Chapter 2

Exterior Calculus

Throughout these notes we assume that the reader is familiar with basic multivariate calculus as well as linear algebra. So what about this **Exterior Calculus**? Why do we need this special flavor of calculus? The main motivation for us is that we want a concise notational apparatus to talk about differential and integral calculus on manifolds. Exterior calculus was designed to deliver on this requirement. For example, vector calculus operators such as ∇, $\nabla\times$, $\nabla\cdot$ all coalesce into a single derivative d. Similarly, important theorems such as the Fundamental Theorem of Calculus, or Gauss' and Green's theorem all become instances of one single general statement, the so called Stokes' theorem

$$\int_M d\,\omega = \int_{\partial M} \omega,$$

i.e., the integral of the differential of ω over the domain M is the same as the integral of ω over the boundary of the domain. Another important feature for our purposes is that we will formulate everything as much as possible in a coordinate free manner. This exposes the underlying structures far better than piles of indices can ever hope to do. Sometimes we need to go to coordinates to actually calculate something or to prove a basic fact before using it, of course. In any case, life will be far easier if we first abstract the usual multivariate calculus into **exterior calculus**.

2.0.1 The Differential as 1-form

Let's start with something we are all very familiar with, the **differential of a function**, and try to really understand the underlying structure to see what the proper generalization is. For now all computations will be done on $\mathbb{R}^m$ to develop the basic apparatus and help you relate everything to things you are familiar with. Then we'll go to the manifold case where all we will have to do is point out how the basic definitions are modified so that everything formally pulls through.

Consider a function $f : \mathbb{R}^m \to \mathbb{R}$ for some $m > 0$. If you like, think of cases such as $m = 1, 2, 3$ to help your intuition here . What is the differential df of f? Consider the Taylor expansion of f around a point p in its domain as a function of a vector X and a step size $h > 0$ along X

$$f(p + Xh) = f(p) + d_p f(X)h + o(h^2).$$

Here we see that for each point in the domain, $p \in \mathbb{R}^m$, $d_p f$ is a linear operator which best approximates the function f at that point. This linear operator maps vectors from the domain to "amounts of change" in the range, *i.e.*, it maps the tangent space of the domain to the tangent space of the range, $d_p f : T_p\mathbb{R}^m \to T_{f(p)}\mathbb{R}$. In the case of a real valued f, the linear map $d_p f$ returns a number ($T_{f(p)}\mathbb{R}$ is isomorphic to $\mathbb{R}$ itself). Such a linear map from a vector space, here the tangent vector space, to the reals is a **linear functional**, also sometimes called a **co-vector** or **co-tangent**. The term which we will use, as it will generalize to higher orders, is **1-form**.

What do we actually measure when we stick a tangent vector into the differential? It measures the change of the function along that vector, in other words, we get the **directional derivative**. By the Riesz representation theorem such a linear functional can always (well, subject to some technical conditions we will not worry about here) be written as the inner product against a suitably chosen representation vector. In fact you are already familiar with this. For the differential of a function it's called the gradient $d_p f(v) = \langle \mathrm{grad}_p f, v\rangle$. This relationship between the differential and its representation vector via the inner product is encoded by what are known as the **musical isomorphisms**, $(df)^\sharp = \mathrm{grad} f$ resp. $(\mathrm{grad} f)^\flat = df$.

Definition 2.1. Whenever we have a vector space V with a bi-linear form $\langle ., .\rangle : V \times V \to \mathbb{R}$ we get the **dual space** V^* of V which consists of all **linear functionals**, *i.e.*, linear maps from V to the reals $\mathbb{R}$, $v \mapsto \langle v, .\rangle$. The most typical example, and the one most relevant for us, would be the inner product. (Later on, when we are on manifolds the bi-linear form of interest will be the **metric**.) The action of taking a vector $v \in V$ and turning it into a linear functional, *i.e.*, a member of the dual space V^*, by sticking it into one of the argument slots of $\langle ., .\rangle$, is one of the **musical isomorphisms** so long as the bi-linear form is non-degenerate. It is denoted $\flat : V \to V^*$ and defined by

$$v^\flat(w) := \langle v, w\rangle,$$

for $w \in V$. The inverse musical isomorphism is then named, you guessed it, $\sharp : V^* \to V$. For $\alpha \in V^*$ it is defined by

$$\langle \alpha^\sharp, v\rangle := \alpha(v).$$

So far this definition is for a single space V, but it carries over just fine if the space V varies with some other quantity. The prime example of course are tangent spaces T_pM to some manifold which vary with the base point $p \in M$. This version with a varying base point is what we used above when we wrote $(df)^\sharp = \operatorname{grad} f$ resp. $(\operatorname{grad} f)^\flat = df$, dropping the explicit mention of the dependence on p.

