Salvaging the useful parts of the cult - a 2-year review of geometric algebra

Vanessa Alexandra Hollmeier

September 2024

TL;DR: A lot of the nice techniques of geometric algebra for doing physics are not exclusive to Clifford algebras, but can be formulated in traditional maths as well. I think that we should focus on porting these features to traditional maths instead of convincing everyone to learn an entirely new mathematical language.

Contents

1 Introduction

In theoretical physics, I have never encountered a more politically charged subject than geometric algebra. The people who have heard of it seem to cleanly split into two camps: The geometric algebra community who thinks that it is the one mathematical language of physics to rule them all, and have made it their mission to translate every single bit of physics ever done into multivectors. On the other hand, outside the closer geometric algebra community, GA has the reputation of making overblown promises and being a cult (see for instance this tweet by mathematical physicist John Baez).

I first encountered geometric algebra about two years ago, and quickly became interested in it. GA significantly advanced my personal understanding of electromagnetism, spinors and multivariate calculus. However, over time, I also came to doubt the claim of David Hestenes and his followers that it is a "unified language for physics", or even a clearly-cut distinct language at all. Over the course of the last year, I developed several techniques to translate common GA techniques into traditional theoretical physics language, and wrote a GA lecture script about it as a uni project. In this article, I will mention a lot of interesting stuff in passing, so if some points of this article interest you, I recommend checking out the script to see the concepts mentioned in here explained in depth.

In this post, I'll go over various examples of GA techniques and discuss their advantages and disadvantages. I'm a theoretical physicist and not a mathematical physicist, so I won't discuss mathematical formalities, but instead focus on intuitive accessibility of the concepts.

1.1 What exactly is "geometric algebra"?

The common narrative both among GA supporters and opponents goes that geometric algebra is just real Clifford algebras applied to various areas of physics. However, I don't think that's true. Geometric algebra is more of an umbrella term for a collection of mathematical and didactical techniques related to:

- exterior algebra (EA),
- Clifford algebra (CA), and
- differential geometry.

I will use the term "geometric algebra" (GA) to refer to said collection, and "Clifford algebra" (CA) to refer to the specific mathematical algebra.

2 The useful parts: non-CA-specific stuff

A lot of the advantages of geometric algebra are not specific to Clifford algebras, but can be expressed in conventional terms perfectly fine. The main reason for that is that there is a huge intersection between Clifford algebras and exterior algebras - an EA is basically just a CA with the scalar product set to zero[1]. In the following couple of subsections, I will explain why they look so different in practice, and what specific aspects make the GA approach more palatable.

2.1 k-vectors and multivectors

The run-of-the-mill exterior algebra that is taught in every physics curriculum is typically based on **covectors** - that is, we consider some vector space V, take its covector space V^*, and then form the exterior algebra

$$\Lambda V^* = \bigoplus_{n=0}^{\infty} \Lambda^n V^*. \tag{1}$$

The grade-n subspace Λ^n is basically just the space of fully antisymmetric covariant rank n tensors over V - for instance, the electromagnetic field tensor $F_{\mu\nu}$ with downstairs indices is an element of $\Lambda^2 V^*$.

I'm sure that there are historical and maths-aesthetic reasons for why people normally do that - but in my opinion, this is horrible both for physics didactics and physical intuition. Some people already have difficulties geometrically visualizing covectors - personally, I've come to imagine them as blocks of lasagna with the layers indicating "equipotential surfaces" - but when it comes to geometrically picturing k-forms, most people (myself included) have zero intuition for them. What on earth is a two-form? I know the mathematical definition, but I have no idea how to mentally picture the combined electric and magnetic fields as some sort of covariant 1+3D entity using this formalism. I think this lack of geometric intuition is a major reason for why people never use k-forms in everyday theoretical physics.

The main didactical innovation of geometric algebra is to define the exterior algebra not over the space of covectors, **but over the underlying vector space itself**:

$$\Lambda V = \bigoplus_{n=0}^{\infty} \Lambda^n V. \tag{2}$$

[1] I once said this in front of a mathematical physicist and she seemed very angry about it. But it is in fact true - if you set the scalar product/quadratic form of a CA to zero, the geometric product $ab = a \cdot b + a \wedge b$ reduces to the wedge product.

