CHAPTER 12

Transformations in 3D

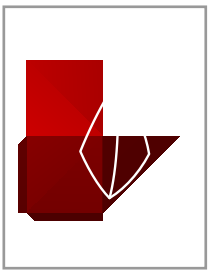

There are no built-in routines for 3D drawing in PostScript. For this purpose we shall have to use a library of PostScript procedures designed especially for the task, as an extension to basic PostScript. In this chapter we shall look at some of the mathematics behind such a library, which is much more complicated than that required in 2D.

We shall examine principally **rigid transformations**, those which affect an object without distorting it, but then at the end look at something related to shadow projections. The importance of rigid transformations is especially great because in order to build an illusion of 3D through 2D images the illusion of motion helps a lot. One point is that motion leads to an illusion of depth through size change, and another is that it allows one to see an object from different sides. The motions used will be mostly rotations and translations, those which occur most commonly in the real world.

There are several reasons why it is a good idea to examine such transformations in dimensions one and two as well as three.

12.1. Rigid transformations

If we move an object around normally, it will not visibly distort—that is to say, to all appearances it will remain rigid. The points of the object themselves will move around, but the relative distances between points of the object will not change. We can formulate this notion more precisely. Suppose we move an object from one position to another. In this process, a point P will be moved to another point P_*. We shall say that the points of the object are **transformed** into other points. A transformation is said to be **rigid** if it preserves relative distances—that is to say, if P and Q are transformed to P_* and Q_* then the distance from P_* to Q_* is the same as that from P to Q.

We shall take for granted something that can actually be proven, but by a somewhat distracting argument:

- *A rigid transformation is* **affine**.

This means that if we have chosen a linear coordinate system in whatever context we are looking at (a line, a plane, or space). then the transformation $P \mapsto P_*$ is calculated in terms of coordinate arrays x and x_* according to the formula

$$x_* = xA + v$$

where A is a matrix and v a vector. Another way of saying this is that first we apply a linear transformation whose matrix is A, then a translation by v. In 3D, for example, we require that

$$[\,x_*\ y_*\ z_*\,] = [\,x\ y\ z\,]\,A + [\,v_x\ v_y\ v_z\,]\ .$$

The matrix A is called the **linear component**, v the **translation component** of the transformation.

It is clear that what we would intuitively call a rigid transformation preserves relative distances, but it might not be so clear that this requirement encapsulates rigidity completely. The following assertion may add conviction:

- *A rigid transformation (in the sense I have defined it) preserves angles as well as distances.*

That is to say, if P, Q and R are three points transformed to P_*, Q_*, and R_*, then the angle θ_* between segments P_*Q_* and P_*R_* is the same as the angle θ between PQ and PR. This is because of the cosine law, which says that

$$\begin{aligned}\cos\theta_* &= \frac{\|Q_*R_*\|^2 - \|P_*Q_*\|^2 - \|P_*R_*\|^2}{\|P_*Q_*\|\,\|P_*R_*\|} \\ &= \frac{\|QR\|^2 - \|PQ\|^2 - \|PR\|^2}{\|PQ\|\,\|PR\|} \\ &= \cos\theta\ .\end{aligned}$$

A few other facts are more elementary:

- *A transformation obtained by performing one rigid transformation followed by another rigid transformation is itself rigid.*
- *The inverse of a rigid transformations is rigid.*

In the second statement, it is implicit that a rigid transformation has an inverse. This is easy to see. First of all, any affine transformation will be invertible if and only if its linear component is. But if a linear transformation does not have an inverse, then it must collapse at least one non-zero vector into the zero vector, and it cannot be rigid.

Exercise 12.1. *Recall exactly why it is that a square matrix with determinant equal to zero must transform at least one non-zero vector into zero.*

An affine transformation is rigid if and only if its linear component is, since translation certainly doesn't affect relative distances. In order to classify rigid transformations, we must thus classify the linear ones. We'll do that in a later section, after some coordinate calculations.

Exercise 12.2. *The inverse of the transformation $x \mapsto Ax + v$ is also affine. What are its linear and translation components?*

12.2. Dot and cross products

In this section I'll recall some basic facts about vector algebra.

• Dot products

In any dimension the dot product of two vectors

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

is defined to be

$$u \bullet v = x_1y_1 + x_2y_2 + \cdots x_ny_n\ .$$

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 $u \bullet v = 0$.

• Parallel projection

One important use of dot products and cross products will be in calculating various 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.

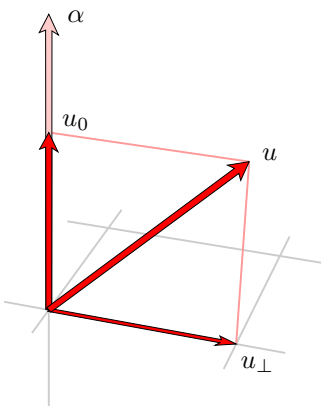

How to calculate this projection? It must be a multiple of α. We can figure out which multiple by using trigonometry. We know three facts: (a) The angle θ between α and u is determined by the formula

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

(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 it 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

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

• Volumes

The basic result about volumes is that in any dimension n, the **signed volume** of the parallelopiped spanned by vectors $v_1, v_2, \ldots, v_n$ is the determinant of the matrix

$$\begin{bmatrix} v_1 \\ v_2 \\ \cdots \\ v_n \end{bmatrix}$$

whose rows are the vectors v_i.