For those readers familiar with raising and lowering indices the musical isomorphisms are easy to remember. $\sharp$ raises indices turning a 1-form into a vector, while $\flat$ lowers them, turning a vector into a 1-form.

Since $V = T_pM$ as well as $V^* = T_p^*M$ are both linear spaces it makes sense to talk about bases for these. For simplicity, let's assume $V = \mathbb{R}^m$ and $V^* = (\mathbb{R}^m)^*$. In principle such a basis could change arbitrarily as the base point p varies, but it makes sense to use a basis which changes smoothly with p. Furthermore computations will be simplified if we allow ourselves the luxury of using a positively oriented orthonormal basis (ONB). Luckily such a basis is always available locally by using orthogonal coordinates . An ONB for T_pM is then given by $\{e_i\}_{i=1\ldots m}$, the unit normal vectors in the Cartesian coordinate directions. For the dual space we will choose a basis $\{dx^i\}_{i=1\ldots m}$ of differentials of the coordinate functions. Because of orthogonality of coordinate directions, the i^{th} directional derivative of the j^{th} coordinate function is given by

$$d_p\, x^j(e_i) = \lim_{t\to 0} \frac{x^j(p+te_i) - x^j(p)}{t} = \frac{\partial x^j}{\partial x^i} = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases} \quad .$$

In other words $dx^i = e_i^\flat$ and $\{dx^i\}_{i=1\ldots m}$ is an ONB for T_p^*M. This establishes **bi-orthogonal** bases for the **primal** and **dual** spaces respectively. Note that even though they vary (smoothly) with p their orthonormality and bi-orthogonality is maintained.

Let's use this basis for df (dropping once again the explicit mention of the dependence on p)

$$df = \sum_i \frac{\partial f}{\partial x^i}\, dx^i \qquad \text{since} \qquad df(e_j) = \langle \operatorname{grad} f, e_j \rangle = \frac{\partial f}{\partial x^j}.$$

Instead of writing $\{e_i\}_{i=1\ldots m}$ for the basis of tangent vectors you will also often find authors using partial derivative symbols $\{\partial_i\}_{i=1\ldots m}$ or $\{\frac{\partial}{\partial x^i}\}_{i=1\ldots m}$. Why? Letting $v = \sum_i v^i \partial_i$ we can write $df(v) = \sum_i v^i \partial_i f = \langle \operatorname{grad} f, v \rangle$. This notation is nice when we later treat vector fields as differentiation operators $v(f) := df(v)$, calling v a **derivation**.

Remark 2.1. So what does this all look like when we are dealing with an m-manifold rather than $\mathbb{R}^m$? Basically an m-manifold looks locally like $\mathbb{R}^m$. What does that mean? For every point $p \in M$ there is an open neighborhood where the manifold can be parameterized by a smooth coordinate system. Think of a surface as an example of an M. A parameterization $\varphi : \mathbb{R}^2 \supset U \to M$ around p is a function from some region in $\mathbb{R}^2$ to a neighborhood of p on the surface. We can visualize this by mapping 2-dimensional "graph" paper to the surface in the vicinity of p . A function on the manifold, $f : M \to \mathbb{R}$ can be composed with this parameterization to yield a function $f \circ \varphi$ from U to $\mathbb{R}$. For such functions we know what differentiation means and we can write