The elements of the individual $\Lambda^k V$ are called k-vectors, and general elements of ΛV are called **multivectors**. Now, we suddenly have a very clear mental picture for everything!

- $\Lambda^1 V$ is the space of ordinary vectors, which are arrows pointing through spacetime.
- $\Lambda^2 V$ is the space of **bivectors**, which are oriented area elements that you can picture as parallelograms spanned by two vectors, $\mathbf{a} \wedge \mathbf{b}$.
- Next, we have the **trivectors** $\Lambda^3 V$, which are oriented volume elements that you can picture as parallelepipeds spanned by three vectors, $\mathbf{a} \wedge \mathbf{b} \wedge \mathbf{b}$.
- etc.

We can even intuitively visualize fully antisymmetric contravariant tensors now. Consider for instance the electromagnetic field with upstairs indices $F^{\mu\nu}$ - this is a bivector in $\Lambda^2 V$. We can visualize it as a parallelogram in a Minkowski diagram. The parallelograms with one timelike and one spacelike axis are called **hyperbolic bivectors**, and represent electric fields. The parallelograms with two spacelike axes are called **circular bivectors**, and represent magnetic fields.

When we're not doing GR, we mostly have a fixed metric $\eta_{\mu\nu}$, so V and V^* are isomorphic and we can seamlessly translate between k-forms and k-vectors[2]. In my opinion, this is a GA technique that should absolutely be adopted into the physics mainstream.

2.2 Interior product between k-vectors

Another common critique of exterior algebra (and perhaps the main reason for why it isn't used more widely) is that it makes practical computations very difficult in its conventional form. Consider for instance the electromagnetic field equations in differential forms:

$$\mathrm{d}F = 0 \tag{3}$$

$$* \, \mathrm{d}(*F) = -j \tag{4}$$

The first one - the exterior derivative - is rather easy to do by hand. When it comes to the double Hodge dual, however, most people I know are pretty lost. Some of them remember the abstract definition of the Hodge dual, and some of them remember the tensor definition with the ϵ symbol and the $(-1)^{\text{something}}$, but no one really has any idea of what is going on, geometrically.

[2]Remember that $\Lambda(V^*) \simeq (\Lambda V)^*$ canonically.

In GA, this construct is replaced with the so-called **interior product** and the **interior derivative**, which gives you a neat geometric picture of what this double Hodge dual is supposed to do. Roughly speaking, while the exterior product with a vector "extends" k-vectors to $k+1$-vectors by adding an axis, the interior product takes away one axis of a k-vector to make it a $(k-1)$-vector. If you want to know the details, I suggest reading the script I linked to.

Crucially, however, Clifford algebra is not needed to define the interior product at all! In GA language, the above equations respecively read

$$\partial \wedge F = 0 \tag{5}$$

$$\partial \cdot F = j. \tag{6}$$

In CA, we define the exterior and interior product as the highest and lowest grade of the geometric product, respectively. However, just as we don't need CA to define the exterior product, we also don't need CA to define the interior product! The two expressions above have the direct tensor equivalent:

$$3\ \partial^{[\mu} F^{\nu\rho]} = 0 \tag{7}$$

$$\partial_\mu F^{\mu\nu} = j^\nu. \tag{8}$$

Basically, in tensor language, the interior product is an index contraction (instead of the antisymmetrization of the tensor product). Again, the details of this are very interesting, and I recommend reading the linked script above, but the relevant point for this post is that the interior product from GA is a very useful concept that does not need CA at all.

(As a corollary, this means that the **reformulation of the cross product** with the exterior and interior product also doesn't need CA.)

