

Elements of Geometric algebra

A computational framework for geometrical applications

Leo Dorst and Stephen Mann

February 21, 2003

This is an extended early version of the double paper [7].
ASCI course Computational Geometry a20, February 2003.

Abstract

Geometric algebra is a consistent computational framework in which to define geometric primitives and their relationships. This algebraic approach contains all geometric operators and permits specification of constructions in a totally coordinate-free manner. Since it contains primitives of any dimensionality (rather than just vectors) it has no special cases: all intersections of primitives are computed with one general incidence operator. We show that the quaternion representation of rotations is also naturally contained within the framework. Models of Euclidean geometry can be made which directly represent the algebra of spheres.

1 Beyond vectors

In the usual way of defining geometrical objects in fields like computer graphics, robotics and computer vision, one uses vectors to characterize the construction. To do this effectively, the basic concept of a vector as an element of a linear space is extended by an inner product and a cross product, and some rather extraneous constructions such as homogeneous coordinates and Grassmann spaces (see [10]) to encode compactly the intersection of, for instance, offset planes in space. Many of these techniques work rather well in 3-dimensional space, although some problems have been pointed out: the difference between vectors and points, and the affine non-covariance of the normal vector as a characterization of a tangent line or tangent plane (i.e. the normal vector of a transformed plane is not the transform of the normal vector). These problems are then traditionally fixed by the introduction of certain data structures with certain combination rules; object-oriented programming can be used to implement this patch tidily.

Yet there are deeper issues in geometric programming which are still accepted as ‘the way things are’. For instance, when you need to intersect linear subspaces, the intersection algorithms are split out in treatment of the various cases: lines and planes, planes and planes, lines and

lines, et cetera, need to be treated in separate pieces of code. The linear algebra of the systems of equations with its vanishing determinants indicates changes in essential degeneracies, and finite and infinite intersections can be nicely unified by using homogeneous coordinates. But there seems no getting away from the necessity of separating the cases. After all, the outcomes themselves can be points, lines or planes, and those are essentially different in their further processing.

Yet this need not be so. If we could see subspaces as basic elements of computation, and do direct algebra with them, then algorithms and their implementation would not need to split their cases on dimensionality. For instance, $A \wedge B$ could be ‘the subspace spanned by the spaces A and B ’, the expression $A \cdot B$ could be ‘the part of B perpendicular to A ’; and then we would always have the computation rule $(A \wedge B) \cdot C = A \cdot (B \cdot C)$ since computing the part of C perpendicular to the span of A and B can be computed in two steps, perpendicularity to B followed by perpendicularity to A . Subspaces therefore have computational rules of their own which can be used immediately, independent of how many vectors were used to span them (i.e. independent of their dimensionality). In this view, the split in cases for the intersection could be avoided, since intersection of subspaces always leads to subspaces. We should consider using this structure, since it would enormously simplify the specification of geometric programs.

This paper intends to convince you that subspaces form an algebra with well-defined products which have direct geometric significance. That algebra can then be used as a language for geometry, and we claim that it is a better choice than a language always reducing everything to vectors (which are just 1-dimensional subspaces). It comes as a bit of a surprise that there is really one basic product between subspaces that forms the basis for such an algebra, namely the *geometric product*. The algebra is then what mathematicians call a Clifford algebra. But for applications, it is often very convenient to consider ‘components’ of this geometric product; this gives us sensible extensions, to subspaces, of the inner product (computing measures of perpendicularity), the cross product (computing measures of parallelness), and the meet and join (computing intersection and union of subspaces). When used in such an obviously geometrical way, the term *geometric algebra* is preferred to describe the field.

In this paper, we will use the basic products of geometric algebra to describe all familiar elementary constructions of basic geometric objects and their quantitative relationships. The goal is to show you that this can be done, and that it is compact, directly computational, and transcends the dimensionality of subspaces. We will not use geometric algebra to develop new algorithms for graphics; but we hope you to convince you that some of the lower level algorithmic aspects can be taken care of in an automatic way, without exceptions or hidden degenerate cases by using geometric algebra as a language – instead of only its vector algebra part as in the usual approach.

2 Subspaces as elements of computation

As in the classical approach, we start with a real vector space R^n which we use to denote 1-dimensional directed magnitudes. Typical usage would be to employ a vector to denote a translation in such a space, to establish the location of a point of interest. (Points are not vectors, but

their locations are.) Another usage is to denote the velocity of a moving point. (Points are not vectors, but their velocities are.) We now want to extend this capability of indicating directed magnitudes to higher-dimensional directions such as facets of objects, or tangent planes. In doing so, we will find that we have automatically encoded the algebraic properties of multi-point objects such as line segments or circles. This is rather surprising, and not at all obvious from the start. For educational reasons, we will start with the simplest subspaces: the ‘proper’ subspaces of a linear vector space which are lines, planes, etcetera through the origin, and develop their algebra of spanning and perpendicularity measures. Only in Section 8 do we show some of the considerable power of the products when used in the context of models of geometries.

2.1 Vectors

So we start with a real m -dimensional linear space V^m , of which the elements are called *vectors*. They can be added, with real coefficients, in the usual way to produce new vectors.

We will always view vectors geometrically: a vector will denote a ‘1-dimensional direction element’, with a certain ‘attitude’ or ‘stance’ in space, and a ‘magnitude’, a measure of length in that direction. These properties are well characterized by calling a vector a ‘directed line element’, as long as we mentally associate an orientation and magnitude with it: \mathbf{v} is not the same as $-\mathbf{v}$ or $2\mathbf{v}$.

2.2 The outer product

In geometric algebra, higher-dimensional oriented subspaces are also basic elements of computation. They are called *blades*, and we use the term *k-blade* for a k -dimensional homogeneous subspace. So a vector is a 1-blade. (Again, we first focus on ‘proper’ linear subspaces, i.e. subspaces which contain the origin: the 1-dimensional homogeneous subspaces are lines through the origin, the 2-dimensional homogeneous subspaces are planes through the origin, etc.)

A common way of constructing a blade is from vectors, using a product that constructs the span of vectors. This product is called the *outer product* (sometimes the *wedge product*) and denoted by \wedge . It is codified by its algebraic properties, which have been chosen to make sure we indeed get m -dimensional space elements with an appropriate magnitude (area element for $m = 2$, volume elements for $m = 3$). As you have seen in linear algebra, such magnitudes are determinants of matrices representing the basis of vectors spanning them. But such a definition would be too specifically dependent on that matrix representation. Mathematically, a determinant is viewed as an anti-symmetric linear scalar-valued function of its vector arguments. That gives the clue to the rather abstract definition of the outer product in geometric algebra:

The outer product of vectors $\mathbf{a}_1, \dots, \mathbf{a}_k$ is anti-symmetric, associative and linear in its arguments. It is denoted as $\mathbf{a}_1 \wedge \dots \wedge \mathbf{a}_k$, and called a k -blade.

The only thing that is different from a determinant is that the outer product is *not* forced to be scalar-valued; and this gives it the capability of representing the ‘attitude’ of a k -dimensional subspace element as well as its magnitude.

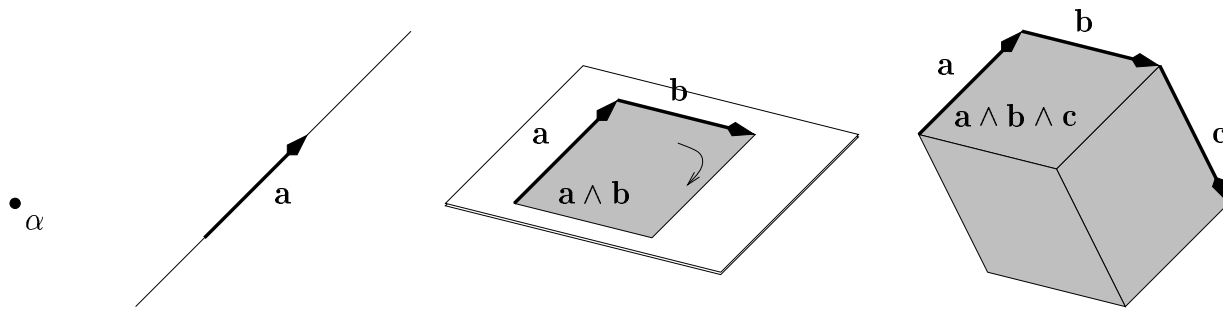


Figure 1: Spanning proper subspaces using the outer product.

2.3 2-blades in 3-dimensional space

Let us see how this works in the geometric algebra of a 3-dimensional space V^3 . For convenience, let us choose a basis $\{e_1, e_2, e_3\}$ in this space, relative to which we denote any vector (there is no need to choose this basis orthonormally – we have not mentioned the inner product yet – but you can think of it as such if you like). Now let us compute $\mathbf{a} \wedge \mathbf{b}$ for $\mathbf{a} = a_1 e_1 + a_2 e_2 + a_3 e_3$ and $\mathbf{b} = b_1 e_1 + b_2 e_2 + b_3 e_3$. By linearity, we can write this as the sum of six terms of the form $a_1 b_2 e_1 \wedge e_2$ or $a_1 b_1 e_1 \wedge e_1$. By anti-symmetry, the outer product of any vector with itself must be zero, so the term with $a_1 b_1 e_1 \wedge e_1$ and other similar terms disappear. Also by anti-symmetry, $e_2 \wedge e_1 = -e_1 \wedge e_2$, so some terms can be grouped. You may verify that the final result is:

$$\begin{aligned} \mathbf{a} \wedge \mathbf{b} &= \\ &= (a_1 e_1 + a_2 e_2 + a_3 e_3) \wedge (b_1 e_1 + b_2 e_2 + b_3 e_3) \\ &= (a_1 b_2 - a_2 b_1) e_1 \wedge e_2 + (a_2 b_3 - a_3 b_2) e_2 \wedge e_3 + (a_3 b_1 - a_1 b_3) e_3 \wedge e_1 \end{aligned} \quad (1)$$

We cannot simplify this further. Apparently, the axioms of the outer product permit us to decompose any 2-blade in 3-dimensional space onto a basis of 3 elements. This ‘2-blade basis’ (also called ‘bivector basis’) $\{e_1 \wedge e_2, e_2 \wedge e_3, e_3 \wedge e_1\}$ consists of 2-blades spanned by the basis vectors. Linearity of the outer product implies that the set of 2-blades forms a linear space on this basis. We will interpret this as the space of all *plane elements* or *area elements*. Let us show that they have indeed the correct magnitude for an area element. That is particularly clear if we choose a particular orthonormal basis $\{e_1, e_2, e_3\}$, chosen such that \mathbf{a} lies in the e_1 -direction, and \mathbf{b} lies in the (e_1, e_2) -plane. Then $\mathbf{a} = a e_1$, $\mathbf{b} = b \cos \phi e_1 + b \sin \phi e_2$ (with ϕ the angle from \mathbf{a} to \mathbf{b}), so that

$$\mathbf{a} \wedge \mathbf{b} = (a b \sin \phi) e_1 \wedge e_2 \quad (2)$$

This single result contains both the correct magnitude of the area $a b \sin \phi$ spanned by \mathbf{a} and \mathbf{b} , and the plane in which it resides – for we should learn to read $e_1 \wedge e_2$ as ‘the unit directed area element of the (e_1, e_2) -plane’. Since we can always adapt our coordinates to vectors in this way, this result is universally valid: $\mathbf{a} \wedge \mathbf{b}$ is an area element of the plane spanned by \mathbf{a} and \mathbf{b} .

You can visualize this as the parallelogram spanned by \mathbf{a} and \mathbf{b} , but you should be a bit careful: the *shape* of the area element is not defined in $\mathbf{a} \wedge \mathbf{b}$. For instance, by the properties of the outer product, $\mathbf{a} \wedge \mathbf{b} = \mathbf{a} \wedge (\mathbf{b} + \lambda \mathbf{a})$, for any λ , so the parallelogram can be sheared.

Also, the area element is free to translate: the sum of the area elements $\frac{1}{4}(\mathbf{a} \wedge \mathbf{b})$, $\frac{1}{4}(\mathbf{b} \wedge (-\mathbf{a}))$, $\frac{1}{4}((-\mathbf{a}) \wedge (-\mathbf{b}))$, $\frac{1}{4}((-\mathbf{b}) \wedge \mathbf{a})$ equals $\mathbf{a} \wedge \mathbf{b}$; drawing this equation shows that we should imagine the area element to have no specific location in its plane. You may also verify that an orthogonal transformation of \mathbf{a} and \mathbf{b} in their common plane (such as a rotation in that plane) leaves $\mathbf{a} \wedge \mathbf{b}$ unchanged. (This is obvious once you know the result for determinants and note that $\mathbf{a} \wedge \mathbf{b}$ can always be expressed as in eq.(1), but we will revisit its deeper meaning in Section 7).

It is important to realize that the 2-blades have an existence of their own, independent of any vectors that one might use to define them; that is reflected in the fact that they are not parallelograms. Planes (or, more precisely, plane elements) are nouns in our computational geometrical language, of the same basic nature as vectors (or line elements).

2.4 Volumes as 3-blades

We can also form the outer product of *three* vectors \mathbf{a} , \mathbf{b} , \mathbf{c} . Considering each of those decomposed onto their 3 components on some basis in our 3-dimensional space (as above), we obtain terms of three different types, depending on how many common components occur: terms like $a_1 b_1 c_1 \mathbf{e}_1 \wedge \mathbf{e}_1 \wedge \mathbf{e}_1$, like $a_1 b_1 c_2 \mathbf{e}_1 \wedge \mathbf{e}_1 \wedge \mathbf{e}_2$, and like $a_1 b_2 c_3 \mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3$. Because of associativity and anti-symmetry, only the last type survives, in all its permutations. The final result is:

$$\mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c} = (a_1 b_2 c_3 - a_1 b_3 c_2 + a_2 b_1 c_3 - a_2 b_3 c_1 + a_3 b_1 c_2 - a_3 b_2 c_1) \mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3.$$

The scalar factor is the determinant of the matrix with columns \mathbf{a} , \mathbf{b} , \mathbf{c} , which is proportional to the signed volume spanned by them (as is well known from linear algebra). The term $\mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3$ is the denotation of which volume is used as unit: that spanned by \mathbf{e}_1 , \mathbf{e}_2 , \mathbf{e}_3 . The order of the vectors gives its orientation, so this is a ‘signed volume’. In 3-dimensional space, there is not really any other choice for the construction of volumes than (possibly negative) multiples of this volume. But in higher dimensional spaces, the attitude of the volume element needs to be indicated just as much as we needed to denote the attitude of planes in 3-space.

2.5 Linear dependence

Note that if the three vectors are linearly dependent, they satisfy:

$$\mathbf{a}, \mathbf{b}, \mathbf{c} \text{ linearly dependent} \iff \mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c} = 0.$$

We interpret the latter immediately as the geometric statement that the vectors span a zero volume. This makes linear dependence a computational property rather than a predicate: three vectors can be ‘almost linearly dependent’. The magnitude of $\mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c}$ obviously involves the determinant of the matrix $(\mathbf{a} \ \mathbf{b} \ \mathbf{c})$, so this view corresponds with the usual computation of determinants to check degeneracy.

2.6 The pseudoscalar as hypervolume

Forming the outer product of *four* vectors $\mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c} \wedge \mathbf{d}$ in 3-dimensional space will always produce zero (since they must be linearly dependent). To see this, just decompose the vectors on

some basis (for instance, the fourth vector on a basis formed by the other 3), and apply the outer product. Since $(\mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c})$ is proportional to $\mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3$, multiplication by \mathbf{d} will always lead to terms like $\mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3 \wedge \mathbf{e}_1$, in which at least two vectors are the same. Associativity and anti-symmetry then makes all terms equal to zero.

The highest order blade which is non-zero in an m -dimensional space is therefore an m -blade. Such a blade, representing an m -dimensional volume element, is called a *pseudoscalar* for that space (for historical reasons); unfortunately a rather abstract term for the elementary geometric concept of ‘hypervolume element’.

2.7 Scalars as subspaces

To make scalars fully admissible elements of the algebra we have so far, we can define the outer product of two scalars, and a scalar and a vector, through identifying it with the familiar scalar product in the vector space we started with:

$$\alpha \wedge \beta = \alpha \beta \quad \text{and} \quad \alpha \wedge \mathbf{v} = \alpha \mathbf{v}$$

This automatically extends (by associativity) to the outer product of scalars with higher order blades.

