

Chapter 1

The Making of GABLE, a Geometric Algebra Learning Environment in Matlab

Stephen Mann, Leo Dorst, and Tim Bouma

1.1 Introduction

Geometric algebra extends Clifford algebra with geometrically meaningful operators, and its purpose is to facilitate geometrical computations. Present textbooks and implementation do not always convey this geometrical flavor or the computational and representational convenience of geometric algebra, so we felt a need for a computer tutorial in which representation, computation and visualization are combined to convey both the intuition and the techniques of geometric algebra. Current software packages are either Clifford algebra only (such as CLICAL [9] and CLIFFORD [1]) or do not include graphics [6], so we decide to build our own. The result is GABLE (Geometric Algebra Learning Environment) a hands-on tutorial on geometric algebra that should be accessible to the second year student in college [3].

The GABLE tutorial explains the basics of Geometric Algebra in an accessible manner. It starts with the outer product (as a constructor of subspaces), then treats the inner product (for perpendicularity), and moves via the geometric product (for invertibility) to the more geometrical operators such as projection, rotors, meet and join, and end with to the homogeneous model of Euclidean space. When the student is done he/she should be able to do simple Euclidean geometry of flats using the geometric algebra of homogeneous blades. For instance, the intersection between lines can be easily expressed in the basic operators:

```
e = e3;           % the homogeneous embedding
P = e+ e1/3+e2;  % point P
Q = e+ e1+e2/2;  % point Q
R = e+ e1/2-e2/4; % point R
PQ = join(P,Q);  % line PQ
QR = join(Q,R);  % line QR
meet(PQ,QR)      % intersection of those lines
```

Drawing the objects of geometric algebra is an important part of GABLE.




FIGURE 1.1. Intersection of lines.

Using them appropriately in this example produces Figure 1.1, clarifying the relationship between the 3-D blade-based computations in the homogeneous model and their geometric semantics in 2-D Euclidean space. Figure 1.2 is an illustration of another exercise in GABLE, requiring interpolation between bivectors. Obviously, with the geometric product we can define and compute spinors simply; then using the n -th root (by logarithm and exponentiation) we may interpolate between orientations, and apply the resulting spinors to the bivectors. That generates the figure.

In Section sec:pappus, we give the a more complete example illustrating Pappus's theorem.

Our desire to visualize meant that we did not need to go beyond 3 dimensions, and our implementation focuses on $\mathcal{C}_{3,0}$ (although our implementation is general enough to handle other signatures in 3-dimensional space as well). Since this software is meant for a tutorial, we did not have great efficiency concerns (though we did have some), and were most interested in ease of implementation and the creation of a software package that could be made widely available for teaching purposes.

These goals led us to implement GABLE in Matlab and to foster distribution of geometric algebra, it was made to work with the student version of Matlab. This paper describes our experiences in developing the package. We ran into some issues that any implementer of geometric algebra will need to decide on (representation, computational efficiency, stability of inverses); but we also encountered insufficiently precisely resolved issues in the very structure of geometric algebra itself, which the need for straightforward practical use brought out. Among these are the various inner products, and the precise definition and semantics of the `meet` and `join` operators in Euclidean geometry (for more general definitions see Chapter 3). This chapter motivates and documents our decisions in these matters.

Our representation of geometric objects is a refinement of the 8×8 representation that has been presented by others, along the lines suggested




FIGURE 1.2. Interpolation of orientations for a bivector.

by Lounesto and Ablamowicz [8]pag. 72), [1]. We compare this representation of geometric algebra to matrix representations of Clifford Algebras in Section 1.2. For the important but potentially expensive operation of inversion (or geometric division), we settled on a variation of a method proposed by Lounesto for $\mathcal{C}_{3,0}$, which we extend in Section 1.3 to work for arbitrary signature (in 3 dimensions). At the higher level, Section 1.4 gives some detail on our implementation of the `meet` and `join` operations, extending them to the non-trivial cases of partially overlapping subspaces. Section 1.5 discusses some of the implementation details, including a brief description of the graphics and numerical aspects of GABLE. The tutorial itself [3] is available on the World Wide Web¹ The webpages contain both the Matlab package GABLE and the tutorial textbook.

1.2 Representation of geometric algebra

The first decision we faced in our implementation was which representation to use for the geometric objects. The most natural representation in Matlab would be a matrix representation. Matrix representations for the geometric product in the Clifford algebras of various signatures are well studied [11]; for each signature a different matrix algebra results. That is slightly un-

¹Available at <http://www.wins.uva.nl/~leo/clifford/gable.html> and <http://www.cgl.uwaterloo.ca/~smann/GABLE/>.

satisfactory. Moreover, our desire to make our algebra a proper *geometric* algebra implies that we should not only represent the geometric product, but also the outer and inner products, and preferably on a par with each other. These issues are discussed in more detail in Section 1.2.4; in brief, we ended up using a modified form of the 8×8 matrix representation.

1.2.1 The matrix representation of GABLE

In GABLE, we represent a multivector \mathbf{A} as an 8×1 column matrix giving its coefficients relative to a basis for the Clifford algebra:

$$\mathbf{A} = [1, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_1 \wedge \mathbf{e}_2, \mathbf{e}_2 \wedge \mathbf{e}_3, \mathbf{e}_3 \wedge \mathbf{e}_1, \mathbf{e}_1 \wedge \mathbf{e}_2 \wedge \mathbf{e}_3] \begin{bmatrix} A_0 \\ A_1 \\ A_2 \\ A_3 \\ A_{12} \\ A_{23} \\ A_{31} \\ A_{123} \end{bmatrix},$$

where $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ form an orthogonal basis for the vector space of our algebra. We will use **bold** font for the multivector, and *math* font for its scalar-valued coefficients. The multivector \mathbf{A} is thus represented by an 8×1 column matrix $[\mathbf{A}]$, which we will denote in shorthand as

$$\mathbf{A} \rightleftharpoons [\mathbf{A}].$$

Now if we need to compute the geometric product $\mathbf{A}\mathbf{B}$, we view this as a linear function of \mathbf{B} determined by \mathbf{A} , i.e., as the linear transformation $\underline{\mathbf{A}}^G(\mathbf{B})$, the ‘ G ’ denoting the geometric product. Such a linear function can be represented by an 8×8 matrix, determined by \mathbf{A} (and the fact that we are doing a geometric product) acting on the 8×1 matrix of $[\mathbf{B}]$. We thus expand the representation $[\mathbf{A}]$ of \mathbf{A} to the 8×8 geometric product matrix $[\underline{\mathbf{A}}^G]$, and apply this to the 8×1 representation $[\mathbf{B}]$ of \mathbf{B} :

$$\mathbf{A}\mathbf{B} \rightleftharpoons [\underline{\mathbf{A}}^G][\mathbf{B}].$$

The result of the matrix product $[\underline{\mathbf{A}}^G][\mathbf{B}]$ is the 8×1 matrix representing the element $\mathbf{A}\mathbf{B}$. The matrix entry $[\underline{\mathbf{A}}^G]_{\alpha,\beta}$ (so in column α and row β , with α and β running through the indices $\{0, 1, 2, 3, 12, 23, 31, 123\}$) can be computed in a straightforward manner from the multiplication table of the geometric product:

$$\mathbf{e}_\alpha \mathbf{e}_\beta = c \mathbf{e}_\gamma \iff [\underline{\mathbf{A}}^G]_{\gamma,\beta} = c A_\alpha, \quad (1.1)$$

where the c are defined in the next section. This 8×8 matrix $[\underline{\mathbf{A}}^G]$ can then be used to evaluate the (bilinear) product $\mathbf{A}\mathbf{B}$ by applying it to the

8×1 column matrix $[\mathbf{B}]$ using the usual matrix multiplication:

$$[\mathbf{AB}]_\gamma = \left([\underline{\mathbf{A}}^G] [\mathbf{B}] \right)_\gamma = \sum_\beta [\underline{\mathbf{A}}^G]_{\gamma,\beta} [\mathbf{B}]_\beta.$$

So for example the identity $\mathbf{e}_1 \mathbf{e}_2 = \mathbf{e}_{12}$ leads to the matrix entry $[\underline{\mathbf{A}}^G]_{12,2} = A_1$; this is the only non-zero entry in column $\beta = 2$. In matrix multiplication between $\mathbf{A} = A_1 \mathbf{e}_1$ and $\mathbf{B} = B_2 \mathbf{e}_2$ this yields $[\mathbf{AB}]_{12} = [\underline{\mathbf{A}}^G]_{12,2} [\mathbf{B}]_2 = A_1 B_2$, which is the correct contribution to the \mathbf{e}_{12} component of the result.

In algebraic parlance, the numbers c defined above (which depend on α , β and γ) are the *structure coefficients* of the algebra determined by the geometric product. In particular, each c will be a positive or negative factor depending on the basis of the algebra, the permutation of the elements, and the signature of the algebra.

Both the outer and the inner product can be implemented as matrix multiplications, since $\mathbf{A} \wedge \mathbf{B}$ and $\mathbf{A} \cdot \mathbf{B}$ are also linear functions of \mathbf{B} , determined by \mathbf{A} . So we implement

$$\mathbf{A} \wedge \mathbf{B} \rightleftharpoons [\underline{\mathbf{A}}^O] [\mathbf{B}] \quad \text{and} \quad \mathbf{A} \cdot \mathbf{B} \rightleftharpoons [\underline{\mathbf{A}}^I] [\mathbf{B}].$$

The 8×8 matrices $[\underline{\mathbf{A}}^O]$ and $[\underline{\mathbf{A}}^I]$ are given below, and they are constructed according to Equation 1.1 for the outer product and inner product, respectively.

