Survey on exterior algebra and differential forms

Daniel Grieser

16. Mai 2013

Inhaltsverzeichnis

1 Exterior algebra for a vector space

Let V be an n-dimensional real vector space. Whenever needed, we let $e_1, \dots, e_n$ be a basis of V and $e^1, \dots, e^n$ its dual basis.
At first reading you may leave out the parts on Hodge $*$ and non-positive definite metrics.

1.1 Alternating forms, wedge and interior product

1. Let $k \in \mathbb{N}_0$. A **k-multilinear form** on V is a map $\omega : V^k \to \mathbb{R}$ which is linear in each entry, i.e.

$$\omega(av_1 + bv_1', v_2, \dots, v_k) = a\,\omega(v_1, v_2, \dots, v_k) + b\,\omega(v_1', v_2, \dots, v_k)$$

for all $v_1, v_1', v_2, \dots, v_k \in V$ and $a, b \in \mathbb{R}$, and similarly for the other entries. The form is called **alternating** if it changes sign under interchange of any two vectors:

$$\omega(v_1, \dots, v_i, \dots, v_j, \dots, v_k) = -\omega(v_1, \dots, v_j, \dots, v_i, \dots, v_k)$$

Equivalent conditions (to alternating) are: $\omega(v_1, \dots, v_k) = 0$ if any two of the v_i are the same. Or:

$$\omega(v_{\sigma(1)}, \dots, v_{\sigma(k)}) = \operatorname{sign}(\sigma)\omega(v_1, v_2, \dots, v_k)$$

for all permutations σ of $\{1, \dots, k\}$.

The space of alternating k-multilinear forms on V is denoted $\Lambda^k V^*$. This is a vector space with basis $\{e^I : |I| = k\}$, where I runs over subsets of $\{1, \dots, n\}$ with k elements and

$$e^I := e^{i_1} \wedge \dots \wedge e^{i_k} \text{ for } I = \{i_1 < i_2 < \dots < i_k\},$$

with $\wedge$ defined below, or explicitly:

$$e^I(e_{j_1}, \dots, e_{j_k}) = \delta^I_J \text{ for } J = \{j_1 < \dots < j_k\}$$

For example, $e^{\{1,2\}} = e^1 \wedge e^2$ satisfies $(e^1 \wedge e^2)(e_1, e_2) = 1$, and this implies that $(e^1 \wedge e^2)(e_2, e_1) = -1$ and all other $(e^1 \wedge e^2)(e_{j_1}, e_{j_2})$ are zero, hence for $v = \sum v^i e_i, w = \sum w^j e_j$ we get $(e^1 \wedge e^2)(v, w) = v^1 w^2 - v^2 w^1$.

It follows that $\Lambda^k V^*$ has dimension $\binom{n}{k}$, in particular $\dim \Lambda^n V^* = 1$, and $\Lambda^k V^* = \{0\}$ if $k > n$. Also $\Lambda^1 V^* = V^*$ and $\Lambda^0 V^* = \mathbb{R}$.[1] We also write

$$k = \deg \omega \text{ if } \omega \in \Lambda^k V^*$$

and call k the **degree** of the form ω.

2. **Wedge (or exterior) product:** For $\omega \in \Lambda^k V^*$, $\nu \in \Lambda^l V^*$ define $\omega \wedge \nu \in \Lambda^{k+l} V^*$ by

$$(\omega \wedge \nu)(v_1, \dots, v_{k+l}) = \sum_\sigma \operatorname{sign}(\sigma)\, \omega(v_{\sigma(1)}, \dots, v_{\sigma(k)}) \nu(v_{\sigma(k+1)}, \dots, v_{\sigma(k+l)})$$

where the sum runs over all permutations σ of $\{1, \dots, k+l\}$ preserving the order in the first and second 'block', i.e. satisfying $\sigma(1) < \dots < \sigma(k)$, $\sigma(k+1) < \dots < \sigma(k+l)$.

For example, for $k = l = 1$

$$(\omega \wedge \nu)(v, w) = \omega(v)\nu(w) - \omega(w)\nu(v)$$

Rule:

$$\boxed{\omega \wedge \nu = (-1)^{\deg \omega \cdot \deg \nu} \nu \wedge \omega}$$

For this property one says that $\wedge$ is 'graded commutative' (in the physics literature also 'super commutative'). Also, $\wedge$ is bilinear and associative.

Remark (Relation to cross product in $\mathbb{R}^3$):
The wedge product generalizes the cross product in the following sense. If $V = \mathbb{R}^3$ then $\dim \Lambda^1 \mathbb{R}^3 = \dim \Lambda^2 \mathbb{R}^3 = 3$. So we have identifications (isomorphisms)

$$\begin{aligned} &\Lambda^1 \mathbb{R}^3 \to \mathbb{R}^3, & \omega = \omega_1 e^1 + \omega_2 e^2 + \omega_3 e^3 \mapsto (\omega_1, \omega_2, \omega_3) \\ &\Lambda^2 \mathbb{R}^3 \to \mathbb{R}^3, & \mu = \mu_1 e^2 \wedge e^3 + \mu_2 e^3 \wedge e^1 + \mu_3 e^1 \wedge e^2 \mapsto (\mu_1, \mu_2, \mu_3) \end{aligned}$$

Now for $\omega, \nu \in \Lambda^1 \mathbb{R}^3$ with $\omega = \sum \omega_i e^i$, $\nu = \sum \nu_j e^j$ we have

$$\omega \wedge \nu = (\omega_2 \nu_3 - \omega_3 \nu_2) e^2 \wedge e^3 + (\omega_3 \nu_1 - \omega_1 \nu_3) e^3 \wedge e^1 + (\omega_1 \nu_2 - \omega_2 \nu_1) e^1 \wedge e^2$$

so if ω, ν are identified with vectors as in the first line, then $\omega \wedge \nu$ corresponds (as in the second line) to the cross product of these vectors.