Let me illustrate this in dimension 2. First of all, what do I mean by the **signed** volume?

A pair of vectors u and v in 2D determine not only a parallelogram but an **orientation**, a sign. The sign is positive if u is rotated in a positive direction through the parallelogram to reach v, and negative in the other case. If u and v happen to lie on the same line, then the parallelogram is degenerate, and the sign is 0. The signed area of the parallelogram they span is this sign multiplied by the area of the parallelogram itself. Thus in the figure on the left the sign is positive while on the right it is negative. Notice that the sign depends on the order in which we list u and v.

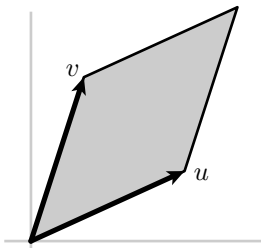

(u, v) oriented $+$

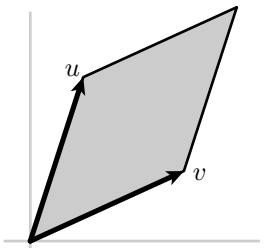

(u, v) oriented $-$

I'll offer two proofs of the fact that in 2D the signed area of the parallelogram spanned by two vectors u and v is the determinant of the matrix whose rows are u and v. The first is the simplest, but it has the drawback that it does not extend to higher dimensions, although it does play an important role.

In the first argument, recall that $u_\perp$ is the vector obtained from u by rotating it positively through $90°$. If $u = [x, y]$ then $u_\perp = [-y, x]$. As the following figure illustrates, the signed area of the paralleogram spanned by u and v is the product of $\|u\|$ and the signed length of the projection of v onto the line through $u_\perp$.

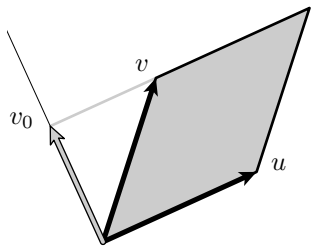

Thus the area is

$$\left(\frac{[-u_y, u_x] \bullet [v_x, v_y]}{\|u\|}\right) \|u\| = u_x v_y - u_y v_x = \det \begin{bmatrix} u_x & u_y \\ v_x & v_y \end{bmatrix} = \begin{vmatrix} u_x & u_y \\ v_x & v_y \end{vmatrix} .$$

The starting observation of the second argument is that shearing one of the rows of the matrix changes neither the determinant nor the area. For the area, this is because shears don't change area, while for the determinant it is a simple calculation:

$$\begin{vmatrix} x_1 + cx_2 & y_1 + cy_2 \\ x_2 & y_2 \end{vmatrix} = \begin{vmatrix} x_1 & y_1 \\ x_2 & y_2 \end{vmatrix} .$$

Thus in each of the transitions below, neither the determinant nor the area changes. Since they agree in the final figure, where the matrix is a diagonal matrix, they must agree in the first.

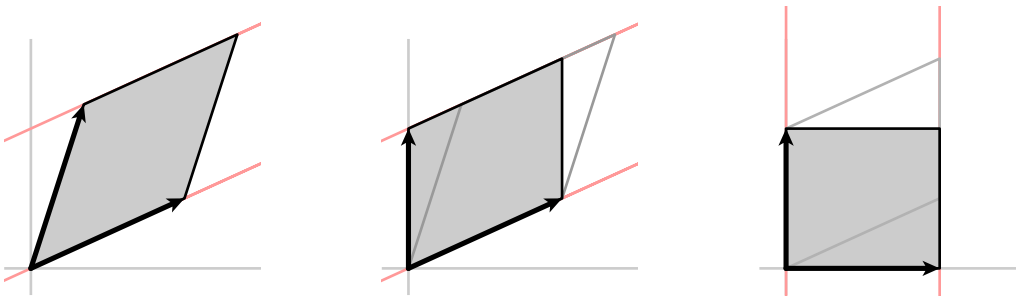

There are exceptional circumstances in which one has to be a bit fussy, but this is nonetheless the core of a valid proof that in dimension two determinants and area agree. It is still, with a little care, valid in dimension three. It is closely related to the Gauss elimination process.

Exercise 12.3. *Suppose A to be the matrix whose rows are the vectors v_1 and v_2. The argument above means that for most A we can write*

$$\begin{bmatrix} 1 & \\ c_2 & 1 \end{bmatrix} \begin{bmatrix} 1 & c_1 \\ & 1 \end{bmatrix} A = D$$

where D is a diagonal matrix. Prove this, being careful about exceptions. For which A in 2D does it fail? Provide an explicit example of an A for which it fails. Show that even for such A (in 2D) the equality of determinant and area remains true.

