

2. Rigid transformations and rigid motions

Suppose (x, y, z) to be the coordinates of a point on the original object. Let (x_*, y_*, z_*) be the coordinates of the point after the object is moved. We shall assume that the object is **rigid**, which means that the distance between points doesn't change when they are transformed. In other words, we are looking at a transformation taking

$P = (x, y, z)$ to $P_* = (x_*, y_*, z_*)$, and it satisfies this condition: if P is moved (or transformed) to P_* and Q to Q_* then the distance between P_* and Q_* is the same as the distance between P and Q. More generally, any transformation of points with this property, in any number of dimensions, is called a **rigid transformation**.

There are a few very general properties of rigid transformations which are useful to know:
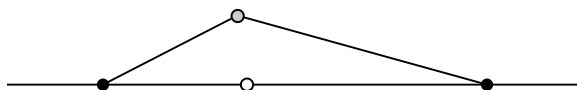
- *A rigid transformation must take line segments into line segments, and lines into lines.*

The segment PQ is the shortest path from P to Q. What this means is that the points R on the segment between P and Q can be characterized as the points for which

$$d(P, Q) = d(P, R) + d(R, Q)$$

whereas in contrast if R does not lie on the segment PQ then

$$d(P, Q) < d(P, R) + d(R, Q) \ .$$

Thus if we are given a set of three points P, Q, R with R on the segment between P and Q, we have

$$d(P, Q) = d(P, R) + d(R, Q) \ .$$

If we apply a rigid transformation T we get

$$d(T(P), T(Q)) = d(T(P), T(R)) + d(T(R), T(Q))$$

which means that $T(R)$ lies on the segment between $T(P)$ and $T(Q)$.

A similar characterization can be given for the points on the line beyond P or beyond Q.

Exercise 2.1. Give such a characterization.

The same argument shows that $T(R)$ is the unique point on the segment PQ which lies at distance PR from $T(P)$.

The points on the line through P and Q are those which can be written $(1 - t)P + tQ$ as t ranges over all of $(-\infty, \infty)$. We can extend the observation made a moment ago to deduce

- *If $R = (1 - t)P + tQ$ then $T(R) = (1 - t)T(P) + tT(Q)$.*

In other words, *if we know what T does to the points P and Q then we know what T does to all points on that line.*

This suggests what we shall later see to be true—that any rigid transformation must be an affine transformation, which is to say it has the form

$$\begin{bmatrix} x_* & y_* & z_* \end{bmatrix} = \begin{bmatrix} x & y & z \end{bmatrix} A + \begin{bmatrix} \tau_x & \tau_y & \tau_z \end{bmatrix}$$

where A is a 3×3 matrix. What the result just stated says is that this is true when a rigid transformation is restricted to $1D$.

- *A rigid transformation must preserve angles as well as distances.*

The angle between two segments PQ and PR is determined by the lengths PQ, PR, QR, according to the cosine law:

$$\cos \theta = \frac{\overline{QR}^2 - \overline{PQ}^2 - \overline{PR}^2}{2\,\overline{PQ}\,\overline{PR}} \ .$$

Apply T to this equation to see that the angle between $T(P)T(Q)$ and $T(P)T(R)$ must be the same as that between PQ and PR.

- *The composition of two rigid transformations (one performed after the other) is a rigid transformation.*

- *If a rigid transformation has an inverse, then that inverse is a rigid transformation.*

The both follow straight from the definition. Incidentally, we shall see later that in actuality every rigid transformation has an inverse, but that is not immediately clear.

The definition of rigid transformations makes sense in any number of dimensions. They are very simple to describe completely in $1D$. Let P and Q be any two points on the line. Let T be a rigid transformation. Given $T(P)$, what can we say about $T(Q)$? The distance between $T(Q)$ and $T(P)$ must be the same as that between P and Q, so there are two possibilities, one on either side of $T(P)$. In one case the direction from $T(P)$ to $T(Q)$ is the same as the direction from P to Q, and in the other it is opposite. In the first case orientation is said to be preserved, in the second reversed. We have seen already that if we know what a rigid transformation does to two points on a line then we know what it does to all of them.
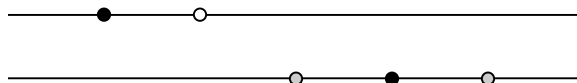
These remarks put together tell us:

- *In $1D$ any rigid transformation is of the form*

$$x \mapsto \epsilon x + \tau$$

where $\epsilon = \pm 1$.