1.2.2 The representation matrices

Using the recipe of Equation 1.1, we can compute the actual matrices. The c values of Equation 1.1 contain signed products of σ_i s. These represent the signature (and metric) through their definition as $\sigma_i \equiv \mathbf{e}_i \mathbf{e}_i$. We further define $\sigma_{ij} \equiv \sigma_i \sigma_j$ and $\sigma_{ijk} \equiv \sigma_i \sigma_j \sigma_k$.

Geometric product matrix:

$$[\underline{\mathbf{A}}^G] = \begin{bmatrix} A_0 & \sigma_1 A_1 & \sigma_2 A_2 & \sigma_3 A_3 & -\sigma_{12} A_{12} & -\sigma_{23} A_{23} & -\sigma_{13} A_{31} & -\sigma_{123} A_{123} \\ A_1 & A_0 & \sigma_2 A_{12} & -\sigma_3 A_{31} & -\sigma_2 A_2 & -\sigma_{23} A_{123} & \sigma_3 A_3 & -\sigma_{23} A_{23} \\ A_2 & -\sigma_1 A_{12} & A_0 & \sigma_3 A_{23} & \sigma_1 A_1 & -\sigma_3 A_3 & -\sigma_{31} A_{123} & -\sigma_{13} A_{31} \\ A_3 & \sigma_1 A_{31} & -\sigma_2 A_{23} & A_0 & -\sigma_{12} A_{123} & \sigma_2 A_2 & -\sigma_1 A_1 & -\sigma_{12} A_{12} \\ A_{12} & -A_2 & A_1 & \sigma_3 A_{123} & A_0 & \sigma_3 A_{31} & -\sigma_3 A_{23} & \sigma_3 A_3 \\ A_{23} & \sigma_1 A_{123} & -A_3 & A_2 & -\sigma_1 A_{31} & A_0 & \sigma_1 A_{12} & \sigma_1 A_1 \\ A_{31} & A_3 & \sigma_2 A_{123} & -A_1 & \sigma_2 A_{23} & -\sigma_2 A_{12} & A_0 & \sigma_2 A_2 \\ A_{123} & A_{23} & A_{31} & A_{12} & A_3 & A_1 & A_2 & A_0 \end{bmatrix} \quad (1.2)$$

Outer product matrix:

$$[\underline{\mathbf{A}}^O] = \begin{bmatrix} A_0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ A_1 & A_0 & 0 & 0 & 0 & 0 & 0 & 0 \\ A_2 & 0 & A_0 & 0 & 0 & 0 & 0 & 0 \\ A_3 & 0 & 0 & A_0 & 0 & 0 & 0 & 0 \\ A_{12} & -A_2 & A_1 & 0 & A_0 & 0 & 0 & 0 \\ A_{23} & 0 & -A_3 & A_2 & 0 & A_0 & 0 & 0 \\ A_{31} & A_3 & 0 & -A_1 & 0 & 0 & A_0 & 0 \\ A_{123} & A_{23} & A_{31} & A_{12} & A_3 & A_1 & A_2 & A_0 \end{bmatrix} \quad (1.3)$$

Inner product matrix:

$$[\underline{\mathbf{A}}^I] = \begin{bmatrix} A_0 & \sigma_1 A_1 & \sigma_2 A_2 & \sigma_3 A_3 & -\sigma_{12} A_{12} & -\sigma_{23} A_{23} & -\sigma_{31} A_{31} & -\sigma_{123} A_{123} \\ 0 & A_0 & 0 & 0 & -\sigma_2 A_2 & 0 & \sigma_3 A_3 & -\sigma_{23} A_{23} \\ 0 & 0 & A_0 & 0 & \sigma_1 A_1 & -\sigma_3 A_3 & 0 & -\sigma_{13} A_{31} \\ 0 & 0 & 0 & A_0 & 0 & \sigma_2 A_2 & -\sigma_1 A_1 & -\sigma_{12} A_{12} \\ 0 & 0 & 0 & 0 & A_0 & 0 & 0 & \sigma_3 A_3 \\ 0 & 0 & 0 & 0 & 0 & A_0 & 0 & \sigma_1 A_1 \\ 0 & 0 & 0 & 0 & 0 & 0 & A_0 & \sigma_2 A_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & A_0 \end{bmatrix} \quad (1.4)$$

Note the relation between these matrices: the inner product matrix and the outer product matrix both have all non-zero elements taken from the geometric product matrix. Note also the lack of signature in the outer product matrix; this is in agreement with the fact that it forms a (non-metric) Grassmann algebra that may be viewed as a geometric algebra of null vectors, for which all σ_i equal 0.

The reader may realize from this matrix description that we have implemented an inner product that differs from Hestenes' inner product [4]: we prefer the *contraction* defined in [7], since we found that its geometric semantics is much more convenient. We will motivate this in Section 1.4.3. We have implemented other inner products as well (see next section), but the contraction is the default.

1.2.3 The derived products

It is common to take the geometric product as basic, and define the other products using it by selecting appropriate grades. This can be the basis for an implementation; the Maple package at Cambridge [6] has been so constructed. For our comparative discussion below, we state the definitions; these can be used in a straightforward manner to derive the matrix representations.

- **Outer product**

$$\mathbf{A} \wedge \mathbf{B} = \sum_{r,s} \langle \langle \mathbf{A} \rangle_r \langle \mathbf{B} \rangle_s \rangle_{s+r}. \quad (1.5)$$

where $\langle \cdot \rangle_r$ is the grade operator taking the part of grade r of a multivector.

- **Contraction inner product**

$$\mathbf{A} \rfloor \mathbf{B} = \sum_{r,s} \langle \langle \mathbf{A} \rangle_r \langle \mathbf{B} \rangle_s \rangle_{s-r}, \quad (1.6)$$

where the grade operator for negative grades is zero (we thank Svenson [12] for this way of writing the contraction). Note that this implies that ‘something of higher grade cannot be contracted onto something of lower grade’. For scalar α and a general non-scalar multivector \mathbf{A} we get $\alpha \rfloor \mathbf{A} = \alpha \mathbf{A}$ and $\mathbf{A} \rfloor \alpha = 0$.

- **Modified Hestenes inner product**

This is a variation of the Hestenes inner product that fixes its odd behavior for scalars (which messes up the `meet` operation, as discussed in Section 1.4.3):

$$\mathbf{A} \cdot_M \mathbf{B} = \sum_{r,s} \langle \langle \mathbf{A} \rangle_r \langle \mathbf{B} \rangle_s \rangle_{|s-r|}. \quad (1.7)$$

For scalar α and a general non-scalar multivector \mathbf{A} we get $\alpha \cdot_M \mathbf{A} = \alpha \mathbf{A} = \mathbf{A} \cdot_M \alpha$.

- **Hestenes inner product**

The original Hestenes product differs from Equation 1.7 in that the contributions of the scalar parts of \mathbf{A} and \mathbf{B} are explicitly set to zero.

Note that mixed-grade multivectors require expansion of a double sum in all these products. We prefer contraction because it simplifies the `meet` and `join` operators, as we will discuss further in Section 1.4.3 after describing `meet` and `join`.

1.2.4 Representational issues in geometric algebra

For Clifford algebras of arbitrary signature (p, q) (which means $p+q$ spatial dimensions, of which p basis vectors have a positive square, and q have a negative square) linear matrix representations have long been known. We repeat part of the table of such representations in Table 1.1, see [11]. In this table, $\mathbb{R}(n)$ are $n \times n$ real matrices, $\mathbb{C}(n)$ are $n \times n$ complex-valued matrices, $\mathbb{H}(n)$ are $n \times n$ quaternion-valued matrices, ${}^2\mathbb{R}(n)$ are ordered pairs of $n \times n$ real matrices (which you may think of as a block-diagonal $2n \times 2n$ matrix containing two real $n \times n$ real matrices on its diagonal and zeroes elsewhere), and similarly for the other number systems.

The various representations are non-equivalent, so the table can be used for arguments on unique representations. Note that the Clifford algebras for

	$q = 0$	$q = 1$	$q = 2$	$q = 3$
$p = 0$	$\mathbb{R}(1)$	$\mathbb{C}(1)$	$\mathbb{H}(1)$	${}^2\mathbb{H}(1)$
$p = 1$	${}^2\mathbb{R}(1)$	$\mathbb{R}(2)$	$\mathbb{C}(2)$	$\mathbb{H}(2)$
$p = 2$	$\mathbb{R}(2)$	${}^2\mathbb{R}(2)$	$\mathbb{R}(4)$	$\mathbb{C}(4)$
$p = 3$	$\mathbb{C}(2)$	$\mathbb{R}(4)$	${}^2\mathbb{R}(4)$	$\mathbb{R}(8)$
$p = 4$	$\mathbb{H}(2)$	$\mathbb{C}(4)$	$\mathbb{R}(8)$	$\mathbb{C}(8)$

TABLE 1.1. Matrix representations of Clifford algebras of signatures (p, q) .

a 3-dimensional space can have many different representations depending on the signature. Though this is not a problem for implementations, it makes it harder to obtain a parametric overview on the metric properties of the various spaces, and a representation that contains the signature as parameters σ_i has our slight preference.