In n dimensions, the Gaussian elimination process finds for any matrix A a permutation matrix w, a lower triangular matrix ℓ, and an upper triangulation matrix u such that

$$A = \ell u w \ .$$

The second argument in 2D shows that the the claim is reduced to the special case of a permutation matrix, in which case it is clear.

• Cross products

In 3D—and essentially only in 3D—there is a kind of product that multiplies two vectors to get another vector. 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

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

It makes a difference in which order we write the terms in the cross product. More precisely

$$u \times v = -v \times u \ .$$

Recall that a vector is completely determined by its direction and its magnitude. The geometrical significance of the cross product is therefore contained in these two rules, which specify the cross product of two vectors uniquely.

- *The length of $w = u \times v$ is the area of the parallelogram spanned in space by u and v.*
- *The vector w lies in the line perpendicular to the plane containing u and v and its direction is determined by the right hand rule—curl the fingers so as to go from u to v and the cross-product will lie along your thumb.*

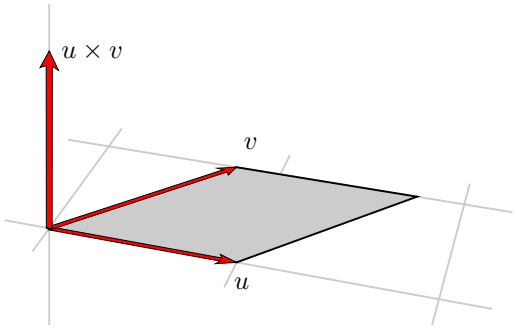

The cross product $u \times v$ will vanish only when the area of this parallelogram vanishes, or when u and v lie in a single line. Equivalently, when they are multiples of one another.

That the length of the cross-product is equal to the area of the parallelogram is a variant of Pythagoras' Theorem applied to areas—it says precisely that the square of the area of a parallelogram is equal to the sum of the squares of the areas of its projections onto the coordinate planes. I do not know a really simple way to see that it is true, however. The simplest argument I know of is to start with the determinant formula for volume. If u, v, and w are three vectors, then on the one hand the volume of the parallelopiped they span is the determinant of the matrix

$$\begin{bmatrix} u \\ v \\ w \end{bmatrix}$$

which can be written as the dot product of u and $v \times w$, while on the other this volume is the product of the projection of u onto the line perpendicular to the plane spanned by v and w and the area of the parallelogram spanned by v and w. A similar formula is valid in all dimensions, where it becomes part of the theory of **exterior products** of vector spaces, a topic beyond the scope of this book.

• Perpendicular projection

Now let $u_\perp$ be the projection of u onto the plane perpendicular to α.

The vector u has the orthogonal decomposition

$$u = u_0 + u_\perp$$

and therefore we can calculate

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

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 $\alpha/\|\alpha\|$.

Any good collection of PostScript procedures for 3D drawing will contain ones to calculate dot products, cross products, u_0 and $u_\perp$.

12.3. Linear transformations and matrices

It is time to recall the precise relationship between linear transformations and matrices. The link is the notion of **frames**. A frame in any number of dimensions d is a set of d linearly independent vectors of dimension d. Thus a frame in 1D is just a single non-zero vector. A frame in 2D is a pair of vectors whose directions do not lie in a single line. A frame in 3D is a set of three vectors not all lying in the same plane.

A frame in dimension d determines a coordinate system of dimension d, and vice-versa. If the frame is made up of $e_1, \ldots, e_d$ then every other vector x of dimension d can be expressed as a linear combination

$$x = x_1 e_1 + \cdots + x_d e_d$$

and the coefficients x_i are its coordinates with respect to that basis. They give rise to the representation of x as a row array

$$[x_1 \ \ldots \ x_d] \ .$$

Conversely, given a coordinate system the vectors

$$e_1 = [1, 0, \ldots, 0], \quad \ldots, e_d = [0, 0, \ldots, 1]$$

are the frame giving rise to it.

I emphasize:

- *A vector is a geometric entity with intrinsic significance, for example the relative position of two points. Or, if a unit of length has been chosen, something with direction and magnitude. In the presence of a frame, and only in the presence of a frame, it can be assigned coordinates.*

In other words, the vector is, if you will, an arrow and it can be assigned coordinates only with respect to a given frame. Change the frame, change the coordinates.

In a context where lengths are important, one usually works with **orthonormal frames**, those made up of a set of vectors e_i with each e_i of length 1 and distinct e_i and e_j perpendicular to each other:

$$e_i \bullet e_j = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

In the presence of any frame, vectors may be assigned coordinates. If the frame is orthonormal, the coordinates of a vector are given simply:

$$v = \sum c_i e_i, \quad c_i = e_i \bullet v \ .$$

If a particular point is fixed as origin, an arbitrary point may be assigned coordinates as well, namely those of the vector by which the origin is displaced to reach that point.