We will denote scalars mostly by Greek lower case letters. Since they are constructed by the outer product of zero vectors, we can interpret the scalars as the representation in geometric algebra of 0-dimensional subspace elements, i.e. as a *weighted points at the origin* – or maybe you prefer ‘charged’, since the weight can be negative. This is indeed consistent, we will get back to that when intersecting subspaces in Section 4.

2.8 The Grassmann algebra of 3-space

Collating what we have so far, we have constructed a geometrically significant algebra containing only two operations: the addition $+$ and the outer multiplication \wedge (subsuming the usual scalar multiplication). Starting from scalars and a 3-dimensional vector space we have generated a 3-dimensional space of 2-blades, and a 1-dimensional space of 3-blades (since all volumes are proportional to each other). In total, therefore, we have a set of elements which naturally group by their dimensionality. Choosing some basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$, we can write what we have as spanned by the set:

$$\left\{ \underbrace{1}_{\text{scalars}}, \underbrace{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3}_{\text{vector space}}, \underbrace{\mathbf{e}_1 \wedge \mathbf{e}_2, \mathbf{e}_2 \wedge \mathbf{e}_3, \mathbf{e}_3 \wedge \mathbf{e}_1}_{\text{bivector space}}, \underbrace{\mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3}_{\text{trivector space}} \right\} \quad (3)$$

Every k -blade formed by \wedge can be decomposed on the k -vector basis using $+$. The ‘dimensionality’ k is often called the *grade* or *step* of the k -blade or k -vector, reserving the term *dimension* for that of the vector space which generated them. A k -blade represents a k -dimensional oriented subspace element.

If we allow the scalar-weighted addition of arbitrary elements in this set of basis blades, we get an 8-dimensional linear space from the original 3-dimensional vector space. This space, with $+$ and \wedge as operations, is called the *Grassmann algebra* of 3-space.

The dimensionality of a k -blade is the number of vector factors that span it; this is usually called the *grade* of the blade. It obeys the simple rule:

$$\text{grade}(\mathbf{A} \wedge \mathbf{B}) = \text{grade}(\mathbf{A}) + \text{grade}(\mathbf{B}). \quad (4)$$

Of course the outcome may be 0, so this zero element of the algebra should be seen as an element of arbitrary grade. There is then no need to distinguish separate zero scalars, zero vectors, zero 2-blades.

We have no interpretation (yet) for mixed-grade terms such as $1 + \mathbf{e}_1$. Actually, even addition of elements of the same grade is hard to interpret in spaces of more than 3 dimensions, since it easily leads to elements that cannot be decomposed using the outer product – so to non-blades, i.e. objects that cannot be ‘spanned’ by vectors. (For instance, $\mathbf{e}_1 \wedge \mathbf{e}_2 + \mathbf{e}_3 \wedge \mathbf{e}_4$ in 4-space cannot be written in the form $\mathbf{a} \wedge \mathbf{b}$ – try it!) The general term for the sum of k -blades (for the same k) is *k-vector*, and the general term for the mixed-grade elements permitted in Grassmann algebra is *multivector*.

2.9 Many blades

From the way it is constructed through the anti-symmetric product, it should be clear that the k -dimensional subspaces of an m -dimensional space have a basis which consists of a number of independent elements equal to the number of ways one can take k distinct indices from a set of m indices. That is

The linear space of k -vectors in m -space is $\binom{m}{k}$ -dimensional.

Adding them all up, we find:

The linear space of all subspaces of an m -dimensional vector space is 2^m -dimensional.

To have a basis for all possible subspaces (through the origin) in 3-dimensional space takes $2^3 = 8$ elements, such as in eq.(3). You can characterize an element X of that space therefore by a 8×1 matrix $[X]$. Since the outer product by another element vector A is linear, $A \wedge X$ can be written as the action of a linear operator A^\wedge on X , and hence be represented as a matrix multiplication $[A^\wedge][X]$, with $[A^\wedge]$ an 8×8 matrix. This is not a particularly efficient representation, but it shows that this algebra of $+$ and \wedge on a vector space is just a special linear algebra; a fact which may give you some confidence that it is at least consistent.

When they just learn about this algebra, most people are put off by how many blades there are, and some have rejected the practical use of geometric algebra because of its exponentially large basis. This is a legitimate concern, and the implementation just sketched obviously does not scale well with dimensionality. For now, a helpful view may be to see this 2^m -dimensional basis as a cabinet in which all relationships which we may care to compute in the course of our

computations in m -dimensional space can be filed properly: k -point relationships in the $\binom{m}{k}$ files in the k -th drawer. And the files themselves have clear computational relationships (we have seen the outer product, more will follow). This should be compared to the usual way in which such k -point relationships are made whenever they are needed, but not preserved in a structural way relating them algebraically to the other relationships of the application. This metaphor suggests that there might be some potential gain in building up the overall structure rather than reinventing it several times along the way, as long as we make sure that this organization does not affect the efficiency of individual computations too much. This paper should provide you with sufficient material to ponder this new possibility.

3 Relative subspaces measures

The outer product gives computational meaning to the notion of ‘spanning subspaces’. It does not use any metric structure which we may have available for our original vector space V^m . The familiar inner product of vectors in a vector space *does* use the metric – in fact, it *defines* the metric, since it gives a bilinear form returning a scalar value $\mathbf{a} \cdot \mathbf{b}$ for each pair of vectors, which can be used to defined the distance measure $\sqrt{(\mathbf{a} - \mathbf{b}) \cdot (\mathbf{a} - \mathbf{b})}$. Now that vectors are viewed as representatives of 1-dimensional subspaces, we of course want to extend this metric capability to arbitrary subspaces. This leads to the *scalar product*, and when meshing with the outer product gives a generalized *inner product* between blades.

3.1 The scalar product: a metric for blades

Between two blades \mathbf{A}_k and \mathbf{B}_k of the same grade k , we can define a metric measure. The most computational way of doing so is to span each of the blades by k vectors: $\mathbf{A}_k = \mathbf{a}_1 \wedge \mathbf{a}_2 \wedge \cdots \wedge \mathbf{a}_k$ and $\mathbf{B}_k = \mathbf{b}_1 \wedge \mathbf{b}_2 \wedge \cdots \wedge \mathbf{b}_k$. Then the scalar product between them is defined as:

$$\mathbf{A}_k * \mathbf{B}_k \equiv \begin{vmatrix} \mathbf{a}_1 \cdot \mathbf{b}_k & \mathbf{a}_1 \cdot \mathbf{b}_{k-1} & \cdots & \mathbf{a}_1 \cdot \mathbf{b}_1 \\ \mathbf{a}_2 \cdot \mathbf{b}_k & \mathbf{a}_2 \cdot \mathbf{b}_{k-1} & \cdots & \mathbf{a}_2 \cdot \mathbf{b}_1 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_k \cdot \mathbf{b}_k & \mathbf{a}_k \cdot \mathbf{b}_{k-1} & \cdots & \mathbf{a}_k \cdot \mathbf{b}_1 \end{vmatrix} \quad (5)$$

The unfortunate order of the factors was chosen historically. We get a nicer form if we introduce an operation that reverses a factorization, for instance $\mathbf{A} = \mathbf{a}_1 \wedge \mathbf{a}_2 \wedge \mathbf{a}_3$ would become $\mathbf{a}_3 \wedge \mathbf{a}_2 \wedge \mathbf{a}_1$. (We need this for other purposes as well, or we would have preferred to fix the scalar product.) Due to the anti-symmetry of the outer product, these differ only by a sign factor, for a k -blade a sign of $(-1)^{\frac{1}{2}k(k-1)}$. We denote it by a tilde, so: $\widetilde{\mathbf{A}} = \mathbf{a}_3 \wedge \mathbf{a}_2 \wedge \mathbf{a}_1 = -\mathbf{A}$. Now $\widetilde{\mathbf{A}} * \mathbf{B}$ has nicely matching coefficients.

The value of $\widetilde{\mathbf{A}} * \mathbf{B}$ is independent of the factorization of \mathbf{A} and \mathbf{B} , as you may verify by the properties of determinants: adding a multiple of, say \mathbf{a}_2 to \mathbf{a}_1 leaves the blade \mathbf{A} unchanged, so it should give the same answer. In $\widetilde{\mathbf{A}} * \mathbf{B}$, it leads to addition of a multiple of the second column

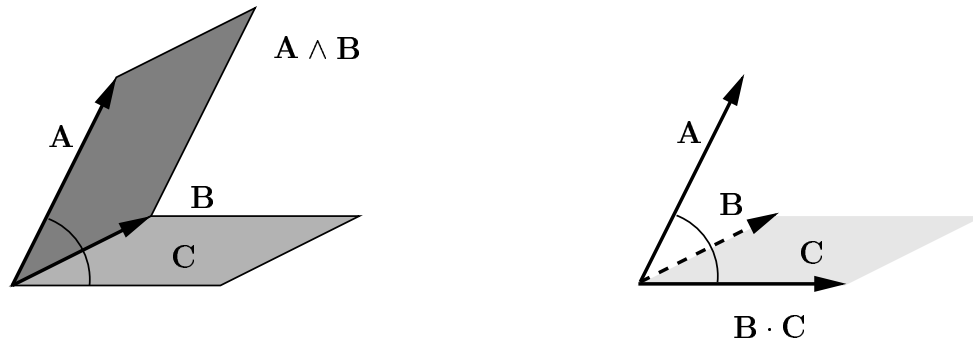


Figure 2: The metric relationship between different spans.

to the first, and this indeed leaves the determinant unchanged – the two anti-symmetries in the definitions of \wedge and $*$ match well. The value of $\widetilde{\mathbf{A}} * \mathbf{B}$ is proportional to the cosine of the angle of the two subspaces – if a rotation exists that rotates one into the other, otherwise it is zero. The definition is extended to blades of different grade by setting $\mathbf{A} * \mathbf{B} = 0$ whenever the grades are different. So no scalar metric comparison is possible between such different subspaces (but for them we have the inner product of the next section).

The scalar product of a subspace with itself gives us the *norm* of the subspace, defined as ¹:

$$|\mathbf{A}| = \sqrt{\widetilde{\mathbf{A}} * \mathbf{A}} \quad (6)$$

For a 2-blade $\mathbf{A} = \mathbf{a}_1 \wedge \mathbf{a}_2$, with an angle of ϕ between \mathbf{a}_1 and \mathbf{a}_2 , you may verify that this gives $|\mathbf{A}| = |\mathbf{a}_1| |\mathbf{a}_2| |\sin \phi|$, the absolute value of the area measure, precisely what one would hope.

3.2 The inner product

The geometric nature of blades means that there are relationships between the metric measures of different grades: for instance, the angle two 2-blades make is related to that of two properly chosen vectors in their planes (see Figure 2). We should therefore be capable of relating those numerically. If a blade is spanned as $\mathbf{A} \wedge \mathbf{B}$, and we are interested in its measure relative to \mathbf{C} we compute $(\mathbf{A} \wedge \mathbf{B}) * \mathbf{C}$; but we should be able to find a similar measure between the subblade \mathbf{A} , and some subblade of \mathbf{C} , which is ‘ \mathbf{C} with \mathbf{B} taken out’’. This can be used to define a new product, through:

$$(\mathbf{A} \wedge \mathbf{B}) * \mathbf{C} = \mathbf{A} * (\mathbf{B} \cdot \mathbf{C}), \quad \text{for all } \mathbf{C} \quad (7)$$

The blade $\mathbf{B} \cdot \mathbf{C}$ is the *inner product* of \mathbf{B} and \mathbf{C} . Its grade is the difference of the grades of \mathbf{C} and \mathbf{B} (since it should equal the grade of \mathbf{A} in the definition). The inner product can be interpreted more directly as

¹This works only in a Euclidean metric in a real vector space; in other metrics one should define the ‘norm squared’ and avoid the square root.

$\mathbf{B} \cdot \mathbf{C}$ is the blade representing the largest subspace which is contained in the subspace \mathbf{C} and which is perpendicular to the subspace \mathbf{B} ; it is linear in \mathbf{B} and \mathbf{C} ; it coincides with the usual inner product $\mathbf{b} \cdot \mathbf{c}$ of V^m when computed for vectors \mathbf{b} and \mathbf{c} .

The above determines the inner product uniquely². It turns out not to be symmetrical (as one would expect since the definition is asymmetrical) and also not associative. But we do demand linearity, to make it computable between any two elements in our linear space (not just blades).

For later use, we just give the rules by which to compute the resulting inner product for arbitrary blades, omitting their derivation. Then we will do some examples to convince you that it does what we want it to do. In the following α, β are scalars, \mathbf{a} and \mathbf{b} vectors and $\mathbf{A}, \mathbf{B}, \mathbf{C}$ blades of arbitrary order. We give the rules in a slightly redundant form, for convenience in evaluating expressions.

scalars	$\alpha \cdot \beta = \alpha \wedge \beta$	(8)
vector and scalar	$\mathbf{a} \cdot \beta = 0$	(9)
scalar and vector	$\alpha \cdot \mathbf{b} = \alpha \wedge \mathbf{b}$	(10)
vectors	$\mathbf{a} \cdot \mathbf{b}$ is the usual inner product in V^m	(11)
vector and blade	$\mathbf{a} \cdot (\mathbf{b} \wedge \mathbf{B}) = (\mathbf{a} \cdot \mathbf{b}) \wedge \mathbf{B} - \mathbf{b} \wedge (\mathbf{a} \cdot \mathbf{B})$	(12)
blades	$(\mathbf{A} \wedge \mathbf{B}) \cdot \mathbf{C} = \mathbf{A} \cdot (\mathbf{B} \cdot \mathbf{C})$	(13)
distributivity 1	$\mathbf{A} \cdot (\mathbf{B} + \mathbf{C}) = \mathbf{A} \cdot \mathbf{B} + \mathbf{A} \cdot \mathbf{C}$	(14)
distributivity 2	$(\mathbf{A} + \mathbf{B}) \cdot \mathbf{C} = \mathbf{A} \cdot \mathbf{C} + \mathbf{B} \cdot \mathbf{C}$	(15)

It should be emphasized that the inner product is not associative. For instance, $\mathbf{a} \cdot (\mathbf{b} \cdot \mathbf{c}) = 0$ since the second argument is a scalar; but $(\mathbf{a} \cdot \mathbf{b}) \cdot \mathbf{c} = \alpha \mathbf{c}$ (with $\alpha = \mathbf{a} \cdot \mathbf{b}$) is a vector. Neither is the inner product symmetrical, as the scalar/vector rules show.

The inner product affects the grades simply (compare eq.(4)):

$$\text{grade}(\mathbf{A} \cdot \mathbf{B}) = \text{grade}(\mathbf{B}) - \text{grade}(\mathbf{A}). \tag{16}$$

with the understanding that all blades of negative grade are 0.

3.3 Perpendicularity and dualization

Having the inner product expands our capabilities in geometric computations. It enables manipulation of expressions involving ‘spanning’ to being about ‘perpendicularity’ and vice versa. Such ‘dual’ formulations turn out to be very convenient. We briefly develop intuition and basic conversion expressions for these manipulations.

²The resulting inner product differs slightly from the inner product commonly used in the geometric algebra literature. Our inner product has a cleaner geometric semantics, and more compact mathematical properties, and that makes it better suited to computer science. It is sometimes called the *contraction*, and denoted as $\mathbf{B} \rfloor \mathbf{C}$ rather than $\mathbf{B} \cdot \mathbf{C}$. The two inner products can be expressed in terms of each other, so this is not a severely divisive issue. They ‘algebraify’ the same geometric concepts, in just slightly different ways [4].

- *perpendicularity*

We define the concept of perpendicularity through the inner product:

$$\mathbf{a} \text{ perpendicular to } \mathbf{A} \iff \mathbf{a} \cdot \mathbf{A} = 0,$$

It is then easy to prove that, for general blades \mathbf{A} , the construction $\mathbf{A} \cdot \mathbf{B}$ is indeed perpendicular to \mathbf{A} , as we suggested in the previous section. For any vector \mathbf{a} satisfies $\mathbf{a} \cdot (\mathbf{A} \cdot \mathbf{B}) = (\mathbf{a} \wedge \mathbf{A}) \cdot \mathbf{B}$. But if \mathbf{a} is in \mathbf{A} it must be linearly dependent on the spanning vectors, so $\mathbf{a} \wedge \mathbf{A} = 0$. Therefore $\mathbf{a} \cdot (\mathbf{A} \cdot \mathbf{B}) = 0$ for any \mathbf{a} in \mathbf{A} . So any vector in \mathbf{A} is perpendicular to $\mathbf{A} \cdot \mathbf{B}$.