3. Let $v \in V$. The **interior product** with v is the linear operator

$$\iota_v : \Lambda^k V^* \to \Lambda^{k-1} V^*, \quad \omega \mapsto \omega(v, \dots)$$

that is, $(\iota_v \omega)(v_2, \dots, v_k) = \omega(v, v_2, \dots, v_k)$ ('plug in v in the first slot'). Here $k \in \mathbb{N}$, but we also define $\iota_v = 0$ on $\Lambda^0 V^*$.

Clearly, ι_v depends linearly on v. With wedge products it behaves as follows (as a 'super derivation'):

$$\boxed{\iota_v(\omega \wedge \nu) = (\iota_v \omega) \wedge \nu + (-1)^{\deg \omega} \omega \wedge (\iota_v \nu)} \tag{1}$$

For example, in $\mathbb{R}^3$, if $v = \sum_i v^i e_i$ then

$$\iota_v(e^1 \wedge e^2 \wedge e^3) = v^1 e^2 \wedge e^3 + v^2 e^3 \wedge e^1 + v^3 e^1 \wedge e^2 \tag{2}$$

(of course one could write $-v^2 e^1 \wedge e^3$ for the middle term)

4. **Behavior under maps:** A linear map $A : V \to W$ defines the pull-back map

$$A^* : \Lambda^k W^* \to \Lambda^k V^*, \quad (A^*\omega)(v_1, \dots, v_k) := \omega(Av_1, \dots, Av_k) \tag{3}$$

[1] By definition, $V^0 = \mathbb{R}$, and a linear map $\mathbb{R} \to \mathbb{R}$ is determined by its value at 1.

where $\omega \in \Lambda^k W^*$, $v_1, \ldots, v_k \in V$. For $k = 1$ this is also called the dual (or transpose) map $A^* : W^* \to V^*$.

Pullback behaves naturally with wedge product

$$A^*(\omega \wedge \nu) = A^*\omega \wedge A^*\nu$$

and with interior product: $\iota_v(A^*\omega) = A^*(\iota_{A(v)}\omega)$, as follows directly from the definitions.

For $V = W$ and $k = n$ pullback relates to the determinant as follows: $A^* = (\det A)\,\mathrm{Id}$ on $\Lambda^n V^*$. Explicitly, this means

$$\omega(Av_1, \ldots, Av_n) = (\det A)\omega(v_1, \ldots, v_n)$$

which follows directly from the facts that ω is multilinear and alternating, and the Leibniz formula for the determinant.

1.2 A scalar product enters the stage

From now on assume that a **scalar product is given** on V, that is, a bilinear, symmetric, positive definite[2] form $g : V \times V \to \mathbb{R}$. We also write $\langle v, w\rangle$ instead of $g(v, w)$. This defines some more structures:

1. Basic geometry: The scalar product allows us to talk about **lenghts** of vectors and **angles** between non-zero vectors:

$$|v| = \sqrt{g(v, v)}, \quad \angle(v, w) = \arccos \frac{g(v, w)}{|v| \cdot |w|}$$

2. Using the scalar product on V we get a map

$$g^\# : V \to V^*, v \mapsto g(v, \cdot)$$

Since g is non-degenerate, this map is injective, hence bijective (since $\dim V = \dim V^* < \infty$). The inverse of $g^\#$ is called

$$g^\flat : V^* \to V$$

Therefore, we may identify vectors and linear forms (but we do this only when necessary).[3]

3. Using this identification, we get a scalar product on V^*, which we also denote by $\langle\ ,\ \rangle$:

$$\langle\alpha, \beta\rangle := \langle g^\flat(\alpha), g^\flat(\beta)\rangle$$

for $\alpha, \beta \in V^*$.

4. More generally, we get a scalar product on $\Lambda^k V^*$ for each k. It is easiest to define it by the property:

If $e_1, \ldots, e_n$ are orthonormal then the basis $\{e^I : |I| = k\}$ of $\Lambda^k V^*$ is orthonormal.

In other words, $\langle \sum_I a_I e^I, \sum_J b_J e^J\rangle := \sum_I a_I b_I$. Then for arbitrary $v^1, \ldots v^k, w^1, \ldots, w^k \in V^*$ one has[4]

$$\boxed{\langle v^1 \wedge \cdots \wedge v^k, w^1 \wedge \cdots \wedge w^k\rangle = \det(\langle v^i, w^j\rangle)} \tag{4}$$

This formula also shows that one obtains the same scalar product if one uses a different orthonormal basis in the definition.

[2]Everything can be done in the more general case that g is only non-degenerate, but one needs to be careful with the signs, see Section 1.5.

[3]The fact that $g^\#$ is surjective, i.e. that every linear form on V can be represented by a vector using the scalar product, is sometimes called the *Riesz lemma*. It holds more generally when (V, g) is a Hilbert space, that is, if V is allowed to be infinite-dimensional but required to be complete with the norm defined by g.

[4]Proof: By definition this holds if all v^i, w^j are taken from the basis vectors $e^1, \ldots, e^n$. Then it holds in general since both sides are multilinear in the $2k$ entries $v^1, \ldots v^k, w^1, \ldots, w^k$.