The outer product forms a Grassmann algebra. We have been unable to find a similar representation table in the literature on Grassmann algebras, but at least for even-dimensional Grassmann algebras this is easily established. A Clifford algebra with signature (p, p) can be converted into a $2p$ -dimensional Grassmann algebra by pairing of the basis vectors with positive and negative signature. The table then shows that the Grassmann algebra of even dimension p is isomorphic to $\mathbb{R}(2^p)$. Odd dimensions can be seen as subalgebras of the next even dimension.

The inner product is not associative, and is therefore not isomorphic to a matrix algebra.

Our initial exclusive interest in $\mathcal{C}_{3,0}$ suggests the representation $\mathbb{C}(2)$, with elements represented as

$$\left[\begin{array}{cc} (A_0 + A_3) + i(A_{12} + A_{123}) & (A_1 + A_{31}) + i(-A_2 + A_{23}) \\ (A_1 - A_{31}) + i(A_2 + A_{23}) & (A_0 - A_3) + i(-A_{12} + A_{123}) \end{array} \right],$$

but this works only for the geometric product; the other products would then have to be implemented using the grade operator. We prefer a representation in which all three products are representable on a par, and in which signatures are parameterized. This desire to represent arbitrary signatures parametrically necessitates viewing $\mathcal{C}_{3,0}$ as a subalgebra of $\mathcal{C}_{3,3}$, and therefore to choose a representation in $\mathbb{R}(8)$.

This algebra $\mathbb{R}(8)$ also contains a representation of the outer product as a certain kind of lower-triangular matrices (in fact, Equation 1.2 works nicely: the matrix product of two such matrices faithfully represents the outer product). For arbitrary signatures, there cannot exist a change of representation in which both the outer product matrices and the geometric product matrices could be reduced to a smaller size (i.e., brought onto a block-diagonal representation of the same kind), since we need the full $\mathbb{R}(8)$ to handle those signatures anyway.

Now the need to represent the inner product as well indicates that we can not represent the elements of the algebra by matrices in their function as both *operator* (i.e., first factor) and *operand* (i.e., second factor). We therefore switch to the view where each product is seen as a linear function of the operand, parameterized by the operator, as detailed in Section 1.2.1. We maintain the $\mathbb{R}(8)$ -representation of these linear functions, but they now operate on 8-dimensional vectors representing the operand (rather than forming an algebra of operators). Thus we arrive at the representation we have chosen (also for the geometric product), with the operator matrices naturally defined as in Equation 1.1.

It should be clear that the same reasoning suggests an $\mathbb{R}(2^n)$ representation of the geometric algebra of n -dimensional space of arbitrary signatures, with matrices defined for the three products in the same way.

1.2.5 Computational efficiency

If we would represent our objects as 8×8 matrices of reals, the resulting matrix multiply to implement the geometric product would cost 512 multiplications and 448 additions. Further, using the 8×8 matrix representation, to compute the outer product and/or inner product, we would have to use the grade operator (or, for the outer product, pay the expansion cost to convert to the outer product 8×8 matrix representation). Addition and scalar multiplication of elements in this form require 64 additions and 64 multiplications respectively. This method is extremely inefficient and we will not discuss it further.

The computational efficiency of the 8×1 format is better and is summarized in Table 1.3, where the notation (a, m) denotes the number of additions and multiplications required. With the 8×1 format, the cost of computing a product is the computational cost of having to expand one of the one of the 8×1 matrices to an 8×8 matrix and then multiply it by an 8×1 matrix at a cost of 64 multiplications, 56 additions, and the cost of expansion. When we include the cost of signatures in the expansion cost, then the total cost is increased by 48 multiplications.

It is of course possible to use the table of Clifford algebra isomorphisms as a literal guide to the implementation. Let us consider the costs of implementing the special case of the 3-dimensional Clifford algebras; Table 1.1 shows that this involves implementation of $\mathbb{C}(2)$, ${}^2\mathbb{R}(2)$ and ${}^2\mathbb{H}(1)$. In all representations the operations of addition and scalar multiplication have take 8 floating point additions and 8 floating point multiplications, respectively (in the $\mathbb{R}(8)$ representation, these operations are performed on the 8×1 matrices representing the objects).

To compute the geometric product we need to multiply elements. The complexity of the basic multiplications is: one complex multiply is $(4, 2)$; one double real multiply takes $(2, 0)$; one quaternion multiply takes $(16, 12)$. For a full matrix implementation to produce the geometric product this

$a \wedge b$	0	1	2	3
0	0	1	2	3
1	1	2	3	(4)
2	2	3	(4)	(5)
3	3	(4)	(5)	(6)

TABLE 1.2. Grade of the outer product of blades. Along the top row and left column are the grades of blades a and b ; the table gives the grade of the outer product of these blades. Note that if the grade is greater than 3, then the result will be 0.

yields for $\mathbb{C}(2)$ a complexity of (32, 16); for ${}^2\mathbb{R}(2)$ a complexity of (16, 8); for ${}^2\mathbb{H}(1)$ a complexity of (32, 24). Depending on the structure of the algebra, one may thus be fortunate by a factor of two. These should be compared to our $\mathbb{R}(8)$ implementation acting on 8×1 matrices, which has a complexity of (110, 56), for general signatures. This is a factor of 3 worse than ${}^2\mathbb{H}(1)$, the most expensive of the other three representations. If we consider only $\mathcal{C}_{3,0}$, then we have no signature cost, and $\mathbb{R}(8)$ costs (64, 56) compared to (32, 16) for $\mathbb{C}(2)$. Table 1.3 compares the cost of several operations of our $\mathbb{R}(8)$ implementation without signature cost and of $\mathbb{C}(2)$.

To implement a full geometric algebra, these specific geometric product implementations need to be augmented with a grade operation to extract elements of the grade desired, according to Equations 1.6 and 1.5. For $\mathbb{C}(2)$ it takes (0, 8) to extract the eight separate elements, and presumably the same for ${}^2\mathbb{H}(1)$ and ${}^2\mathbb{R}(2)$. For simplicity of discussion, when extracting a single grade, we will assume that it costs 3 additions (although for scalars and trivectors, the cost is only 1 addition).

This process of a geometric product followed by grade extraction is simple if the objects to be combined are *blades* (rather than general multivectors). Such an operation requires a geometric product followed by grade extraction, which for $\mathbb{C}(2)$ has a total worst case cost of (32, 19), although there may be some additional cost to test for the grade of the blade, etc., which would add (0, 16) to the cost ((32, 35) total) if we need to perform a full grade extraction of each operand.

When taking the outer or inner product of multivectors that are *not* blades, the use of the geometric product and grade extraction becomes quite expensive, since we must implement a double sum (see Equations 1.5, 1.6, and 1.7). A naive implementation of this formula would require 16 geometric products and grade extractions, an additional 12 additions to combine the results for each grade, and 8 additions to reconstruct the result, for a total cost of (512, 324). However, looking at Table 1.2, we see that that six of these geometric products will always be 0, and we can easily rewrite our code to take advantage of this. This modification to the code reduces the cost to (320, 210).

By unbundling the loop and simplifying the scalar cases (i.e., multiplying

Operation	8×1	$\mathbb{C}(2)$
addition	(0, 8)	(0, 8)
scalar multiplication	(8, 0)	(8, 0)
grade extraction	(0, 0)	(0, 8)
geometric product	(64, 56)	(32, 16)
other products of blades	(64, 56)	(32, 19)
other products of multivectors	(64, 56)	(111, 79)

TABLE 1.3. Comparison of costs for 8×1 and $\mathbb{C}(2)$.

B by the scalar portion of A reduces 4 geometric products to one floating point addition (to extract the scalar) and 8 floating point multiplies, and multiplying A by the scalar portion of B reduces 3 more geometric products to one addition and 7 floating point multiplies) we can get the cost down to $(32 * 3 + 15, 19 * 3 + 2 + 12 + 8) = (111, 79)$. Further special casing of the vector and bivector terms can reduce this cost to $(33, 45)$ (details of the analysis can be found in [10]), but note that in doing this (a) we have left the complex representation for computing these products and (b) each product will need its own special case code.

Note that the above discussion is on the cost of writing special case code for the outer product only. If we choose this route, we would also need to write special case code for each of the inner products and possibly for each dimensional space in which we wish to work. A reasonable compromise of special cases versus general code for the complex representation would be to handle the scalars as special cases and write the loops to avoid the combinations that will always give zero. Table 1.3 compares the costs of using the 8×1 representation and the $\mathbb{C}(2)$ representation, assuming we do the these optimizations for the $\mathbb{C}(2)$ products of multivectors.

Asymptotic costs

If we are interested in arbitrary dimensional spaces, then we need to look at the asymptotic costs. Table 1.4 summarizes the costs of the complex and of the $n \times 1$ representation (where $n = p + q$ is the dimension of the underlying vector space) for the geometric product and for the other products (e.g, inner and outer products) on blades and for the other products on general multivectors. In this table, we only give the top term in the cost expression, ignoring grade extraction, etc., for the complex representation of other products. Use of only this top order term also ignores the savings achieved for the complex representation by not computing the products whose grade is higher than n and special casing the scalar products; such optimizations roughly equate to a factor of two savings. Note that we use the complex representation as a coarse representative of the other representations; in the other cases we would use the quaternion or double-real representation, which cost roughly a factor of 2 less than the complex representation.

