Where do alternating multilinear maps come from?

Aaron Fenyes

May 24, 2012

Abstract

Alternating multilinear maps, as the pointwise constituents of differential forms, play a fundamental role in differential geometry, while other kinds of multilinear maps—for example, symmetric ones—hardly show up at all. [1] In these notes, I'll try to illuminate the geometric nature of alternating maps by linking them to our intuitive concepts of volume and area.

1 What is volume?

A parallelepiped is one of the simplest kinds of shapes. You can specify a k-dimensional parallelepiped by giving k vectors, like this:

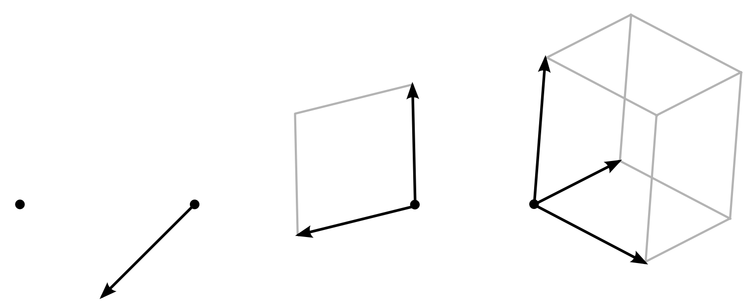

Figure 1: 0-, 1-, 2-, and 3-dimensional parallelepipeds.

I'll call these k vectors the *legs* of the parallelepiped.

Any engineer can tell you that in an n-dimensional vector space, every n-dimensional parallelepiped is associated with a non-negative real number, called its *volume*. I don't know exactly what this "volume" is, but I do know it has the following properties:

1. If you *stretch* a parallelepiped by multiplying one of its legs by a scalar λ, its volume gets multiplied by $|\lambda|$.

2. If you *shear* a parallelepiped by adding a multiple of one of its legs to another leg, its volume stays the same.

More formally, in an n-dimensional vector space V over an absolute value field[1] K, the function $\varepsilon\colon V^n \to [0,\infty)$ that sends the legs of an n-parallelepiped to its volume has the following properties:

1. For any $\lambda \in K$,
$$\varepsilon(\ldots \underset{i}{\lambda v} \ldots) = |\lambda| \varepsilon(\ldots \underset{i}{v} \ldots).$$

2. For any $\lambda \in K$ and $i \neq j$,
$$\varepsilon(\ldots \underset{i}{v} \ldots \underset{j}{w} \ldots) = \varepsilon(\ldots \underset{i}{v + \lambda w} \ldots \underset{j}{w} \ldots).$$

As we'll see in the next section, these properties are surprisingly strong for their size. They're strong enough, in fact, to determine the volume function ε uniquely, up to multiplication by a constant. Moreover, they imply that ε is very nice algebraically: in Section 4, we'll find that ε is the absolute value of a multilinear function, which is also unique up to scaling.

2 What are volume functions like?

Let's say a *volume function* is any function that fits the description of the function ε above. If you start playing with the definition, you'll soon discover some general features of volume functions.

Proposition 1. *If ε is a volume function,*

$$\varepsilon(\ldots \underset{i}{v} \ldots \underset{j}{v} \ldots) = 0$$

for all $i \neq j$. In other words, ε is alternating.

Proof. This follows from the shear property and a blindingly obvious consequence of the stretch property, which I'll let you find for yourself. □

Proposition 2. *If ε is a volume function,*

$$\varepsilon(\ldots \underset{i}{v} \ldots \underset{j}{w} \ldots) = \varepsilon(\ldots \underset{i}{w} \ldots \underset{j}{v} \ldots)$$

for all $i \neq j$. In other words, ε is symmetric.

[1] That is, a field equipped with a map $|\cdot|\colon K \to [0,\infty)$ such that

1. $|\alpha| = 0$ if and only if $\alpha = 0$.
2. $|\alpha\beta| = |\alpha||\beta|$.
3. $|\alpha + \beta| \leq |\alpha| + |\beta|$.

Proof. Just keep pounding away with the shear property. At the end, you'll need the fact that in an absolute value field, $|\alpha| = |-\alpha|$. □

More complicated manipulations uncover features of deeper linear algebraic significance.