1.3 Now add an orientation: Volume element, Hodge $*$

Now assume that on V a scalar product and an orientation is given.

1. The **Hodge $*$ operator** is the unique linear map (for each k)
$$*: \Lambda^k V^* \to \Lambda^{n-k} V^*$$
with the property that[5]
$$*\,(e^1 \wedge \cdots \wedge e^k) = e^{k+1} \wedge \cdots \wedge e^n \quad \text{for any oONB,} \tag{5}$$
that is, for any oriented orthonormal basis (oONB) $e_1, \dots, e_n$ with dual basis $e^1, \dots, e^n$.
Intuition: k-forms $e^1 \wedge \cdots \wedge e^k$ correspond to k-dimensional subspaces $W = \text{span}\{e^1, \dots, e^k\}$ of V^*. Then $*(e^1 \wedge \cdots \wedge e^k)$ corresponds to the orthogonal complement of W.
Of course not every form can be written in this way, but using linearity $*$ is defined when it is defined on forms of this type.
So one can say:[6]
    - Alternating multilinear forms are a 'linear extension' of the notion of vector subspace.
    - Then $*$ corresponds to orthogonal complement.

2. Define the **volume element** (or volume form) of V as
$$\text{dvol} = *1, \qquad \text{dvol} \in \Lambda^n V^*.$$
Why is this reasonable? Because for any oONB $e_1, \dots, e_n$ we have, by definition of $*$, $\text{dvol} = e^1 \wedge \cdots \wedge e^n$ and therefore
$$\text{dvol}(e_1, \dots, e_n) = 1 \quad \text{for any oONB.} \tag{6}$$
So the volume of a 'unit cube' is one, as it should be.

3. Properties of $*$:
$$\boxed{\omega \wedge *\nu = \langle \omega, \nu \rangle\, \text{dvol} \quad \text{for } \omega, \nu \in \Lambda^k V^*} \tag{7}$$
Also, if ν is fixed then the validity of (7) for all ω defines $*\nu$.
(7) can easily be checked on basis elements, and then extends by linearity.[7]
$$*\,(e^{i_1} \wedge \cdots \wedge e^{i_k}) = \text{sign}(\sigma) e^{j_1} \wedge \cdots \wedge e^{j_{n-k}} \quad \text{for oONB} \tag{8}$$
where $\{j_1, \dots, j_{n-k}\} = \{1, \dots, n\} \setminus \{i_1, \dots, i_k\}$ and σ is the permutation sending $(1, \dots, n) \mapsto (i_1, \dots, i_k, j_1, \dots, j_{n-k})$. From this one gets easily[8]
$$** = (-1)^{k(n-k)} \quad \text{on } \Lambda^k V^* \tag{9}$$
That is, if $\omega \in \Lambda^k V^*$ then $*\omega \in \Lambda^{n-k} V^*$, and $*(*\omega) \in \Lambda^k V^*$ equals $(-1)^{k(n-k)}\omega$.

---

[5] Uniqueness of such a linear map is clear, existence is less obvious. See footnote 7.

[6] As an exercise, you might try to make these somewhat vague ideas more precise. For example: To what extent does a subspace of dimension k determine a 'pure' form of degree k (i.e. one which can be written as wedge product of one-forms) uniquely?

[7] This assumes we know the existence of the linear map $*$. A logically more sound way of introducing $*$ is this:

(a) Define $\text{dvol} \in \Lambda^n V^*$ by equation (6) for a fixed oriented ONB, and check that (6) must then hold for any oriented ONB (this follows from (11)). Since $\dim \Lambda^n V^* = 1$, $\{\text{dvol}\}$ is a basis of $\dim \Lambda^n V^*$.

(b) Consider the map $P : \Lambda^k V^* \times \Lambda^{n-k} V^* \to \mathbb{R}$, $(\omega, \mu) \mapsto$ (the coefficient a in $\omega \wedge \mu = a\,\text{dvol}$). This is easily seen to be bilinear and (e.g. using a basis) non-degenerate. Therefore, by Riesz' lemma, for any linear form $q : \Lambda^k V^* \to \mathbb{R}$ there is a unique element $\mu \in \Lambda^{n-k} V^*$ so that $P(\omega, \mu) = q(\omega)$ for all $\omega \in \Lambda^k V^*$.

(c) Now given $\nu \in \Lambda^k V^*$, apply the Riesz lemma to the form $q(\omega) = \langle \omega, \nu \rangle$. This determines an element $\mu \in \Lambda^{n-k} V^*$. Define $*\nu := \mu$. Then (7) holds by definition, and from this (5) follows.

[8] From $\text{sign}(k+1, \dots, n, 1, \dots, k) = (-1)^{k(n-k)}$

4. As an exercise, use the previous properties to prove: If $v \in V$ then

$$* g^{\#}(v) = \iota_v \mathrm{dvol} \tag{10}$$

Also check this in the example (2), where e_1, e_2, e_3 is the standard basis und g the standard scalar product.[9]

1.4 Formulas in an arbitrary basis

For the application in the manifold setting we need formulas in terms of *any* basis $e_1, \dots, e_n$ of V (not necessarily orthonormal), for the objects defined by a scalar product.