	Geometric Product	Other products on blades	Other products on multivectors
Complex	$2^{(3n+1)/2}$	$2^{(3n+1)/2}$	$n^2 2^{(3n+1)/2}$
$n \times 1$	2^{2n}	2^{2n}	2^{2n}

TABLE 1.4. Comparison of costs of various methods, with n being the dimension of the underlying vector space.

From the table, we see that asymptotically the complex representation is always best. However, substituting number in these equations shows that for small n , the $n \times 1$ representation is best when performing inner or outer products of general multivectors, with the cross-over point being around $n = 14$. But when n is 14, the cost of even the geometric product in the complex representation is extremely large, requiring roughly 3×10^6 multiplications.

For smaller n , the complex representation is better than the $n \times 1$ representation for the geometric product and the products of blades, while the $n \times 1$ representation is computationally less expensive than the complex representation for the other products of general multivectors. However the other products of general multivectors are rarely (if ever) performed in our present understanding of what constitute geometrically significant combinations. Thus, in general the complex/quaternion/double-real representation will be more efficient than the $n \times 1$ representation by a factor of $2^{n/2}$. The conclusion must be that once one has decided on a particular geometry for one's application, reflected in a particular signature, it makes sense to implement it literally using the isomorphism of Table 1.1.

For our tutorial in 3 dimensional spaces, the cost of the 8×1 representation is only a factor of three more expensive than the complex representation. Since we were writing tutorial code, we felt this cost was more than offset by the explicitness of the signature and ease of implementation.

1.3 Inverses

In Matlab, the obvious way to compute the inverse of a geometric object \mathbf{M} is to express it in the 8×8 geometric product matrix representation, $[\mathbf{M}]$. Then inversion of $[\mathbf{M}]$ may be done using the Matlab matrix inverse routine, and the first column of $[\mathbf{M}]^{-1}$ will be the representation of the inverse of \mathbf{M} . However, when we implemented this method for computing the inverse, we found that it introduced small numerical errors on rather simple data, and thus was less stable than we would like. We investigated a method of Lounesto's that was more stable in our testing, and is computationally considerably more efficient than a matrix inverse.

Lounesto [8] (pag. 57) proposes a method to compute inverses in Clifford algebras of 3-dimensional spaces. We discuss it now, and extend it slightly.

Lounesto's trick is based on the observation that in three dimensions (and that is essential!) the product of a multivector \mathbf{M} and its Clifford conjugate $\overline{\mathbf{M}}$ only has two grades, a scalar and a pseudoscalar (the Clifford conjugate is the grade involution of the reverse of a multivector). Let \mathbf{M}_i denote the part of \mathbf{M} of grade i , though we will write M_0 for the scalar part. Then we compute

$$\begin{aligned}\mathbf{M}\overline{\mathbf{M}} &= (M_0 + \mathbf{M}_1 + \mathbf{M}_2 + \mathbf{M}_3)(M_0 - \mathbf{M}_1 - \mathbf{M}_2 + \mathbf{M}_3) \\ &= (M_0^2 - \mathbf{M}_1^2 - \mathbf{M}_2^2 + \mathbf{M}_3^2) + 2(M_0\mathbf{M}_3 - \mathbf{M}_1 \wedge \mathbf{M}_2),\end{aligned}$$

and the first bracketed term is a scalar, the second a trivector.

Further, at least in Euclidean 3-space, if such an object of the form 'scalar plus trivector' $N_0 + \mathbf{N}_3$ is non-zero, then it has an inverse that is easily computed:

$$(N_0 + \mathbf{N}_3)^{-1} = \frac{N_0 - \mathbf{N}_3}{N_0^2 - \mathbf{N}_3^2}.$$

Please note that not all multivectors have an inverse, not even in a Euclidean space: for instance $\mathbf{M} = 1 + \mathbf{e}_1$ leads to $\mathbf{M}\overline{\mathbf{M}} = 0$, so this \mathbf{M} is non-invertible. In a non-Euclidean space, the denominator may become zero even when N_0 and \mathbf{N}_3 are not, and we need to demand at least that $N_0^2 \neq \mathbf{N}_3^2$. (When it exists, the inverse is unique. This follows using the associativity of the geometric product: if \mathbf{A} and \mathbf{A}' are left and right inverses of \mathbf{B} , respectively, then $\mathbf{A} = \mathbf{A}(\mathbf{B}\mathbf{A}') = (\mathbf{A}\mathbf{B})\mathbf{A}' = \mathbf{A}'$. Therefore any left inverse is a right inverse, and both are identical to *the* inverse.)

These two facts can be combined to construct an inverse for an arbitrary multivector \mathbf{M} (still in Euclidean 3-space) as follows:

$$\mathbf{M}^{-1} = \overline{\mathbf{M}}\overline{\mathbf{M}}^{-1}\mathbf{M}^{-1} = \overline{\mathbf{M}}(\mathbf{M}\overline{\mathbf{M}})^{-1} = \frac{\overline{\mathbf{M}}((\mathbf{M}\overline{\mathbf{M}})_0 - (\mathbf{M}\overline{\mathbf{M}})_3)}{(\mathbf{M}\overline{\mathbf{M}})_0^2 - (\mathbf{M}\overline{\mathbf{M}})_3^2}$$

The following two lemmas and their proofs demonstrate the correctness of Lounesto's method in 3-dimensional spaces of arbitrary signature.

Lemma 1.1. \mathbf{M}^{-1} exists if and only if $(\mathbf{M}\overline{\mathbf{M}})^{-1}$ exists.

Proof: First, assume that \mathbf{M}^{-1} exists. Then $1 = \overline{\mathbf{M}^{-1}}\overline{\mathbf{M}} = (\overline{\mathbf{M}^{-1}}\mathbf{M}^{-1})(\mathbf{M}\overline{\mathbf{M}})$, so that $(\mathbf{M}\overline{\mathbf{M}})^{-1} = \overline{\mathbf{M}^{-1}}\mathbf{M}^{-1}$, which exists.

Secondly, assume that $(\mathbf{M}\overline{\mathbf{M}})^{-1}$ exists. Then we have $1 = (\mathbf{M}\overline{\mathbf{M}})(\mathbf{M}\overline{\mathbf{M}})^{-1} = \mathbf{M}(\overline{\mathbf{M}}(\mathbf{M}\overline{\mathbf{M}})^{-1})$, so that $\mathbf{M}^{-1} = \overline{\mathbf{M}}(\mathbf{M}\overline{\mathbf{M}})^{-1}$, which exists. \square

Lemma 1.2. Let $\mathbf{N} = N_0 + \mathbf{N}_3$. Then iff $N_0^2 \neq \mathbf{N}_3^2$, \mathbf{N}^{-1} exists and equals

$$(N_0 + \mathbf{N}_3)^{-1} = \frac{N_0 - \mathbf{N}_3}{N_0^2 - \mathbf{N}_3^2}$$

Proof: Assume $N_0^2 \neq \mathbf{N}_3^2$, then $(N_0 + \mathbf{N}_3)(N_0 - \mathbf{N}_3)/(N_0^2 - \mathbf{N}_3^2) = (N_0^2 +$

Term	8×1			$\mathcal{C}(2)$		
	Naive	Good	Scalar	Naive	Good	Scalar
\mathbf{M}	(0,0)	(0,0)	(0,0)	(0,8)	(0,8)	(0,8)
$\overline{\mathbf{M}\mathbf{M}}$	(64,56)	(64,56)	(64,56)	(32,16)	(32,16)	(32,16)
$\overline{\mathbf{M}\mathbf{M}}^{-1}$	(4,1)	(4,1)	(0,0)	(4,1)	(4,1)	(0,0)
$\mathbf{M}(\overline{\mathbf{M}\mathbf{M}})^{-1}$	(64,56)	(16,8)	(8,0)	(32,16)	(16,8)	(8,0)
Total	(132,113)	(84,65)	(72,56)	(64,41)	(52,33)	(40,24)

TABLE 1.5. Cost of Lounesto's inverse.

$\mathbf{N}_3 N_0 - N_0 \mathbf{N}_3 - \mathbf{N}_3^2)/(N_0^2 - \mathbf{N}_3^2) = (N_0^2 - \mathbf{N}_3^2)/(N_0^2 - \mathbf{N}_3^2) = 1$, so \mathbf{N}^{-1} is as stated.