Proposition 3. *Say ε is a volume function. If $v_1, \dots, v_n$ are linearly dependent, $\varepsilon(v_1, \dots, v_n) = 0$.*

Proof. Say $\alpha_1 v_1 + \dots + \alpha_n v_n = 0$, with at least one of the α_i nonzero. Since ε is symmetric, we can assume without loss of generality that α_1 is nonzero. Then,

$$
\begin{aligned}
\varepsilon(v_1, \dots, v_n) &= \tfrac{1}{|\alpha_1|}\varepsilon(\alpha_1 v_1, v_2, \dots, v_n) \\
&= \tfrac{1}{|\alpha_1|}\varepsilon(\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n, v_2, \dots, v_n) \\
&= \tfrac{1}{|\alpha_1|}\varepsilon(0, v_2, \dots, v_n) \\
&= 0.
\end{aligned}
$$

□

The next feature—a partial converse of Proposition 3—depends crucially on the fact that volumes are only defined for parallelepipeds of the same dimension as V.

Lemma 1. *Say ε is a volume function. If there are linearly independent vectors $v_1, \dots, v_n$ for which $\varepsilon(v_1, \dots, v_n) = 0$, then $\varepsilon = 0$.*

Proof. Suppose the proposition holds when $n = m$, for some $m \geq 1$, and consider the case $n = m + 1$.

Suppose there are linearly independent vectors $v_1, \dots, v_{m+1} \in V$ for which $\varepsilon(v_1, \dots, v_{m+1}) = 0$. Pick any vectors $w_1, \dots, w_{m+1} \in V$. Our goal is to show that $\varepsilon(w_1, \dots, w_{m+1}) = 0$. If $w_1, \dots, w_{m+1}$ are linearly dependent, this follows immediately from Proposition 3, so we only need to consider the case where $w_1, \dots, w_{m+1}$ are linearly independent.

Since $n = m + 1$ is the dimension of V, the vectors $v_1, \dots, v_{m+1}$ form a basis for V. Linear independence guarantees that $w_{m+1} \neq 0$, so at least one component of w_{m+1} in this basis is nonzero. Since ε is symmetric, we can reorder the vectors $v_1, \dots, v_{m+1}$ however we like and still have $\varepsilon(v_1, \dots, v_{m+1}) = 0$, so we might as well choose an order for which the v_{m+1} component of w_{m+1} is nonzero.

Let U be the subspace of V spanned by $v_1, \dots, v_m$. Since the v_{m+1} component of w_{m+1} is nonzero, we can find $\alpha_1, \dots, \alpha_m$ such that $\tilde{w}_i = w_i + \alpha_i w_{m+1}$ lies in U for all $i \in \{1 \dots m\}$. Since $w_1, \dots, w_{m+1}$ are linearly independent, so are $\tilde{w}_1, \dots, \tilde{w}_m$. As a result, $\tilde{w}_1, \dots, \tilde{w}_m$ span U. We can therefore find $\beta_1, \dots, \beta_m$ such that $w_{m+1} + \beta_1 \tilde{w}_1 + \dots + \beta_m \tilde{w}_m = \lambda v_{m+1}$.

By the shear property,

$$
\begin{aligned}
\varepsilon(w_1, \dots, w_m, w_{m+1}) &= \varepsilon(\tilde{w}_1, \dots, \tilde{w}_m, w_{m+1}) \\
&= \varepsilon(\tilde{w}_1, \dots, \tilde{w}_m, \lambda v_{m+1}) \\
&= |\lambda|\varepsilon(\tilde{w}_1, \dots, \tilde{w}_m, v_{m+1}).
\end{aligned}
$$

Define a function $\eta\colon U^m \to K$ by $\eta(x_1, \ldots, x_m) = \varepsilon(x_1, \ldots, x_m, v_{m+1})$. Because $m \geq 1$, it's simple to verify that η is a volume function on U. Observe that $v_1, \ldots, v_m$ are linearly independent vectors for which $\eta(v_1, \ldots, v_m) = 0$. Since we're assuming the proposition holds in dimension m, it follows that $\eta = 0$, which means

$$\begin{aligned}\varepsilon(w_1, \ldots, w_m, w_{m+1}) &= |\lambda|\eta(\tilde{w}_1, \ldots, \tilde{w}_m) \\ &= 0,\end{aligned}$$

which is what we wanted to show.

We have now proven that if the proposition holds when $n = m$, for some $m \geq 1$, it also holds when $n = m + 1$. When $n = 1$, the proposition follows easily from the stretch property, so the proposition holds for all $n \geq 1$.