- *orthogonal complement and dual*

If we take the inner product of a blade relative to the volume element of the space it resides in (i.e. relative to the pseudoscalar of the space), we get the whole subspace perpendicular to it. This is how *duality* sits in geometric algebra: it is simply taking an orthogonal complement. A good example in a 3-dimensional Euclidean space is the dual of a 2-blade (or bivector). Using an orthonormal basis $\{\mathbf{e}_i\}_{i=1}^3$ and the corresponding bivector basis, we write: $\mathbf{B} = b_1\mathbf{e}_2 \wedge \mathbf{e}_3 + b_2\mathbf{e}_3 \wedge \mathbf{e}_1 + b_3\mathbf{e}_1 \wedge \mathbf{e}_2$. We take the dual relative to the space with volume element $\mathbf{I}_3 \equiv \mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3$ (i.e. the ‘right-handed volume’ formed by using a right-handed basis). Any scalar multiple would do, but it turns out that the best definition is to use the *reverse* of \mathbf{I}_3 to define the dual (since that generalizes to higher dimensions; here $\tilde{\mathbf{I}}_3 = -\mathbf{I}_3$). The subspace of \mathbf{I}_3 dual to \mathbf{B} is then:

$$\begin{aligned} \mathbf{B} \cdot \tilde{\mathbf{I}}_3 &= (b_1\mathbf{e}_2 \wedge \mathbf{e}_3 + b_2\mathbf{e}_3 \wedge \mathbf{e}_1 + b_3\mathbf{e}_1 \wedge \mathbf{e}_2) \cdot (\mathbf{e}_3 \wedge \mathbf{e}_2 \wedge \mathbf{e}_1) \\ &= b_1\mathbf{e}_1 + b_2\mathbf{e}_2 + b_3\mathbf{e}_3. \end{aligned} \tag{17}$$

This is a vector, and we recognize it (in this Euclidean space) as the *normal vector* to the planar subspace represented by \mathbf{B} . So we have normal vectors in geometric algebra as the duals of 2-blades, if we would want them (but we will see in Section 7.3 why we prefer the direct representation of a planar subspace by a 2-blade rather than the indirect representation by normal vectors).

If it is clear from context relative to which pseudoscalar \mathbf{I} the dual is taken, we will use the convenient shorthand \mathbf{B}^* for $\mathbf{B} \cdot \tilde{\mathbf{I}}$.

- *duality relationships*

Going over to a dual representation involves translating formulas given in terms of spanning to formulas using perpendicularity. An example is the specification of a plane in 3-space given its 2-blade \mathbf{B} . On the one hand, all vectors in the plane satisfy $\mathbf{x} \wedge \mathbf{B} = 0$ (zero volume spanned with the 2-blade); but dually they satisfy $\mathbf{x} \cdot \mathbf{B}^* = 0$ (perpendicular to the normal vector). This is an example of a more general duality relationship between blades, which we state without proof. Let \mathbf{A} , \mathbf{B} and \mathbf{I} be blades, with \mathbf{A} contained in \mathbf{I} (this is essential). Then:

$$(\mathbf{A} \cdot \mathbf{B}) \cdot \mathbf{I} = \mathbf{A} \wedge (\mathbf{B} \cdot \mathbf{I}) \quad \text{if } \mathbf{A} \subseteq \mathbf{I}. \tag{18}$$

Remember also the universally valid eq.(13)

$$(\mathbf{A} \wedge \mathbf{B}) \cdot \mathbf{I} = \mathbf{A} \cdot (\mathbf{B} \cdot \mathbf{I}). \quad (19)$$

Together, these equations allow the change to a ‘dual perspective’ converting spanning to orthogonality and *vice versa*, permitting more flexible interpretation of equations.

Let us use these to verify the motivating example above in full detail. In a 3-dimensional space with pseudoscalar \mathbf{I}_3 , the equation $\mathbf{x} \wedge \mathbf{B} = 0$ (meaning that \mathbf{x} is in the 2-dimensional subspace determined by \mathbf{B}) can be dualized to $0 = (\mathbf{x} \wedge \mathbf{B}) \cdot \tilde{\mathbf{I}}_3 = \mathbf{x} \cdot (\mathbf{B} \cdot \tilde{\mathbf{I}}_3)$. This characterizes the vectors in the \mathbf{B} -plane through its normal vector $\mathbf{n} \equiv \mathbf{B} \cdot \tilde{\mathbf{I}}_3 = \mathbf{B}^*$. It is the familiar ‘normal equation’ of the plane, and identical to the common way to represent a plane by its normal vector \mathbf{n} .

In general, we will say that a blade \mathbf{B} represents a subspace \mathcal{B} of vectors \mathbf{x} if

$$\mathbf{x} \in \mathcal{B} \iff \mathbf{x} \wedge \mathbf{B} = 0 \quad (20)$$

and that a blade \mathbf{B}^* *dually represents* the subspace \mathcal{B} if

$$\mathbf{x} \in \mathcal{B} \iff \mathbf{x} \cdot \mathbf{B}^* = 0. \quad (21)$$

Switching between the two standpoints is done by the duality relations above.

- *the cross product*

Classical computations with vectors in Euclidean 3-space often use the cross product, which produces from two vectors \mathbf{a} and \mathbf{b} a new vector $\mathbf{a} \times \mathbf{b}$ perpendicular to both (by the right-hand rule), proportional to the area they span. We can make this in geometric algebra as the dual of the 2-blade spanned by the vectors:

$$\mathbf{a} \times \mathbf{b} \equiv (\mathbf{a} \wedge \mathbf{b}) \cdot \tilde{\mathbf{I}}_3. \quad (22)$$

This shows a number of things explicitly which one always needs to remember about the cross product: there is a convention involved on handedness (this is coded in the sign of \mathbf{I}_3); there are metric aspects since it is perpendicular to a plane (this is coded in the usage of the inner product ‘ \cdot ’); and the construction really only works in three dimensions, since only then is the dual of a 2-blade a vector (this is coded in the 3-gradedness of \mathbf{I}_3). The vector relationship $\mathbf{a} \wedge \mathbf{b}$ does not depend on any of these embedding properties, yet characterizes the (\mathbf{a}, \mathbf{b}) -plane just as well.

You may verify that computing eq.(22) explicitly using eq.(1) and eq.(17) indeed retrieves the usual expression:

$$\mathbf{a} \times \mathbf{b} = (a_2 b_3 - a_3 b_2) \mathbf{e}_1 + (a_3 b_1 - a_1 b_3) \mathbf{e}_2 + (a_1 b_2 - a_2 b_1) \mathbf{e}_3 \quad (23)$$

In geometric algebra, we have the possibility of replacing the cross product by more elementary constructions. In Section 7.3 we discuss the advantages of doing so.

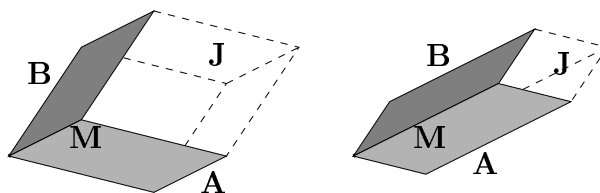


Figure 3: The ambiguity of scale for meet M and join J of two blades A and B . Both figures are examples of acceptable solutions.

4 Intersecting subspaces

So far, we can span subspaces and consider their containment and orthogonality. Geometric algebra also contains operations to determine the *union* and *intersection* of subspaces. These are the join and meet operations. Several notations exist for these in literature, causing some confusion. For this paper, we will simply use the set notations \cup and \cap to make the formulas more easily readable.³

4.1 Union of subspaces

The join of two subspaces is their smallest superspace, i.e. the smallest space containing them both. Representing the spaces by blades A and B , the join is denoted $A \cup B$. If the subspaces of A and B are disjoint, their join is obviously proportional to $A \wedge B$. But a problem is that if A and B are not disjoint (which is precisely the case we are interested in), then $A \cup B$ contains an unknown scaling factor which is fundamentally unresolvable due to the reshapable nature of the blades discussed in Section 2.3 (see Figure 3; this ambiguity was also observed by [17][Stolfi]). Fortunately, it appears that in all geometrically relevant entities which we compute this scalar ambiguity cancels.

The join is a more complicated product of subspaces than the outer product and inner product; we can give no simple formula for the grade of the result (like eq.(4)), and it cannot be characterized by a list of algebraic computation rules. Although computation of the join may appear to require some optimization process, finding the *smallest* superspace can actually be done in virtually constant time [1].

4.2 Intersection of subspaces

The meet of two subspaces A and B is their largest common subspace. If this is the blade M , then A can be factorized as $A = A' \wedge M$ and B as $B = M \wedge B'$, and their join is a multiple of $A' \wedge M \wedge B' = A \wedge B' = A' \wedge B$. This gives the relationship between meet and join.

³We should also say that there are some issues currently being resolved to make meet and join a properly embedded part of geometric algebra since they produces blades modulo a multiplicative scaling factor rather than actual blades. Most literature now uses them only in projective geometry, in which there is no problem. For the general case we refer to the current work of Bouma [1].

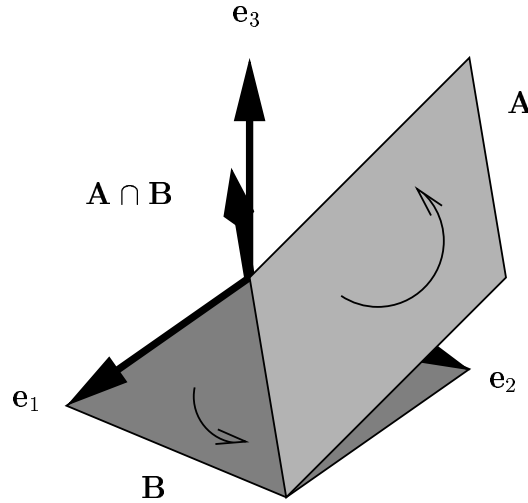


Figure 4: An example of the meet

Given the join $\mathbf{J} \equiv \mathbf{A} \cup \mathbf{B}$ of \mathbf{A} and \mathbf{B} , we can compute their meet by the property that its dual (with respect to the join) is the outer product of their duals (this is a not-so-obvious consequence of the required ‘containment in both’). In formula, this is:

$$(\mathbf{A} \cap \mathbf{B}) \cdot \tilde{\mathbf{J}} = (\mathbf{B} \cdot \tilde{\mathbf{J}}) \wedge (\mathbf{A} \cdot \tilde{\mathbf{J}}) \quad \text{or} \quad (\mathbf{A} \cap \mathbf{B})^* = \mathbf{B}^* \wedge \mathbf{A}^*$$

with the dual taken with respect to the join \mathbf{J} . (The somewhat strange order is a consequence of the factorization chosen above, and it corresponds to [17] for vectors). This leads to a formula for the meet of \mathbf{A} and \mathbf{B} relative to the chosen join (use eq.(19)) :

$$\mathbf{A} \cap \mathbf{B} = (\mathbf{B} \cdot \tilde{\mathbf{J}}) \cdot \mathbf{A}. \quad (24)$$

Let us do an example: the intersection of two planes represented by the 2-blades $\mathbf{A} = \frac{1}{2}(\mathbf{e}_1 + \mathbf{e}_2) \wedge (\mathbf{e}_2 + \mathbf{e}_3)$ and $\mathbf{B} = \mathbf{e}_1 \wedge \mathbf{e}_2$. Note that we have normalized them (this is not necessary, but convenient for a point we want to make later). These are planes in general position in 3-dimensional space, so their join is proportional to \mathbf{I}_3 . It makes sense to take $\mathbf{J} = \mathbf{I}_3$. This gives for the meet:

$$\begin{aligned} \mathbf{A} \cap \mathbf{B} &= \frac{1}{2}((\mathbf{e}_1 \wedge \mathbf{e}_2) \cdot (\mathbf{e}_3 \wedge \mathbf{e}_2 \wedge \mathbf{e}_1)) \cdot ((\mathbf{e}_1 + \mathbf{e}_2) \wedge (\mathbf{e}_2 + \mathbf{e}_3)) \\ &= \frac{1}{2}\mathbf{e}_3 \cdot ((\mathbf{e}_1 + \mathbf{e}_2) \wedge \mathbf{e}_3) \\ &= -\frac{1}{2}(\mathbf{e}_1 + \mathbf{e}_2) = -\frac{1}{\sqrt{2}}\left(\frac{\mathbf{e}_1 + \mathbf{e}_2}{\sqrt{2}}\right) \end{aligned} \quad (25)$$

(the last step expresses the result in normalized form, for a point we make below). Figure 4 shows the answer; as in [17] the sign of $\mathbf{A} \cap \mathbf{B}$ is the right-hand rule applied to the turn required to make \mathbf{A} coincide with \mathbf{B} , in the correct orientation.

Classically, one computes the intersection of two planes in 3-space by first converting them to normal vectors, and then taking the cross product. We can see that this gives the same answer in this non-degenerate case in 3-space, using our previous equations eq.(18), eq.(19), and noting that $\tilde{\mathbf{I}}_3 = -\mathbf{I}_3$:

$$\begin{aligned}
 (\mathbf{A} \cdot \tilde{\mathbf{I}}_3) \times (\mathbf{B} \cdot \tilde{\mathbf{I}}_3) &= ((\mathbf{A} \cdot \tilde{\mathbf{I}}_3) \wedge (\mathbf{B} \cdot \tilde{\mathbf{I}}_3)) \cdot \tilde{\mathbf{I}}_3 \\
 &= ((\mathbf{B} \cdot \tilde{\mathbf{I}}_3) \wedge (\mathbf{A} \cdot \tilde{\mathbf{I}}_3)) \cdot \mathbf{I}_3 \\
 &= (\mathbf{B} \cdot \tilde{\mathbf{I}}_3) \cdot ((\mathbf{A} \cdot \tilde{\mathbf{I}}_3) \cdot \mathbf{I}_3) \\
 &= (\mathbf{B} \cdot \tilde{\mathbf{I}}_3) \cdot (\mathbf{A} \wedge (\tilde{\mathbf{I}}_3 \cdot \mathbf{I}_3)) \\
 &= (\mathbf{B} \cdot \tilde{\mathbf{I}}_3) \cdot \mathbf{A}.
 \end{aligned}$$

So the classical result is a special case of eq.(24), but that formula is much more general: it applies to the intersection of subspaces of *any* grade, within a space of *any* dimension. With it, we begin to see some of the potential power of geometric algebra.

When the meet is a scalar, the two subspaces intersect in the point at the origin. This is in agreement with our geometrical interpretation in Section 2.7 of scalars as the weighted point at the origin. Scalars are geometrical objects, too!

The norm of the meet gives an impression of the ‘strength’ of the intersection. Between normalized subspaces in Euclidean space, the magnitude of the meet is the sine of the angle between them. From numerical analysis, this is a well-known measure for the ‘distance’ between subspaces in terms of their orthogonality: it is 1 if the spaces are orthogonal, and decays gracefully to 0 as the spaces get more parallel, before changing sign. This numerical significance is very useful in applications.

5 Dividing by subspaces

With subspaces as basic elements of computation, we would really like to complete our algebra by the ability to solve equations in similarity problems such as indicated in Figure 5:

Given two vectors \mathbf{a} and \mathbf{b} , and a third vector \mathbf{c} , determine \mathbf{x} so that \mathbf{x} is to \mathbf{c} as \mathbf{b} is to \mathbf{a} , i.e. solve (in a symbolic notation which we will soon make exact):

$$\frac{\mathbf{x}}{\mathbf{c}} = \frac{\mathbf{b}}{\mathbf{a}} \quad (26)$$

Such equations require a *division* of subspaces (here vectors), and so, really, an invertible product of subspaces. This *geometric product* is at the core of geometric algebra, and it is a rather amazing construction, at first sight.