2.3 Differential geometry

In standard differential geometry and general relativity, it is common to use metrics, tensors and Christoffel symbols. On the other hand, the GA community advocates the use of **tetrads, k-vectors** and **spin connections**. Not that there is anything wrong with it - I think this is the better approach both conceptually and practically - but what really drives me up the wall is how the GA community acts like they *invented these things*. It may not be common, but mainstream GR absolutely has these concepts and uses them without any reference to GA/CA - see for instance Appendix J of Carrol's *Spacetime and Geometry*. Of course, they become more palatable with GA techniques - imagining the Riemann tensor as a map from the parallel transport bivector to the bivector generating a rotation is a lot more easy than imagining an abstract map between two 2-forms. Similarly, the spin connection is no longer an eldritch complex spin-1/2 representation thing, but a map taking a vector and returning a bivector (more on that later). But there is nothing inherently Clifford about these techniques, and it is certainly possible to adopt the GA style of teaching them without actually having to use Clifford algebras.

Similarly, the techiques called "**geometric calculus**" and "**directed integration theory**" are pretty much the same as standard exterior integrals - although again, they are made a lot more accessible by the use of k-vectors and the interior product. For instance, the generalized Stokes theorem in exterior calculus reads:

$$\oint_{\partial\Omega} \omega = \int_{\Omega} \mathrm{d}\omega. \tag{9}$$

where Ω is a k-dimensional volume. In order to formulate a generalized Gauss theorem, we'd have to do something weird with two Hodge duals again. In geometric calculus notation[3] however, we can just write

$$\oint_{\partial\Omega} \mathrm{d}Y \cdot \omega = \int_{\Omega} \mathrm{d}X \cdot (\partial \wedge \omega) \tag{10}$$

for the generalized Stokes theorem, and

$$\oint_{\partial\Omega} \mathrm{d}Y \wedge \omega = \int_{\Omega} \mathrm{d}X \wedge (\partial \cdot \omega) \tag{11}$$

for the generalized Gauss theorem. In both cases, $\mathrm{d}X$ stands for the k-vector-valued measure over the volume Ω, and $\mathrm{d}Y$ for the $k-1$-vector-valued measure over its surface $\partial\Omega$.

However, this way of handling k-vector integrals is, in principle, **completely independent of Clifford algebra**. For instance, let's say we want a Gauss theorem for 1+3D spacetime, a 3D submanifold Ω and a bivector field ω. The tensor notation translation of 11 for this case would be:

$$\frac{(2+2)!}{2!\ 2!} \oint_{\partial\Omega} \mathrm{d}Y^{[\mu\nu}\omega^{\rho\sigma]} = \frac{(3+1)!}{3!\ 1!} \int_{\Omega} \mathrm{d}X^{[\mu\nu\rho|} \left(\partial_\lambda \omega^{\lambda|\sigma]}\right) \tag{12}$$

where the $[\mu\nu\rho|\cdots|\sigma]$ notation means that the antisymmetrization is only performed over these indices (i.e. the contraction over λ is performed first).

Note how every single tensor in the above integral has a direct geometric interpretation as a vector, bivector or trivector. In my opinion, we should definitely adopt this technique into mainstream physics - but again, we don't need Clifford algebras to do so.

[3] It's not strictly correct to call this way of writing the integral a different "notation", because this integral is scalar-valued as opposed to the differential-form-valued exterior integral. The relevant difference is that we contract the integrand with the the measure in geometric calculus instead of taking the pullback along the coordinate chart.

2.4 Bivector algebra

The bivectors $\Lambda^2 V$ can be represented as antisymmetric matrices acting on V. To be precise: When we single out a metric, the bivectors $\Lambda^2 V$ become isomorphic to a subspace of the endomorphisms $\mathrm{End}(V)$. In physics vernacular, this means that we can pull down one index of the fully antisymmetric rank (2, 0) tensor $B^{\mu\nu}$ to obtain the rank (1, 1) tensor $B^{\mu}{}_{\nu}$.

These matrices $B^{\mu}{}_{\nu}$ form the Lie algebra $\mathrm{so}(p, q) \subset \mathrm{End}(V)$, where (p, q) is the signature of the metric. Hence, when we exponentiate these generators, we get the rotation matrices $\mathrm{SO}(p, q)$. The huge advantage of the GA picture is that the generators directly correspond to the bivectors describing the respective plane of rotation. For instance, the rotation around the z axis is generated by the matrix corresponding to the bivector lying in the xy plane. In GA language, we also say that this is a rotation **along the xy bivector**.