Exercise 2.2. Suppose that $P = 0$, $Q = 1$. If we are given $P_* = x_0$ and $Q_* = x_1$ on the one-dimensional line with $|x_1 - x_0| = 1$ there exists a unique rigid transformation of the line taking P, Q to P_*, Q_*. Find a formula for ϵ and τ. (Hint: as a first step try a few explicit examples.)

In $2D$ and $3D$ things are a bit more complicated. Instead of trying to deal with things geometrically right off, I shall introduce some matrix algebra first. I call a rigid transformation affine if it is given by a formula

$$v \mapsto Av + \tau$$

where A is a square matrix of the proper size (2×2 in $2D$, 3×3 in $3D$).

As I have already mentioned, every rigid transformation turns out to be affine, but we shall postpone that problem for a while and instead restrict our attention to the affine ones. The first question we must settle is this: *Under what circumstances is an affine transformation*

$$v \mapsto Av + \tau$$

a rigid transformation?

If we are to believe that every rigid transformation is affine, then we can classify them according to whether the determinant of A is positive or negative. In the first case I shall call the transformation a rigid **motion**. The idea is

that in any continuous motion of an object from rest it cannot suddenly reverse its orientation, so that the motions are what we expect to see.

Any translation is rigid. If we follow the transformation above by a translation by $-\tau$, we get the transformation

$$v \mapsto Av$$

which according to the composition principle will be rigid precisely when the original is. There is in fact no condition on τ, and it is only on the matrix A that it is necessary to impose one. So we now ask *under what circumstances is a linear transformation (in either $2D$ or $3D$)*

$$v \mapsto Av$$

a rigid transformation?

I begin by looking at things geometrically. Work in $2D$ for convenience–mostly because I can use pictures in $2D$ to illustrate how things are going. Define a **unit frame** to be a set of three points P_0, P_1, P_2 that fit into the following picture:

$$\bullet P_2$$
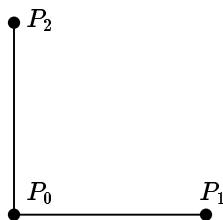
$$P_0 \qquad P_1$$

The distance between P_0 and P_1, as well as between P_0 and P_2, is required to be 1. The angle between the segments $P_0 P_1$ and $P_0 P_2$ is required to be $90°$.

The **standard unit frame** is made up of $(0,0)$, $(1,0)$, and $(0,1)$. Because a rigid transformation preserves distances and angles and takes line segments to line segments:

- *A rigid transformation takes any unit frame into another unit frame.*

The definition of a unit frame can be expressed in terms of vectors. Let P_0, P_1, P_2 be any three points, let v_1 be the vectors from P_0 to P_1, v_2 that from P_0 to P_2. The three points make up a unit frame precisely when v_1 and v_2 each have unit length, and the dot product $v_1 \bullet v_2$ is equal to 0. Another way of expressing this: make up a matrix A whose columns are v_1 and v_2. If $^t A\, A$ is the transpose of A then its rows are the columns of A and vice-versa. The product of $^t A$ and A has the expression

$$^t A\, A = \begin{bmatrix} v_1 \bullet v_1 & v_1 \bullet v_2 \\ v_2 \bullet v_1 & v_2 \bullet v_2 \end{bmatrix} \; .$$

The condition for a unit frame therefore becomes

$$^t A\, A = I \, .$$

in which case the matrix A is said to be **orthogonal**.

One property of such matrices is simple. If we take determinants of both sides we get

$$\det(^t A\, A) = \det(^t A) \det(A) = \det(A) \det(A) = \det(A)^2 = \det(I) = 1$$

so that

$$\det(A) = \pm 1$$

if A is orthogonal.