5.1 The geometric product

For vectors, the geometric product is defined in terms of the inner and outer product as:

$$\mathbf{a} \mathbf{b} \equiv \mathbf{a} \cdot \mathbf{b} + \mathbf{a} \wedge \mathbf{b} \quad (27)$$

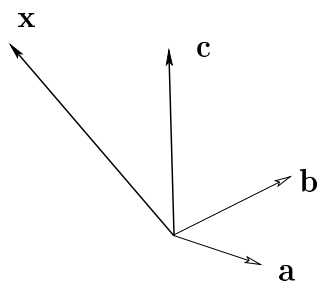


Figure 5: Ratios of vectors

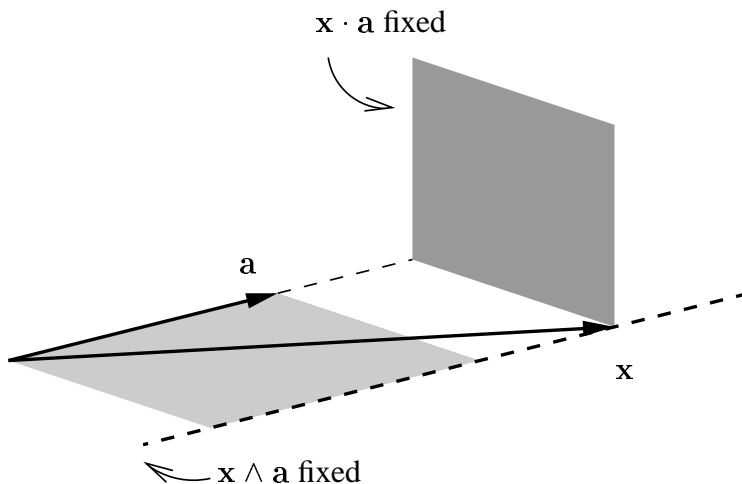


Figure 6: Invertibility of the geometric products.

So the geometric product of two vectors is an element of mixed grade: it has a scalar (0-blade) part $a \cdot b$ and a 2-blade part $a \wedge b$. It is therefore *not* a blade; rather, it is an operator on blades (as we will soon show). Changing the order of a and b gives:

$$b a \equiv b \cdot a + b \wedge a = a \cdot b - a \wedge b$$

The geometric product of two vectors is therefore neither fully symmetric (or rather: commutative), nor fully anti-symmetric.

A simple drawing may convince you that the geometric product is indeed invertible, whereas the inner and outer product separately are not. In Figure 6, we have a given vector a . We denote the set of vectors x with the same value of the inner product $x \cdot a$ – this is a plane perpendicular to a . The set of all vectors with the same value of the outer product $x \wedge a$ is also denoted – this is the line of all points which span the same directed area with a . Neither of these sets is a singleton (in spaces of more than 1 dimension), so the inner and outer products are not fully invertible. The geometric product provides both the plane and the line, and therefore permits determining their unique intersection x , as illustrated in the figure. Therefore it is invertible.

Note that the geometric product is sensitive to the relative directions of the vectors: for parallel vectors \mathbf{a} and \mathbf{b} , the outer product contribution is zero, and $\mathbf{a}\mathbf{b}$ is a scalar and commutative in its factors; for perpendicular vectors, $\mathbf{a}\mathbf{b}$ is a 2-blade, and anti-commutative. In general, if the angle between \mathbf{a} and \mathbf{b} is ϕ in their common plane with unit 2-blade \mathbf{I} , we can write (in a Euclidean space):

$$\mathbf{a}\mathbf{b} = |\mathbf{a}||\mathbf{b}|(\cos\phi + \mathbf{I}\sin\phi) \quad (28)$$

We will see below that $\mathbf{I}\mathbf{I} = -1$, so this is very reminiscent of complex numbers. More about that later, we mention it here to make the construction of the different grade elements in eq.(27) somewhat less outrageous than it may appear at first.

Eq.(27) defines the geometric product *only for vectors*. For arbitrary elements of our algebra it is defined using linearity and associativity, and making it coincide with the usual scalar product in the vector space, as the notation already suggests. That gives the following axioms (where α and β are scalars, \mathbf{x} is a vector, A is a general element of the algebra):

$$\text{scalars} \quad \alpha\beta \text{ and } \alpha\mathbf{x} \text{ have their usual meaning in } V^m \quad (29)$$

$$\text{scalars commute} \quad \alpha A = A\alpha \quad (30)$$

$$\text{vectors} \quad \mathbf{x}A = \mathbf{x} \cdot A + \mathbf{x} \wedge A \quad (31)$$

$$\text{associativity} \quad A(BC) = (AB)C \quad (32)$$

$$\text{distributivity 1} \quad A(B+C) = AB+AC \quad (33)$$

$$\text{distributivity 2} \quad (A+B)C = AC+BC \quad (34)$$

(One can avoid the reference to the inner and outer product through replacing eq.(31) by ‘the square of a vector \mathbf{x} must be equal to the scalar $Q(\mathbf{x}, \mathbf{x})$ ’, with Q the bilinear form of the vector space. Then one can re-introduce inner and outer product through the commutative properties of the geometric product:

$$\mathbf{a} \cdot \mathbf{b} = \frac{1}{2}(\mathbf{a}\mathbf{b} + \mathbf{b}\mathbf{a}) \quad \text{and} \quad \mathbf{a} \wedge \mathbf{b} = \frac{1}{2}(\mathbf{a}\mathbf{b} - \mathbf{b}\mathbf{a}). \quad (35)$$

This is mathematically cleaner, but too indirect for our purpose here.)

It may not be obvious that these equations give enough information to compute the geometric product of arbitrary elements. Rather than show this abstractly, let us show by example how the rules can be used to develop the geometric algebra of 3-dimensional Euclidean space. We introduce, for convenience only, an orthonormal basis $\{\mathbf{e}_i\}_{i=1}^3$. Since this implies that $\mathbf{e}_i \cdot \mathbf{e}_j = \delta_{ij}$, we get the commutation rules:

$$\mathbf{e}_i\mathbf{e}_j = \begin{cases} -\mathbf{e}_j\mathbf{e}_i & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases} \quad (36)$$

In fact, the former is equal to $\mathbf{e}_i \wedge \mathbf{e}_j$, whereas the latter equals $\mathbf{e}_i \cdot \mathbf{e}_i$. Considering the unit 2-blade $\mathbf{e}_1 \wedge \mathbf{e}_2$, we find for its square:

$$\begin{aligned} (\mathbf{e}_i \wedge \mathbf{e}_j)^2 &= (\mathbf{e}_i \wedge \mathbf{e}_j)(\mathbf{e}_i \wedge \mathbf{e}_j) = (\mathbf{e}_i \mathbf{e}_j)(\mathbf{e}_i \mathbf{e}_j) \\ &= \mathbf{e}_i \mathbf{e}_j \mathbf{e}_i \mathbf{e}_j = -\mathbf{e}_i \mathbf{e}_i \mathbf{e}_j \mathbf{e}_j = -1 \end{aligned} \quad (37)$$

\mathcal{C}_3	1	\mathbf{e}_1	\mathbf{e}_2	\mathbf{e}_3	\mathbf{e}_{12}	\mathbf{e}_{31}	\mathbf{e}_{23}	\mathbf{e}_{123}
1	1	\mathbf{e}_1	\mathbf{e}_2	\mathbf{e}_3	\mathbf{e}_{12}	\mathbf{e}_{31}	\mathbf{e}_{23}	\mathbf{e}_{123}
\mathbf{e}_1	\mathbf{e}_1	1	\mathbf{e}_{12}	$-\mathbf{e}_{31}$	\mathbf{e}_2	$-\mathbf{e}_3$	\mathbf{e}_{123}	\mathbf{e}_{23}
\mathbf{e}_2	\mathbf{e}_2	$-\mathbf{e}_{12}$	1	\mathbf{e}_{23}	$-\mathbf{e}_1$	\mathbf{e}_{123}	\mathbf{e}_3	\mathbf{e}_{31}
\mathbf{e}_3	\mathbf{e}_3	\mathbf{e}_{31}	$-\mathbf{e}_{23}$	1	\mathbf{e}_{123}	\mathbf{e}_1	$-\mathbf{e}_2$	\mathbf{e}_{12}
\mathbf{e}_{12}	\mathbf{e}_{12}	$-\mathbf{e}_2$	\mathbf{e}_1	\mathbf{e}_{123}	-1	\mathbf{e}_{23}	$-\mathbf{e}_{31}$	$-\mathbf{e}_3$
\mathbf{e}_{31}	\mathbf{e}_{31}	\mathbf{e}_3	\mathbf{e}_{123}	$-\mathbf{e}_1$	$-\mathbf{e}_{23}$	-1	\mathbf{e}_{12}	$-\mathbf{e}_2$
\mathbf{e}_{23}	\mathbf{e}_{23}	\mathbf{e}_{123}	$-\mathbf{e}_3$	\mathbf{e}_2	\mathbf{e}_{31}	$-\mathbf{e}_{12}$	-1	$-\mathbf{e}_1$
\mathbf{e}_{123}	\mathbf{e}_{123}	\mathbf{e}_{23}	\mathbf{e}_{31}	\mathbf{e}_{12}	$-\mathbf{e}_3$	$-\mathbf{e}_2$	$-\mathbf{e}_1$	-1

Figure 7: The multiplication table of the geometric algebra of 3-dimensional Euclidean space, on an orthonormal basis. Shorthand: $\mathbf{e}_{12} = \mathbf{e}_1 \wedge \mathbf{e}_2$, etcetera.

So a unit 2-blade squares to -1 (we just computed for $\mathbf{e}_1 \wedge \mathbf{e}_2$ for convenience, but there is nothing exceptional about that particular unit 2-blade, since the basis was arbitrary). Continued application of eq.(36) gives the full multiplication for all basis elements in the Clifford algebra of 3-dimensional Euclidean space. The resulting multiplication table is given in Figure 7. Arbitrary elements are expressible as a linear combination of these basis elements, so this table determines the full algebra.

5.2 Solving geometric equations

The geometric product is invertible, so ‘dividing by a vector’ has a unique meaning. We usually do this through ‘multiplication by the inverse of the vector’. Since multiplication is not necessarily commutative, we have to be a bit careful: there is a ‘left division’ and a ‘right division’. As you may verify, the unique inverse of a vector \mathbf{a} is

$$\mathbf{a}^{-1} = \frac{\mathbf{a}}{\mathbf{a} \cdot \mathbf{a}}$$

since that is the unique element that satisfies: $\mathbf{a}^{-1} \mathbf{a} = 1 = \mathbf{a} \mathbf{a}^{-1}$. In general, a blade \mathbf{A} has the inverse

$$\mathbf{A}^{-1} = \frac{\widetilde{\mathbf{A}}}{\mathbf{A} \cdot \widetilde{\mathbf{A}}}$$

where $\widetilde{\mathbf{A}}$ is the *reverse* of \mathbf{A} , obtained by switching its spanning factors: if $\mathbf{A} = \mathbf{a}_1 \wedge \mathbf{a}_2 \wedge \cdots \wedge \mathbf{a}_k$, then $\widetilde{\mathbf{A}} = \mathbf{a}_k \wedge \cdots \wedge \mathbf{a}_2 \wedge \mathbf{a}_1$. The reverse of \mathbf{A} differs from \mathbf{A} by a sign $(-1)^{\frac{1}{2}k(k-1)}$. You may verify that $\mathbf{A} \cdot \widetilde{\mathbf{A}}$ is a scalar (and in Euclidean space, even a positive scalar, which can be considered as the ‘norm squared’ of \mathbf{A} ; if it is zero, the blade \mathbf{A} has no inverse, but this does not happen in Euclidean vector spaces).

Invertibility is a great help in solving geometric problems in a closed coordinate-free computational form. The common procedure is as follows: we know certain defining properties of

objects in the usual terms of perpendicularity, spanning, rotations etcetera. These give equations typically expressed using the derived products. You combine these equations algebraically, with the goal of finding an expression for the unknown object involving only the geometric product; then division (permitted by the invertibility of the geometric product) should provide the result.

Let us illustrate this by an example. Suppose we want to find the component \mathbf{x}_\perp of a vector \mathbf{x} perpendicular to a vector \mathbf{a} . The perpendicularity demand is clearly

$$\mathbf{x}_\perp \cdot \mathbf{a} = 0.$$

A second demand is required to relate the magnitude of \mathbf{x}_\perp to that of \mathbf{x} . Some practice in ‘seeing subspaces’ in geometrical problems reveals that the area spanned by \mathbf{x} and \mathbf{a} is the same as the area spanned by \mathbf{x}_\perp and \mathbf{a} , see Figure 8a. This is expressed using the outer product:

$$\mathbf{x}_\perp \wedge \mathbf{a} = \mathbf{x} \wedge \mathbf{a}.$$

These two equations should be combined to form a geometric product. In this example, it is clear that just adding them works, yielding

$$\mathbf{x}_\perp \cdot \mathbf{a} + \mathbf{x}_\perp \wedge \mathbf{a} = \mathbf{x}_\perp \mathbf{a} = \mathbf{x} \wedge \mathbf{a}.$$

This one equation contains the full geometric relationship between \mathbf{x} , \mathbf{a} and the unknown \mathbf{x}_\perp . Geometric algebra solves this equation through division on the right by \mathbf{a} :

$$\mathbf{x}_\perp = (\mathbf{x} \wedge \mathbf{a}) / \mathbf{a} = (\mathbf{x} \wedge \mathbf{a}) \mathbf{a}^{-1}. \quad (38)$$

We rewrote the division by \mathbf{a} as multiplication by the subspace \mathbf{a}^{-1} to show clearly that we mean ‘division on the right’.

This is an example of how the indefinite shape $\mathbf{x} \wedge \mathbf{a}$ spanned by the outer product is just the right element to generate a perpendicular to a vector \mathbf{a} in its plane, through the geometric product. Note that eq.(38) agrees with the well-known expression of \mathbf{x}_\perp using the inner product of vectors:

$$\mathbf{x}_\perp = (\mathbf{x} \wedge \mathbf{a}) \mathbf{a}^{-1} = (\mathbf{x} \mathbf{a} - \mathbf{x} \cdot \mathbf{a}) \mathbf{a}^{-1} = \mathbf{x} - (\mathbf{x} \cdot \frac{\mathbf{a}}{\|\mathbf{a}\|}) \frac{\mathbf{a}}{\|\mathbf{a}\|}. \quad (39)$$

The geometric algebra expression using outer product and inverse generalizes immediately to arbitrary subspaces \mathbf{A} .

5.3 Projection of subspaces

The availability of an inverse thus gives us an interesting way of decomposing a vector \mathbf{x} relative to a given blade \mathbf{A} using the geometric product:

$$\mathbf{x} = (\mathbf{x} \mathbf{A}) \mathbf{A}^{-1} = (\mathbf{x} \cdot \mathbf{A}) \mathbf{A}^{-1} + (\mathbf{x} \wedge \mathbf{A}) \mathbf{A}^{-1} \quad (40)$$

The first term is a blade fully inside \mathbf{A} : it is the *projection* of \mathbf{x} onto \mathbf{A} . The second term is a vector perpendicular to \mathbf{A} , sometimes called the *rejection* of \mathbf{x} by \mathbf{A} . The projection of a blade \mathbf{X} onto a blade \mathbf{A} is given by the extension of the above, as:

$$\text{projection of } \mathbf{X} \text{ onto } \mathbf{A}: \quad \mathbf{X} \mapsto (\mathbf{X} \cdot \mathbf{A}) \mathbf{A}^{-1}$$

Again geometric algebra has allowed a straightforward extension to arbitrary dimensions of subspaces, without additional computational complexity.

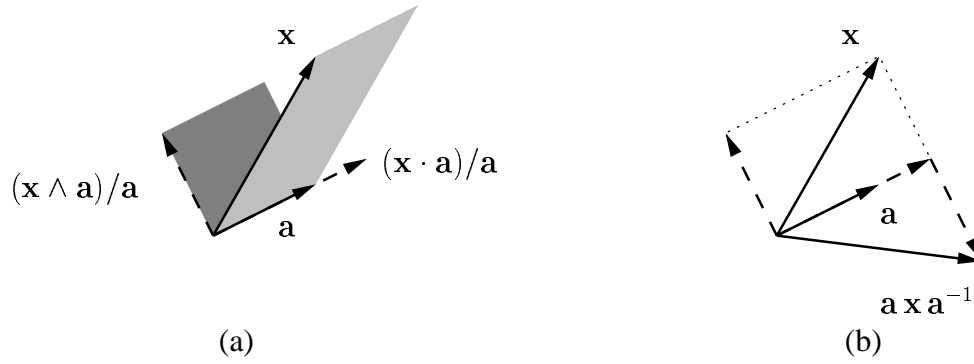


Figure 8: (a) Projection and rejection of x relative to a . (b) Reflection of x in a .