The huge advantage of this approach (as opposed to Euler vectors with the cross product as the Lie bracket) is that it generalizes to arbitrary dimensions and signatures. For instance, in special relativity, the hyperbolic bivectors generate boosts, and the circular bivectors generate spatial rotations. In my opinion, this approach should definitely be adopted into the mainstream - but again, the CA approach of rotors isn't really needed here. The tensor formulation of bivectors and the bivector algebra works perfectly fine on its own.

3 The useful parts: CA-specific stuff

While Clifford algebra is not strictly required for a lot of useful GA stuff, it is indispensable for some parts. For example:

3.1 Electromagnetic field equation and U(1) EM duality

The two electromagnetic field equations mentioned above already have a succinct formulation in terms of the interior and exterior derivative:

$$\partial \cdot F = j \tag{13}$$
$$\partial \wedge F = 0 \tag{14}$$

With the geometric product, they can be merged into a single equation:

$$\partial F = j. \tag{15}$$

This short and simple form of Maxwell's equations is one of the main triumphs of GA. On top of that, this form can be used to perform actual, practical computations. The main reason for this is that 1+3D CA provides a significantly better way of describing electromagnetic duality.

As a short reminder: If a specific field F is a solution of the electromagnetic equations in vacuum, its Hodge dual $*F$ is a solution too. This seemingly discrete symmetry can be made a continuous $U(1)$ symmetry by linearly combining F and $*F$ the following way:

$$F' = \cos(\theta)F + \sin(\theta)(*F) \tag{16}$$

For every θ, F' is a solution too. The analogous construct in CA reads:

$$F' = F\exp(I\theta), \tag{17}$$

where I is the pseudoscalar for which $I^2 = -1$. This allows us to make arbitary duality rotations - in contrast, the Hodge dual only allows $\tau/4$ turns[4]. The CA approach allows us to write a plane-wave solution as:

$$F(x) = F_0 \exp(Ik \cdot x). \tag{18}$$

We can check that this is a solution by:

$$\partial F(x) = \partial F_0 \exp(Ik \cdot x) = kF_0 I \exp(Ik \cdot x) = 0, \tag{19}$$

where we have used $kF_0 = 0$ in the last step. The huge advantage of this approach is that we have completely eliminated the need for complex numbers, such that only physical degrees of freedom remain.

Again, if the details interest you, read the linked script - this extreme ease of handling electrodynamics extends to the gauge-field formulation and quantum electrodynamics. However, Clifford algebra is an indispensable part of rewriting the Hodge dual with the pseudoscalar and the geometric product, so there's no easy way out here.

3.2 Rotors

In the section about the bivector algebra, I have glossed over a very important question: The question of representations.

When we represent the bivectors as matrices/(1, 1) tensors over V and use the matrix product to exponentiate them, we get the standard rotation matrices we are all used to, aka the **spin-1 representation of** $SO(p, q)$. However, when we formulate bivectors as elements of a Clifford algebra and use the geometric product to exponentiate them, we get the **rotors**, also known as **spin-1/2 representation of** $SO(p, q)$.

The rotation matrices M are used to rotate vectors (aka spin-1 objects),

$$\mathbf{v} \to M\mathbf{v}, \tag{20}$$

[4] $\tau = 2\pi$.

and the rotors are used to rotate spinors (aka spin-1/2 objects),

$$|\psi\rangle \to R\,|\psi\rangle\,. \tag{21}$$

If we want to use a rotor to rotate a vector, we need to use a double-sided transformation law,

$$\mathbf{v} \to R\mathbf{v}\tilde{R}, \tag{22}$$

analogously to how we need to use a double-sided transformation law to rotate bilinear forms (aka spin-2 objects) with rotation matrices. This double-sided rotor law is somewhat tedious to handle in practice, but didactically speaking, the rotor approach to spinors completely wipes the floor with the common approach of deriving spinors by complexifying the Lorentz algebra, because suddenly, we have a clear geometric picture of what a spin-1/2 representation is supposed to be.