1. The scalar product determines (and is determined by) the $n \times n$ matrix (g_{ij}) where

$$g_{ij} := \langle e_i, e_j \rangle$$

2. The maps $g^{\#}$, $g^{\flat}$ are given as follows: Suppose $v \in V$, $\alpha \in V^*$ satisfy $\alpha = g^{\#}(v)$, or equivalently $v = g^{\flat}(\alpha)$. Then

$$\alpha_j = \sum_i g_{ij} v^i, \quad v^i = \sum_j g^{ij} \alpha_j$$

Here (g^{ij}) is the inverse matrix of (g_{ij}). These operations (going from the coefficients v^i to the α_j, and vice versa) are called **lowering and raising indices** using the scalar product g.[10]

3. From this it easily follows that the scalar product on V^* is given by the matrix (g^{ij}):

$$\langle e^i, e^j \rangle = g^{ij}$$

More generally, (4) gives for k-forms

$$\langle e^I, e^J \rangle = \det(g^{ij})_{i \in I, j \in J}$$

(where the indices on the right are listed in increasing order).

4. Now assume that V is oriented with oriented basis $e_1, \dots, e_n$ (still not necessarily orthonormal). Then[11]

$$\mathrm{dvol} = \sqrt{\det(g_{ij})} e^1 \wedge \dots \wedge e^n \tag{11}$$

---

[9] Hint for (10): By the statement after (7) this follows if we show that for all $\omega \in \Lambda^1 V$

$$\omega \wedge (\iota_v \mathrm{dvol}) = \langle \omega, g^{\#}(v) \rangle \mathrm{dvol}$$

Now by definition of $g^{\#}$, we have $\langle \omega, g^{\#}(v) \rangle = \omega(v) = \iota_v \omega$. Now use the product rule (1) for ι_v.

Alternative proof of (10): By linearity it suffices to prove this for unit vectors v. Set $e_1 = v$ and extend to an oONB $e_1, \dots, e_n$. Then check equality of both sides when applied to any $(n-1)$-tuple out of $e_1, \dots, e_n$.

Explicitly in an oONB, both sides are $\sum v^i (-1)^{i-1} e^1 \wedge \dots \wedge \widehat{e^i} \wedge \dots \wedge e^n$, where the hat means omission.

[10] Conventions often used in physics:

- A vector is denoted by its compoents: (v^i), or simply v^i (rather than $\sum v^i e_i$). Similarly a covector (element of V^*) is denoted v_i (instead of $\sum v_i e^i$).
- The summation sign is omitted (Einstein summation convention).
- The same letter is used for a vector and the corresponding covector (i.e. element of V^*). Thus, one writes $v_i = g_{ij} v^j$.

[11] Proof: Choose an oONB $E_1, \dots E_n$ and write $e_i = \sum_k a_i^k E_k$. Then, using $\langle E_k, E_l \rangle = \delta_{kl}$ we get

$$g_{ij} = \langle \sum_k a_i^k E_k, \sum_l a_j^l E_l \rangle = \sum_{k,l} a_i^k a_j^l \langle E_k, E_l \rangle = \sum_k a_i^k a_j^k$$

which is the ij entry of the matrix $A^t A$, where A is the matrix (a_i^k). Therefore $\det(g_{ij}) = (\det A)^2$, so $\det A = \sqrt{\det(g_{ij})}$ since $\det A > 0$ (both bases $e_1, \dots, e_n$ and $E_1, \dots, E_n$ are oriented). Therefore, $\mathrm{dvol}(e_1, \dots, e_n) = \det A \, \mathrm{dvol}(E_1, \dots, E_n) = \det A$, and $\mathrm{dvol} = \mathrm{dvol}(e_1, \dots, e_n) e^1 \wedge \dots \wedge e^n$ gives the claim.

5. For the Hodge $*$ operator we have: Let $\omega = \sum_I \omega_I e^I$, then $*\omega = \sum_J (*\omega)_J e^J$ with

$$(*\omega)_J = \omega^I \sqrt{\det(g_{ij})}\,\mathrm{sign}(\sigma)$$

where σ is the permutation $(1, \dots, n) \mapsto (I, J)$ with I, J listed in increasing order. Here, ω^I is obtained by raising indices from the ω_I, that is

$$\omega^{i_1,\dots,i_k} = \sum g^{i_1 l_1} \cdots g^{i_k l_k} \omega_{l_1,\dots,l_k}$$

where $\omega_{l_1,\dots,l_k} := \omega(e_{l_1}, \dots, e_{l_k})$.

1.5 Modifications for not positive definite inner product

If the bilinear form g on V is not positive definite (but still symmetric and non-degenerate) then we need to modify some of the formulas slightly.

Define the **index** of g as the dimension of a maximal subspace on which g is negative definite. Equivalently[12], it is the number of negative eigenvalues of the matrix of g with respect to any basis. We denote the index of g by ν.

1. First, $g(v, v)$ may be negative, so the length of a vector is defined as

$$|v| := \sqrt{|g(v, v)|}$$

2. A **standard basis** of V is a basis $e_1, \dots, e_n$ for which

$$\langle e_i, e_j \rangle = \varepsilon_i \delta_{ij}$$

where[13]

$$\varepsilon_1 = \cdots = \varepsilon_\nu = -1, \varepsilon_{\nu+1} = \cdots = \varepsilon_n = 1$$

Standard bases replace orthonormal bases in this context.

3. The scalar product on $\Lambda^k V^*$ is still characterized by property (4).

4. The volume form is still defined by the property (6) (for an oriented standard basis), so that[14]

$$\mathrm{dvol} = \sqrt{|\det(g_{ij})|}\, e^1 \wedge \cdots \wedge e^n \text{ (any oriented basis)}$$

5. The Hodge $*$ operator is defined by property (7). Then in (5) and (8) there is an extra factor $(-1)^{\nu'}$ on the right, where ν' is the number of vectors e^i, $i \in \{i_1, \dots, i_k\}$, with $\langle e^i, e^i \rangle = -1$.