Another geometric object is a **linear function**. A linear function ℓ assigns to very vector a number, with the property, called **linearity**, that

$$\ell(au + bv) = a\ell(u) + b\ell(v) \ .$$

A linear function ℓ may also be assigned coordinates, namely the coefficients ℓ_i in its expression in terms of coordinates:

$$\ell(x) = \ell_1 x_1 + \cdots + \ell_d x_d \ .$$

A linear function is represented as a column array

$$\begin{bmatrix} \ell_1 \\ \ldots \\ \ell_d \end{bmatrix} \ .$$

Then $\ell(x)$ is the matrix product

$$[x_1 \ \ldots \ x_d] \begin{bmatrix} \ell_1 \\ \ldots \\ \ell_d \end{bmatrix} \ .$$

Another point to emphasize:

- *A linear function is a geometric entity with intrinsic significance, which assigns a number to any vector. In the presence of a frame, and only in the presence of a frame, it can be assigned coordinates, the coefficients of its expression with respect to the coordinates of vectors.*

In physics, what I call vectors are called contravariant vectors and linear functions are called covariant ones. There is often some confusion about these notions, because often one has at hand an orthonormal frame and a notion of length, and given those each vector u determines a linear function f_u:

$$f_u(v) = u \bullet v \ .$$

A **linear transformation** T assigns vectors to vectors, also with a property of linearity. In the presence of a coordinate system it, too, may be assigned coordinates, the array of coefficients $t_{i,j}$ such that

$$(xT)_j = \sum_{i=1}^{d} x_i t_{i,j} \ .$$

For example, in 2D we have

$$xT = [x_1 \ \ x_2] \begin{bmatrix} t_{1,1} & t_{1,2} \\ t_{2,1} & t_{2,2} \end{bmatrix} \ .$$

The rectangular array $t_{i,j}$ is called the of a linear **matrix** assigned to T in the presence of the coordinate system. In particular, *the i-th row of the matrix is the coordinate array of the image of the frame element e_i with respect to T.* For example,

$$[1 \ \ 0] \begin{bmatrix} t_{1,1} & t_{1,2} \\ t_{2,1} & t_{2,2} \end{bmatrix} = [t_{1,1} \ t_{1,2}]$$
$$[0 \ \ 1] \begin{bmatrix} t_{1,1} & t_{1,2} \\ t_{2,1} & t_{2,2} \end{bmatrix} = [t_{2,1} \ t_{2,2}] \ .$$

A final point, then, to emphasize:

- *A linear transformation is a geometric entity with intrinsic significance, transforming a vector into another vector. In the presence of a coordinate system (i.e. a frame), and only in the presence of a coordinate system, it can be assigned a matrix.*

In summary, *it is important to distinguish between intrinsic properties of a geometric object and properties of its coordinates.*

If A is any square matrix, one can calculate its **determinant**. It is perhaps surprising to know that this has intrinsic significance.

- *Suppose T to be a linear transformation in dimension d and A the matrix associated to T with respect to some coordinate system. The determinant of A is the factor by which T scales all d-dimensional volumes.*

This explains why, for example, the determinant of a matrix product AB is just the same as $\det(A)\det(B)$. If A changes volumes by a factor $\det(A)$ and B by the factor $\det(B)$ then the composite changes them by the factor $\det(A)\det(B)$.

If this determinant is 0, for example, it means that T is degenerate: it collapses d-dimensional objects down to something of lower dimension.

What is the significance of the sign of the determinant? Let's look at an example, where T amounts to 2D reflection in the y-axis. This takes $(1, 0)$ to $(-1, 0)$ and $(0, 1)$ to itself, so the matrix is

$$\begin{bmatrix} -1 & \\ & 1 \end{bmatrix} \ ,$$

which has determinant -1. Here is its effect on a familiar shape:

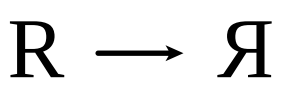

Now the transformed letter **R** is qualitatively different from the original—there is no continuous way to deform one into the other without some kind of one-dimensional degeneration. In effect, reflection in the y-axis changes **orientation** in the plane. This is always the case:

- *A linear transformation with negative determinant changes orientation.*

12.4. Changing coordinate systems

Vectors, linear functions, linear transformations all acquire coordinates in the presence of a frame. What happens if we change frames?

The first question to be answered is, *how can we describe the relationship between two frames?*

The answer is, by a matrix. To see how this works, let's look at the 2D case first. Suppose e_1, e_2 and f_1, f_2 are two frames in 2D. Make up column matrices E and F whose entries are vectors instead of numbers, vectors chosen from the frames:

$$e = \begin{bmatrix} e_1 \\ e_2 \end{bmatrix}, \quad f = \begin{bmatrix} f_1 \\ f_2 \end{bmatrix} .$$

We can write

$$f_1 = f_{1,1}e_1 + f_{1,2}e_2$$
$$f_2 = f_{2,1}e_1 + f_{2,2}e_2$$

and it seems natural to define the 2×2 matrix

$$T_f^e = \begin{bmatrix} f_{1,1} & f_{1,2} \\ f_{2,1} & f_{2,2} \end{bmatrix} .$$

Then

$$f = T_f^e e ,$$

so the matrix T_f^e relates the two frames. This use of a matrix, relating two different coordinate systems, is conceptually quite different from that in which it describes a linear transformation.