Now assume that \mathbf{N}^{-1} exists. Then if $\mathbf{N}_3 = 0$ the result is trivial. If $\mathbf{N}_3 \neq 0$ and $N_0 = 0$ the result is trivial. So take $\mathbf{N}_3 \neq 0$ and $N_0 \neq 0$. Let \mathbf{K} be the inverse of $\mathbf{N} = N_0 + \mathbf{N}_3$. Then it needs to satisfy

$$(N_0 + \mathbf{N}_3)(K_0 + \mathbf{K}_1 + \mathbf{K}_2 + \mathbf{K}_3) = 1,$$

so, written out in the different grades

$$(N_0 K_0 + \mathbf{N}_3 \mathbf{K}_3) + (N_0 \mathbf{K}_1 + \mathbf{N}_3 \mathbf{K}_2) + (N_0 \mathbf{K}_2 + \mathbf{N}_3 \mathbf{K}_1) + (N_0 \mathbf{K}_3 + \mathbf{N}_3 K_0) = 1$$

Straightforward algebra on the terms of grade 0 and 3 yields $(N_0^2 - \mathbf{N}_3^2)\mathbf{K}_3 + \mathbf{N}_3 = 0$, and since $\mathbf{N}_3 \neq 0$ this gives $N_0^2 \neq \mathbf{N}_3^2$. Then the case above shows that the inverse is $\mathbf{N}^{-1} = (N_0 - \mathbf{N}_3)/(N_0^2 - \mathbf{N}_3^2)$. \square

Table 1.5 summarizes the costs to compute the inverse for both the 8×1 representation and for the $\mathcal{C}(2)$ representation. In this table, we give three algorithms for each representation: a naive algorithm, that does not try to exploit any extra knowledge we have about the terms we are manipulating; a good algorithm, that exploits the structure of $(\overline{\mathbf{M}\mathbf{M}})^{-1}$, which is a scalar plus a pseudo-scalar, and thus does not require a full product when multiplied by \mathbf{M} ; and a scalar version that can be used when $(\overline{\mathbf{M}\mathbf{M}})^{-1}$ is a scalar. This last case occurs when \mathbf{M} is a blade, a scalar plus a bivector, or a vector plus the pseudo-scalar, which covers most of the geometrically significant objects we manipulate.

Note that in this table we have omitted the cost of the six negations needed to compute the Clifford conjugate. Also note that the complex representation requires 8 additions when computing the Clifford conjugate because it has to separate and recombine the scalar and pseudo-scalar part of the geometric object.

Lounesto's method is computationally much cheaper than the matrix inverse method, with a good implementation of Lounesto's method requiring 149 Matlab floating point operations for the 8×1 representation, while the Matlab matrix inverse routine on 8×8 matrices requires 1440 Matlab

floating point operations. Lounesto's method really makes convincing use of the special structure of our matrices. While a faster matrix inversion routine may be available, it is unlikely that there will be a general routine capable of inverting our special 8×8 matrix in fewer than 149 floating point operations (which is after all little more than twice the number of matrix elements!). Further, in practice we found our modified Lounesto inverse to compute a more numerically stable inverse than the matrix inverse routine provided by Matlab (perhaps not surprising, since it involves fewer operations).

Had we used the $\mathbb{C}(2)$ representation of elements in our geometric algebra, the cost of matrix inversion would have dropped dramatically, with Matlab requiring only 260 floating point operations to invert a 2×2 complex matrix. However, Lounesto's method using the complex representation only requires 75 floating point operations. Thus Lounesto's inversion method is also less expensive in the $\mathbb{C}(2)$ representation.

1.4 Meet and Join

The geometric intersection and union of subspaces is done by the `meet` and `join` operations. These have mostly been used by others in the context of *projective geometry*, which has led to the neglect of some scalar factors and signs (since they do not matter in that application). This issue was partly treated in [2], but the development of the tutorial required some more investigation of those scalar factors. This section reports on that.

1.4.1 Definition

The `meet` and `join` operations are geometrical 'products' of a higher order than the elementary products treated before. They are intended as geometrical intersection and union operators on (sub)spaces of the algebra. Since subspaces are represented by pure blades, *these operations should only be applied to blades*.

Let blades \mathbf{A} and \mathbf{B} contain as a common factor a blade \mathbf{C} of maximum grade (this is like a 'largest common divisor' in the sense of the geometric product), so that we can write

$$\mathbf{A} = \mathbf{A}' \wedge \mathbf{C} \quad \text{and} \quad \mathbf{B} = \mathbf{C} \wedge \mathbf{B}'$$

(note the order!). We will actually choose \mathbf{A}' and \mathbf{B}' to be perpendicular to \mathbf{C} , so that we may also write the factorization in terms of the geometric product: $\mathbf{A} = \mathbf{A}'\mathbf{C}$ and $\mathbf{B} = \mathbf{C}\mathbf{B}'$ (but note that \mathbf{A}' and \mathbf{B}' are in general *not* mutually perpendicular!). If \mathbf{A} and \mathbf{B} are disjoint, then \mathbf{C} is a scalar (a 0-blade). We now define `meet` and `join` as

$$\text{join}(\mathbf{A}, \mathbf{B}) = \mathbf{A}' \wedge \mathbf{C} \wedge \mathbf{B}' \quad \text{and} \quad \text{meet}(\mathbf{A}, \mathbf{B}) = \mathbf{C}.$$

Note that the factorization is *not* unique: we may multiply \mathbf{C} by a scalar γ . This affects the `join` result by $1/\gamma$ and the `meet` by γ , so `meet` and `join` are *not well-defined*. (Since γ may be negative, not even the orientation of the results is defined unambiguously.) So these operations are hard to define in a Clifford algebra; but for a Geometric Algebra, they are definitely desired. Many geometric constructions are actually insensitive to the magnitude and/or sign of the blade representing the subspace. A prime example is the projection $(\mathbf{x} \rfloor \mathbf{A}) / \mathbf{A}$ onto the subspace represented by \mathbf{A} – there is not problem using for \mathbf{A} the outcome of a `meet` or `join`.

In our implementation, we do want to guarantee that `meet` and `join` of the same subspaces can be used consistently, so we do need to base both on the same factorization. We can make the computational relationships between `meet` and `join` explicit. The definition gives for the `join`, given the `meet` (where the fraction denotes right-division):

$$\text{join}(\mathbf{A}, \mathbf{B}) = \frac{\mathbf{A}}{\text{meet}(\mathbf{A}, \mathbf{B})} \wedge \mathbf{B}. \quad (1.8)$$

Note that this is only valid if the `meet` is an invertible blade. In non-Euclidean spaces there may therefore be a problem with this equation and the factorization on which it is built: if \mathbf{C} is a null blade (i.e., a blade with norm 0, non-invertible) then we cannot compute \mathbf{A}' in terms of \mathbf{A} from the factorization equation $\mathbf{A} = \mathbf{A}'\mathbf{C}$, and therefore not compute $\text{join}(\mathbf{A}, \mathbf{B}) = \mathbf{A}' \wedge \mathbf{B}$ from the `meet` (or vice versa, by a similar argument). We thus have to limit `join` and `meet` to non-null blades; which means that we restrict ourselves to Euclidean spaces only. (Actually, anti-Euclidean spaces in which all signatures are -1 would obviously be permissible as well.) Since no blades are now null-blades, we can agree to make the common factor \mathbf{C} a *unit blade* (so that $|\mathbf{C}| = 1$) leaving only the sign of its orientation undetermined. But please be aware that this is a rather arbitrary partial fixing of the scalar factor!

By duality relative to $\text{join}(\mathbf{A}, \mathbf{B})$ and symmetry of a scalar-valued contraction (or inner product) it follows from Equation 1.8 that

$$1 = \frac{\mathbf{A}}{\text{meet}(\mathbf{A}, \mathbf{B})} \rfloor \frac{\mathbf{B}}{\text{join}(\mathbf{A}, \mathbf{B})} = \frac{\mathbf{B}}{\text{join}(\mathbf{A}, \mathbf{B})} \rfloor \frac{\mathbf{A}}{\text{meet}(\mathbf{A}, \mathbf{B})}.$$

The division by $\text{meet}(\mathbf{A}, \mathbf{B})$ can be factored out (this is due to the containment relationship of the factors of the contraction and easy to prove using the techniques in [2]) and we obtain

$$\text{meet}(\mathbf{A}, \mathbf{B}) = \frac{\mathbf{B}}{\text{join}(\mathbf{A}, \mathbf{B})} \rfloor \mathbf{A}. \quad (1.9)$$

Thus we can start from either `meet` or `join` and compute the other in a consistent manner. The symmetry of the equations means that either way

$\text{join}(\mathbf{A}, \mathbf{B})$	0	1	2	3
0	0	1	2	3
1	1	2(1)	3(2)	3
2	2	3(2)	3(2)	3
3	3	3	3	3

$\mathbf{A} \wedge \mathbf{B}$	0	1	2	3
0	0	1	2	3
1	1	2(\emptyset)	3(\emptyset)	\emptyset
2	2	3(\emptyset)	\emptyset	\emptyset
3	3	\emptyset	\emptyset	\emptyset

$(\mathbf{B}/\mathbf{I}_3) \rfloor \mathbf{A}$	0	1	2	3
0	3	2	1	0
1	2	1(\emptyset)	0(\emptyset)	1
2	1	0(\emptyset)	1(\emptyset)	2
3	0	1	2	3

TABLE 1.6. The result of the `join` in 3-space can often be computed using the wedge product.

is equally feasible.² For a more detailed discussion on algebra of incidence see chapter 3 and chapter 7.

1.4.2 Implementation

We saw that the three issues, factorization of \mathbf{A} and \mathbf{B} , computing their `join` (smallest containing superspace) and computing their `meet` (largest common subspace) are intertwined; giving any one determines the other two (at least in Euclidean signatures).

We have chosen to use the `join` (i.e., the smallest common space of \mathbf{A} and \mathbf{B}) as the one to implement, and to base the `meet` on it using Equation 1.9. In principle, this determination of the smallest common space is a minimization problem, which may be solved by starting with a space that is too big and reducing it, or by growing one that is too small. In either case, the general case will involve some administration of polynomial time in the number of blades, and therefore exponential in the dimensionality of the space. We have not solved this general issue; in the 3-dimensional Euclidean space of interest in the tutorial the `join` is fairly easy to implement case by case.