5.4 Reflection of subspaces

The *reflection* of a vector x relative to a fixed vector a can be constructed from the decomposition of eq.(40) (used for a vector a), by changing the sign of the rejection (see Figure 8b). This can be rewritten in terms of the geometric product:

$$(x \cdot a) a^{-1} - (x \wedge a) a^{-1} = (a \cdot x + a \wedge x) a^{-1} = a x a^{-1}.$$

So the reflection of x in a is the expression $a x a^{-1}$, see Figure 8b; the reflection in a plane perpendicular to a is then $-a x a^{-1}$,

We can extend this formula to the reflection of a blade X relative to the vector a , this is simply:

$$\text{reflection in vector } a: X \mapsto a X a^{-1}.$$

and even to the reflection of a blade X in a k -blade A , which turns out to be:

$$\text{general reflection: } X \mapsto -(-1)^k A X A^{-1}.$$

Note that these formulas permit you to do reflections of subspaces without first decomposing them in constituent vectors. It gives the possibility of reflection a polyhedral object by directly using a facet representation, rather than acting on individual vertices.

5.5 Angles as geometrical objects

We have found in eq.(37) that any unit 2-blade I in a Euclidean space satisfies $I^2 = -1$, so this is also true for the unit 2-blade occurring in eq.(28). Therefore, using the usual definition of the exponential as a converging series of terms, we are actually permitted to write the geometric product in an exponential form:

$$a b = |a| |b| (\cos \phi + I \sin \phi) = |a| |b| e^{I\phi} \quad (41)$$

with I the unit 2-blade containing a and b , oriented from a to b . This exponential form will be very convenient when we do rotations. Note that all elements occurring in this equation have a

straightforward geometrical interpretation, we are not doing complex numbers here! (Really, we aren't: \mathbf{I} is not a complex scalar, since then it would have to commute with *all* elements of the algebra by eq.(30), but it instead satisfies $\mathbf{a} \mathbf{I} = -\mathbf{I} \mathbf{a}$ for vectors \mathbf{a} in the \mathbf{I} -plane.)

The combination $\mathbf{I}\phi$ is a full indication of the angle between the two vectors: it denotes not only the magnitude, but also the plane in which the angle is measured, and even the orientation of the angle. If you ask for the scalar magnitude of the geometrical quantity $\mathbf{I}\phi$ in the plane $-\mathbf{I}$ (the plane 'from \mathbf{b} to \mathbf{a} ' rather than 'from \mathbf{a} to \mathbf{b} '), it is $-\phi$; so the scalar value of the angle automatically gets the right sign. The fact that the angle as expressed by $\mathbf{I}\phi$ is now a geometrical quantity independent of the convention used in its definition removes a major headache from many geometrical computations involving angles. We call this true geometric quantity the *bivector angle* (it is just a 2-blade, of course, not a new kind of element – but we use it as an angle, hence the name).

5.6 Rotations in the plane

Using the inverse of a vector, we can now solve the motivating problem of eq.(26), to find a vector \mathbf{x} that is to \mathbf{c} as \mathbf{b} is to \mathbf{a} . Denoting the 2-blade of the $(\mathbf{a} \wedge \mathbf{b})$ -plane by \mathbf{I} , we obtain:

$$\mathbf{x} \mathbf{c}^{-1} = \mathbf{b} \mathbf{a}^{-1}$$

so that

$$\mathbf{x} = (\mathbf{b} \mathbf{a}^{-1}) \mathbf{c} = \frac{|\mathbf{b}|}{|\mathbf{a}|} e^{-\mathbf{I}\phi} \mathbf{c} \quad (42)$$

Here $\mathbf{I}\phi$ is the angle in the \mathbf{I} plane from \mathbf{a} to \mathbf{b} , as in eq.(41), so $-\mathbf{I}\phi$ is the angle from \mathbf{b} to \mathbf{a} . If we happen to have $|\mathbf{a}| = |\mathbf{b}|$, we get $\mathbf{x} = e^{-\mathbf{I}\phi} \mathbf{c}$; apparently we should interpret 'pre-multiplying by $e^{-\mathbf{I}\phi}$ ' as a *rotation operator* in the \mathbf{I} -plane. The full expression of eq.(42) denotes a rotation/dilation in the \mathbf{I} -plane.

Let us write this out, to get familiar with the geometric algebra way of looking at rotations:

$$e^{-\mathbf{I}\phi} \mathbf{c} = \mathbf{c} \cos \phi - \mathbf{I} \mathbf{c} \sin \phi = \mathbf{c} \cos \phi + \mathbf{c} \mathbf{I} \sin \phi$$

What is $\mathbf{c} \mathbf{I}$? Introduce orthonormal coordinates $\{\mathbf{e}_1, \mathbf{e}_2\}$ in the \mathbf{I} -plane, with \mathbf{e}_1 along \mathbf{c} , so that $\mathbf{c} \equiv c \mathbf{e}_1$. Then $\mathbf{I} = \mathbf{e}_1 \wedge \mathbf{e}_2 = \mathbf{e}_1 \mathbf{e}_2$. Therefore $\mathbf{c} \mathbf{I} = c \mathbf{e}_1 \mathbf{e}_1 \mathbf{e}_2 = c \mathbf{e}_2$: it is \mathbf{c} *turned over a right angle*, following the orientation of the 2-blade \mathbf{I} (here anti-clockwise). So $\mathbf{c} \cos \phi + \mathbf{c} \mathbf{I} \sin \phi$ is 'a bit of \mathbf{c} plus a bit of its anti-clockwise perpendicular' – and those amounts are precisely right to make it equal to the rotation by ϕ , see Figure 9.

If you use a classical rotation matrix in 2 dimensions, it does precisely this construction, but in a coordinate system that is adapted to an arbitrary basis $\{\mathbf{e}_1, \mathbf{e}_2\}$, rather than to \mathbf{c} . That is why you then need 4 coefficients, to describe how each of those 2 basis vectors turns. Geometric algebra is coordinate-free in this sense: orthogonal directions can be made from the vectors for which you need them in a coordinate-free manner. Then a specification of the rotation requires only 2 trigonometric functions, just for the scaling of those 2 components.

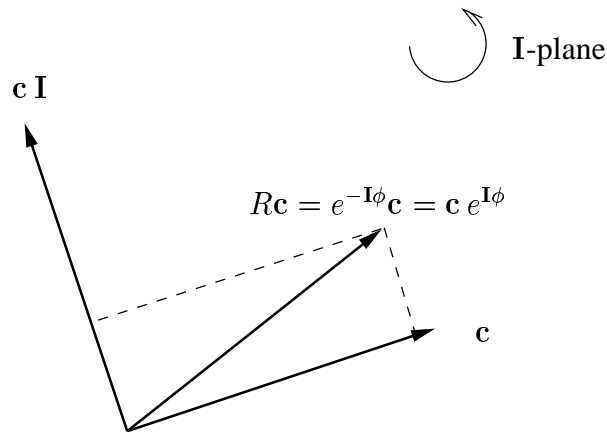


Figure 9: Coordinate-free specification of rotation.

5.7 Rotations in 3 dimensions

Two subsequent reflections in lines which make an angle of $\phi/2$ in a plane with unit 2-blade I constitute a rotation over ϕ in the I -plane. In 2-dimensional space, this is obvious, but it also works in 3-dimensional space, see Figure 10 (and even in m -dimensional space). It gives us the way to express general rotations in geometric algebra.

Two successive reflections of a vector x in vectors u and v give

$$v(u x u^{-1})v^{-1} = \frac{v}{|v|} \frac{u}{|u|} x \frac{u}{|u|} \frac{v}{|v|} = e^{-I\phi/2} x e^{I\phi/2}$$

where we used the exponential notation for the geometric product of two unit vectors (I is the unit 2-blade from u to v). The expression for the rotation is therefore directly given by the bivector angle, i.e. by angle and rotation plane. An operator $e^{-I\phi/2}$, used in this way, is called a *rotor*. Writing out this expression in terms of the perpendicular component x_{\perp} (rejection) and the parallel component x_{\parallel} (projection) of x relative to the I plane gives

$$\text{rotation over } I\phi: x \mapsto e^{-I\phi/2} x e^{I\phi/2} = x_{\perp} + e^{-I\phi} x_{\parallel} \quad (43)$$

(this is a good exercise, it requires $I x_{\perp} = x_{\perp} I$ and $I x_{\parallel} = -x_{\parallel} I$; why do these hold?). So the perpendicular component to the rotation plane is unchanged (as it should!), and the parallel component becomes pre-multiplied by $e^{-I\phi}$. We have seen in eq.(42) that this is a rotation in the I -plane. (In fact, we could have defined the higher dimensional rotation by the right hand side of eq.(43) and then derived the left hand side.)

5.8 Combining rotations

Two successive rotations R_1 and R_2 are equivalent to a single new rotation R of which the rotor R is the geometric product of the rotors R_1 and R_2 , since

$$R_2 R_1 x R_1^{-1} R_2^{-1} = (R_2 R_1) x (R_2 R_1)^{-1} \equiv R x R^{-1}.$$

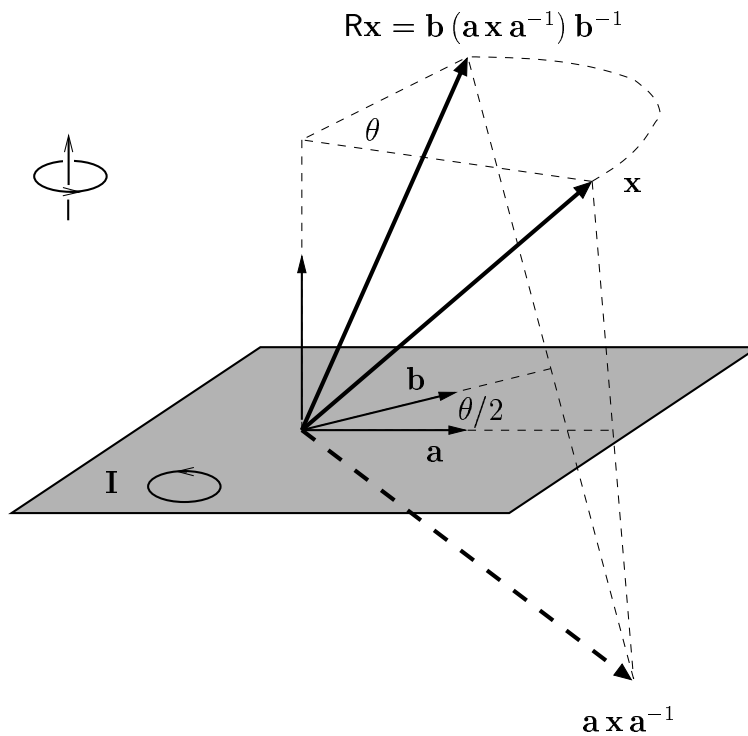


Figure 10: A rotation as 2 reflections in vectors \mathbf{u} and \mathbf{v} , making an angle of $\mathbf{I}\phi/2$.

This applies in 3-dimensional space as well as in 2-dimensional space. Therefore the combination of rotations is a simple consequence of the definition of the geometric product on rotors, i.e. elements of the form $e^{-\mathbf{I}\phi/2} = \cos \phi/2 - \mathbf{I} \sin \phi/2$, with $\mathbf{I}^2 = -1$. (We could allow a scalar factor in the rotor, since the inverse divides it out; yet it is common to restrict rotor to be normalized to unity – then one can replace R^{-1} by \tilde{R} , defining the rotation by $R \mathbf{x} \tilde{R}$. Reversion is a simpler (cheaper) operation than inversion, though the normalization may add some additional computational cost.)

Let's see how it works in 3-space. In 3 dimensions, we are used to specifying rotations by a *rotation axis* \mathbf{a} rather than by a *rotation plane* \mathbf{I} . The relationship between axis and plane is given by duality: $\mathbf{a} \equiv \mathbf{I} \cdot \tilde{\mathbf{I}}_3 = -\mathbf{I} \mathbf{I}_3$ (check that this indeed gives the correct orientation). Given the axis \mathbf{a} , we therefore find the plane as the 2-blade $\mathbf{I} = -\mathbf{a} \mathbf{I}_3^{-1} = \mathbf{a} \mathbf{I}_3 = \mathbf{I}_3 \mathbf{a}$. A rotation over an angle ϕ around an axis with unit vector \mathbf{a} is therefore represented by the rotor $e^{-\mathbf{I}_3 \mathbf{a} \phi/2}$.

To compose a rotation R_1 around the \mathbf{e}_1 axis of $\pi/2$ with a subsequent rotation R_2 over the \mathbf{e}_2 axis over $\pi/2$, we write out their rotors:

$$R_1 = e^{-\mathbf{I}_3 \mathbf{e}_1 \pi/4} = \frac{1 - \mathbf{e}_{23}}{\sqrt{2}} \quad \text{and} \quad R_2 = e^{-\mathbf{I}_3 \mathbf{e}_2 \pi/4} = \frac{1 - \mathbf{e}_{31}}{\sqrt{2}}$$

The total rotor is their product, and we rewrite it back to the exponential form to find the axis:

$$R \equiv R_2 R_1 = \frac{1}{2}(1 - \mathbf{e}_{31})(1 - \mathbf{e}_{23}) = \frac{1}{2}(1 - \mathbf{e}_{23} - \mathbf{e}_{31} + \mathbf{e}_{12})$$

$$= \frac{1}{2} - \frac{1}{2}\sqrt{3} \mathbf{I}_3 \frac{\mathbf{e}_1 + \mathbf{e}_2 - \mathbf{e}_3}{\sqrt{3}} = e^{-\mathbf{I}_3 \mathbf{a} \pi / 3}$$

Therefore the total rotation is over the axis $\mathbf{a} = (\mathbf{e}_1 + \mathbf{e}_2 - \mathbf{e}_3)/\sqrt{3}$, over the angle $2\pi/3$. But of course you do not need to decompose the resulting rotor into those geometrical constituents: you can apply it immediately to a vector \mathbf{x} as $R \mathbf{x} R^{-1}$, or even to an arbitrary blade through the formula:

$$\text{general rotation: } \mathbf{X} \mapsto R \mathbf{X} R^{-1}$$

This enables you to rotate a plane in one operation, for instance:

$$R(\mathbf{e}_1 \wedge \mathbf{e}_2)R^{-1} = \frac{1}{4}(1 - \mathbf{e}_{23} - \mathbf{e}_{31} + \mathbf{e}_{12}) \mathbf{e}_{12} (1 + \mathbf{e}_{23} + \mathbf{e}_{31} - \mathbf{e}_{12}) = -\mathbf{e}_{31}$$

No need to decompose the plane into its spanning vectors first!

5.9 Quaternions: based on bivectors

You may have recognized the example above as strongly similar to quaternion computations. Quaternions are indeed part of geometric algebra, in the following straightforward manner.

Choose an orthonormal basis $\{\mathbf{e}_i\}_{i=1}^3$. Construct out of that a bivector basis with elements $\mathbf{e}_{12} \equiv \mathbf{e}_1 \wedge \mathbf{e}_2 (= \mathbf{e}_1 \mathbf{e}_2)$ and cyclic. Note that these elements satisfy: $\mathbf{e}_{12}^2 = \mathbf{e}_{23}^2 = \mathbf{e}_{31}^2 = -1$, and $\mathbf{e}_{12} \mathbf{e}_{23} = \mathbf{e}_{13}$ (and cyclic) and also $\mathbf{e}_{12} \mathbf{e}_{23} \mathbf{e}_{31} = 1$. In fact, setting $i \equiv \mathbf{e}_{23}$, $j \equiv -\mathbf{e}_{31}$ and $k \equiv \mathbf{e}_{12}$, we find $i^2 = j^2 = k^2 = i j k = -1$ and $j i = k$ and cyclic. Algebraically these objects are the quaternions obeying the quaternion product, commonly interpreted as some kind of ‘4-D complex number system’. There is nothing ‘complex’ about quaternions; but they are not really vectors either (as some still think) – they are just real 2-blades in 3-space, denoting elementary rotation planes, and multiplying through the geometric product. Visualizing quaternions is therefore straightforward: each is just a rotation plane with a rotation angle, and the ‘bivector angle’ concept represents that well (the corresponding quaternion is simply its exponential, elevating the bivector angle to a rotation operator).

5.10 Constructing rotors

For a 2-dimensional rotation, if you know for certain that a (unit) vector \mathbf{e} has been rotated to become a vector \mathbf{f} (which therefore necessarily has the same norm) by a rotation in the $\mathbf{e} \wedge \mathbf{f}$ -plane, it is easy to find a rotor that does that:

$$R = 1 + \mathbf{f} \mathbf{e}$$

(if you want the unit rotor, you need to normalize this). For a 3-dimensional rotation, if you know an orthonormal frame $\{\mathbf{e}_i\}_{i=1}^3$ which has rotated to the frame $\{\mathbf{f}_i\}_{i=1}^3$, then a rotor doing that is:

$$R = 1 + \mathbf{f}_1 \mathbf{e}_1 + \mathbf{f}_2 \mathbf{e}_2 + \mathbf{f}_3 \mathbf{e}_3$$

(which needs to be normalized if you want a unit rotor). This formula can be generalized simply to non-orthonormal frames, see [15]. Warning: the formulas do not work for rotations over π (there is then no unique rotation plane!) – but are very useful elsewhere.