In $3D$ also a 3×3 matrix is called orthogonal if $^tA\,A = I$.

What we have said so far amounts to:

- *If T is an affine rigid transformation of the form*

$$v \mapsto Av + \tau$$

    then A is orthogonal.

I now claim that the condition that A be orthogonal is sufficient as well, or in other words the converse holds:

- *If A is orthogonal then*

$$v \mapsto Av$$

    is a rigid linear transformation.

This is simple algebra. Let P and Q be any two points, and let u and v be vectors whose ends are P and Q. Let A be an orthogonal transformation, $T\colon v \mapsto Av$. We want to show that the distance from $T(P)$ to $T(Q)$ is the same as that from P to Q. The distance from P to Q is $\|v - u\|$, and that from $T(P)$ to $T(Q)$ is $\|Av - Au\|$. Since A is linear

$$\|Av - Au\| = \|A(v - u)\|$$

so that what we must show is that if w is any vector then $\|Aw\| = \|w\|$. Let

$$w = c_1 e_1 + c_2 e_2, \quad Aw = c_1 v_1 + c_2 v_2$$

where $e_1 = (1, 0)$ and $e_2 = (0, 1)$, then

$$\|Aw\|^2 = (c_1 v_1 + c_2 v_2) \bullet (c_1 v_1 + c_2 v_2) = c_1^2(v_1 \bullet v_1) + 2c_1 c_2(v_1 \bullet v_2) + c_2^2(v_2 \bullet v_2) = c_1^2 + c_2^2 = \|Au\|^2$$

because of simple algebraic properties of the dot product.

In $3D$ the calculation is a bit more complicated, but otherwise very similar.

We can now see why the following fact in $2D$ is true:

- *Given any unit frame, there exists at least one rigid transformation transforming the standard unit frame into it.*

3. Orthogonal transformations in two dimensions

We now have an algebraic characterization of rigid transformations in any number of dimensions in terms of unit frames.

In two dimensions, we can in fact describe all rigid linear transformations (no translation component) geometrically. Start with the particular triple of points $P_0 = (0, 0)$, $P_1 = (1, 0)$, $P_2 = (0, 1)$. Suppose we are given a rigid linear transformation T in $2D$, and let $Q_i = T(P_i)$. Since the transformation is linear, $Q_0 = (0, 0)$. The point Q_1 must lie at distance 1 from it, that is to say at any point on the unit circle centred at the origin. If we fix Q_1 then we have exactly two choices for Q_2, on opposite sides of the unit circle around Q_0 on the diameter perpendicular to the radius $Q_0 Q_1$.
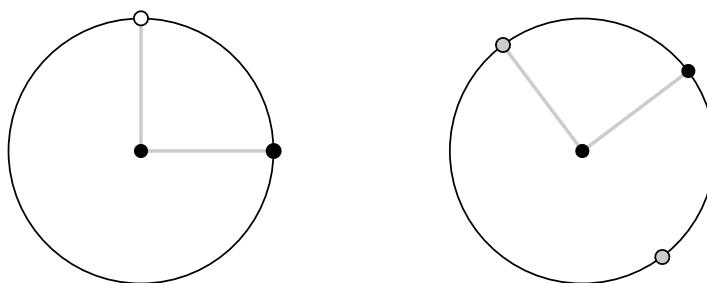
If we choose Q_2 so that it is obtained from Q_1 by a rotation of $90°$, then the whole transformation must be a rotation through the angle between $P_0 P_1$ and $Q_0 Q_1$. This is the case where the orientation of frames is preserved, and the matrix A has determinant 1.

In the other case, the determinant is -1. Our reasoning and our final description are a bit more complicated. A **reflection** in the plane is a transformation mirror-imaging points through some line in the plane. In other words, it fixes points on the mirror-line and reverses points on the line perpendicular to that one. A more mathematical description is that it is a scale change by factors 1 and -1. I claim that *any rigid linear transformation with determinant -1 is a reflection*.