Then (9) gets replaced by

$$** = (-1)^{k(n-k)+\nu} \quad \text{on } \Lambda^k V^*$$

Example: Minkowski space is $\mathbb{R}^4$ with the standard scalar product of index 1 and standard orientation. Coordinates are usually denoted t, x, y, z (in this order), so[15]

$$\langle \partial_t, \partial_t \rangle = -1, \langle \partial_x, \partial_x \rangle = \langle \partial_y, \partial_y \rangle = \langle \partial_z, \partial_z \rangle = 1.$$

Then $\mathrm{dvol} = dt \wedge dx \wedge dy \wedge dz$ and

$$\begin{aligned} *dt &= -dx \wedge dy \wedge dz & *(dx \wedge dy \wedge dz) &= -dt \\ *dx &= -dt \wedge dy \wedge dz & *(dt \wedge dy \wedge dz) &= -dx \\ *(dt \wedge dx) &= -dy \wedge dz & *(dy \wedge dz) &= dt \wedge dx \end{aligned}$$

etc. (cyclically permute x, y, z). Note $** = 1$ on Ω^1 and Ω^3 and $** = -1$ on Ω^2.

[12] This fact is called Sylvester's law of inertia.

[13] Sometimes a different convention is used, where the last ν elements are negative.

[14] Sometimes a different convention is used, where dvol gets an extra factor $(-1)^\nu$, so that $\mathrm{dvol} = (-1)^\nu \sqrt{\det(g_{ij})}\, e^1 \wedge \cdots \wedge e^n$.

[15] This is one of the common conventions, mostly used by mathematicians and graviational physicists. Particle physicists mostly use a different convention, where all the signs are turned around.

2 Differential forms

Let M be a manifold of dimension n.

2.1 Pointwise ('tensorial') constructions

The constructions of the previous section can be done on each tangent space $V = T_pM$. In this way we obtain, for example:

- A differential form ω of degree k (or differential k-form, or k-form) is given by $\omega_p \in \Lambda^k T_p^* M$ for each p, smoothly depending on p. The space of differential forms of degree k is denoted $\Omega^k(M)$. In particular, 0-forms are functions, $\Omega^0(M) = C^\infty(M, \mathbb{R})$.
- Wedge product defines a bilinear map $\wedge : \Omega^k(M) \times \Omega^l(M) \to \Omega^{k+l}(M)$. (For $k = 0$ this is simply multiplying a form by a function.)
- Interior product with a vector field $X \in \mathcal{X}(M)$ defines a linear map $\iota_X : \Omega^k(M) \to \Omega^{k-1}(M)$,[16] more precisely a $C^\infty(M,\mathbb{R})$-bilinear map $\iota : \mathcal{X}(M) \times \Omega^k(M) \to \Omega^{k-1}(M)$.
- Any smooth map $F : M \to N$ defines a pullback map

  $$F^* : \Omega^k(N) \to \Omega^k(M), \quad (F^*\omega)_p(v_1, \dots, v_k) := \omega_{F(p)}(dF_{|p}(v_1), \dots dF_{|p}(v_k))$$

  for $\omega \in \Omega^k(N)$, $p \in M$ and any vectors $v_1, \dots, v_k \in T_pM$ (apply (3) with $A = dF_{|p}$).
- A Riemannian metric on M is given by a scalar product g_p on T_pM for each p. It defines linear maps $g^\# : \mathcal{X}(M) \to \Omega^1(M)$, $g^\flat : \Omega^1(M) \to \mathcal{X}(M)$ and a scalar product on $\Lambda^k T_p^* M$ for each p.
- An orientation of M is given by an orientation on each T_pM, varying continuously with p. Given a scalar product and an orientation, we get the Hodge $*$ operator

  $$* : \Omega^k(M) \to \Omega^{n-k}(M)$$

  and the volume form $\mathrm{dvol} \in \Omega^n(M)$. [17]

All the rules from before still hold since they hold pointwise at each p.

Formulas in local coordinates

Given local coordinates $x^1, \dots, x^n$ on a coordinate patch $U \subset M$, one can express all these concepts and operations in terms of the basis $\partial_1, \dots, \partial_n$ of T_pM and its dual basis $dx^1, \dots, dx^n$ of T_pM^*, for $p \in U$. That is, in the formulas of Section 1 (especially 1.4)[18] one sets[19]

$$e_i = \partial_i, \; e^i = dx^i, \; i = 1, \dots, n$$

Some examples of this are:

---

[16] For the case $k = 0$, i.e. functions f, we define $\iota_X f = 0$. In this way the iddi formula below, see (18), holds on forms of any degree, including functions. Also $\Omega^{-1}(M) := \{0\}$.

[17] Note that dvol is *not* d applied to an $(n-1)$-form – at least not globally. Locally it is (by the Poincaré Lemma).

[18] Note that when considering (semi-)RIemannian manifolds, one should use the formulas for an arbitrary basis, not for an ON basis. Why? Because usually one cannot choose local coordinates for which the ∂_i form an ONB at each $p \in U$. To be precise:

- Fix $p \in M$. Then local coordinates can be chosen near p so that $\partial_1, \dots, \partial_n$ form an ONB *at* p.
- Local coordinates can be chosen with $\partial_1, \dots, \partial_n$ an ONB for *each* $p \in U$ if and only if (U, g) is locally isometric to $\mathbb{R}^n$ with the Euclidean metric (or, equivalently, if the curvature of g is identically zero on U).

Proof as exercise. (The statement about curvature is harder, will be proved in lecture.)