In Table 1.6 we show the dimension of the 3-space `join` operator on operands of different dimensions. In this table, we have indicated the grade of the results, with \emptyset indicating a zero result, and results in brackets indicating alternative outcomes in degenerate cases. First observe that the definition implies that for *disjoint* spaces \mathbf{A} and \mathbf{B} , factored by a scalar

²Equation 1.9 is frequently extended to provide a 3-argument `meet` function relative to a general blade \mathbf{I} : $\text{meet}(\mathbf{A}, \mathbf{B}, \mathbf{I}) \equiv (\mathbf{B}/\mathbf{I}) \cdot \mathbf{A}$. However, since the geometric significance of using anything but `join`(\mathbf{A}, \mathbf{B}) as third argument is unclear, we will not use it. Also, beware that some writers may switch the order of the arguments in this formula!

$\mathbf{C} = 1$, $\text{join}(\mathbf{A}, \mathbf{B})$ equals $\mathbf{A} \wedge \mathbf{B}$. In particular, we see in Table 1.6 that $\mathbf{A} \wedge \mathbf{B}$ equals $\text{join}(\mathbf{A}, \mathbf{B})$ unless the dimensions of \mathbf{A} and \mathbf{B} are too high (sum exceeds 3), with some exceptional degeneracies when the grades are 1 and 2. So we may use the outer product as a basis for an algorithm. The table shows that of $(\mathbf{B}/\mathbf{I}_3)]\mathbf{A}$ may aid in treating some of the non-outer-product cases, where \mathbf{I}_3 is the pseudoscalar of our 3-dimensional Euclidean space (details below).

This has led us to consider two algorithms for the computation of the `join`:

- *Algorithm A*

For non-degenerate arguments, we can implement $\text{join}(\mathbf{A}, \mathbf{B})$ by computing the quantity $\mathbf{J} = \mathbf{A} \wedge \mathbf{B}$. If \mathbf{J} is non-zero, then it is the result we want. Otherwise, if the grade of one of the arguments equals 3, the result is proportional to this argument. For instance, let $\text{grade}(\mathbf{A}) = 3$, then a possible factorization of \mathbf{B} is through $\mathbf{B} = \mathbf{C}\beta$ (with $\mathbf{B}' = \beta$ a scalar) which yields $\text{join}(\mathbf{A}, \mathbf{B}) = \mathbf{A} \wedge \mathbf{B}' = \beta\mathbf{A}$. If we choose the common factor to be a unit blade, then $\beta = \pm|\mathbf{B}|$, so that the result is $\text{join}(\mathbf{A}, \mathbf{B}) = \pm|\mathbf{B}|\mathbf{A}$. We choose, arbitrarily, the positive sign.

That leaves the exceptions. When not both grades are 2, the result is proportional to the argument of highest grade, by a scalar factor depending on the other argument (by the same reasoning as above, taking that scalar factor equal to the norm implies considering the common factor to be a unit blade). When both grades are 2, we need to find whether these 2-blades are coincident or not. If they are not, then their `join` is proportional to \mathbf{I}_3 , so we may use Equation 1.8 to compute a carrier for this common subspace: $\mathbf{M} = (\mathbf{B}/\mathbf{I}_3)]\mathbf{A}$. We normalize this to a unit blade: $\mathbf{C} = \mathbf{M}/|\mathbf{M}|$, and then return $\text{join}(\mathbf{A}, \mathbf{B}) = (\mathbf{A}/\mathbf{C}) \wedge \mathbf{B}$ as the proper multiple of \mathbf{I}_3 . If they are coincident, the computation of \mathbf{M} yields zero (which is how we may detect it) and we return $|\mathbf{A}|\mathbf{B}$ or, equivalently up to the undetermined sign, $|\mathbf{B}|\mathbf{A}$.

- *Algorithm B*

For non-degenerate arguments, we can implement $\text{join}(\mathbf{A}, \mathbf{B})$ by computing the quantity $\mathbf{J} = \mathbf{A} \wedge \mathbf{B}$. If \mathbf{J} is non-zero, then it is the result we want. Otherwise, we compute $\mathbf{M} = (\mathbf{B}/\mathbf{I}_3)]\mathbf{A}$. If \mathbf{M} is non-zero, then it is proportional to the `meet` (since precisely in those cases, the `join` is proportional to the blade \mathbf{I}_3 of grade 3). The common factor \mathbf{C} is then the unit blade $\mathbf{C} = \mathbf{M}/|\mathbf{M}|$, so the `join` is then $\mathbf{A}' \wedge \mathbf{C} \wedge \mathbf{B}' = (\mathbf{A}/\mathbf{C}) \wedge \mathbf{B}$.

In the degenerate cases indicated in parentheses in the table for the `join`, both \mathbf{J} and \mathbf{M} of the previous paragraph are zero. Which degenerate case we have is readily determined by testing the grade of \mathbf{A} and

B. If both are vectors or bivectors, then they must be parallel. The factoring is thus $\mathbf{A} = \alpha\mathbf{C}$ and $\mathbf{B} = \mathbf{C}\beta$, with both α and β scalars. The result of the `join` is then $\text{join}(\mathbf{A}, \mathbf{B}) = \alpha\beta\mathbf{C} = \alpha\mathbf{B} = \beta\mathbf{A}$. We can implement this as $|\mathbf{A}|\mathbf{B}$ or $|\mathbf{B}|\mathbf{A}$, if we agree to factor out a unit blade \mathbf{C} . If exactly one of \mathbf{A} and \mathbf{B} is a vector \mathbf{c} , then the other must be a bivector containing this vector as a factor. The factorization is now $\mathbf{A} = \mathbf{a}\mathbf{c}$ and $\mathbf{B} = \beta\mathbf{c}$, so $\text{join}(\mathbf{A}, \mathbf{B}) = \beta\mathbf{a}\mathbf{c} = \beta\mathbf{A}$ (if \mathbf{A} is the bivector) or $\mathbf{A} = \alpha\mathbf{c}$ and $\mathbf{B} = \mathbf{c}\mathbf{b}$, so $\text{join}(\mathbf{A}, \mathbf{B}) = \alpha\mathbf{B}$ (if \mathbf{B} is the bivector). If we fix the common blade \mathbf{c} to be a unit blade, this may be implemented as $|\mathbf{B}|\mathbf{A}$ or $|\mathbf{A}|\mathbf{B}$, respectively.

Algorithm A is computationally faster since it mostly does testing of grades to establish the exceptional cases. Algorithm B has a simpler conditional structure, leading to simpler code. It is the one we implemented.

With the `join` found, the `meet` is computed from Equation 1.9. Although either is only determined only up to a scalar, they are consistent in the sense of those equations, and their relative magnitudes may therefore be used to derive geometrically meaningful results.

1.4.3 Why we use a contraction as inner product

We gave three different inner product definitions in Section 1.2.3, and we still owe the explanation on why we prefer the contraction, which has been used so rarely in geometric algebra. The main reason is that Hestenes' original inner product (abbreviated as HIP) has some features that make it less suitable for straightforward geometric interpretations. This shows up rather clearly when it is used in the `meet` operation, and in projection operations. The former can be fixed by treating the scalars differently, the latter requires more and leads to the contraction.

Suppose we take the `meet` of \mathbf{I}_3 and a vector \mathbf{a} , in the 3-dimensional space with pseudoscalar \mathbf{I}_3 . We would obviously expect the outcome to be a multiple of \mathbf{a} , since the `meet` should have the semantics of geometric intersection, and the intersection of the subspace spanned by \mathbf{I}_3 and the subspace spanned by \mathbf{a} should be the same subspace spanned by \mathbf{a} . The `join` of \mathbf{I}_3 with any subspace is \mathbf{I}_3 , so we may use Equation 1.9 to compute the `meet`. Using the Hestenes inner product, denoted as \cdot_H , we obtain:

$$\text{meet}_H(\mathbf{a}, \mathbf{I}_3) = (\mathbf{I}_3/\mathbf{I}_3) \cdot_H \mathbf{a} = 1 \cdot_H \mathbf{a} = 0,$$

since the HIP with a scalar is zero. On the other hand

$$\text{meet}_H(\mathbf{I}_3, \mathbf{a}) = (\mathbf{a}/\mathbf{I}_3) \cdot_H \mathbf{I}_3 = \mathbf{a}.$$

So the `meetH` is severely asymmetrical, which is unexpected for an operation that should be geometric intersection. In this case, it is due to the awkward properties of scalars (which [5] page 20 notes, but does not fix).

We can fix this by modifying the Hestenes inner product to a new inner product denoted \cdot_M , the same as \cdot_H except for scalars. For scalars, we demand that $\alpha \cdot_M u = \alpha u$ for scalar α , and any multivector u . This leads to the modified Hestenes inner product defined in Equation 1.7. We will abbreviate it as modified HIP.

For non-scalars this modified HIP has a certain symmetry in grades of the arguments: the inner product of a blade of grade r with one of grade s , or *vice versa* is a blade of grade $|r-s|$. The contraction of Equation 1.6 does *not* have this property: it actually ‘contracts’ the first argument inside the second, one cannot contract something of a bigger grade onto something of a smaller grade. For its use in the `meet`, this is a distinction without a difference, since in the evaluation of Equation 1.9 as $(\mathbf{B}/\text{join}(\mathbf{A}, \mathbf{B})) \rfloor \mathbf{A}$, the first argument blade of the inner product has a grade that never exceeds that of the second argument blade.