Let A be an orthogonal linear transformation in $2D$ with determinant -1. Its characteristic equation will be

$$\lambda^2 - t\lambda - 1 = 0$$

where t is the trace of A. Its roots are

$$\frac{t}{2} \pm \frac{\sqrt{t^2 + 4}}{2}$$

and are in particular must be real.

We now make the following remark: *If λ is a real eigenvalue of any orthogonal matrix then it is either 1 or -1.* This is true because if $Av = \lambda v$ then $\|v\| = \|\lambda v\|$, so $|\lambda| = 1$.

The determinant of A is the product of its eigenvalues, which in our case is assumed to be -1. Therefore one is 1 and the other -1. The transformation associated to A is reflection in the line of eigenvectors for the eigenvalue 1.

Exercise 3.1. Find a formula for the reflection of v in the line along u.

Exercise 3.2. Find the matrix associated to reflection in the line through the origin which makes an angle of θ with respect to the x axis. (Hint: do this directly by geometry.)

4. Rigid motions in three dimensions

To each affine rigid transformation is associated its linear component, determined by a matrix. It is called a **rigid motion** if the determinant of this matrix is 1, rather than -1. In $2D$ a rigid linear motion is a rotation, while a rigid linear transformation which is not a motion is a reflection. In $3D$ things are more complicated.

Geometrically, there is one simple way to obtain rigid motions in $3D$. We choose an axis in space, and rotate points around it through some fixed angle.
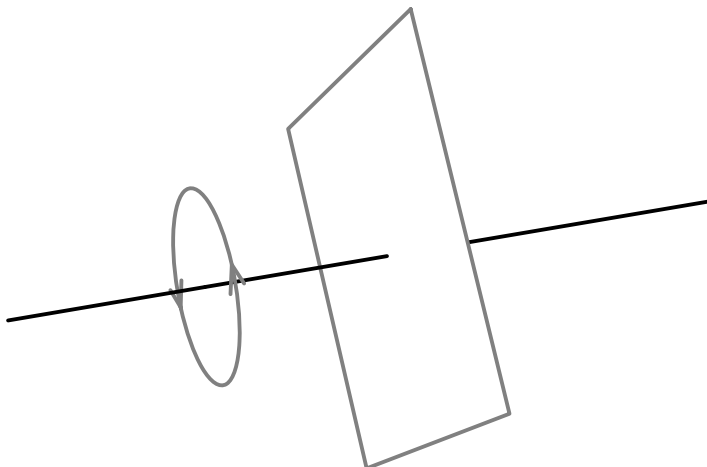
I call such a linear transformation an **axial rotation**. We shall see a bit later how to describe its effect in terms of coordinates.

- *Every rigid linear motion in $3D$ is an axial rotation.*

This is not quite an intuitive fact. For example, if we combine two rotations around distinct axes it is not at all simple to see why this composition is also a rotation.

The proof of this claim is similar to that where we saw that every rigid linear transformation in $2D$ is a reflection. Let A be an orthogonal matrix with determinant 1.

Step (1). If A were an axial rotation, then it would fix the vectors on its axis. Our first step will be to see that there at least exist non-zero vectors fixed by A. A vector fixed by A satisfies the condition

$$Av = v$$

which means that v is an eigenvector of A with eigenvalue 1. Therefore we must show that A has 1 as an eigenvalue. The characteristic polynomial of A has degree 3, and has an expression

$$\det(A - \lambda I) = -\lambda^3 + \cdots$$

If λ is large and negative, this expression will be positive, while if λ is large and positive this expression will be negative. There must exist some place in between where its graph crosses the x-axis, so in other words A has at least one real eigenvalue λ. If

$$Av = \lambda v$$

then

$$\|Av\| = \|v\| = |\lambda| \|v\|$$

so that necessarily $|\lambda| = 1$. If $\lambda = 1$ we are done in this step. Otherwise, suppose $\lambda = -1$, so we have a vector $v \neq 0$ with

$$Av = -v .$$

Then A takes the whole line through v into itself, although reversing directions on that line. Since A preserves angles, it must take points on the plane through the origin perpendicular to v into itself. It acts on this plane as a $2D$ rigid transformation. On that plane it must have determinant 1 or -1. If it had determinant 1 there then it would preserve orientation on that plane, but then since it reverses v it would have to reverse orientation in space, which is a contradiction. Therefore on the plane it must reverse orientation, in which case we know that it is a reflection there. Hence there exists a non-zero vector in that plane fixed by the transformation.