$$df(v) = \sum_{i=1}^{m} \frac{\partial(f \circ \varphi)}{\partial x^i} v^i$$

for $v = \sum_{i=1}^{m} v^i \frac{\partial}{\partial x^i}$. Here the x^i are the cartesian coordinates in the domain of φ. At first sight the expressions appear to depend on the parameterization chosen. This is certainly true for $\frac{\partial}{\partial x^i}$ and v^i individually, but their combination into v is independent of the parameterization. The fact that the ultimate expressions are independent of the parameterization, though the coefficients do depend on the parameterization, *i.e.*, the basis chosen, is true also for the expression for df. Let $\varphi(x^1, \ldots, x^m) = (y^1, \ldots, y^m) \in M$ and write

$$df = \sum_{i=1}^{m} \frac{\partial(f \circ \varphi)}{\partial x^i} dx^i = \sum_{i,j=1}^{m} \frac{\partial f}{\partial y^j} \frac{\partial y^j}{\partial x^i} dx^i = \sum_{j=1}^{m} \frac{\partial f}{\partial y^j} dy^j.$$

Since x and y were just formal names we will from now on write $df = \sum_i \frac{\partial f}{\partial x^i} dx^i$ no matter whether we are working with f defined on $\mathbb{R}^m$ or on some m-manifold M.

2.0.2 What are Forms good for?

So now we have this fancy way of talking about the meaning of df. But what is it good for? One of the fundamental uses for forms is in integration which corresponds to taking a physical measurement (we'll see plenty of examples soon). Consider $\mathbb{R}^m$ with a function f defined on it and a curve $\gamma\colon [0,1] \to \mathbb{R}^m$ with $\gamma(0) = p$ and $\gamma(1) = q$ then

$$\int_\gamma df = \int_0^1 df(\dot\gamma) = f(q) - f(p).$$

To get some intuition as to what is going on, think of this integral for the moment as an infinite Riemann sum over infinitesimal steps along $[0,1]$. $\dot\gamma$ is a tangent vector to γ and its length is proportional to the speed with which γ is traversed as we take little "Riemann steps." Sticking it into df tells us the amount of change of f along this direction weighted by the "step size" induced by $\dot\gamma$. In particular the result is independent of the particular (smooth) parameterization of γ used in the computation and hence $\int_\gamma df$ is well defined. For this particular example the final result is of course the content of the fundamental theorem of calculus. More interesting is the case

$$\int_\gamma \omega$$

for *some* 1-form ω, not necessarily the differential of a function. (Can you think of an example?) Physically, for example, $\omega(X)$ might represent the intensity of a field along the vector X and we want to know the aggregate circulation along γ. Incidentally the result of this integration being independent of the parameterization of γ is essential if this computation is to represent a physical measurement.

It is these features that we now want to generalize. So far we know how to do integrals along a 1-dimensional domain (γ in the example above). What about things such as the flux across a surface? In that case we need an object which can be integrated along a surface, something that acts like the dA often seen underneath a 2-dimensional integral fence. Or perhaps the dV sometimes used to denote the local volume element to integrate some density of "stuff" over a 3-dimensional domain. In essence the object underneath the integral fence takes an infinitesimal parallelipiped and asks for its volume modulated by whatever function we are integrating.

What we need then to go to higher dimensions are objects which act like k-dimensional volume measures. The tricky part is that we have to deal with k-dimensional volumes in m dimensions, so writing down a $k \times k$ determinant, with (maybe) some spatially varying factors, won't be totally obvious. What we will discover though is that the most general of such measurement functions can be written as a (varying) linear combination of $k \times k$ minors of the "surrounding" $m \times m$ determinant.

For starters though, to understand what the right abstraction is, we make some observations about the usual area and volume measures. Consider the area $A(v_1, v_2)$ of a parallelogram spanned by two vectors $v_1, v_2 \in \mathbb{R}^n$. The function A is bi-linear and anti-symmetric, $A(v_1, v_2) = -A(v_2, v_1)$. Now consider the volume $V(v_1, v_2, v_3)$ of the parallelipiped spanned by three vectors $v_1, v_2, v_3 \in \mathbb{R}^n$. The function V is linear in each entry and once again anti-symmetric in its arguments. More precisely, if $\sigma\colon \{1,2,3\} \to \{1,2,3\}$ is a permutation then $V(v_{\sigma(1)}, v_{\sigma(2)}, v_{\sigma(3)}) = \mathrm{sgn}(\sigma)V(v_1, v_2, v_3)$. Here sgn returns ± 1 if the permutation is even resp. odd, *i.e.*, it takes an even or odd number of pairwise exchanges to map σ to sorted order. So, such alternating multi-linear forms then is what we are after.