But the selection of the inner product also has an effect on the evaluation of the *projection*. For blades, Hestenes and Sobczyk [5] define the projection of \mathbf{A} into \mathbf{B} as $(\mathbf{A} \cdot_H \mathbf{B}) \cdot_H \mathbf{B}^{-1}$. Some problems with scalars are noted (see [5], page 20) which we can fix by using \cdot_M instead. Using that but following the reasoning of [5], we can then show that the projection can be simplified to $(\mathbf{A} \cdot_M \mathbf{B})/\mathbf{B}$ if $\text{grade}(\mathbf{A}) \leq \text{grade}(\mathbf{B})$, and zero otherwise. Rightly, [5] prefers the new algebraic form since it makes proofs easier. Yet there is still this conditional split into cases; when treating blades of grades not known *a priori* this may lead to lots of cases and still make work hard.

Using the *contraction*, the projection onto the blade \mathbf{B} can be defined as $(\mathbf{A} \rfloor \mathbf{B})/\mathbf{B}$, for *all* arguments \mathbf{A} . This is automatically zero when \mathbf{A} exceeds \mathbf{B} ; since the algebraic properties of the contraction are similar to those of the (extended) HIP, most proofs still work unchanged in their mechanics, but now do not require careful conditional splits dependent on the grades of the arguments complicating the reasoning. Geometrically interpreted, the contraction implicitly contains subspace relationships, in the sense that the blades resulting from the simple formula $\mathbf{A} \rfloor \mathbf{B}$ *must* be contained in \mathbf{B} , and so must the result of the division by \mathbf{B} to produce the actual projection. This therefore encodes something that must be added as a separate concept when using the (modified) HIP: that subspaces may be contained inside each other; Hestenes’ conditions on the grade impose this explicitly; the contraction does it implicitly without additional structure, and thus provides a simpler algebra without sacrificing any geometry.

In summary, the choice between HIP and modified HIP is clear: use the modified HIP or you will get a much more complicated geometric intersection operation as `meet`. It can probably be fixed with some grade testing, but this is not as elegant as fixing the inner product instead. Our preference for the contraction is based on the algebra and geometrical semantics that

permits a simpler projection operator.³ Again, this can be fixed with appropriate grade testing, but we prefer the more straightforward modification of the inner product.

The power of the `meet` and `join` defined in this way – making essential usage of the contraction – shows in the `connection` function in GABLE (see [3]): it is possible to give an algorithm *without cases* to compute the translation to make the `meet` between two affine subspaces non-trivial. Here ‘case-less’ means: no internal separate treatment of situations, all situations are computed using the same formula (whether point-to-point, line-to-line, line-to-point, parallel lines). Moreover, this formula is also valid in arbitrarily dimensional (Euclidean) space. The fact that we can do this shows that we are beginning to have the right primitive operations at our disposal in a computational language for geometry.

1.5 Implementation details

As with any software package, we encountered a number of implementation details. Here we highlight the more interesting and important details; further discussion on the implementation can be found in our technical report [10].

1.5.1 Matlab objects

The Matlab language has objects. For our implementation of the geometric algebra, we created a `GA` object, which stores the 8×1 matrix representing a geometric object. No other information is stored in the object. The benefit of using an object is that we were able to overload the ‘*’ and ‘^’ operators to perform the geometric product and the inner product respectively. We also overloaded ‘+’ and ‘-’ (both binary and unary) to add, subtract, and negate geometric objects.

Operations such as dual, inverse, and exponentiation we left as named routines (`dual`, `inverse`, and `gexp` respectively), although we did overload the ‘/’ operator to allow for right-division by a geometric object. I.e., the expression `A/B` is computed as `A*inverse(B)`.

³It is of some concern that the contraction *combines* the notions of perpendicularity and containment in one operation (for `A|B` is contained in `B` and perpendicular to `A`), and we need to investigate whether the remainder of the structure of geometric algebra enables their disentanglement; the projection operation suggests that it does. We should also mention an alternative definition of the contraction, as the adjoint of the outer product in terms of an extension of the bilinear form, which demonstrates its basic algebraic tidiness (see [7]) and its nicely ‘dual’ relationship to the outer product. This makes for improved duality relationships, see [2].




FIGURE 1.3. Graphical representation of vector, bivector, trivector

We also overloaded ‘==’ and ‘~=’ to allow for the comparison of geometric objects.

1.5.2 Graphics

Since we wanted a visual tutorial, we created graphical representations for all blades, and used Matlab rendering commands to draw them. The following table summarizes our representations:

Type	Representation	Orientation
scalar	Text above window	Sign
vector	Line from origin	Arrow head
bivector	Disk centered at origin	Arrows along edge
trivector	Line drawn sphere	Line segments going out or in

Figure 1.3 illustrates the vector, bivector, and trivector; the axes are put in automatically by Matlab.

We chose the disk as our representation for bivectors since with our matrix representation of the geometric objects, we do not necessarily have the defining vectors for the bivector (which may not even exist, as is the case if the bivector was created as the dual of a vector). Without such vectors, we can not use the standard parallelogram representation of the bivector. There is a similar issue with the trivector (i.e., we were unable to use a parallelepiped as its representation) and thus we used the sphere. However, we also provide demonstration routines to illustrate the more standard representations of bivectors and trivectors; the user must then provide the basis on which to decompose them.

Objects of mixed grade presented a more difficult problem. While it is easy to draw the scalar, vector, bivector, and trivector components independently, this is not particularly illustrative. In particular, we needed to find a way to illustrate the operations of the inner, outer, and geometric products. The first two are fairly easy to demonstrate: we have two sub-windows, in the former we draw the operands and in the latter we draw the result. The geometric product is more difficult to illustrate. So in addition

to providing a routine to show the operands and result of the geometric product, we presented examples of using the geometric product as an *operator* to perform rotations and interpolation between orientations, rather than as a (composite) *object* by itself.

1.5.3 Dealing with numerical issues

Numerically, some routines (particularly the inverse routine) may create small error terms. For example, we might get a Geometric Object that should be a vector, but has a small (on the order of 10^{-16}) bivector term. Several of our routines (e.g., the drawing routines) check to make sure that the arguments are blades, and such numerical errors, though small, will cause these routines to fail. For example, while the numerical error causes no particular computational problems, some routines will reject such geometric objects as not being blades. Thus, we wrote `gazv`, which sets all small terms of a geometric object to zero.

The routine `gazv` will set to zero all terms of a GA that are smaller in absolute value than `1e-15`, giving a warning when it does so. When developing code, it is a good idea to use `gazv` to overcome small numerical problems, and once the code is debugged switch to `grade` (since presumably you know the grade you want). Although you could use `grade` from the beginning, its use might hide some bugs that the system would otherwise automatically catch for you. We additionally wrote two similar routines: `GAZ`, which is identical to `gazv` except that it does not produce a warning message, and `blade` which converts a geometric object into a blade.

Since this software is meant for a tutorial, our `==` and `~=` operators compare to within a numerical tolerance. Thus, vectors, etc., differing by only small amounts will be considered equal. If an exact equality is desired, one may use the `eeq` function. Any further testing will require extracting the coordinates using `inner`.

Other than these small round-off errors (usually introduced by `inverse`), we encountered no numerical problems in developing GABLE. However, it should be noted that we did not perform extremely complex computations, and further testing of our software is needed to better assess its numerical stability.

1.6 Example: Pappus's theorem

As a more complete example of GABLE, we present an illustration of Pappus's theorem, which says take any two lines and three points on each line (`P1 P2 P3` and `Q1 Q2 Q3`), cross-join the point (i.e., build the line segments `P1Q2`, `P1Q3`, `P2Q1`, `P2Q3`, `P3Q1`, and `P3Q2`) and compute the intersection of the three cross-joined pairs of segments (i.e., intersect `P1Q2` with `P2Q1`, `P2Q3`




FIGURE 1.4. Illustration of Pappus's theorem created by GABLE.

with P_3Q_2 , and P_3Q_1 with P_1Q_3), and then these three points of intersection will be collinear.

To illustrate this theorem, we first need to construct six points and draw the relevant line segments. For this example, we have chosen to let e_3 be the homogeneous coordinate:

```
>> P1 = e3+e1; P2 = e3+2*e1; P3 = e3+4*e1;
>> Q1 = e3+e2; Q2 = e3+e1+2*e2; Q3 = e3+2*e1+3*e2;
>> DrawPolyline({P1,P3}, 'r'); DrawPolyline({Q1,Q3}, 'r');
>> DrawPolyline({P1,Q2}, 'k'); DrawPolyline({P1,Q3}, 'k');
>> DrawPolyline({P2,Q1}, 'k'); DrawPolyline({P2,Q3}, 'k');
>> DrawPolyline({P3,Q1}, 'k'); DrawPolyline({P3,Q2}, 'k');
>> GAvew([0,90]);
```

The `DrawPolyline` calls draw the line segments of Pappus's theorem. Next we want to compute the intersection of corresponding line segments. As a first step, we need to compute each line segment (the join of two points on the segment) and as a second step we need to intersect pairs of line segment (the meet of the two segments). Note that the meet will give us a homogeneous point, and we need to normalize its coordinates to put the point back in the homogeneous plane:

```
>> %...
>> H3 = meet(join(P1,Q2),join(P2,Q1)); A3 = H3/inner(H3,e3);
>> H2 = meet(join(P1,Q3),join(P3,Q1)); A2 = H2/inner(H2,e3);
>> H1 = meet(join(P2,Q3),join(P3,Q2)); A1 = H1/inner(H1,e3);
>> DrawHomogeneous(e3,H1,'n','g');
>> DrawHomogeneous(e3,H2,'n','g');
>> DrawHomogeneous(e3,H3,'n','g');
>> DrawPolyline({A1,A3}, 'b')
```

The resulting GABLE drawing appears in Figure 1.4. Two other examples that appear in our tutorial are an illustration of Napoleon's theorem and an illustration of Morley's triangle.