Step (2). We now have $v \neq 0$ fixed by A. The transformation acts, as we have just discussed, on the plane perpendicular to v, where (by the same reasoning as above) it must preserve orientation. Hence on that plane it is a rotation. We are through.

5. How axial rotations act

To begin this section, I remark that to determine an axial rotation we must specify not only an axis but a direction on that axis. This is because the sign of a rotation in $3D$ is only determined if we know whether it is assigned by a left hand or right hand rule. At any rate if choosing a vector along an axis fixes a direction on it. Given a direction on an axis we shall adopt the convention that the direction of positive rotation follows the right hand rule.

So now the question we want to answer is this: *Given a vector $\alpha \neq 0$ and an angle θ. If u is any vector in space and we rotate u around the axis through α by θ, what new point v do we get?* This is one of the main calculations we will make to draw moved or moving objects in $3D$.

There are some cases which are simple. If u lies on the axis, it is fixed by the rotation. If it lies on the plane perpendicular to α it is rotated by θ in that plane (with the direction of positive rotation determined by the right hand rule).

If u is an arbitrary vector, we express it as a sum of two vectors, one along the axis and one perpendicular to it, and then use linearity to find the effect of the rotation on it.

To be precise, let R be the rotation we are considering. Given u we can find its **projection** onto the axis along α to be

$$u_0 = \left(\frac{\alpha \bullet u}{\alpha \bullet \alpha} \right) \alpha$$

Its projection $u_\perp$ is then $u - u_0$. We write

$$u = u_0 + u_\perp$$
$$Ru = Ru_0 + Ru_\perp$$
$$= u_0 + Ru_\perp .$$

How can we find $Ru_\perp$?

Normalize α so $\|\alpha\| = 1$, in effect replacing α by $\alpha/\|\alpha\|$. This normalized vector has the same direction and axis as α. The vector $u_* = \alpha \times u_\perp$ will then be perpendicular to both α and to $u_\perp$ and will have the same length as $u_\perp$. The plane perpendicular to α is spanned by $u_\perp$ and u_*, which are perpendicular to each other and have the same length. The following picture shows what we are looking at from on top of α.
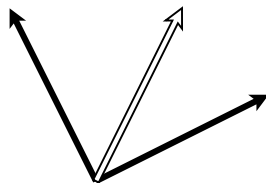
It shows:

- *The rotation by θ takes $u_\perp$ to*

$$Ru_\perp = (\cos\theta)\, u_\perp + (\sin\theta)\, u_* \;.$$

In summary:

(1) Normalize α, replacing α by $\alpha/\|\alpha\|$.

(2) Calculate

$$u_0 = \left(\frac{\alpha \bullet u}{\alpha \bullet \alpha}\right) \alpha \;.$$

(3) Calculate

$$u_\perp = u - u_0 \;.$$

(4) Calculate

$$u_* = \alpha \times u_\perp \;.$$

(5) Finally set

$$Ru = u_0 + (\cos\theta)\, u_\perp + (\sin\theta)\, u_* \;.$$

Exercise 5.1. What do we get if we rotate the vector $(1, 0, 0)$ around the axis through $(1, 1, 0)$ by $36°$?

Exercise 5.2. Write a PostScript procedure with α and θ as arguments and returns the matrix associated to rotation by θ around α.