6 Differentiation with respect to blades

Geometric algebra also has a much extended operation of differentiation, which contains the classical vector calculus, and much more. It is possible to differentiate with respect to a scalar or a vector, as before, but now also with respect to k -blades. This enables efficient encoding of differential geometry, in a coordinate-free manner, and gives an alternative look at differential shape descriptors like the ‘second fundamental form’ (it becomes an immediate indication of how the tangent plane changes when we slide along the surface).

Somebody should rewrite classical differential geometry texts into geometric algebra; but this has not been done yet and it would lead too far to do so in this introductory paper. Let us just briefly show the scalar differentiation of a rotor, to demonstrate how the commutation rules of geometric algebra naturally group to a well-known classical result, which is then automatically extended beyond vectors.

So, suppose we have a rotor $R = e^{-\mathbf{I}\phi/2}$, and use it to produce a rotated version $\mathbf{X} = R \mathbf{X}_0 \tilde{R}$ of some constant blade \mathbf{X}_0 . Scalar differentiation with respect to time gives (using chain rule and commutation rules):

$$\begin{aligned} \frac{d}{dt} \mathbf{X} &= \frac{d}{dt} (e^{-\mathbf{I}\phi/2} \mathbf{X}_0 e^{\mathbf{I}\phi/2}) \\ &= -\frac{1}{2} \frac{d}{dt} (\mathbf{I}\phi) (e^{-\mathbf{I}\phi/2} \mathbf{X}_0 e^{\mathbf{I}\phi/2}) + \frac{1}{2} (e^{-\mathbf{I}\phi/2} \mathbf{X}_0 e^{\mathbf{I}\phi/2}) \frac{d}{dt} (\mathbf{I}\phi) \\ &= \frac{1}{2} (\mathbf{X} \frac{d}{dt} (\mathbf{I}\phi) - \frac{d}{dt} (\mathbf{I}\phi) \mathbf{X}) \\ &= \mathbf{X} \times \frac{d}{dt} (\mathbf{I}\phi) \end{aligned}$$

using the commutator product \times defined in geometric algebra as the shorthand $A \times B \equiv \frac{1}{2}(AB - BA)$; this product often crops up in computations with Lie groups such as the rotations. This simple expression which results assumes a more familiar form when \mathbf{X} is a vector \mathbf{x} in 3-space, the rotation plane is fixed so that $\frac{d}{dt} \mathbf{I} = 0$, and we introduce a scalar angular velocity $\omega \equiv \frac{d}{dt} \phi$. It is then common practice to introduce the vector dual to the plane as the angular velocity vector $\boldsymbol{\omega}$, so $\boldsymbol{\omega} \equiv \omega \mathbf{I} \cdot \tilde{\mathbf{I}}_3 = \omega \mathbf{I} / \mathbf{I}_3$. We then obtain:

$$\frac{d}{dt} \mathbf{x} = \mathbf{x} \cdot \frac{d}{dt} (\mathbf{I}\phi) = \mathbf{x} \cdot (\boldsymbol{\omega} \mathbf{I}_3) = (\mathbf{x} \wedge \boldsymbol{\omega}) \mathbf{I}_3 = \boldsymbol{\omega} \times \mathbf{x}$$

where \times is the vector cross product. As before when we treated the meet and other operations, we find that an equally simple geometric algebra expression is much more general; here it describes the differential rotation of k -dimensional subspaces in n -dimensional space, rather than merely of vectors in 3-D.

But in geometric algebra we can also differentiate with respect to vectors. Suppose that we project a vector \mathbf{x} on the unit sphere by the function $\mathbf{x} \mapsto P(\mathbf{x}) = \mathbf{x}/|\mathbf{x}|$. We compute its derivative in the \mathbf{a} direction, denoted as $(\mathbf{a} \cdot \partial_{\mathbf{x}})P(\mathbf{x})$ or $P_{\mathbf{a}}(\mathbf{x})$, as a standard differential quotient and using Taylor series expansion. Note how geometric algebra permits compact expression of the result, with geometrical significance:

$$(\mathbf{a} \cdot \partial_{\mathbf{x}}) \frac{\mathbf{x}}{|\mathbf{x}|} \equiv \lim_{\lambda \rightarrow 0} \frac{1}{\lambda} \left(\frac{\mathbf{x} + \lambda \mathbf{a}}{|\mathbf{x} + \lambda \mathbf{a}|} - \frac{\mathbf{x}}{|\mathbf{x}|} \right) = \lim_{\lambda \rightarrow 0} \frac{1}{\lambda} \left(\frac{\mathbf{x} + \lambda \mathbf{a}}{|\mathbf{x}| \sqrt{1 + 2\lambda \mathbf{a} \cdot \mathbf{x}^{-1}}} - \frac{\mathbf{x}}{|\mathbf{x}|} \right)$$

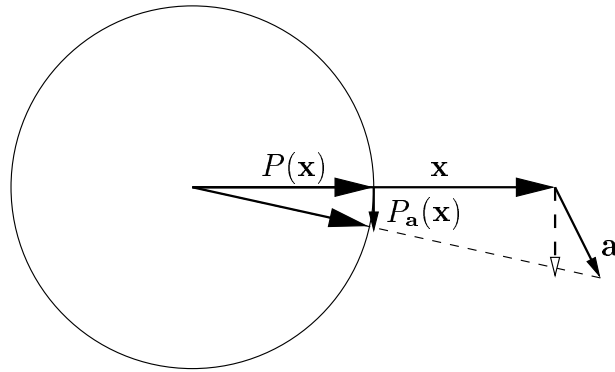


Figure 11: The derivative of the spherical projection.

$$\begin{aligned}
 &= \lim_{\lambda \rightarrow 0} \frac{(\mathbf{x} + \lambda \mathbf{a})(1 - \lambda \mathbf{a} \cdot \mathbf{x}^{-1}) - \mathbf{x}}{\lambda |\mathbf{x}|} = \frac{\mathbf{a} - \mathbf{x}(\mathbf{a} \cdot \mathbf{x}^{-1})}{|\mathbf{x}|} \\
 &= \frac{(\mathbf{a} \wedge \mathbf{x}) \mathbf{x}^{-1}}{|\mathbf{x}|}
 \end{aligned}$$

We recognize the result as the rejection of \mathbf{a} by \mathbf{x} (Section ?? from [?]), scaled appropriately. The sketch of Figure 11 confirms the outcome. You may verify in a similar manner that $(\mathbf{a} \cdot \partial_{\mathbf{x}}) \mathbf{x}^{-1} = -\mathbf{x}^{-1} \mathbf{a} \mathbf{x}^{-1}$, and interpret geometrically.

Similar generalizations result for differentiation relative to general blades or even rotors or linear mappings; the interested reader is referred to the tutorial of [2], which introduces these differentiations using examples from physics, and the application paper [15].

7 Linear algebra

In the classical ways of using vector spaces, linear algebra is an important tool. In geometric algebra, this remains true: linear transformations are of interest in their own right, or as first order approximations to more complicated mappings. Indeed, linear algebra is an integral part of geometric algebra, and acquires much extended coordinate-free methods through this inclusion. We show some of the basic principles; much more may be found in [2] or [13].

7.1 Outermorphisms: spanning is linear

When vectors are transformed by a linear transformation on the vector space, the blades they span can be viewed to transform as well, simply by the rule: ‘the transform of a span of vectors is the span of the transformed vectors’. This means that a linear transformation $f : V^n \rightarrow V^n$ on a vector space has a natural extension to the whole geometric algebra of that vector space, as an *outermorphism*, i.e. a mapping that preserves the outer product structure:

$$f(\mathbf{a}_1 \wedge \mathbf{a}_2 \wedge \cdots \wedge \mathbf{a}_k) \equiv f(\mathbf{a}_1) \wedge f(\mathbf{a}_2) \wedge \cdots \wedge f(\mathbf{a}_k).$$

Note that this is grade-preserving: a k -blade transforms to a k -blade. To this we have to add what the extension does to scalars, which is simply: $f(\alpha) = \alpha$.

This outermorphism definition has immediate consequences. Apply it to a pseudoscalar \mathbf{I}_m , which is an m -blade: it must produce another m -blade. But the linear space of m -blades in m -dimensional vector space is 1-dimensional, so this must again be a multiple of \mathbf{I}_m . That multiple is precisely the determinant of f in m -dimensional space:

$$\det(f) = f(\mathbf{I}_m)\mathbf{I}_m^{-1}.$$

The determinant is thus simply the change of hypervolume under f . This is nothing new, but it is satisfying that all the usual properties of the determinant, including its expression in terms of coordinates, follow immediately from this straightforward, coordinate-free definition.

7.2 Linear transformation of the inner product

The transformation rule for the inner product now follows automatically from the definition through eq.(7), and is found to be rather more involved:

$$f(A \cdot B) = \bar{f}^{-1}(A) \cdot f(B),$$

where \bar{f} is the *adjoint* of f , defined by

$$f(A) * B = A * \bar{f}(B) \quad \text{for all } A \text{ and } B.$$

(In terms of matrices on an orthonormal basis, \bar{f} is the mapping represented by the *transpose* of the matrix representing f .)

7.3 No normal vectors or cross products!

Since the inner product transformation under a linear mapping is so involved, one should steer clear of any constructions that involve the inner product, especially in the characterization of basic properties of one's objects. Therefore the practice of characterizing a plane by its normal vector – which contains the inner product in its duality, see Section 3.3 – should be avoided. Under linear transformations, *the normal vector of a transformed plane is not the transform of the normal vector of the plane!* (this is a well known fact, but always a shock to novices). The normal vector is in fact a cross product of vectors, which (as you may verify from eq.(22) and the above) transforms as:

$$f(\mathbf{a} \times \mathbf{b}) = \bar{f}^{-1}(\mathbf{a}) \times \bar{f}^{-1}(\mathbf{b}) \det(f)$$

and that is usually not equal to $f(\mathbf{a}) \times f(\mathbf{b})$. It is therefore much better to characterize the plane by a 2-blade, now that we can. *The 2-blade of the transformed plane is the transform of the 2-blade of the plane*, since linear transformations are outermorphisms preserving the 2-blade construction. Especially when the planes are tangent planes constructed by differentiation, 2-blades are appropriate: under *any* transformation f , the construction of the tangent plane is

only dependent on the first order linear approximation mapping f of f . Therefore a tangent plane represented as a 2-blade transforms simply under *any* transformation (and the same applies of course to tangent k -blades in higher dimensions). Using blades for those tangent spaces should enormously simplify the treatment of object through differential geometry, especially in the context of affine transformations – but this has not yet been done.

8 All you need is blades: models of geometries

So far we have been treating only homogeneous subspaces of the vector spaces, i.e. subspaces containing the origin. We have spanned them, projected them, and rotated them, but we have not moved them out of the origin to make more interesting geometrical structures such as lines floating in space.

There is a very nice way of making such basic primitives in geometric algebra. At first it looks like a straightforward embedding of the classical ideas behind ‘homogeneous coordinates’, but it rapidly becomes much more powerful than that. It creates an *algebra of points* (rather than vectors). We present three models of Euclidean space, all useful to computer graphics, and show how the geometric algebra of those models implements totally different semantics using the same basic products (but in different spaces). This goes much beyond resolving the issues raised in the classical papers by Goldman [9, 10].

8.1 The vector space model

The most straightforward model of Euclidean space represents its points by the translation vectors required to get there. We call those *position vectors*. This representation strongly depends on the location of the origin. It is well known [9] that this easily leads to bad representations and software which depend heavily on the chosen origin. It is inappropriate to take the position vectors \mathbf{a} and \mathbf{b} as ‘being’ the points A and B , and then form new points by addition of their vectors. The construction $\mathbf{a} + \mathbf{b}$ cannot represent a geometrical point, for its value changes as the origin changes, and no geometrically relevant objects should depend on that.

Still, the vector space model of a Euclidean space is appropriate for translation vectors (the null translation *is* special: it is the identity operation) and for tangent planes to a manifold (again, the origin is special since it is where the tangent space is attached to the manifold). For those, $\mathbf{a} + \mathbf{b}$ has a clear meaning: it is the resultant translation or resultant velocity, of a point. Beyond these applications, one has to be careful with the vector space model.

The products between vectors are just as much part of the model as the embedding of the points themselves (this is a point which Goldman [9, 10] neglects somewhat in his discussion of representations). In the vector space model, they simply have the meaning we have used throughout this paper: the *outer product* constructs the higher-dimensional proper subspaces; the *inner product* constructs the orthogonal complement of subspaces; and the *geometric product* gives us the rotation/dilation operator between subspaces. Elementary combinations of these give us projection and reflection. Note that all these operations are origin-centered in this model: rotations are around an axis through the origin, reflections are in planes through the origin, etcetera. It is

simple to shift them out of the origin of course, but algebraically, that is a ‘hack’ – it would be much more tidy if we could find a representation in which those operations are all elementary relationships between blades (and we will). Even an basic concept like the Euclidean distance between two points P and Q is a fairly involved expression – we have to form $\sqrt{(\mathbf{p} - \mathbf{q}) \cdot (\mathbf{p} - \mathbf{q})}$ to obtain this geometric invariant. It would be much nicer if this elementary concept were one of the elementary products.

The vector space model, then, contains a lot of the basic elements to do Euclidean geometry, especially when we consider its full geometric algebra of higher dimensional subspaces. But we can do better, tidying up the algebra by embedding Euclidean geometry of E^m in a space of more than m dimensions and using the geometric algebra of that space to describe the Euclidean objects and operators of interest.

8.2 The homogeneous model

So far we have been treating only homogeneous subspaces of the vector space V^m , i.e. subspaces containing the origin. We have spanned them, projected them, and rotated them, but we have not moved them out of the origin to make more interesting geometrical structures such as lines floating in space. We construct those now, by extending the thoughts behind ‘homogeneous coordinates’ to geometric algebra. It turns out that these elements of geometry can also be represented by blades, in a representational space with an extra dimension. The geometric algebra of this space then turns out to give us precisely what we need. In this view, more complicated geometrical objects do not require new operations or techniques in geometric algebra, merely the standard computations in a higher dimensional space.

The *homogeneous model* in geometric algebra extends the usual ‘homogeneous coordinates’ model from linear algebra, which works merely on vectors. That model is often described as augmenting a 3-dimensional vector \mathbf{v} with coordinates $[v_1, v_2, v_3]$ to a 4-vector $[v_1, v_2, v_3, 1]$. That extension makes non-linear operations such as translations implementable as linear mappings.

Following the approach of this paper, we have to give the $(m+1)$ -dimensional ‘homogeneous space’ into which we embed our m -dimensional Euclidean space a full geometric algebra. Let the unit vector for the extra dimension be denoted by e . This vector must be perpendicular to all regular vectors in the Euclidean space E^m , so $e \cdot \mathbf{x} = 0$ for all $\mathbf{x} \in E^m$. We also need to define $e \cdot e$ to make our algebra complete. This involves some dilemmas that are only fully resolvable in the conformal model of Section ???. For now we can take $e \cdot e = 1$. We let e denote ‘the Euclidean point at the origin’.

Let us now represent the subspaces of interest into this model, simply by using the structure of its geometric algebra.

- *Points:* A point at a location \mathbf{p} is made by translation of the point at the origin over the Euclidean vector \mathbf{p} . This is done by adding \mathbf{p} to e . This construction therefore gives the representation of the point P at location \mathbf{p} as the vector p in $(m + 1)$ -dimensional space:

$$p = e + \mathbf{p}$$

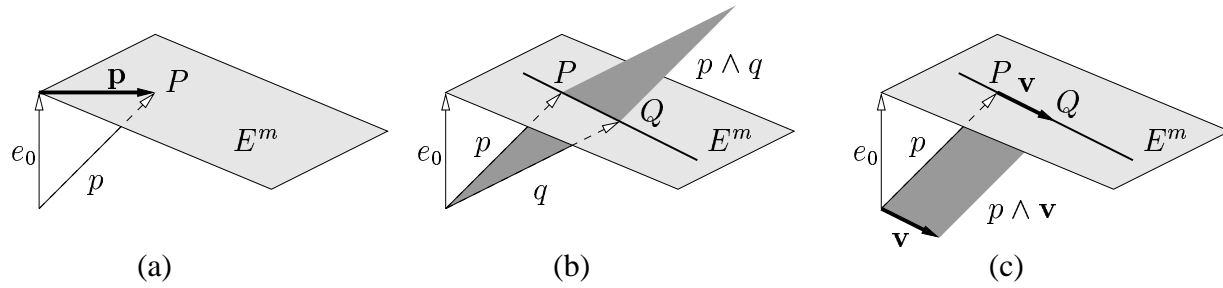


Figure 12: Representing offset subspaces of E^m in $m + 1$ -dimensional space.