[19] More precise notation would be $\partial_{i|p}$, $dx^i_{|p}$, but often the p is left out for better readability.

- A differential k-form in local coordinates is of the form

$$\omega = \sum_I a_I(x)\, dx^I, \quad dx^I := dx^{i_1} \wedge \cdots \wedge dx^{i_k} \text{ if } I = \{i_1 < \cdots < i_k\}$$

with smooth functions $a_I : U \to \mathbb{R}$.

- **Pull-back is just plugging in:** Let $F : M \to N$ be a smooth map. Suppose F is given in local coordinates $x^1, \ldots, x^n$ for M and $y^1, \ldots, y^m$ for N as $y(x) = (y^1(x), \ldots, y^m(x))$.[20] Then for $\omega = \sum a_I(y)\, dy^I \in \Omega^k(N)$, we have

$$F^*\omega = \sum_{I=\{i_1<\cdots<i_k\}} a_I(y(x))\, dy^{i_1} \wedge \cdots \wedge dy^{i_k}$$

where each y^{i_j} is considered as function of x, so one should write $dy^{i_j} = \sum_l \frac{\partial y^{i_j}}{\partial x^l} dx^l$ and then multiply out.

- The volume form on an oriented Riemannian manifold is

$$\boxed{\mathrm{dvol} = \sqrt{\det(g_{ij})}\, dx^1 \ldots dx^n} \tag{12}$$

in oriented local coordinates, where $g_{ij} = g(\partial_i, \partial_j)$.

2.2 Integration

One of the motivations for considering differential forms is that they are the objects that can be integrated invariantly over a manifold. More precisely, if $(M, \mathcal{O})$ is an oriented manifold and $\omega \in \Omega^n_0(M)$,[21] where $n = \dim M$, then

$$\int_{(M,\mathcal{O})} \omega \tag{13}$$

is well-defined[22]. Instead of (13) one usually writes $\int\limits_M \omega$, if $\mathcal{O}$ is fixed in the context. The definition proceeds in two steps:

1. First assume $\operatorname{supp}\omega \subset U$ for an orientation preserving local chart $\varphi : \tilde{U} \to U$. The local coordinate representation $\varphi^*\omega$ can be written as $\varphi^*\omega = a(x)\, dx^1 \wedge \cdots \wedge dx^n$ for some function a on $\tilde{U}$. Define

$$\int_M \omega = \int_{\tilde{U}} a(x)\, dx \tag{14}$$

Note that dx here stands for n-dimensional Lebesgue measure.

One then checks that the result is independent of the choice of coordinates. This is due to the way that differential forms transform under coordinate transformations: A $\det d\kappa$ factor appears, and this corresponds precisely to the $|\det d\kappa|$ factor in the transformation formula for integrals – if the determinant is positive, which is true if both charts are orientation preserving.

2. Any $\omega \in \Omega^n_0(M)$ can be integrated by summing over coordinate patches and applying the first part. In practice, often one or two coordinate systems suffice[23]. For theoretical purposes

[20] That is, for any $p \in M$, if p has coordinates x_0 and $F(p)$ has coordinates y_0 then $y_0 = y(x_0)$.

[21] The 0 in $\Omega^n_0(M)$ means compact support, i.e. elements of $\Omega^n_0(M)$ are zero outside of some compact set. This is assumed for simplicity to avoid problems with integrability. Of course weaker conditions would suffice.

[22] In contrast, $\int_M f$ would not be well-defined for a function f. Naively, one might try to define this, if f is supported in a coordinate patch $U \subset M$ with coordinates $x : U \to \tilde{U}$, as $\int_{\tilde{U}} \tilde{f}(x)dx$, where $\tilde{f}$ is f in coordinates x; however, this would depend on the choice of coordinates: If $y : V \to \tilde{V}$ is a different coordinate system then $\int_{\tilde{V}} \tilde{\tilde{f}}(y)\, dy = \int_{\tilde{U}} \tilde{f}(x)\, |\det d\kappa| dx$ with $\kappa = y \circ x^{-1}$ the coordinate change.
A different way to overcome this difficulty is to choose a measure μ on M and consider $\int_M f\, \mu$. The advantage of n-forms over measures is that they are part of the exterior calculus (i.e. $\wedge, d$ etc.).

[23] However, the restriction to the patch will usually not be compactly supported, and one possibly misses a set of measure zero, which does not affect the integral

(for example the proof of Stokes' theorem) it is better to use a partition of unity, which means splitting up ω smoothly into pieces which are compactly supported in coordinate patches. That is, choose a cover of M by orientation preserving charts $(U_i)_{i\in I}$ and a corresponding partition of unity $(\rho_i)_{i\in I}$. Then set

$$\int_M \omega := \sum \int_{U_i} \omega_i, \quad \omega_i := \rho_i \omega$$

This makes sense since $\sum_i \rho_i = 1$, so $\sum_i \omega_i = \omega$. Also $\omega_i \in \Omega_0^n(U_i)$.

One then checks that the result is independent of the choice of the U_i and of the ρ_i.

2.3 Derivative operations

There are several different operations in which derivatives are taken: Exterior derivative and Lie derivative (and later also covariant derivative).

The exterior derivative is defined only on differential forms (alternating $\mathcal{T}_k^0$-tensors). Lie derivative and covariant derivative are defined for all tensors.

Both d and Lie derivative are defined for a manifold, without scalar product.