In general, suppose e and f are two frames. Each of them is a set of d vectors. Let T_f^e be the matrix whose i-th row is the array of coordinates of f_i with respect to e. We can write this relationship in an equation $f = T_f^e e$. This may seem a bit cryptic. In this equation, we think of the frames e and f themselves as column matrices

$$e = \begin{bmatrix} e_1 \\ \dots \\ e_d \end{bmatrix}, \quad f = \begin{bmatrix} f_1 \\ \dots \\ f_d \end{bmatrix}$$

whose entries are vectors, not numbers. One odd thing to keep in mind is that when we apply a matrix to a vector we write the matrix on the right, but when we apply it to a frame we put it on the left.

Now suppose x to be a vector whose array of coordinates with respect to e is x_e. We can write this relationship as

$$x = x_e\, e = [x_1 \ \dots \ x_d] \begin{bmatrix} e_1 \\ \dots \\ e_d \end{bmatrix} ,$$

again expressing the frame as a column of vectors.

With respect to the frame f we can write

$$x = x_f\, f\ .$$

We want to know how to compute x_e in terms of x_f and vice-versa. But we can write

$$x = x_f\, f = x_f(T_f^e e) = (x_f T_f^e)e = x_e\, e$$

which implies:

- *If e and f are two frames with $f = T_f^e e$ then for any vector x*

$$x_e = x_f T_f^e\ .$$

A similar result is this:

- *If e and f are two frames with $f = T_f^e e$ then for any linear function ℓ*

$$\ell_e = (T_f^e)^{-1}\ell_f\ .$$

This is what has to happen, since the function evaluation

$$x_f \ell_f = x_e \ell_e = x_e T_f^e \ell_e\ ,$$

is intrinsic.

Finally, we deal with linear transformations. We get from the first frame e a matrix A_e associated to A, and from the second a matrix A_f. *What is the relationship between the matrices A_e and A_f?* We start out by recalling that the meaning of A_e can be encapsulated in the formula

$$(xA)_e = x_e A_e$$

for any vector x. In other words, the matrix A_e calculates the coordinates of xA with respect to the frame e. Similarly

$$(xA)_f = x_f A_f$$

and then we deduce

$$\begin{aligned} x_e A_e &= (xA)_e \\ &= (xA)_f T_f^e \\ &= (x_f A_f) A_f^e \\ &= x_f (A_f T_f^e) \\ &= (x_e (T_f^e)^{-1}) A_f T_f^e \\ &= x_e ((T_f^e)^{-1} A_f T_f^e)\ . \end{aligned}$$

Hence

- *If e and f are two frames with $f = T_f^e e$ then for any linear transformation T*

$$T_e = (T_f^e)^{-1} A_f T_f^e\ .$$

This is an extremely important formula. For example, it allows us to see immediately that the determinant of a linear transformation, which is defined in terms of a matrix associated to it, is independent of the coordinate system which gives rise to the matrix. That's because

$$A_e = (T_f^e)^{-1} A_f T_f^e, \qquad \det(A_e) = \det((T_f^e)^{-1} A_f T_f^e) = \det((T_f^e)^{-1})\det(A_f)\det(T_f^e) = \det(A_f)\ .$$

Two matrices A and $T^{-1}AT$ are said to be **similar**. They are actually equivalent, in the sense that they are the matrices of the same linear transformation, but with respect to different coordinate systems.

12.5. Rigid linear transformations

Let A be an $d \times d$ matrix, representing a rigid linear transformation T with respect to the frame e.

$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,d} \\ a_{2,1} & a_{2,2} & \dots & a_{2,d} \\ & & \dots & \\ a_{d,1} & a_{d,2} & \dots & a_{d,d} \end{bmatrix} .$$

Then $e_i A$ is equal to the i-th row of A. Therefore the length of the i-th row of A is also 1. The angle between e_i and e_j is 90° if $i \neq j$, and therefore the angle between the i-th and j-th rows of A is also 90°, and the two rows must be perpendicular to each other. In other words, the rows of A must be an orthogonal frame. In fact:

- *A linear transformation is rigid precisely when its rows make up an orthogonal frame.*

Any such matrix A is said to be **orthogonal**.

The **transpose** ${}^tA$ of a matrix A has as its rows the columns of A, and vice-versa. By definition of the matrix product $A\,{}^tA$, its entries are the various dot products of the rows of A with the columns of ${}^tA$. Therefore:

- *A matrix A is orthogonal if and only if*
$$A\,{}^tA = I\ ,$$
or equivalently if and only if its transpose is the same as as its inverse.