Since an empty set of vectors is linearly independent, the proposition also holds for $n = 0$. Its proof in this case is left as a somewhat mind-bending exercise for the reader. (There's actually no need to consider the $n = 0$ case separately, but doing so allows a much clearer treatment of the other cases.) □

With Lemma 1 in hand, the main result of this section follows by a clever flick of the wrist.

Theorem 1. *If ε and η are volume functions, one is a constant multiple of the other.*

Proof. The theorem is obviously true if $\varepsilon = 0$, so we only have to consider the case $\varepsilon \neq 0$. Choose a basis $v_1, \ldots, v_n$ for V. Because $\varepsilon \neq 0$, Lemma 1 guarantees that $\varepsilon(v_1, \ldots, v_n) \neq 0$, so

$$C = \frac{\eta(v_1, \ldots, v_n)}{\varepsilon(v_1, \ldots, v_n)}$$

is well-defined.

Let $\delta\colon V^n \to [0, \infty)$ be the function

$$\delta(w_1, \ldots, w_n) = |\eta(w_1, \ldots, w_n) - C\varepsilon(w_1, \ldots, w_n)|_{\mathbb{R}},$$

where $|\cdot|_{\mathbb{R}}$ is the standard absolute value on $\mathbb{R}$. It's not hard to show that δ is a volume function. By construction, $\delta(v_1, \ldots, v_n) = 0$, so $\delta = 0$ by Lemma 1. □

3 Where does volume come from?

The notion of volume is an essential part of our everyday experience—what one might call the real world. A volume function, however, can be defined on a vector space over any absolute value field K. The absolute value serves as an intermediary between the K world and the real world, converting the K scale factor λ into the real scale factor $|\lambda|$ in the statement of the stretch property. Metaphorically, a volume function looks something like this:

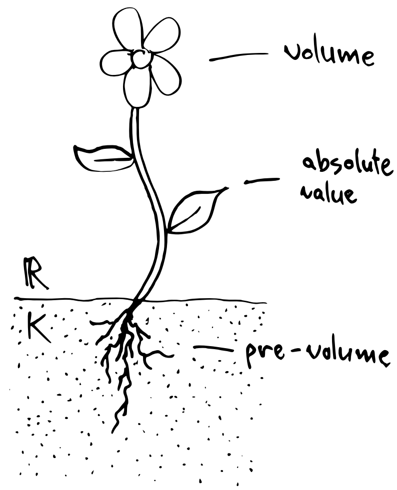

Figure 2: A volume function as a flower.

Like a flower, a volume function has part of its structure hidden underground, in the K world. In the next section, we'll explore volume's K-valued roots.

4 Pre-volume functions

A volume function associates every n-parallelepiped in an n-dimensional vector space with a non-negative real number. Let's try associating every n-parallelepiped with a scalar instead. We'll call our association a *pre-volume* function if it has the following properties:

1. If you stretch a parallelepiped by multiplying one of its legs by a scalar λ, its pre-volume gets multiplied by λ.

2. If you shear a parallelepiped by adding a multiple of one of its legs to another leg, its pre-volume stays the same.

More formally, a pre-volume on an n-dimensional vector space V over a field K is a function $E\colon V^n \to K$ with the following properties:

1. For any $\lambda \in K$,
$$E(\ldots \underset{i}{\lambda v} \ldots) = \lambda E(\ldots \underset{i}{v} \ldots).$$

2. For any $\lambda \in K$ and $i \neq j$,
$$E(\ldots \underset{i}{v} \ldots \underset{j}{w} \ldots) = E(\ldots \underset{i}{v + \lambda w} \ldots \underset{j}{w} \ldots).$$

It's immediately apparent that...

Proposition 4. *If K is an absolute value field, the absolute value of a pre-volume function is a volume function.*

We'll soon learn that the converse is also true: every volume function is the absolute value of a pre-volume function.

Pre-volume functions act a lot like volume functions. The following propositions, for example, are almost identical to our first three propositions about volume functions, and they're proven in essentially the same way.

Proposition 5. *If E is a pre-volume function,*

$$E(\dots \underset{i}{v} \dots \underset{j}{v} \dots) = 0$$

for all $i \neq j$. In other words, E is alternating.

Proposition 6. *If E is a pre-volume function,*

$$E(\dots \underset{i}{v} \dots \underset{j}{w} \dots) = -E(\dots \underset{i}{w} \dots \underset{j}{v} \dots)$$

for all $i \neq j$. In other words, E is skew-symmetric.