2.0.3 The Exterior Algebra

The pattern of volume measures of different dimension is indeed general and defines the notion of a **k-form**. We begin with a fixed vector space V which will later correspond to a tangent space T_pM. Eventually we'll let the base point $p \in M$ vary smoothly and talk about integration over some domain again.

A k-form $\omega^k \in \bigwedge^k(V;\mathbb{R})^*$ is an alternating k-linear mapping, $\omega^k\colon V^k \to \mathbb{R}$ for V some m-dimensional vector space. (The range space can be something more general such as $\mathbb{R}^n$ or the complex numbers $\mathbb{C}$, but we'll stick with $\mathbb{R}$ for now.) Being k-linear just means that ω^k is linear in each slot. Alternating means

$$\omega^k(v_{\sigma(1)},\ldots,v_{\sigma(k)}) = \operatorname{sgn}(\sigma)\omega^k(v_1,\ldots,v_k)$$

where $\sigma\colon \{1,\ldots,k\} \to \{1,\ldots,k\}$ denotes a permutation of the integers 1 through k and $\operatorname{sgn}(\sigma) = \pm 1$ if the permutation is even resp. odd. If $\dim(V) = m$ then $0 \leq k \leq m$ is the valid range for k. The case of 0-forms, objects that take zero vectors, are just the reals.

Exercise 2.1. Show that for $k > m$, $\omega^k = 0$.

So how then do we construct k-forms? For this we will use the **wedge product**. The simplest case is given by the wedge product of two 1-forms $\alpha, \beta \in \bigwedge^1(V;\mathbb{R})^*$ which is defined as

$$\alpha \wedge \beta(v,w) := \alpha(v)\beta(w) - \alpha(w)\beta(v) = \det\begin{pmatrix} \alpha(v) & \beta(v) \\ \alpha(w) & \beta(w) \end{pmatrix}$$

Obviously $\alpha \wedge \beta$ is a 2-form since it is multi-linear (bi-linear) and alternating. These properties follow directly from its definition as a determinant. More generally we have a k-fold wedge product to produce a k-form

$$\alpha_1 \wedge \ldots \wedge \alpha_k(v_1,\ldots,v_k) = \det(\alpha_j(v_i)).$$

Remark 2.2. Notice one small subtlety here. We wrote $\det(\alpha_j(v_i))$ rather than $\det(\alpha_i(v_j))$ as many authors do. For real valued forms this makes no difference, but we will later see cases where the forms take values in an algebra which is not commutative. For example, the forms may be $\mathbb{R}^3$ vector valued and multiplication be given by the cross product. In that case index order matters and we follow the usual convention that for matrices the index i denotes rows while j denotes columns.

Ok. So this gets us *some* k-forms, but does it get all of them? (No.) Since $\bigwedge^k(V;\mathbb{R})^*$ is itself a linear vector space with the obvious addition and scalar multiplication of k-forms we can talk about a basis for it. For 1-forms we have already seen this. Let $\{e_i\}_{i=1\ldots m}$ be an ONB for V and $\{e_i^\flat\}_{i=1\ldots m}$ the corresponding basis for 1-forms. How then do we construct bases for k-forms with $k > 1$? You guessed it: we'll build them up through wedge products of the basis for 1-forms. To figure out what these bases look like observe that a multi-linear transformation, just as a linear transformation, is fully determined if we know its action on the basis vectors. For example, consider $k = 2$, *i.e.*, a bi-linear, alternating form on $\mathbb{R}^m$. For $a = \sum_{i=1}^m a^i e_i$, resp. $b = \sum_{i=1}^m b^i e_i$ we have

$$\omega(a,b) = \sum_{i,j=1}^m a^i b^j \omega(e_i,e_j) = \sum_{1\leq i<j\leq m} (a^i b^j - a^j b^i)\omega(e_i,e_j)$$

which demonstrates that knowledge of $\omega(e_i, e_j)$ for $1 \le i < j \le m$ is sufficient to evaluate ω on any pair of vectors from $\mathbb{R}^m$. A similar, though more messy expression holds for a k-form and we'll need some additional notation to manage this.