If A and B are two $n \times n$ matrices, then

$$\det(AB) = \det(A)\det(B)\ .$$

The determinant of A is the same as that of its transpose. If A is an orthogonal matrix, then

$$\det(I) = \det(A)\det({}^tA) = \det(A)^2$$

so that:

- *The determinant of an orthogonal matrix is ± 1.*

If $\det(A) = 1$, A preserves orientation, otherwise it reverses orientation. As we have already seen, there is a serious qualitative difference between the two types. If we start with an object in one position and move it continuously, then the transformation describing its motion will be a continuous family of rigid transformations. The linear component at the beginning is the identity matrix, with determinant 1. Since the family varies continuously, the linear component can never change the sign of its determinant, and must therefore always be orientation preserving. A way to change orientation would be to reflect the object, as if in a mirror.

12.6. Orthogonal transformations in 2D

In 2D the classification of orthogonal transformations is very simple. First of all, we can rotate an object through some angle θ (possibly 0°).

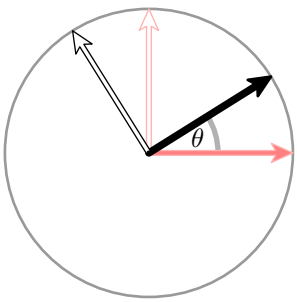

This preserves orientation. The matrix of this transformation is, as we saw much earlier,

$$\begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}.$$

Second, we can reflect things in an arbitrary line.

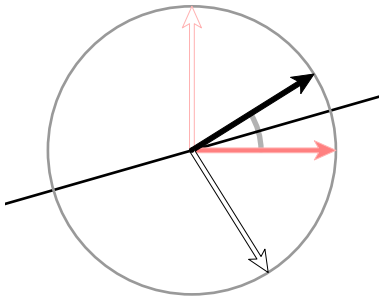

That is to say, given a line ℓ, we can transform points on ℓ into themselves, points in the line through the origin perpendicular to ℓ into their negatives. This reverses orientation.

Exercise 12.4. *If ℓ is the line at angle θ with respect to the positive x-axis, what is the matrix of this reflection?*

It turns out there are no more possibilities.

- *Every linear rigid transformation in 2D is either a rotation or a reflection.*

Let $e_1 = [1, 0]$, $e_2 = [0, 1]$, and let T be a linear rigid transformation. Since e_1 and e_2 both have length 1, both Te_1 and Te_1 also have length 1. All of these lie on the unit circle. Since the angle between e_1 and e_2 is $90°$, so is that between Te_1 and Te_2. There are two distinct possibilities, however. Either we rotate in the positive direction from Te_1 to Te_2, or in the negative direction.

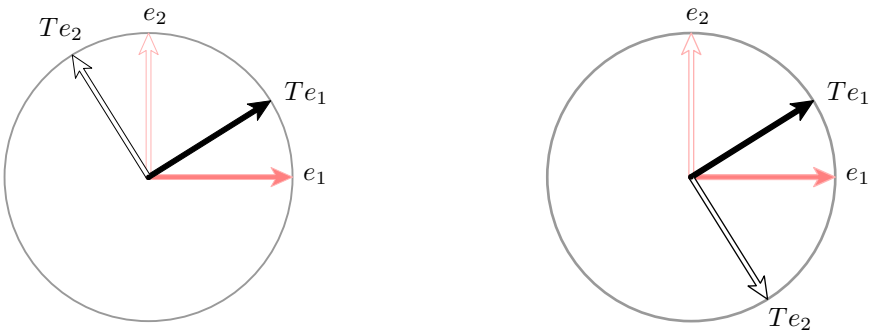

In the first case, we obtain Te_1 and Te_2 from e_1 and e_2 by a rotation. In the second case, something more complicated is going on. Here, as we move a vector u from e_1 towards e_2 and all the way around again to e_1, Tu moves along the arc from Te_1 to Te_2 all the way around again to Te_1, and *in the opposite direction*. Now if we start with two points anywhere on the unit circle and move them around in opposite directions, sooner or later they will meet. At that point we have $Tu = u$. Since T fixes u it fixes the line through u, hence takes points on the line through the origin perpendicular to it into itself. It cannot fix the points on that line, so it must negate them. In other words, T amounts to reflection in the line through u.

Exercise 12.5. *Explain why we can take u to be either of the points half way between e_1 and Te_1.*

Exercise 12.6. *Find a formula for the reflection of v in the line through u.*

12.7. Orthogonal transformations in 3D

There is one natural way to construct rigid linear motions in 3D. Choose an axis, and choose on it a direction. Equivalently choose a unit vector u, and the axis to be the line through u with direction that of u.