Proposition 7. *Say E is a pre-volume function. If $v_1, \dots, v_n$ are linearly dependent, $E(v_1, \dots, v_n) = 0$.*

For volume functions, our next step was to prove a partial converse of the third proposition, using the fact that volumes are only defined for parallelepipeds of the same dimension as V. For pre-volume functions, we can get a much stronger result.

Lemma 2. *Every pre-volume function is multilinear.*

Proof. Say $E \colon V^n \to K$ is a pre-volume. Our goal is to show that E has the following properties:

1. $E(\dots \lambda \underset{i}{v} \dots) = \lambda E(\dots \underset{i}{v} \dots)$
2. $E(\dots v \underset{i}{+} w \dots) = E(\dots \underset{i}{v} \dots) + E(\dots \underset{i}{w} \dots)$

The first property is just the stretch property, so it holds by definition. Since E is skew-symmetric, we only have to prove the second property for $i = 1$.

Let

$$\begin{aligned} A &= E(u, v_2, \dots, v_n) \\ B &= E(\tilde{u}, v_2, \dots, v_n) \\ C &= E(u + \tilde{u}, v_2, \dots, v_n). \end{aligned}$$

We want to show that $A + B = C$. If $v_2, \dots, v_n$ are linearly dependent, this follows immediately from Proposition 7, so we only need to consider the case where $v_2, \dots, v_n$ are linearly independent.

If u lies in the subspace spanned by $v_2, \dots, v_n$, the shear property can be used to show that $A = 0$ and $C = B$, and the desired result follows. Hence, we only need to consider the case where $u, v_2, \dots, v_n$ are linearly independent.

Since n is the dimension of V, the vectors $u, v_2, \ldots, v_n$ span V. That means we can write $\tilde{u}$ as $\lambda u + x$, where x lies in the subspace spanned by $v_2, \ldots, v_n$. Notice that $u + \tilde{u} = (\lambda + 1)u + x$. By the shear property,

$$\begin{aligned} B &= E(\lambda u, v_2, \ldots, v_n) \\ C &= E((\lambda + 1)u, v_2, \ldots, v_n), \end{aligned}$$

so the stretch property gives $B = \lambda A$ and $C = (\lambda + 1)A$. The desired result follows. □

We now have an extremely useful characterization of pre-volume functions.

Theorem 2. *A function $V^n \to K$ is a pre-volume if and only if it's multilinear and alternating.*

Proof. The "only if" direction is given by Proposition 5 and Lemma 2, and the "if" direction is a straightforward exercise in definition-chasing. □

This characterization makes it easy to prove that pre-volume functions always exist, and are unique up to scaling. In fact...

Corollary 1. *The pre-volume functions on V form a one-dimensional vector space.*

Proof. The alternating multilinear functions $V^n \to K$ form a vector space under the usual addition and scalar multiplication of functions. If $v_1, \ldots, v_n$ is a basis for V, an alternating multilinear function $E \colon V^n \to K$ is completely determined by the value $E(v_1, \ldots, v_n)$, and any choice of $E(v_1, \ldots, v_n)$ gives an alternating multilinear function. Hence, the space of alternating multilinear functions $V^n \to K$ is one-dimensional. □

Together, the existence of pre-volume functions and the uniqueness of volume functions imply the converse of Proposition 4.

Theorem 3. *Every volume function is the absolute value of a pre-volume function.*

Proof. Say K is an absolute value field, and ε is a volume function on V. Corollary 1 guarantees the existence of a nonzero pre-volume function H on V, and Proposition 4 tells us that the absolute value η of H is a volume function. By Theorem 1, one of the volume functions ε and η is a multiple of the other. Since η is nonzero, it follows that ε is a multiple of η. □

Using this result, we can turn Theorem 2 into a nice characterization of volume functions.

Corollary 2. *If K is an absolute value field, a function $V^n \to [0, \infty)$ is a volume function if and only if it's the absolute value of an alternating multilinear function.*

5 What is area?

Imagine you wake up one morning to find a two-dimensional parallelepiped floating in the air outside your window. Naturally, you'd like to know more about it. Your first thought is to measure its volume using a volume function, but of course that doesn't make sense: in a three-dimensional space, volumes are only defined for three-dimensional parallelepipeds. If you could somehow wrestle the parallelepiped to the ground, you could pin it flat against a two-dimensional board and measure its volume there. Unfortunately, the parallelepiped is far out of reach.