1.7 Conclusions

In GABLE, our Matlab package for the geometric algebra tutorial, we have chosen an 8×1 representation of multivectors, to be expanded to an 8×8 matrix representation when they are used as operands in the elementary products (geometric product, inner product, outer product). In our detailed comparison of the complexity of this representation with representations based on the isomorphisms of Clifford algebras with matrix algebras, this choice appeared not always the most efficient for software used in an actual application (rather than a mere tutorial), especially if the signature of the space required could be known beforehand, and if one would deal mostly with pure blades. Further developments in the practical use of geometric algebra should show whether blades are indeed sufficient for our needs. If applications would require many inner and outer product of multivectors of mixed grade, then our explicit representation of these products by matrices should be considered.

GABLE is an implementation of 3-dimensional Clifford algebras with arbitrary signature. The generalization to arbitrary dimensions is readily obtained from Equation 1.1. However, as noted earlier in this paper, high dimensional Clifford algebras are computationally expensive, and rather than use our $n \times 1$ representation, a specific complex, etc., representation may be preferred (although you would then lose the arbitrary signature property).

For the geometric division, we have extended Lounesto's method to computer inverses to work in 3-dimensional spaces of arbitrary signatures; but it should be emphasized that the method does *not* work in spaces of higher dimensions since it is based on properties of the Clifford conjugation that do not generalize to such spaces. In those spaces, an inversion of the geometric product matrix will be required.

The need to make geometrical macros for intersection and connection of geometrical objects 'without case statements' necessitated a detailed study of the `join` and `meet` operations and their relationship. We have now embedded them properly into the geometric algebra of blades, even though each is only determined up to a scalar factor; the key is to realize that both are based on the same factorization of blades. The tutorial shows that despite this unknown scalar, geometrically significant quantities based on them are unambiguously determined. This explicit realization appears to be new. Further, the GABLE `meet` and `join` operators only work for Euclidean signature; further research is needed to extend these operations to arbitrary signature. Likewise the GABLE implementation of `meet` and `join` exploit simplifications that occur in three dimensions; the implementation of higher dimensional `meet` and `join` operations will require more complex code.

At the start of this project, we thought it would be straight-forward to implement this software using results in the literature. However, we found the literature lacking in several areas, which we have partly addressed in

this paper. As a result of our work, we now have GABLE, a Matlab package and tutorial that should ease the learning of geometric algebra for people new to the subject. Moreover, we have found the package useful for testing out ideas and results in our own research.

1.8 Acknowledgments

We gratefully acknowledge the funding of Stephen Mann's sabbatical by the Dutch Organization for Scientific Research (NWO), and the funding of Tim Bouma's summer stay by the University of Amsterdam.

A.1 Inner product matrices

Hestenes Inner Product matrix:

$$[\underline{\mathbf{A}}^I] = \begin{bmatrix} 0 & \sigma_1 A_1 & \sigma_2 A_2 & \sigma_3 A_3 & -\sigma_{12} A_{12} & -\sigma_{23} A_{23} & -\sigma_{13} A_{31} & -\sigma_{123} A_{123} \\ 0 & 0 & \sigma_2 A_{12} & -\sigma_3 A_{31} & -\sigma_2 A_2 & -\sigma_{23} A_{123} & \sigma_3 A_3 & -\sigma_{23} A_{23} \\ 0 & -\sigma_1 A_{12} & 0 & \sigma_3 A_{23} & \sigma_1 A_1 & -\sigma_3 A_3 & -\sigma_{31} A_{123} & -\sigma_{13} A_{31} \\ 0 & \sigma_1 A_{31} & -\sigma_2 A_{23} & 0 & -\sigma_{12} A_{123} & \sigma_2 A_2 & -\sigma_1 A_1 & -\sigma_{12} A_{12} \\ 0 & 0 & 0 & \sigma_3 A_{123} & 0 & 0 & 0 & \sigma_3 A_3 \\ 0 & \sigma_1 A_{123} & 0 & 0 & 0 & 0 & 0 & \sigma_1 A_1 \\ 0 & 0 & \sigma_2 A_{123} & 0 & 0 & 0 & 0 & \sigma_2 A_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Modified Hestenes Inner Product matrix:

$$[\underline{\mathbf{A}}^I] = \begin{bmatrix} A_0 & \sigma_1 A_1 & \sigma_2 A_2 & \sigma_3 A_3 & -\sigma_{12} A_{12} & -\sigma_{23} A_{23} & -\sigma_{13} A_{31} & -\sigma_{123} A_{123} \\ A_0 & A_0 & \sigma_2 A_{12} & -\sigma_3 A_{31} & -\sigma_2 A_2 & -\sigma_{23} A_{123} & \sigma_3 A_3 & -\sigma_{23} A_{23} \\ A_0 & -\sigma_1 A_{12} & A_0 & \sigma_3 A_{23} & \sigma_1 A_1 & -\sigma_3 A_3 & -\sigma_{31} A_{123} & -\sigma_{13} A_{31} \\ A_0 & \sigma_1 A_{31} & -\sigma_2 A_{23} & A_0 & -\sigma_{12} A_{123} & \sigma_2 A_2 & -\sigma_1 A_1 & -\sigma_{12} A_{12} \\ A_0 & 0 & 0 & \sigma_3 A_{123} & A_0 & 0 & 0 & \sigma_3 A_3 \\ A_0 & \sigma_1 A_{123} & 0 & 0 & 0 & A_0 & 0 & \sigma_1 A_1 \\ A_0 & 0 & \sigma_2 A_{123} & 0 & 0 & 0 & A_0 & \sigma_2 A_2 \\ A_0 & 0 & 0 & 0 & 0 & 0 & 0 & A_0 \end{bmatrix}$$

Contraction Inner Product Matrix:

$$[\underline{\mathbf{A}}^C] = \begin{bmatrix} A_0 & \sigma_1 A_1 & \sigma_2 A_2 & \sigma_3 A_3 & -\sigma_{12} A_{12} & -\sigma_{23} A_{23} & -\sigma_{31} A_{31} & -\sigma_{123} A_{123} \\ 0 & A_0 & 0 & 0 & -\sigma_2 A_2 & 0 & \sigma_3 A_3 & -\sigma_{23} A_{23} \\ 0 & 0 & A_0 & 0 & \sigma_1 A_1 & -\sigma_3 A_3 & 0 & -\sigma_{13} A_{31} \\ 0 & 0 & 0 & A_0 & 0 & \sigma_2 A_2 & -\sigma_1 A_1 & -\sigma_{12} A_{12} \\ 0 & 0 & 0 & 0 & A_0 & 0 & 0 & \sigma_3 A_3 \\ 0 & 0 & 0 & 0 & 0 & A_0 & 0 & \sigma_1 A_1 \\ 0 & 0 & 0 & 0 & 0 & 0 & A_0 & \sigma_2 A_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & A_0 \end{bmatrix}$$

- [1] R. Ablamowicz. Clifford algebra computations with Maple. In W.E. Baylis, editor, *Clifford (Geometric) Algebras with Applications to Physics, Mathematics and Engineering*. Birkhäuser, 1996.
- [2] Leo Dorst. Honing geometric algebra for its use in the computer sciences. In G. Sommer, editor, *Geometric Computing with Clifford Algebra*. Springer, expected 2000.
- [3] Leo Dorst, Stephen Mann, and Tim Bouma. GABLE: A matlab tutorial for geometric algebra. Available at <http://www.wins.uva.nl/~leo/clifford/gable.html>, 1999.
- [4] David Hestenes. *New Foundations for Classical Mechanics*. Reidel, 1986.
- [5] David Hestenes and Garrett Sobczyk. *Clifford Algebra to Geometric Calculus*. Reidel, 1984.
- [6] Anthony Lasenby and M. Ashdown et al. GA package for Maple V. 1999.
- [7] Pertti Lounesto. Marcel riesz's work on clifford algebras. In *Clifford numbers and spinors*, pages 119–241. Kluwer Academic, 1993.
- [8] Pertti Lounesto. *Clifford Algebras and Spinors*. London Mathematical Society Lecture Note Series 239. Cambridge University Press, 1997.
- [9] Pertti Lounesto, Risto Mikkola, and Vesa Vierros. Clical user manual: Complex number, vector space and clifford algebra calculator for ms-dos personal computers. Technical Report A248, Institute of Mathematics, Helsinki University of Technology, 1987.
- [10] Stephen Mann, Leo Dorst, and Tim Bouma. The making of a Matlab geometric algebra. Technical Report CS-99-27, University of Waterloo, December 1999. Available at <ftp://cs-archive.uwaterloo.ca/cs-archive/CS-99-27/>.
- [11] Ian R Porteous. *Topological Geometry*. Cambridge university Press, Cambridge, 1981.
- [12] Lars Svenson. Personal communication at ACACSE'99, Ixtapa, Mexico, 1999.