This is just a regular vector in the $(m+1)$ -space, now interpretable as a Euclidean point. It is of course no more than the usual ‘homogeneous coordinates’ method in disguise: p has coordinates $[p_1, p_2, p_3, 1]$ on the orthonormal basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, e\}$. We will denote points of the m -dimensional Euclidean space in script, the vectors and blades in the corresponding vector space in **bold**, and vectors and blades in the $(m + 1)$ -dimensional homogeneous space in *italic*. You can visualize this construction as in Figure 12a (necessarily drawn for $m = 2$).

These vectors in $(m+1)$ -dimensional space can be multiplied using the products in geometric algebra. Let us consider in particular the outer product, and form blades.

- *Lines*: To represent a line, we compute the 2-blade spanned by the representative vectors of two points:

$$p \wedge q = (e + \mathbf{p}) \wedge (e + \mathbf{q}) = e \wedge (\mathbf{q} - \mathbf{p}) + \mathbf{p} \wedge \mathbf{q} \quad (44)$$

We recognize the vector $\mathbf{q} - \mathbf{p}$, and the area spanned by \mathbf{p} and \mathbf{q} . Both are elements that we need to describe an element of the directed line through the points P and Q . The former is the *direction vector* of the directed line, the latter is an area that we will call the *moment* of the line through p and q . It specifies the distance to the origin, for we can rewrite it to a rectangle spanned by the direction $(\mathbf{q} - \mathbf{p})$ and the *perpendicular support vector* \mathbf{d} :

$$\mathbf{p} \wedge \mathbf{q} = \mathbf{p} \wedge (\mathbf{q} - \mathbf{p}) = \mathbf{d} (\mathbf{q} - \mathbf{p}) \quad (45)$$

where $\mathbf{d} \equiv (\mathbf{p} \wedge (\mathbf{q} - \mathbf{p}))(\mathbf{q} - \mathbf{p})^{-1} = (\mathbf{p} \wedge \mathbf{q})(\mathbf{q} - \mathbf{p})^{-1}$ is the rejection of \mathbf{p} by $\mathbf{q} - \mathbf{p}$. We can therefore rewrite the same 2-blade $p \wedge q$ in various ways, such as $p \wedge q = p \wedge (\mathbf{q} - \mathbf{p}) = d (\mathbf{q} - \mathbf{p})$ (with $de + \mathbf{d}$), separating the positional part p or d and the purely Euclidean directional part $\mathbf{v} \equiv \mathbf{q} - \mathbf{p}$, see Figure 12c. The element $p \wedge q$ contains all these potential interpretations in one data structure, which may be constructed in all of these ways.

Temporarily reverting to the cross product to make a connection with a classical representation, we can rewrite eq.(44) as

$$p \wedge q = e \wedge (\mathbf{q} - \mathbf{p}) + \mathbf{p} \wedge \mathbf{q} = (\mathbf{p} - \mathbf{q}) e + (\mathbf{p} \times \mathbf{q}) \mathbf{I}_3 \quad (46)$$

We recognize the six Plücker coefficients $[\mathbf{p} - \mathbf{q}; \mathbf{p} \times \mathbf{q}]$ characterizing the line by its direction vector $\mathbf{v} = \mathbf{q} - \mathbf{p}$ and its ‘moment’ vector $\mathbf{m} = \mathbf{p} \times \mathbf{q}$ as

$$\ell = -\mathbf{v}e + \mathbf{m}\mathbf{I}_3$$

The six coefficients $[-v_1, -v_2, -v_3, m_1, m_2, m_3]$ of the line in this representation are in fact the coefficients of a 2-blade on the bivector basis $\{\mathbf{e}_1e, \mathbf{e}_2e, \mathbf{e}_3e, \mathbf{e}_2\mathbf{e}_3, \mathbf{e}_3\mathbf{e}_1, \mathbf{e}_1\mathbf{e}_2\}$. This integrates the Plücker representation fully into the homogeneous model (for which it was indeed historically designed). We will see that the compact and efficient Plücker intersection formulas are now straightforward consequences of the meet operation in geometric algebra.

- *Hyperplanes:* If we have an $(m - 1)$ -dimensional hyperplane characterized as $\mathbf{x} \cdot \mathbf{n} = \delta$, this can be written as $(e + \mathbf{x}) \cdot (\mathbf{n} - \delta e) = 0$, so as $x \cdot (\mathbf{n} - \delta e) = 0$. Therefore $\mathbf{n} - \delta e$ is the dual of the blade representing the hyperplane. The dual A^* of a blade A in the homogeneous model is obtained relative to the pseudoscalar $e \wedge \mathbf{I}_3 = e \mathbf{I}_3$ of the full space as $A^* = A \cdot (e\mathbf{I}_3)^{-1} = A \cdot (e\mathbf{I}_3) = A(e\mathbf{I}_3)$, so we get

$$(\mathbf{n} - \delta e)(e\mathbf{I}_3) = \mathbf{n}e\mathbf{I}_3 - \delta\mathbf{I}_3$$

This has the usual four Plücker coefficients $[n_1, n_2, n_3, -\delta]$, but on the trivector basis $\{-\mathbf{e}_2\mathbf{e}_3e, -\mathbf{e}_3\mathbf{e}_1e, -\mathbf{e}_1\mathbf{e}_2e, \mathbf{e}_1\mathbf{e}_2\mathbf{e}_3\}$, clearly different from the vector basis for points.

We can of course also construct the blade representing the hyperplane *directly* (rather than dually), given m points on it, namely as the outer product of the vectors representing those points.

- *And beyond:* These ways of making offset planar subspaces extend easily. An element of the oriented plane through the points P, Q and R is represented by the 3-blade $p \wedge q \wedge r$, and so on for higher dimensional ‘offset’ subspaces – if the space has enough dimensions to accommodate them. The blades we construct in this way can always be rewritten in the form $A = d\mathbf{A}$, where \mathbf{A} is a purely Euclidean blade, and d is a vector of the form $e + \mathbf{d}$, with \mathbf{d} a Euclidean vector. We should interpret \mathbf{A} as the direction element, and its grade therefore denotes the dimensionality of the flat subspace represented by A . The vector d represents the closest point to the origin (so that \mathbf{d} is the perpendicular support vector).
- *Scalar distances:* A small surprise is that even a 0-blade (i.e. a scalar) is useful: it is the representation of a scalar distance in the Euclidean space (with a sign, but without a direction), as we will see in the next section. Such distances are of course regular elements of geometry, so it is satisfying to find them on a par with position vectors, direction vectors and other elements of higher dimensionality as just another case of a representing blade in the homogeneous model of a flat Euclidean space.

Having such a unified representation for the various geometrical elements implies that computations using them are unified as well: they have just become operations on blades in $(m+1)$ -space, blissfully ignorant of what different geometrical situations these computations might represent. This opens the way to caseless computation in geometrical algorithms.

8.2.1 Equivalence of alternative characterizations

A special and rather satisfying property of this construction is its insensitivity to the kind of objects we use to construct the subspace. Of course the element of the line through p and q is determined by two points, or by a point and a direction. We would normally think of those as different constructions. However, in geometric algebra

$$p \wedge q = p \wedge (\mathbf{q} - \mathbf{p}) \quad (47)$$

(verify this!). So the two are exactly equal, they produce the same element by the same operation of ‘taking the outer product’. Moreover, the intrinsic ‘sliding’ symmetry of the support vector (any of $\mathbf{p} + \lambda(\mathbf{q} - \mathbf{p})$ can be used) is also automatically absorbed in the representation $p \wedge q$ due to the ‘sliding’ symmetry of the outer product term $\mathbf{p} \wedge \mathbf{q}$ in it. For instance, we may rewrite it as $\mathbf{p} \wedge \mathbf{q} = \frac{1}{2}(\mathbf{p} + \mathbf{q}) \wedge (\mathbf{q} - \mathbf{p})$, showing that the midpoint $\frac{1}{2}(\mathbf{p} + \mathbf{q})$ is on the carrier line. We have in $p \wedge q$ just the right mixture of specificity and freedom to denote the desired geometric entity.

You may verify that in general, a k -dimensional subspace element B determined by the points at locations $\mathbf{p}_0, \dots, \mathbf{p}_k$ is represented in the homogeneous model by the $(k + 1)$ -blade

$$B = p_0 \wedge \dots \wedge p_k$$

and that this is equivalent, by the rules of computation for the outer product, to specifying it by a point and k directions

$$B = p_0 \wedge (\mathbf{p}_1 - \mathbf{p}_0) \wedge \dots \wedge (\mathbf{p}_k - \mathbf{p}_0)$$

or any intermediate form specifying some positions and some directions. It is satisfying not to have to make different data-structures for those many ways of specifying this single geometrical object; the ‘constructor’ \wedge takes care of it automatically. Testing of equivalence of various objects is therefore much simplified. The paper [16] goes on to use this to develop a complete ‘simplicial calculus’ for simplices specified in this manner, deriving advanced results in a highly compact algebraic and computational manner.

8.3 Caseless subspace interactions in the homogeneous model

The `meet` and `join` in the homogeneous model function just as you would expect, providing the intersection of lines, planes, etc. Writing these out in their (Plücker) coordinates retrieves the familiar compact formulas for the coefficients of the result, but now accompanied by automatic evaluation of the basis on which these coefficients should be interpreted. That provides immediate identification of the kind of intersection, in a manner so well integrated that it suggests we might compute on without intermediate interpretation. This leads to caseless geometrical algorithms in which the dimensionality of intermediate results does not affect the data flow.

Even though the actual computation is a caseless `meet` applied to blades, let us see what is going on in detail in some typical situations. Following the internal computational combination of the Plücker-like coefficients shows how the algebra of the basis elements takes care of the proper intersection computation, at a small additional expense compared to the usual implementation using ‘pre-compiled’ tables of intersection formulas.

1. *Line and plane*: The meet of a line ℓ and a plane $\pi^* = \mathbf{n} - \delta e$ in general position is computed as:

$$\begin{aligned} \ell \cap \pi &= \pi^* \cdot \ell = (\mathbf{n} - \delta e) \cdot (-\mathbf{v}e + \mathbf{m}\mathbf{I}_3) \\ &= -(\mathbf{n} \cdot \mathbf{v})e + \mathbf{n} \cdot (\mathbf{m}\mathbf{I}_3) - \delta \mathbf{v} + 0 = -(\mathbf{n} \cdot \mathbf{v})e + (\mathbf{n} \wedge \mathbf{m})\mathbf{I}_3 - \delta \mathbf{v} \\ &= -(\mathbf{n} \cdot \mathbf{v})e + (\mathbf{m} \times \mathbf{n} - \delta \mathbf{v}) \end{aligned}$$

This is the correct result, representing a point at the location $(\delta \mathbf{v} + \mathbf{n} \times \mathbf{m})/(\mathbf{n} \cdot \mathbf{v})$ in its homogeneous (Plücker) coordinates. Note how the orthogonality relationships between the basis elements automatically kill the potential term involving δ and \mathbf{m} . But the fact that this term is zero *is* computed, and that is a slight inefficiency relative to the direct implementation of the same result from a table with Plücker formulas. It is the computational price we pay for the membership of the full geometric algebra.

2. *Two lines*: The meet of two lines in general position is a measure of their signed distance (remember that in this model a dual is made through right-multiply by $e\mathbf{I}_3$, so the dual of the line $-\mathbf{v}_2e + \mathbf{m}_2\mathbf{I}_3$ is $-\mathbf{v}_2\mathbf{I}_3 + \mathbf{m}_2e$):

$$\ell_1 \cap \ell_2 = \ell_2^* \cdot \ell_1 = (-\mathbf{v}_2\mathbf{I}_3 + \mathbf{m}_2e) \cdot (-\mathbf{v}_1e + \mathbf{m}_1\mathbf{I}_3) = \mathbf{m}_2 \cdot \mathbf{v}_1 + \mathbf{v}_2 \cdot \mathbf{m}_1,$$

retrieving the well-known compact Plücker way of determining how lines pass each other in space; three tests on the signs of such quantities representing the edges of a triangle determine efficiently whether a ray hits the triangle. Again the basis orthogonality relationships have made terms containing $\mathbf{m}_1 \cdot \mathbf{m}_2$ or $\mathbf{v}_1 \cdot \mathbf{v}_2$ equal to zero.

The directional outcomes are accompanied by numerical factors (such as $\mathbf{n} \cdot \mathbf{v}$ in the first example) relating to the numerical significance of the computation. These are an intrinsic part of the computation of the object, not just secondary aspects that need to be thought of separately (with the danger of being *ad hoc*) or that need be computed separately (costing time).

8.4 The conformal model

A recently developed model of Euclidean space E^m is the *conformal model* $V^{n+1,1}$. This is a true *algebra of points*, or rather, *an algebra of spheres* (with points being spheres of zero radius). Again, points at locations \mathbf{p} and \mathbf{q} are represented by vectors p and q in the model, but now in a manner such that *the inner product represents their Euclidean distance*:

$$p \cdot q = -\frac{1}{2}(\mathbf{p} - \mathbf{q})^2 \tag{48}$$

In particular, $p \cdot p = 0$, so that points are represented by vectors which have – in their representative space – a zero norm! To do this and still have a complete geometric algebra requires *two* extra dimensions, so an m -dimensional Euclidean space is now represented using the geometric algebra of an $(m+2)$ -dimensional space. Moreover, one of these extra dimensions is represented by a basis vector which squares to -1 (such spaces are known as Minkowski spaces).

A useful basis for this space is: an orthonormal basis for the Euclidean space embedded in it, and the vectors e_0 and e_∞ to represent the *point at the origin*, and the *point at infinity*, respectively. The two satisfy: $e_0 \cdot e_\infty = 1$, they are null vectors: $e_0 \cdot e_0 = 0$ and $e_\infty \cdot e_\infty = 0$, and they are orthogonal to the Euclidean subspace, so that $e_0 \cdot \mathbf{x} = 0$ and $e_\infty \cdot \mathbf{x} = 0$ for any $\mathbf{x} \in E^m$. The representation of a point p of Euclidean space in this conformal model is the vector:

$$p = e_0 + \mathbf{p} - \frac{1}{2}\mathbf{p}^2 e_\infty$$

(or a scalar multiple). You may verify that $p^2 = 0$, and that

$$p \cdot q = (e_0 + \mathbf{p} - \frac{1}{2}\mathbf{p}^2 e_\infty) \cdot (e_0 + \mathbf{q} - \frac{1}{2}\mathbf{q}^2 e_\infty) = -\frac{1}{2}\mathbf{q}^2 + \mathbf{p} \cdot \mathbf{q} - \frac{1}{2}\mathbf{p}^2 = -\frac{1}{2}(\mathbf{q} - \mathbf{p})^2$$

as desired.

Any point \mathbf{x} on the hyperplane perpendicularly bisecting the line segment pq satisfies $(\mathbf{x} - \mathbf{p})^2 = (\mathbf{x} - \mathbf{q})^2$, and therefore:

$$x \cdot (q - p) = 0.$$

It follows that $q - p = (\mathbf{q} - \mathbf{p}) - \frac{1}{2}(\mathbf{q}^2 - \mathbf{p}^2)e_\infty$ *dually represents the midplane of p and q* , see eq.(21). In general, a hyperplane with orthogonal support vector \mathbf{d} is (dually) represented by the vector

$$d = \mathbf{d}^{-1} - e_\infty$$

or any multiple of it, such as $\mathbf{n} - \delta e_\infty$ with \mathbf{n} its normal vector and δ the support along \mathbf{n} of the hyperplane. You may verify that the equation $x \cdot d = 0$ is indeed equivalent to the normal hyperplane equation $\mathbf{x} \cdot \mathbf{n} = \delta$.