Let $I = \{i_1, \ldots, i_k\}$ be a sorted index set $1 \le i_1 < \ldots < i_k \le m$ with $|I| = k$ elements. With this e_I will denote the correspondingly indexed set of basis vectors. Now consider some $\omega \in \bigwedge^k(V;\mathbb{R})^*$ and record its value on the basis vectors $\{e_i\}_{i=1\ldots m}$. To this end define the $\binom{m}{k}$ coefficients

$$w_I := \omega(e_I),$$

for all sorted k-element subsets of $\{1 \ldots m\}$. These are sufficient to completely characterize ω since a repeated index yields no additional information (why?), nor does an unsorted index set yield any additional information (why?). In other words, instead of the m^k different choices for bases in the k slots we really only need $\binom{m}{k}$ many. The claim now is that

$$\omega = \sum_I w_I e_I^\flat,$$

where $e_I^\flat = e_{i_1}^\flat \wedge \ldots \wedge e_{i_k}^\flat$ denotes the k-fold wedge product of the correspondingly indexed elements of $\{e_i^\flat\}_{i=1\ldots m}$. That ω can be written in this way follows from e_I and $e_I^\flat$ being dual to one another ($e_I^\flat(e_{\hat{I}}) = \delta_{I\hat{I}}$) and multi-linearity. With this we conclude that a basis for $\bigwedge^k(V;\mathbb{R})^*$ is given by $\{e_I^\flat\}_{I\subset\{1\ldots m\},|I|=k}$.

Definition 2.2. Let $\alpha = \sum_I a_I e_I^\flat$ and $\beta = \sum_J b_J e_J^\flat$, where $I = (i_1, \ldots, i_k), i_1 < \ldots < i_k$ and $J = (j_1, \ldots, j_l), j_1 < \ldots < j_l$. Define

$$\alpha \wedge \beta = \sum_{IJ} a_I b_J e_I^\flat \wedge e_J^\flat.$$

Note that this definition agrees with our previous definition $\alpha_1 \wedge \ldots \wedge \alpha_k(v_1, \ldots, v_k) = \det(\alpha_j(v_i))$ when $\alpha_1, \ldots, \alpha_k$ are 1-forms.

Exercise 2.2. Let $V = \mathbb{R}^4$ and define the 2-form $\alpha = u_{12} e_1^\flat \wedge e_2^\flat + u_{24} e_2^\flat \wedge e_4^\flat + u_{34} e_3^\flat \wedge e_4^\flat$ and the 1-form $\beta = w_2 e_2^\flat + w_3 e_3^\flat$. Compute $\alpha \wedge \beta$ and $\alpha \wedge \alpha$.

Exercise 2.3. Prove that $e_I^\flat(e_{\hat{I}}) = \delta_{I\hat{I}}$, *i.e.*, takes on the value 1 when $I = \hat{I}$ and 0 otherwise.

Exercise 2.4. Let α, β, and γ be k-, l-, and r-forms respectively. Show that the wedge product is associative, $(\alpha \wedge \beta) \wedge \gamma = \alpha \wedge (\beta \wedge \gamma)$, distributive over addition (for $l = r$), $\alpha \wedge (\beta + \gamma) = \alpha \wedge \beta + \alpha \wedge \gamma$, and anti-commutative, $\alpha \wedge \beta = (-1)^{kl} \beta \wedge \alpha$.

2.0.4 Inner Product on k-forms

So we constructed a basis for $\bigwedge^k(V;\mathbb{R})^*$ starting with an ONB $\{e_i\}_{i=1\dots m}$ for V. Are the bases for k-forms also ONB? We haven't talked about the inner product for even just 1-forms. We can certainly *declare* these bases to be ONB. By linear extension this then would define an inner product for the underlying spaces. Alternatively we can define an inner product more directly and then show that the bases we constructed starting with an ONB for V are all orthonormal. We'll take that path. To start with we define the inner product on 1-forms via the musical isomorphism.