1. The **exterior derivative** $d : \Omega^k(M) \to \Omega^{k+1}(M)$ is defined for $k = 0$ (functions) as the usual differential $d : f \mapsto df$ (in coordinates $df = \sum_i \frac{\partial f}{\partial x^i}\, dx^i$) and for any k in local coordinates by the formula

   $$\boxed{d\left(\sum_I a_I\, dx^I\right) = \sum_I da_I \wedge dx^I}$$

   Rules for d: d is linear, obeys the product rule[24]

   $$\boxed{(\omega \wedge \nu) = (d\omega) \wedge \nu + (-1)^{\deg \omega} \omega \wedge (d\nu)}$$

   commutes with pullback by a smooth map $F : M \to N$:

   $$\boxed{F^* \circ d = d \circ F^*}$$

   (this implies that d is well-defined on a manifold, independent of the choice of coordinates) and

   $$\boxed{d^2 = 0} \tag{15}$$

   (this will be essential for cohomology).

   One of the main reasons to consider the exterior derivative is that the general **Stokes' theorem** holds (see below for more on this): If M is an oriented manifold with boundary and ∂M is equipped with the induced orientation and $\omega \in \Omega_0^{n-1}(M)$ where $n = \dim M$ then

   $$\boxed{\int_M d\omega = \int_{\partial M} \omega} \tag{16}$$

2. The **Lie derivative** along a vector field $X \in \mathcal{X}(M)$. As for general tensors this is defined as

   $$L_X : \Omega^k(M) \to \Omega^k(M), \quad L_X\omega = \frac{d}{dt}_{|t=0} \Phi_t^* \omega$$

   where Φ is the flow of X. So L_X measures how ω changes ('deforms') under the flow of X. In particular,

   $$L_X \omega = 0 \iff \Phi_t^* \omega = \omega \quad \forall t \tag{17}$$

[24] So d is a 'graded derivation', just like the interior product ι_v, see (1). The $(-1)^{\deg \omega}$ factor comes from 'pulling d past ω' in the second summand, and similarly for ι_v. If ω is a product of 1-forms, then pulling d past each 1-form produces a -1 factor. A general ω is a sum of such products.
Note that both d and ι_v change the degree of a form by one. The Lie derivative does not, and its product rule has no $\pm$ in front of the second term.

The right side expresses a symmetry (invariance) property of ω.

Rules for the Lie derivative: Most importantly the 'iddi-formula':

$$\boxed{L = \iota d + d\iota} \tag{18}$$

that is, $L_X = \iota_X d + d\iota_X$, that is $L_X\omega = \iota_X(d\omega) + d(\iota_X\omega)$. This makes calculating $L_X\omega$ much easier than the original definition.[25]

Also there is a product rule[26][27]

$$L_X(\omega \wedge \nu) = (L_X\omega) \wedge \nu + \omega \wedge (L_X\nu) \tag{19}$$

and L_X commutes with d:[28]

$$L_X \circ d = d \circ L_X$$

3. **Comparison of Lie derivative and exterior derivative.**

   Recall that Lie derivative and exterior derivative agree on functions, in the sense that

$$L_X f = df(X) \tag{20}$$

   For forms of higher degree this is no longer true[29].

   Lie derivative and exterior derivative generalize two different ideas connected to the derivative of a function. To see this, consider for simplicity a function of one variable $f : \mathbb{R} \to \mathbb{R}$.

   - *Derivative as rate of change* $\rightsquigarrow$ *Lie derivative:*
     The formula $f'(x) = \lim_{t\to 0} \frac{f(x+t)-f(x)}{t}$ can be written $f'(x) = \frac{d}{dt}_{|t=0}(\Phi_t^* f)(x)$ for $\Phi_t(x) = x + t$ the flow of the unit vector field ∂_x.
   - *Derivative as inverse of integration* $\rightsquigarrow$ *exterior derivative:*
     The fundamental theorem of calculus

$$\int_a^b f'(x)\,dx = f(b) - f(a)$$

     is the special case $M = [a, b]$ (where $\partial M = \{a, b\}$ and standard orientation is used) of Stokes' theorem since the left side is $\int_M df$ and the right is $\int_{\partial M} f$. The exterior derivative is defined in such a way that this generalizes to higher dimensions. More precisely: There is a unique way to define linear maps $d : \Omega^{k-1}(M) \to \Omega^k(M)$ for any $k \in \mathbb{N}$ and any manifold M which is natural (i.e. commutes with pull-back by smooth maps) and so that Stokes' theorem (16) holds for all oriented manifolds M with boundary and all compactly supported forms $\omega \in \Omega_0^{\dim M-1}(M)$.[30][31]

4. **grad, div, rot**. These are really special cases of the exterior derivative d. But to define them on a manifold, *one needs a (semi-)Riemannian metric* (for d one doesn't). In this sense d is the more basic (and more general) operation.

[25] The formula is also the central piece in proving homotopy invariance of de Rham cohomology.

[26] So L_X is a 'derivation'. Note that this is different from the product rule for d because there is no $\pm$ sign in front of the second summand.

[27] Proof: Use $\Phi_t^*(\omega \wedge \nu) = (\Phi_t^*\omega) \wedge (\Phi_t^*\nu)$ and differentiate both sides in t.

[28] Follows directly from the definition of L_X and $\Phi_t^* \circ d = d \circ \Phi_t^*$.

[29] More precisely, one could write $df(X) = \iota_X f$ and then ask if $L_X f = \iota_X df$ holds with f replaced by a k-form. The iddi formula $L_X\omega = \iota_X d\omega + d(\iota_X\omega)$ shows that this is not the case, and shows that the correction term is $d(\iota_X\omega)$. Recall that $\iota_X f = 0$ by definition, so this term disappears for functions.