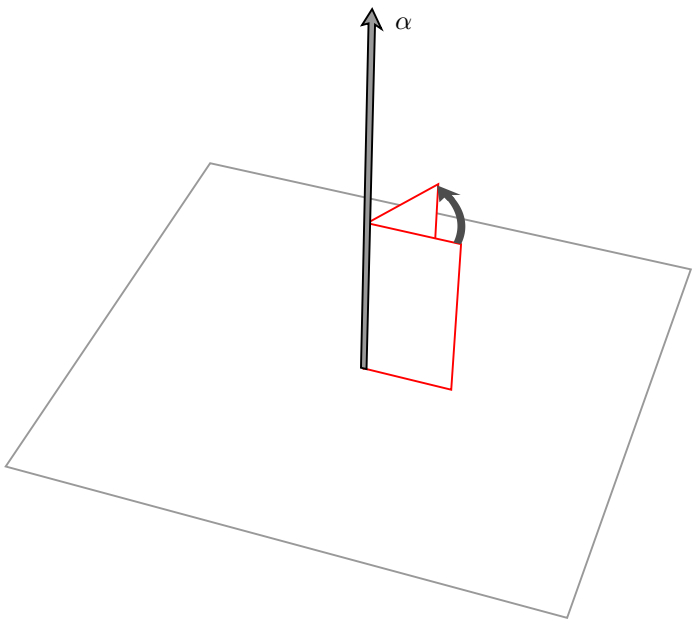

Choose an angle θ. Rotate things around the axis through angle θ, in the positive direction as seen along the axis from the positive direction. This is called a **axial rotation**.

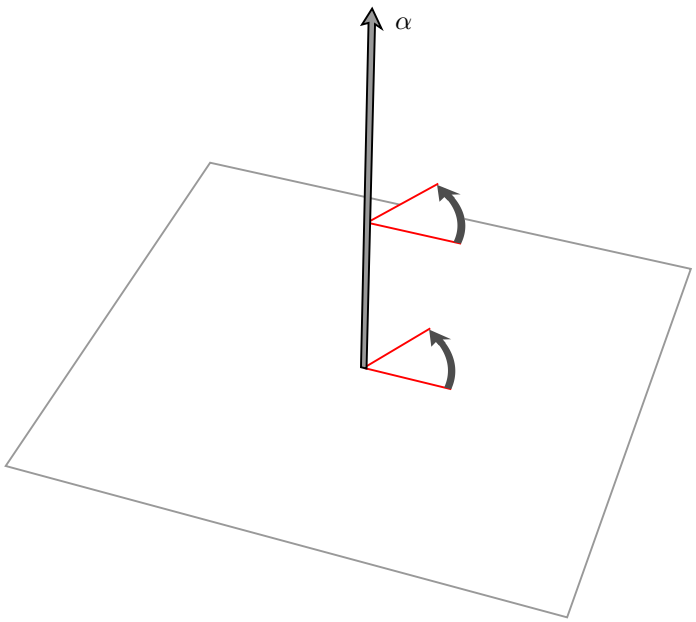

The motion can in some sense be decomposed into two parts. The plane through the origin perpendicular to he axis is rotated into itself, and points on the axis remain fixed. Therefore the height of a point above the plane remains constant, and the projection of the motion onto this plane is just an ordinary 2D rotation.

- *Every orientation-preserving linear rigid transformation in 3D is an axial rotation.*

I'll show this in one way here, and in a slightly different way later on.

Recall that an **eigenvector** v for a linear transformation T is a non-zero vector v taken into a multiple of itself by T:

$$Tv = cv$$

for some constant c, which is called the associated **eigenvalue**. This equation can be rewritten

$$Tv = cv = (T - cI)v = 0 \, .$$

If $T - cI$ were invertible, then we would deduce from this that

$$v = (T - cI)^{-1}0 = 0$$

which contradicts the assumption that $v \neq 0$. Therefore $T - cI$ is not invertible, and $\det(T - cI) = 0$. In other words, c is a root of the **characteristic polynomial**

$$\det(A - xI)$$

where A is a matrix representing T and x is a variable. For a 3×3 matrix

$$A - xI = \begin{bmatrix} a_{1,1} - x & a_{1,2} & a_{1,3} \\ a_{2,1} & a_{2,2} - x & a_{2,3} \\ a_{3,1} & a_{3,2} & a_{3,3} - x \end{bmatrix}$$

and the characteristic polynomial is a cubic polynomial which starts out

$$-x^3 + \cdots \, .$$

For $x < 0$ and $|x|$ large, this expression is positive, and for $x > 0$ and $|x|$ large it is negative. It must cross the x-axis somewhere, which means that it must have at least one real root. Therefore A has at least one real eigenvalue. In 2D this reasoning fails—there may be two conjugate complex eigenvalues instead.

Let c be a real eigenvalue of T, v a corresponding eigenvector. Since T is a rigid transformation, $\|Tv\| = \|v\|$, or $\|cv\| = \|v\|$. Since $\|cv\| = |c|\|v\|$ and $\|v\| \neq 0$, $|c| = 1$ and $c = \pm 1$.

If $c = 1$, then we have a vector fixed by T. Since T preserves angles, it takes all vectors in the plane through the origin perpendicular to v into itself. Since T reserves orientation and $Tv = v$, the restriction of T on this plane also preserves orientation. Therefore T rotates vectors in this plane, and must be a rotation around the axis through v.