Definition 2.3. For $\alpha, \beta \in \bigwedge^1(V;\mathbb{R})^*$ we define

$$\langle \alpha, \beta \rangle := \langle \alpha^\sharp, \beta^\sharp \rangle.$$

This obviously makes $\{e_i^\flat\}_{i=1\dots m}$ into an ONB for $\bigwedge^1(V;\mathbb{R})^*$.

With this in hand we can immediately define the inner product for k-forms which are k-fold wedge products.

Definition 2.4. For $\alpha_i, \beta_i \in \bigwedge^1(V;\mathbb{R})^*$, $i = 1\dots k$ we define

$$\langle \alpha_1 \wedge \dots \wedge \alpha_k, \beta_1 \wedge \dots \wedge \beta_k \rangle = \det(\langle \alpha_j, \beta_i \rangle).$$

By linearity this then extends to an inner product to all of $\bigwedge^k(V;\mathbb{R})^*$ and once again the bases we constructed above for k-forms turn out to be ONB.

So what does that look like in practice? Let's try an example. Let $V = \mathbb{R}^4$ and consider 2-forms. That makes $\dim(\bigwedge^2(\mathbb{R}^4;\mathbb{R})^*) = 6$. Let $\{\omega^i\}_{1=1\dots 6}$ be some basis, not necessarily orthogonal, for this space. Any 2-forms α, β are thus given by coefficient vectors $(\alpha_i)_{i=1\dots 6}$ and $(\beta_i)_{i=1\dots 6}$ in $\mathbb{R}^6$ and the inner product is given by

$$\langle \alpha, \beta \rangle = \sum_{i,j\in\{1\dots 6\}} \alpha_i \beta_j \langle \omega^i, \omega^j \rangle$$

Notice that the symmetric tableaux $\langle \omega^i, \omega^j \rangle_{i,j\in\{1\dots 6\}}$ needs to be computed only once when the basis is chosen.

Exercise 2.5. Show that for $\{e_i\}_{i=1\dots m}$ an ONB for V and $\alpha, \beta \in \bigwedge^k(V;\mathbb{R})^*$ we have

$$\langle \alpha, \beta \rangle = \sum_I \alpha(e_I)\beta(e_I)$$

where I ranges over all sorted k-element subsets of $\{1 \dots m\}$. The version for 1-forms is often very handy in computations

These inner products are defined for V a single vector space. As we go to differential forms V will become a tangent space T_pM to a manifold M at a point p and coefficients of forms will become functions (of $p \in M$). The inner products we just constructed then carry over naturally as functions of p and may be used, for example, to define integral norms in the usual way. Since nothing really changes except for the dependency on p we won't say anything more about the extension of these inner products to differential k-forms.

Ok. What are these differential k-forms and what can you do with them?

2.0.5 Differential k-forms

Back to integration. We have seen that k-forms have the right structure for integration, but now the vector space V becomes the tangent space T_pM with the base point $p \in M$ varying over some region. Consequently we want all the machinery we worked out pointwise to vary smoothly as well. Let $p \in M$ be inside an open neighborhood $p \in U_p \subset M$ and let this neighborhood be parameterized by a system of orthogonal coordinate functions $\{x^i\}_{i=1\dots m}$. The elementary coordinate differentials $\{dx^i\}_{i=1\dots m}$ then yield a smoothly (in p) varying ONB for the space of 1-forms which is dual to the ONB $\{e_i\}_{i=1\dots m}$ of unit vectors along the cartesian coordinate directions. What was V above is now T_pM and everything gets a dependence on the point $p \in M$. We thus consider **differential k-forms, $\Omega^k(M;\mathbb{R})$**, whose coefficients with respect to a basis are now differentiable functions $w_I : M \to \mathbb{R}$

$$\omega = \sum_I w_I \, dx^I$$

where the dependence of w_I and dx^I on p is understood.

Now we are ready for the differential as it applies to k-forms, $d\colon \Omega^k(M;\mathbb{R}) \to \Omega^{k+1}(M;\mathbb{R})$. This **exterior derivative** d acts on functions f (0-forms) just as the usual differential df, which we already saw. It also satisfies $d(df) = 0$ and more generally $d \circ d = 0$. Finally it satisfies a product rule with respect to the wedge product

$$d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^k \alpha \wedge d\beta$$

where $\alpha \in \Omega^k(M;\mathbb{R})$ and $\beta \in \Omega^l(M;\mathbb{R})$.