[30] Proof of uniqueness: Let $\omega \in \Omega^{k-1}(M)$. First, if $\dim M = k$ then $d\omega$ is uniquely determined since (16) must also hold for any open subset of M with smooth boundary – then use (14) and the corresponding fact for the Lebesgue integral. Next, if $\dim M = n$ with $n > k$ arbitrary then apply this argument for any k-dimensional submanifold N of M. It shows that $d(i_N^*\omega)$, with $i_N : N \hookrightarrow M$ the inclusion, is uniquely determined. By naturality $d(i_N^*\omega) = i_N^* d\omega$. Finally, a k-form is uniquely determined by its pull-backs to arbitrary k-dimensional submanifolds (use coordinate subspaces in a local coordinate system), so $d\omega$ is determined.

[31] As an exercise, try to derive the formula for $d\omega$ from this condition!

grad : $C^\infty(M) \to \mathcal{X}(M)$ and div : $\mathcal{X}(M) \to C^\infty(M)$ are defined on semi-Riemannian manifolds of any dimension, but rot : $\mathcal{X}(M) \to \mathcal{X}(M)$ is defined only in three dimensions.

Let g be a (semi-)Riemannian metric on M.

grad: The map $g^\# : \mathcal{X}(M) \to \Omega^1(M)$ identifies vector fields with one-forms. We define the gradient of a function f to be the vector field corresponding to the one-form df. That is, $g^\#(\operatorname{grad} f) = df$,[32] or explicitly, for $p \in M$,

$$\langle \operatorname{grad} f(p), w\rangle = df_{|p}(w) \quad \text{for all } w \in T_pM \tag{21}$$

div: On the other hand, a vector field can also be identified with an $(n-1)$-form, by first applying $g^\#$ and then $*$. A function, i.e. 0-form, can be identified with an n-form using $*$. Explicitly, the function f corresponds to the n-form $f\,\mathrm{dvol}$. Then the divergence of a vector field is defined by first transforming the vector field to an $(n-1)$-form, applying d, then transforming the resulting n-form to a function. That is

$$\operatorname{div} X = *^{-1} d(*g^\#(X)) \text{ or equivalently } \;(\operatorname{div} X)\,\mathrm{dvol} = d(*g^\#(X))$$

rot: If $n = 3$ then $n-1 = 2$, so using the identifications above we can translate $d : \Omega^1(M) \to \Omega^2(M)$ into a map $\mathcal{X}(M) \to \mathcal{X}(M)$. This is the 'rotation' rot.[33]

It is easiest to understand and remember this if we put it all in a diagram:

$$\begin{array}{ccccccc} \Omega^0(M) & \xrightarrow{d} & \Omega^1(M) & \overset{d}{\dashrightarrow} & \Omega^{n-1}(M) & \xrightarrow{d} & \Omega^n(M) \\ = \uparrow & & g^\# \uparrow & & *\circ g^\# \uparrow & & * \uparrow \\ C^\infty(M) & \xrightarrow{\text{grad}} & \mathcal{X}(M) & \overset{\text{rot}}{\dashrightarrow} & \mathcal{X}(M) & \xrightarrow{\text{div}} & C^\infty(M) \end{array} \tag{22}$$

(the dashed arrows only make sense if $n = 3$). The identity $d^2 = 0$ then translates into

$$\operatorname{div}\operatorname{rot} = 0, \operatorname{rot}\operatorname{grad} = 0 \quad (n = 3)$$

There are two useful identities for the divergence. First[34]

$$(\operatorname{div} X)\,\mathrm{dvol} = d(\iota_X \mathrm{dvol}) \tag{23}$$

The **geometric meaning of the divergence** is 'volume change under the flow'

$$L_X \mathrm{dvol} = (\operatorname{div} X)\,\mathrm{dvol}$$

(proof: use iddi-formula and (23)). The meaning of this may become clearer after integration over any open set U[35]

$$\frac{d}{dt}_{|t=0} \operatorname{vol} \Phi_t(U) = \int_U \operatorname{div} X\,\mathrm{dvol}$$

Then (17) says in this context

$$\operatorname{div} X = 0 \iff \text{the flow of } X \text{ preserves volume}$$

i.e. $\operatorname{vol} \Phi_t(U) = \operatorname{vol} U \;\forall t \forall U$.

The **geometric meaning of the gradient** is (for $df_{|p} \neq 0$):

- $\operatorname{grad} f(p)$ points in the direction of steepest increase of f
- $|\operatorname{grad} f(p)|$ is the rate of that increase

This follows easily from (21).

[32]Sometimes we write $\nabla f = \operatorname{grad} f$.
[33]Sometimes this is called curl.
[34]Use (10).
[35]Use $\int_U \Phi_t^*(\mathrm{dvol}) = \int_{\Phi_t(U)} \mathrm{dvol} = \operatorname{vol} \Phi_t(U)$.

Local coordinate formulas for $\operatorname{grad}, \operatorname{div}$

Since $g^\flat$ is pulling up indices, we have

$$\operatorname{grad} f = \sum (\operatorname{grad} f)^i \partial_i \quad \text{with} \quad (\operatorname{grad} f)^i = \sum g^{ij} \frac{\partial f}{\partial x^j} \tag{24}$$

Also, using (23) and (12) one gets

$$\operatorname{div} X = \frac{1}{\sqrt{\det(g_{jk})}} \sum_i \frac{\partial \left(X^i \sqrt{\det(g_{jk})} \right)}{\partial x^i} \quad \text{for } X = \sum X^i \partial_i \tag{25}$$