6. Finding the axis and angle

If we are given a matrix R which we calculate to be orthogonal and with determinant 1, how do we find its axis and rotation angle? (1) *How do we find its axis?* If e_i is the i-th standard basis vector (one of $\mathbf{i}$, $\mathbf{j}$, or $\mathbf{k}$) the i-th column of R is Re_i. Now for any vector u the difference $Ru - u$ is perpendicular to the rotation axis. Therefore we can find the axis by calculating a cross product $(Re_i - e_i) \times (Re_j - e_j)$ for one of the three possible distinct pairs from the set of indices 1, 2, 3–unless it happens that this cross-product vanishes. Usually all three of these cross products will be non-zero vectors on the rotation axis, but in exceptional circumstances it can happen that one or more will vanish. It can even happen that all three vanish! But this only when A is the identity matrix, in which case we are dealing with the trivial rotation, whose axis isn't well defined anyway.

At any rate, any of the three which is not zero will tell us what the axis is.

(2) *How do we find the rotation angle?*

As a result of part (1), we have a vector α on the rotation axis. Normalize it to have length 1. Choose one of the e_i so that α is not a multiple of e_i. Let $u = e_i$. Then Ru is the i-th column of R.

Find the projection u_0 of u along α, set $u_\perp = u - u_0$. Calculate $Ru_\perp = Ru - u_0$. Next calculate

$$u_* = \alpha \times u_\perp \ .$$

and let θ be the angle between $u_\perp$ and $Ru_\perp$. The rotation angle is θ if $Ru_\perp \bullet u_* >= 0$ otherwise $-\theta$.

Exercise 6.1. If

$$R = \begin{bmatrix} 0.899318 & -0.425548 & 0.100682 \\ 0.425548 & 0.798635 & -0.425548 \\ 0.100682 & 0.425548 & 0.899318 \end{bmatrix}$$

find the axis and angle.

7. Euler's construction

The fact that every orthogonal matrix with determinant 1 is an axial rotation may seem quite reasonable, after some thought about what else might such a linear transformation be, but I claim that it is not quite intuitive. To demonstrate this, let me point out that it implies that the combination of rotations around distinct axes is again a rotation. This is not at all obvious, and in particular it is difficult to see what the axis of the combination should be. This axis was constructed geometrically by Euler.
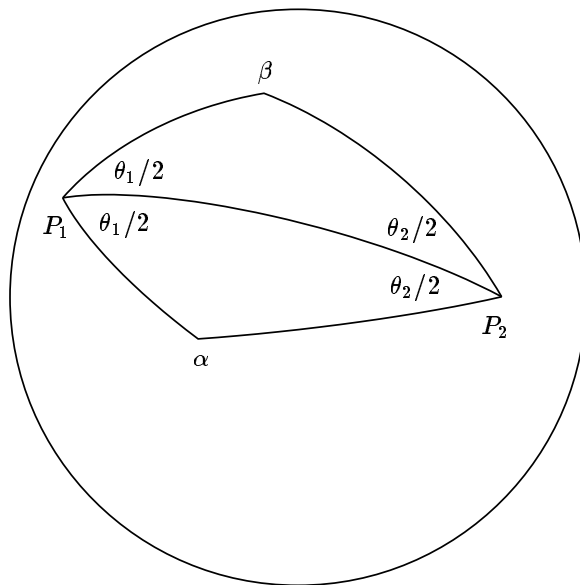
Let P_1 and P_2 be points on the unit sphere. Suppose P_1 is on the axis of a rotation of angle θ_1, P_2 that of a rotation of angle θ_2. Draw the spherical arc from P_1 to P_2. On either side of this arc, at P_1 draw arcs making an angle of $\theta_1/2$ and at P_2 draw arcs making an angle of $\theta_2/2$. Let these side arcs intersect at α and β on the unit sphere. The the rotation R_1 around P_1 rotates α to β, and the rotation R_2 around P_2 moves β back to α. Therefore α is fixed by the composition $R_2 R_1$, and must be on its axis.

Exercise 7.1. What is the axis of $R_1 R_2$? Prove geometrically that generally

$$R_1 R_2 \neq R_2 R_1 \ .$$