A few minutes of furious pondering bring you no closer to satisfying your curiosity, so you decide to forget about the apparition altogether. As you're leaving for work, however, you happen to glance up and notice that the morning sun is casting a shadow of the parallelepiped against the wall of your building. Now there's an idea! The shadow is a two-dimensional parallelepiped in a two-dimensional space, so you can measure its volume with a volume function. As the sun moves across the sky, it'll illuminate the parallelepiped from other directions and cast its shadow against other surfaces. You can use these multiple perspectives to build up a more complete picture of your mysterious visitor.

The story I've just told may not hold up under philosophical scrutiny, but it suggests a reasonably natural way of generalizing the notion of volume to lower-dimensional parallelepipeds. This generalization, which I'll call *area*, will be the subject of the next section.

6 Area functions

Let's say a *k-area function* on a vector space V is a function $\alpha\colon V^k \to [0, \infty)$ of the form

$$\alpha(v_1, \ldots, v_k) = \varepsilon(Lv_1, \ldots, Lv_k),$$

where $L\colon V \to W$ is a linear map, W is a k-dimensional vector space, and $\varepsilon\colon W^k \to [0, \infty)$ is a volume function. (For this to make sense, the base field K has to be an absolute value field.) Similarly, we'll define a *k-pre-area function* on V as a function $A\colon V^k \to K$ of the form

$$A(v_1, \ldots, v_k) = E(Lv_1, \ldots, Lv_k),$$

where L and W are as before, and $E\colon W^k \to K$ is a pre-volume function.

If n is the dimension of V, an n-area is the same thing as a volume, and an n-pre-area is the same thing as a pre-volume. Hence, area really is a generalization of volume.

In Section 4 (Theorem 2), we used the notion of pre-volume to characterize the alternating multilinear functions $V^n \to K$. Now, more generally, we can use the notion of pre-area to characterize the alternating multilinear functions $V^k \to K$ for any k.

Theorem 4. *A map $V^k \to K$ is multilinear and alternating if and only if it can be written as a linear combination of k-pre-areas.*

Proof. The "if" direction is easy. We'll start the "only if" direction by picking a basis $v_1, \ldots, v_n$ for V. For each function $I\colon \{1 \ldots k\} \to \{1 \ldots n\}$, define $\theta_I\colon V^k \to K$ as the unique multilinear function with the property that

$$\theta_I(v_{J(1)}, \ldots, v_{J(k)}) = \begin{cases} 1 & J = I \\ 0 & J \neq I \end{cases}$$

for all $J\colon \{1 \ldots k\} \to \{1 \ldots n\}$. For each strictly increasing function $I\colon \{1 \ldots k\} \to \{1 \ldots n\}$, define

$$\chi_I(w_1, \ldots, w_k) = \sum_\sigma \operatorname{sgn}(\sigma)\theta_I(w_{\sigma(1)}, \ldots, w_{\sigma(k)}),$$

where σ runs over all the permutations of $\{1 \ldots k\}$.

Any multilinear function can be written as a linear combination of the θ_I. Since every alternating multilinear function is skew-symmetric, any alternating multilinear function can be written as a linear combination of the χ_I.

It's apparent from the definition that each χ_I is alternating and skew-symmetric when its arguments are restricted to the set of basis vectors. A short calculation then proves that χ_I is alternating for all arguments. It follows that each χ_I is a k-pre-area, because χ_I is a pre-volume when its arguments are restricted to the subspace spanned by $v_{I(1)}, \ldots, v_{I(k)}$. □

This result links alternating multilinear maps to the geometric concept of area, using the notion of pre-area as an intermediary. In other words, it does exactly what I set out to do in the abstract of these notes. Unfortunately, the connection I've established is a lot more subtle than I'd hoped it would be. The subtlety is that, although Proposition 4 and Theorem 3 let us pass from pre-areas to areas and back in a straightforward way, they don't tell us what to make of a linear combination of pre-areas. In the case of n-pre-areas, the point is moot, because a linear combination of n-pre-areas is itself an n-pre-area (Corollary 1). This fact gives us a direct connection between n-linear maps and n-areas, enshrined in Corollary 2. For general k-pre-areas, however, the subtlety remains; until it can be resolved, the story told in these notes will be incomplete.

References

[1] Matt Noonan. Antisymmetry in geometry, I. In *The Everything Seminar*. `http://cornellmath.wordpress.com/2007/09/13/antisymmetry-in-geometry-i/`, 2007.