Also, rotors are a way better picture than quaternions, but I won't get into that here because Hestenes' followers have already beaten this point into everyone's head (and I fully agree with them on that count).

3.3 Pauli and Dirac matrices

When describing spinors, physicists conventionally use the Pauli and Dirac matrices - Pauli matrices for non-relativistic spinors, and Dirac matrices for relativistic spinors. In the usual aproach to spinors, they basically just fall out of the sky, and most people just accept them as something abstract with no physical intuition.

However, Clifford algebra tells us that they can also be interpreted as the basis vectors of the vector space we're working on - the Pauli matrices $\sigma_1, \sigma_2, \sigma_3$ are a matrix representation of the three basis vectors of the 3D geometric algebra, and the Dirac gamma matrices γ^μ are a matrix representation of the four spacetime basis vectors. The matrix product between them is equivalent to the geometric product between multivectors. In order to introduce this isomorphism, we absolutely need Clifford algebra.

4 The shadow side

David Hestenes and his followers frequently make the claim that geometric algebra is not just another mathematical tool, but an entire unified language for physics. I strongly disagree with this claim - a lot of things just are better left formulated in traditional maths, and the GA reformulation of some concepts seems significantly worse than the original in some cases. IMO, the most glaring examples of this are:

4.1 Complex numbers

While reformulating quaterions as rotors, we replaced the three imaginary units i, k, k with the three bivectors in three dimensions, and the EM duality. Similarly, in the context of EM plane waves, we replaced the i in the phasor $\exp(ik \cdot x)$ with the pseudoscalar $I = \gamma^{0123}$.

This way of replacing imaginary elements with multivectors is common in geometric algebra - in fact, it is one of the stated aims of the GA community to replace *all* complex numbers with real multivectors. There are some cases were this works really well - see quaternions and EM duality - but in other cases, it just makes things more difficult. Consider for instance the standard replacement of complex numbers by 2D even multivectors. Apart from slightly different terminology, this adds literally nothing. Complex numbers and the even 2D Clifford algebra are isomorphic to each other. It's a nice toy example to explain the working principles of Clifford algebras in an informal way, but I really don't see what I could gain by writing $z = a + e_{12}b$ and $\exp(e_{12}\theta)$ instead of $z = a + bi$ and $\exp(i\theta)$.

4.2 Matrices and tensors

With exterior algebra, we can reformulate antisymmetric matrices as bivectors - or more generally, fully antisymmetric rank-k tensors as k-vectors. This does a lot of work in terms of geometric intuition, and I think that this should be done wherever possible.

However, the GA community has made it its mission to replace literally all matrices and tensors with some sort of Clifford algebra construct - mainly in the name of "beautiful non-coordinate descriptions"[5]. This often results in unyieldy constructs. For instance, the most prominent example of a non-antisymmetric tensor in classical mechanics is the **inertia tensor** I_{ij}. If we want to calculate e.g. the rotational energy of an object, we'd write

$$E = \frac{1}{2} I_{ij} \omega^i \omega^j. \tag{23}$$

In geometric algebra, this tensor is expressed as a linear map from bivectors to bivectors, i.e. $L = I(\omega)$. This forces us to use significantly less idiomatic constructs like

$$E = -\frac{1}{2} \langle \omega I(\omega) \rangle, \tag{24}$$

where we had to insert the minus sign because bivectors square to -1. This problem gets especially bad in GR, where it is vital to keep track of the index types of the objects we are using.

[5] I'm surprised that there aren't more people who know about abstract tensor index notation. Basically, it's a way of reinterpreting tensor index notation such that becomes coordinate-free, with essentially zero changes.