Exercise 2.6. Instead of stating the properties that define the exterior derivative we could also give a working definition. Let $\omega = \sum_I w_I \, dx^I$ and define d by

$$d\omega = \sum_I d(w_I) \wedge dx^I$$

and d applied on a 0-form f is defined by

$$df = \sum_{i=1}^{n} \frac{\partial f}{\partial x^i} dx^i.$$

Now one needs to show that this amounts to the usual differential for functions, is a linear operator, satisfies $d^2 = 0$, and the product rule. Show that this is true. To show $d^2 = 0$ first show that this is true for functions using the fact that for multiple partial derivatives their order does not matter.

A very important property that differential forms have, which is key to the parameterization independence of their integrals, is their behavior under **pullback**.

Definition 2.5. Let $\varphi : N \to M$ be a smooth map between an n-manifold and an m-manifold. Once again, if you like, think $\mathbb{R}^n$ and $\mathbb{R}^m$. For a k-form $\omega \in \Omega^k(M;\mathbb{R})$, $k \geq 1$ define the pullback

$$(\varphi^*\omega)_q(v_1, \ldots, v_k) := \omega_{p=\varphi(q)}(d\varphi(v_1), \ldots, d\varphi(v_k)).$$

For a 0-form f we simply have composition $\varphi^* f = f \circ \varphi$.

Pullback, as the name suggests, allows us to take a k-form on M, thought of as the range of some φ, and bring it "back" to N, the domain of φ. What makes pullback useful is that it distributes over the wedge product

$$\varphi^*(\alpha \wedge \beta) = (\varphi^*\alpha) \wedge (\varphi^*\beta).$$

Even more useful is the fact that pullback commutes with the exterior derivative

$$d(\varphi^*\alpha) = \varphi^*(d\alpha).$$

Exercise 2.7. Prove the above two assertions and use them to show that for $\omega = w\, dy^I$ and $\varphi : N \to M$, $\varphi(x_1, \ldots, x_n) = (y_1, \ldots, y_m)$, *i.e.*, a map from an n-manifold to an m-manifold we get

$$\varphi^*\omega = (w \circ \varphi)\, d\varphi^I.$$

Finally, let's talk about integration of k-forms. Let $\omega \in \Omega^k(M;\mathbb{R})$ and M^k a k-dimensional submanifold of M. To make life simpler here, suppose further that M^k is the image of the k-dimensional unit cube under the parameterization $\varphi(x_1, \ldots, x_k) = (y_1, \ldots, y^k)$, *i.e.*, $M^k = \varphi([0,1]^k)$ and $\omega = dy^1 \wedge \ldots \wedge dy^k$ (why is this enough?). Then

$$\int_{M^k} \omega = \int_{[0,1]^k} \varphi^*\omega = \int_{[0,1]^k} \det(d\varphi)\, dx^1 \wedge \ldots \wedge dx^k = \int_0^1 \cdots \int_0^1 \det(d\varphi)\, dx^1 \ldots dx^k$$

where the last integral is the usual iterated multi-dimensional integral from standard calculus. This all follows from the above assertions. In particular

$$\int_{M^k} \omega$$

is parameterization independent and hence well defined so long as any re-parameterization uses only smooth, regular functions. This follows from the change of variable formula for integrals.

2.0.6 Hodge Duality

Whenever we have an inner product (bi-linear form) we get a duality relation. This is no different with the inner product on differential k-forms. This duality carries the name of **Hodge star** and maps k-forms to $m-k$-forms, $*: \Omega^k(M;\mathbb{R}) \to \Omega^{m-k}(M;\mathbb{R})$. For $\alpha, \beta \in \Omega^k(M;\mathbb{R})$ it is defined by

$$\alpha \wedge *\beta := \langle \alpha, \beta \rangle \mu$$

where $\mu \in \Omega^m(M;\mathbb{R})$ is the volume form for M. Since $\Omega^m(M;\mathbb{R})$ is one dimensional, μ is unique up to a scale factor, giving rise to a Hodge star which is unique only up to scale. Of course the natural choice is the normalization which gives the ONB of T_pM unit volume. Ok. But what does this mean in practice? If you actually have to work it out, how do you do it? The easiest direct calculation follows from working it out for an ONB and then have it work for any k-form by linear extension. Assuming that the inner product on T_pM has an all positive signature (the standard case) we get