8.4.1 Spheres are blades

The direct expression of the Euclidean distance by the inner product in eq.(48) implies that the equation

$$x \cdot c = -\frac{1}{2}\rho^2$$

is the equation of a sphere with radius ρ and center c . We rewrite this to

$$x \text{ on sphere with radius } \rho \text{ and center } c \iff x \cdot (c + \frac{1}{2}\rho^2 e_\infty) = 0,$$

so this shows that *the vector $c + \frac{1}{2}\rho^2 e_\infty$ dually represents a sphere*. Where the homogeneous model can be used to code a hyperplane by a homogeneous normal vector, the conformal model (dually) represents a complete sphere by a single representative vector! In the conformal model, *(dual) spheres are basic elements of computation*. We get an algebra of spheres; a point is just a (dual) sphere of radius zero.

The direct (rather than dual) representation of a sphere is through the wedge product: *spheres are blades* in the conformal model. This is obvious since the dual of the vector $c + \frac{1}{2}\rho^2 e_\infty$ is an $(m - 1)$ -blade in the $(m + 2)$ -dimensional representation space.

$$x \text{ on sphere through } p, q, r, s \iff x \wedge (p \wedge q \wedge r \wedge s) = 0.$$

Moreover, the two representations are exactly dual in the conformal representation, so we can compute the center and radius of a sphere given by four points immediately through using:

$$p \wedge q \wedge r \wedge s = (c + \frac{1}{2}\rho^2 e_\infty)^*$$

We now give a derivation using the midplane equation above, as an example of the kind of computational and reasoning power we have available with the conformal model. The sphere through the points p , q , r and s must be on the intersection of the midplanes dually represented by $(q - p)$, $(r - p)$ and $(s - p)$. Assuming these in general position so that their join is the full common space, the element dually representing their intersection is $(q - p) \wedge (r - p) \wedge (s - p)$ (by eq.(??)). This is a 3-blade, so its dual in the $(3 + 2)$ -dimensional representation space is a 2-blade: it represents the two intersection points e_∞ and c , so $(q - p) \wedge (r - p) \wedge (s - p) = (e_\infty \wedge c)^*$ (the dual is relative to the pseudoscalar $e_\infty \wedge e_0 \wedge \mathbf{I}_m$ of the representation space). Now the point p must be on the sphere, so $p \cdot c = -\frac{1}{2}\rho^2$. Using this we obtain:

$$\begin{aligned} p \wedge q \wedge r \wedge s &= p \wedge (q - p) \wedge (r - p) \wedge (s - p) = p \wedge (e_\infty \wedge c)^* \\ &= (p \cdot (e_\infty \wedge c))^* = (c - e_\infty(p \cdot c))^* = (c + \frac{1}{2}\rho^2 e_\infty)^* \end{aligned} \quad (49)$$

Done. Note how easy it is to perform these intersections!

It is very satisfying that these two totally different specifications of a sphere should be literally duals of each other, i.e. perpendicular to each other in the representative space of the conformal model. It is also a very pleasant surprise that the very complicated symmetries of four points determining the same sphere are simply reduced to the anti-symmetry of the outer product (as were the symmetries of the support vectors of hyperplanes in the homogeneous model). Spheres are not really new objects requiring totally new products – as long as you treat them in their own algebra, they behave just like subspaces.

We note that $p \wedge q$ is a 1-dimensional sphere, i.e. the computational representation of a point pair. In contrast to the homogeneous model, $p \wedge q$ now really has the semantics of a localized line *segment* rather than merely a line *element*.

8.4.2 Intersection of spheres

In the homogeneous model we saw that a factorization like eq.(47) gave literal equivalence of the same geometric object specified in different ways. Such simplifications also occur in the conformal model. Indeed, the dual equivalence of the sphere specifications just treated can be used in this way. Another example is the intersection of two spheres, which should produce a circle in a well-defined plane. Let us take a simple example, equal sized spheres of radius ρ at opposite sides $\pm c$ of the origin. The dual of their intersection is computed as the outer product of their duals, which can then be rewritten in more convenient form:

$$(e_0 - \mathbf{c} - \frac{1}{2}(\mathbf{c}^2 - \rho^2)e_\infty) \wedge (e_0 + \mathbf{c} - \frac{1}{2}(\mathbf{c}^2 - \rho^2)e_\infty) = 2\mathbf{c} \wedge (e_0 + \frac{1}{2}(\mathbf{c}^2 - \rho^2)e_\infty)$$

The right hand side is immediately recognizable as the dual of the intersection of a hyperplane with normal \mathbf{c} through the origin (its dual representation is \mathbf{c}) with a sphere at the origin of radius $c^2 - \rho^2$. So these two alternative representations of the intersection circle are just two factorizations of the element of geometric algebra representing it (many other factorizations exist). Note how we can *compute directly with spheres and planes* rather than with equations asserting properties of points on it.

8.4.3 Homogeneous model subsumed

The direct (rather than dual) representation of a linear ('flat') k -dimensional subspace is obtained as a blade using the wedge product with e_∞ ; you can see them all as k -spheres containing the point at infinity. Since

$$e_\infty \wedge p = e_\infty \wedge (e_0 + \mathbf{p}),$$

the point at infinity 'factors out', leaving essentially the homogeneous model we had before. So the homogeneous model is contained in the conformal model.

8.4.4 Points and point pairs

In the conformal model, we get a rather subtle distinction between the point at \mathbf{p} viewed as a 0-dimensional subspace (represented by $e_\infty \wedge p = e_\infty \wedge (e_0 + \mathbf{p})$) and the point at \mathbf{p} viewed as a sphere with zero radius (*dually* represented by $p = e_0 + \mathbf{p} - \frac{1}{2}\mathbf{p}^2 e_\infty$).

The above shows that the outer product of k terms in the conformal model constructs a k -sphere. In particular, the outer product of two points should be a 1-dimensional sphere, which is essentially a point pair; the anti-symmetry of the outer product actually makes $p \wedge q$ the *ordered point pair*. A point pair is therefore not an arbitrary construction; in the conformal model, it is automatically endowed with all the algebraic properties to make it part of a consistent system of computation. Actually computing it leads to a rather frightful expression:

$$p \wedge q = e_0(\mathbf{q} - \mathbf{p}) + \mathbf{p} \wedge \mathbf{q} + \frac{1}{2}(\mathbf{q}^2 - \mathbf{p}^2)e_0 \wedge e_\infty + \frac{1}{2}(\mathbf{q}^2\mathbf{p} - \mathbf{p}^2\mathbf{q})e_\infty$$

The terms without e_∞ are the same as in the homogeneous model and denote the line element. The rest is apparently needed for the considerable symmetry properties required to make $p \wedge q \wedge r$ a circle, $p \wedge q \wedge r \wedge s$ a sphere, etcetera.

8.4.5 Unification of translations and rotations

The conformal model unites rotations and translations in a satisfying manner: both are representable as the exponent of a 2-blade. We have seen that the rotations require a 2-blade $\mathbf{I}\phi/2$ denoting a plane in the Euclidean space, and that a rotation can then be represented as

$$\text{rotation: } \mathbf{x} \mapsto e^{-\mathbf{I}\phi/2}\mathbf{x}e^{\mathbf{I}\phi/2}.$$

A translation turns out to be representable as the exponent of a 2-blade $e_\infty \wedge \mathbf{p}/2$ containing the point at infinity and the translation vector \mathbf{p} . Because e_∞ squares to zero and commutes with \mathbf{p} ,

we obtain

$$e^{e_\infty \mathbf{p}/2} = 1 + e_\infty \mathbf{p}/2.$$

You can now use this to verify that the translation of the point at the origin (represented by e_0) indeed gives the point at \mathbf{p} :

$$p = e^{-e_\infty \mathbf{p}/2} e_0 e^{e_\infty \mathbf{p}/2} = e_0 + \mathbf{p} - \frac{1}{2} \mathbf{p}^2 e_\infty.$$

Having rotations and translations in the same form permits a concise treatment of rigid body motions, presenting new unifying insights in traditional representations such as screws [12]. This may well transfer them from theoretical mechanics to practical computational geometry, as the next refinement after quaternions.

9 Some remarks on implementation

The geometric algebra of an m -dimensional vector space contains nicely linear objects (the blades) that can be represented on a basis which should contain 2^m elements (since we need $\binom{m}{k}$ for each k -blade). The various products are all linear, and can be implemented using matrix products of $2^m \times 2^m$ matrices. That may be straightforward, but it is obviously inefficient in both space and time. This is even more urgent when we use the homogeneous model in which an m -dimensional Euclidean space requires an $(m+1)$ -dimensional geometric algebra in the homogeneous model, or the $(m+2)$ -dimensional conformal model. It seems a lost cause.

However, the recognition that the important elements are blades and their products suggests a more efficient implementation. When two blades multiply by inner or outer product, a blade of unique grade results. This suggests designing a data representation for the elements of geometric algebra that permits easy retrieval of their grades, and also to automatically generate optimized code for the inner and outer multiplication of blades of specific grades k and ℓ . The geometric product generates a more general element of a mixed, but still limited, set of grades $|\ell - k|, |\ell - k| + 2, \dots, k + \ell$. Division requires inversion, but this is closely related to the much simpler operation of reversion, simply switching the signs of certain grades. The structural membership of an element to the larger geometric algebra, with its benefit of unified relationships between the various operations, is then merely paid for by a grade-dependent jump to a piece of code; when processing data in a batch mode with many similar operations, this would not slow down things significantly.

Work is underway on an efficient implementation which capitalizes on these structural properties of geometric algebra [8]. First results look very promising indeed, getting close to the usual efficiency of the geometrical computations, but in a simpler code without exceptions or *ad hoc* data structures, and naturally integrating computational techniques which classically belong to different realms than vector/matrix algebra (such as quaternions and Plücker coordinates).

10 Conclusion

This introduction of geometric algebra intends to alert you to the existence of a limited set of products that appears to generate all geometric constructions in one consistent framework. Using this framework can simplify the set of data structures representing objects since it inherently encodes all relationships and symmetries of the geometrical primitives in those operators (an example was eq.(45)). Also, it could serve as a straight-jacket for the specification of geometric algorithms, preventing the unbridled invention of new operations and objects without clear and clean geometrical meaning, or well-defined relationships to other objects in the application. If our hopes are correct, this straight-jacket would actually *not* be a limitation on what one can construct; rather it contains precisely the right set of operations to provide a precise language for arbitrary constructions. The basic operations even have the power to model the geometry of spheres and their interactions; thus the same syntax admits of varied semantics.

That such a system exists is a happy surprise to all learning about it. Whether it is also the way we should structure our programming is at the moment an open question. Use of the conformal model would require representing the computations on the Euclidean geometry of a 3-dimensional space on a basis of $2^{3+2} = 32$ elements, rather than just 3 basis vectors (plus 1 scalar basis). It seems a hard sell. But you often have to construct objects representing higher order relationships between points (such as lines, planes and spheres) anyway, even if you do not encode them on such a ‘basis’. Also, our investigations show that perhaps all one needs to do all of geometry are blades and operators composed of products of vectors; the product combinations of this limited subset can be optimized in time and space requirements, with very little overhead for their membership of the full geometric algebra. That automatic membership would enable us to compute directly with lines, planes, circles and spheres and their intersections without needing to worry about special or degenerate cases, which should eliminate major headaches and bugs. We also find the coordinate-free specification of the operations between objects very attractive; relegating the use of coordinates purely to the input and output of geometric objects banishes them from the body of the programs and frees the specification of algorithms from details of the data structures used to implement them. Such properties makes geometric programs so much more easy to verify, and – once we have learned to express ourselves fluently in this new language – to construct.

We are currently investigating these possibilities, doing our best to make the geometric algebra approach a reasonable alternative. The main delay now is that the algebra dictates a new way of thinking about geometry which requires one to revisit many old constructions. This takes time, but is worthwhile since it appears to simplify the whole structure of geometric programming. At the very least, we would hope geometric algebra to be a useful meta-language in which to specify geometric programs; but the proven efficiency of quaternions, which are such a natural part of geometric algebra, suggests that we might even want to do our low-level computations in this new computational framework.

11 Further reading

There is a growing body of literature on geometric algebra. Unfortunately much of the more readable writing is not very accessible, being found in books rather than journals. Little has been written with computer science in mind, since the initial applications have been to physics. No practical implementations in the form of libraries with algorithms yet exist (though there are packages for Maple [14] and Matlab [6] which can be used as a study-aid or for algorithm design). We would recommend the following as natural follow-ups on this paper:

- GABLE: a Matlab package for geometric algebra, accompanied by a tutorial [6].
- The introductory chapters of ‘New Foundations of Classical Mechanics’ [11].
- An introductory course intended for physicists [2].
- An application to a basic but involved geometry problem in computer vision, with a brief introduction into geometric algebra [15].
- A paper showing how linear algebra becomes enriched by viewing it as a part of geometric algebra: [13].

If you read them in approximately this order, you should be alright. We are working on texts more specifically suited for a computer graphics audience; these will probably first appear as SIGGRAPH courses.

If you want to write application programs in geometric algebra, GABLE will not keep you satisfied for long. We try to keep links to available software up to date at the site <http://www.wins.uva.nl/~leo/ga/>.

References

- [1] Bouma, T.A., Dorst, L. and Pijls, H.G.J., *Geometric Algebra for Subspace Operations*, Acta Mathematicae Applicandae, pp. 285-300, 2002, available on-line at series <http://xxx.lanl.gov/abs/math.LA/0104102>.
- [2] C. Doran and A. Lasenby, *Physical Applications of Geometric Algebra*, 2001, available at <http://www.mrao.cam.ac.uk/~clifford/ptIIIcourse/>
- [3] C. Doran, A. Lasenby, S. Gull, *Chapter 6: Linear Algebra*, in: Clifford (Geometric) Algebras with applications in physics, mathematics and engineering, W.E. Baylis (ed.), Birkhäuser, Boston, 1996.
- [4] L. Dorst, *The Inner Products of Geometric Algebra*, in: *Applications of Geometric Algebra in Computer Science and Engineering*, L. Dorst, C. Doran, J. Lasenby (eds.), Birkhauser, Boston, pp. 35-46, 2002

- [5] L. Dorst, *Honing geometric algebra for its use in the computer sciences*, in: G. Sommer ed.: *Geometric Computing with Clifford Algebras: Theoretical Foundations and Applications in Vision and Robotics*, Springer ISBN 3-540-41198-4, Chapter 6, 2001, pp. 127-152. also at <http://www.wins.uva.nl/~leo/ga/>
- [6] L.Dorst, S.Mann, T.A.Bouma, GABLE: a Geometric AlgeBra Learning Environment, www.science.uva.nl/~leo/GABLE/
- [7] L. Dorst and S. Mann, *Geometric algebra: a computational framework for geometrical applications (part I: algebra)*, IEEE Computer Graphics and Applications, vol.22. no.3, May/June 2002, pp.24-31.
S. Mann and L. Dorst, *Geometric algebra: a computational framework for geometrical applications (part II: applications)*, IEEE Computer Graphics and Applications, vol.22 no.4, July/August 2002, pp.58-67.
- [8] D. Fontijne, T. Bouma, L. Dorst, *GAIGEN: a Geometric Algebra Implementation Generator*, <http://www.science.uva.nl/~fontijne/raytracer>
- [9] R. Goldman, *Illicit Expressions in Vector Algebra*, ACM Transactions of Graphics, vol.4, no.3, 1985, pp. 223—243.
- [10] R. Goldman, *The Ambient Spaces of Computer Graphics and Geometric Modeling*, IEEE Computer Graphics and Applications, vo.20, pp. 76–84, 2000.
- [11] D. Hestenes, *New foundations for classical mechanics*, 2nd edition, D. Reidel, Dordrecht, 2000.
- [12] D. Hestenes, *Old wine in new bottles*, in: *Geometric Algebra: A Geometric Approach to Computer Vision, Quantum and Neural Computing, Robotics and Engineering*, Bayro-Corrochano, Sobczyk, eds, Birkhäuser, ISBN 0-8176-4199-8, 2001, Chapter 24, pp. 498-520.
- [13] D. Hestenes, *The design of linear algebra and geometry*, Acta Applicandae Mathematicae 23: 65-93, 1991.
- [14] A. Lasenby, M. Ashdown et al., *GA package for Maple V*, 1999, available at <http://www.mrao.cam.ac.uk/~clifford/software/GA/>
- [15] J. Lasenby, W. J. Fitzgerald, C. J. L. Doran and A. N. Lasenby. *New Geometric Methods for Computer Vision* Int. J. Comp. Vision 36(3), p. 191-213 (1998).
- [16] G. Sobczyk, *Simplicial Calculus with Geometric Algebra*, in: *Clifford Algebras and their Applications in Mathematical Physics*, A. Micali et al. eds., Kluwer 1993.
- [17] J. Stolfi, *Oriented projective geometry*, Academic Press, 1991.