4.3 "Dirac-Hestenes" equation

As explained above, Clifford algebras are very useful to give the Pauli and Dirac matrices geometric meaning. However, the standard geometric algebra approach to spinors goes further - instead of just translating the involved matrices to multivectors, people also commonly **translate spinors to multivectors**. Roughly speaking, the multivector M_ψ representing some spinor $|\psi\rangle$ represents the matrix one needs to apply to a given reference spinor $|psi_0\rangle$ (conventionally spin-up) to obtain $|\psi\rangle$, i.e.

$$M_\psi |\psi_0\rangle = |\psi\rangle . \tag{25}$$

The multivector M_ψ then transforms under a one-sided (i.e. spin-1/2) rotor transformation law,

$$M_\psi \to RM_\psi . \tag{26}$$

If we single out zero-3-momentum spin-up as the reference spinor, the Dirac equation reads

$$\partial M_\psi \gamma_{12} + mM_\psi \gamma_0 = 0 \tag{27}$$

in terms of the multivector M_ψ. This reformulation is commonly called the "Dirac-Hestenes equation" in the GA community (+10 Cultishness points).

I think this reformulation is a horrible idea for several reasons:

- Spinors aren't multivectors, in the sense that these two have completely different transformation properties. Multivectors are invariant under 360° rotations, but spinors aren't. Imagining spinors as multivectors in the geometric sense actively hurts your physical intuition of them, and it serves little to brute-force postulate the one-sided rotor transformation law (26) for them.

- The fact that we had to single out a reference spinor $|\psi_0\rangle$ breaks manifest Lorentz invariance. In fact, this is glaringly obvious from the fact that random multivectors like γ_{12} and γ_0 appear in (27). That alone would be a dealbreaker for me.

- If we want to perform operations on the multivector M_ψ, we have to use modified multiplication rules depending on the operation in question. For instance, the operation $\gamma_\mu |\psi\rangle$ corresponds to $\gamma_\mu M_\psi \gamma_0$. This significantly complicates the handling of these multivectors in the context of e.g. Feynman diagrams (on top of breaking manifest Lorentz covariance all over again).

- Rotating the reference spinor by $\tau/2$ and multiplying it by i has the same effect:

$$\gamma_{21} |\psi\rangle = i |\psi\rangle . \tag{28}$$

> The standard GA approach to spinors uses this fact to eliminate complex numbers from the Dirac equation by replacing all imaginary units with a right-hand γ_{12} - in fact, this is where the γ_{12} in (27) comes from. However, this conceptually mixes up two different symmetries of Dirac spinors - Spin(1,3) and U(1) symmetry respectively, which causes confusion at best (in fact, it completely breaks down if we are dealing with multiple Dirac fields with different electric charges).

This reformulation of spinors probably is the aspect of GA I dislike the most.

5 Conclusion

As shown above, I think that a lot of the mathematical techniques that make geometric algebra so great in the eyes of many physicists are not exclusive to Clifford algebras at all. If we want to adopt them into the physics mainstream, I think we're better off not branding them as a part of a "new mathematical language", but instead preferentially introduce them with conventional maths, as outlined above. Clifford algebras should be used iff the specific technique in question absolutely requires it - and that is the case a lot less frequently than most GA advocates claim. In any case, I think it is better to drop the "geometric algebra" banner, because it has become too strongly associated with the exaggerated claims of the GA community.

5.1 Bonus content: Is GA a cult?

I've called GA a cult somewhat provocatively in the title, but I think that the question whether the GA community deserves to be called a cult is a bit more nuanced. Strictly speaking, this question comes down to how you define the word "cult": Everyone agrees that GA is not a new religious movement whose members dress up in black robes, gather in dark rooms and chant "All hail the Dark Lord Hestenes". On the other hand, I'd wager that everyone also agrees that the GA community has both a higher degree of social cohesion and homogenity of ideas than other parts of the maths/physics community. Whether or not you choose to call this a "cult" is down to where you draw the border between "normal group" and "cult".

In practice, however, the matter of calling GA a cult or not really just is a social signaling issue - by using a definition of "cult" that includes the GA community, you signal your disapproval of the concept of geometric algebra and/or the GA community, and by using a definition that doesn't include it, you signal your approval (or indifference) to it. So this whole discussion really is pointless.