If $c = -1$, then we have $Tv = -v$. The transformation T still takes the complementary plane into itself. Since T preserves orientation in 3D but reverses orientation on the line through v, T reverses orientation on this plane. But then T must be a reflection on this plane, according to the results of the previous section. We can find u such that $Tu = u$, and w perpendicular to u and v such that $Tw = -w$. In this case, T is rotation through $180°$ around the axis through u.

12.8. Calculating the effect of an axial rotation

To begin this section, I remark again that to determine an axial rotation we must specify not only an axis and an angle but a direction on the 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 choosing a vector along an axis fixes a direction on it. Given a direction on an axis I'll 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 basic 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

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

Let $u_\perp$ be the projection of u perpendicular to α. It is equal to $u - u_0$. We write

$$\begin{aligned} u &= u_0 + u_\perp \\ Ru &= Ru_0 + Ru_\perp \\ &= u_0 + Ru_\perp \ . \end{aligned}$$

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 rotation R acts as a 2D rotation in the plane perpendicular to α, so:

- *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
$$u_0 = \left(\frac{\alpha \bullet u}{\alpha \bullet \alpha}\right) \alpha = (\alpha \bullet u)\, \alpha \ .$$

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

Exercise 12.8. *Write a PostScript procedure with α and θ as arguments that returns the matrix associated to rotation by θ around α.*

12.9. Finding the axis and angle

If we are given a matrix R which we know to be orthogonal and with determinant 1, how do we find its axis and rotation angle? As we have seen, it is a special case of the problem of finding eigenvalues and eigenvectors. But the situation is rather special, and can be done in a more elementary manner. In stages:

(1) *How do we find its axis?* If e_i is the i-th standard basis vector (one of **i**, **j**, or **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 the dot-product $Ru_\perp \bullet u_* \geq 0$ otherwise $-\theta$.

Exercise 12.9. *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.

12.10. Euler's Theorem

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 don't think it is 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.

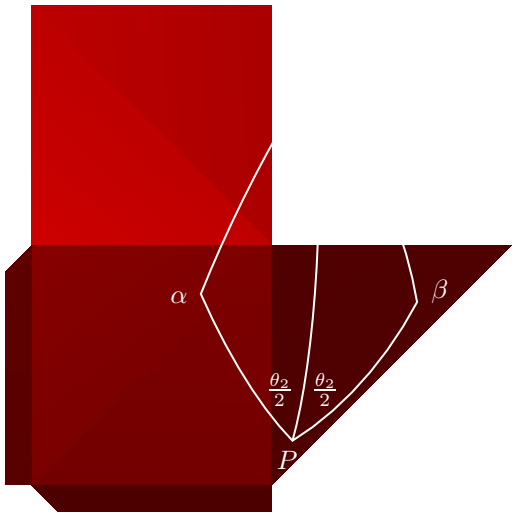

Let P_1 and P_2 be points on the unit sphere. Suppose P_1 to be on the axis of a rotation of angle θ_1, P_2 on that of a rotation of angle θ_2. Draw the spherical arc from P_1 to P_2. On either side of this arc, at P_1 draw arcs making an angle of $\theta_1/2$ and at P_2 draw arcs making an angle of $\theta_2/2$. 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_2R_1, and must be on its axis.

Exercise 12.10. *What is the axis of R_1R_2 in the diagram above? Deduce from this result under what circumstances $R_1R_2 = R_2R_1$.*

12.11. More about projections

If P is a point in space and $f(x, y, z) = Ax + By + Cz + D = 0$ a plane, then just about any point other than P can be projected from P onto the plane. The formula for this is very simple. Suppose the point being projected is Q. The projection of Q onto the plane will be the point of the line through P and Q lying in the plane. The points of the line through P and Q are those of the form $R = (1 - t)P + tQ$, so we must solve

$$f((1 - t)P + tQ) = (1 - t)f(P) + tf(Q) = 0$$

to get

$$R = \frac{f(P)\,Q - f(Q)\,P}{f(P) - f(Q)}\ .$$

The explicit formula for this is ugly, unless we use 4D homogeneous coordinates. We embed 3D into 4D by setting the last coordinate 1, making

$$P = (a, b, c, 1), \quad Q = (x, y, z, 1)$$

and thus in homogeneous coordinates the projection formula becomes simply

$$R = f(P)Q - f(Q)P\ .$$

Projection in homogeneous coordinates is a linear transformation whose matrix is

$$f(P)I - \begin{bmatrix} A \\ B \\ C \\ D \end{bmatrix} [x_P\ y_P\ z_P\ w_P] = f(P)I - \begin{bmatrix} Ax_P & Ay_P & Az_P & Aw_P \\ Bx_P & By_P & Bz_P & Bw_P \\ Cx_P & Cy_P & Cz_P & Cw_P \\ Dx_P & Dy_P & Dz_P & Dw_P \end{bmatrix}\ .$$

This is an especially pleasant formula because it continues to make sense even if P is at infinity. Also, it respects our interpretation of points as row vectors and affine functions as column vectors.