$$dx^I \wedge *dx^I = dx^I \wedge dx^J = \operatorname{sgn}(\sigma)\, dx^1 \wedge \ldots \wedge dx^m$$

where $J = \{1 \ldots m\} \setminus I$ and σ is the permutation of $I \cup J$. For example, let $M = \mathbb{R}^4$ with the usual inner product then

$$*(dx^1 \wedge dx^2) = dx^3 \wedge dx^4 \qquad *(dx^1 \wedge dx^3) = -dx^2 \wedge dx^4$$
$$*dx^2 = -dx^1 \wedge dx^3 \wedge dx^4 \qquad *(dx^1 \wedge dx^3 \wedge dx^4) = dx^2.$$

Exercise 2.8. Prove that $** \omega^k = (-1)^{k(m-k)} \omega^k$ and thus $*^{-1} = (-1)^{k(m-k)} *$. Conclude that for odd m no minus sign appears and for even m a minus sign appears for odd k.

Exercise 2.9. The Hodge star depends on the inner product and so far we have only seen Euclidean inner products. An important non-trivial example of a different inner product arises from space-time and its treatment as a Minkowski space. We will see it in action when talking about electromagnetism. We say that this inner product has signature $(+,+,+,-)$, *i.e.*, $|e_1|^2 = |e_2|^2 = |e_3|^2 = 1$ and $|e_4|^2 = -1$, sometimes calling the coordinates (x, y, z, t)

for the three spatial dimensions followed by the time dimension. Work out

$$\begin{array}{ccc} *dx & *dy & *dt \\ *(dx\wedge dz) & *(dt\wedge dx) & *(dy\wedge dt) \end{array}$$

and prove that $*^{-1} = (-1)^{k(m-k)}\mathrm{sgn}(\langle e^\flat_{1\ldots m}, e^\flat_{1\ldots m}\rangle)$ in general.

2.0.7 Vector Calculus

Now we have all the tools to talk about the relationship between vector calculus operators and differential forms. Let $f:\mathbb{R}^3\to\mathbb{R}$ be a 0-form and $v:\mathbb{R}^3\to\mathbb{R}^3$ a vector field, then

$$\mathrm{grad}\, f = \nabla f = (df)^\sharp \qquad \mathrm{curl}\, v = \nabla\times v = (*d(v^\flat))^\sharp \qquad \mathrm{div}\, v = \nabla\cdot v = *d*(v^\flat)$$

That's supposed to be easier? Well, no. Here we have used exterior calculus operators to express vector calculus operators. Hence we have a lot of $\flat$ and $\sharp$ and Hodge stars to go to and from vector fields. If instead we stay with k-forms all the way things become very simple. Using subscripts on d to denote the degree of the form it takes as an argument we have

$$d_0 \triangleq \mathrm{grad} \qquad d_1 \triangleq \mathrm{curl} \qquad d_2 \triangleq \mathrm{div}\,.$$

This is what we meant with the initial claim that all of vector calculus coalesces into a single notion of d. You may have faint memories of things such as $\mathrm{curl}\circ\mathrm{grad} = 0$ or $\mathrm{div}\circ\mathrm{curl} = 0$. Both are just $d^2 = 0$. Easy.

Exercise 2.10. Let $v = \sum_{i=1}^3 v^i e_i$ for $\{e_i\}_{i=1\ldots3}$ the standard ONB and verify $\mathrm{curl}\, v = (*d(v^\flat))^\sharp$ and $\mathrm{div}\, v = *d*(v^\flat)$ by unwinding each of the exterior calculus operators.

Exercise 2.11. In a later chapter we will meet the Hodge Laplacian for k-forms. Here we just write the Laplacian for functions, *i.e.*, 0-forms, and 1- resp. 2-forms using exterior calculus. Let $f:\mathbb{R}^m\to\mathbb{R}$ then

$$\Delta f = \mathrm{div}(\mathrm{grad}\, f) = *d*df.$$

This follows immediately from the definitions of div and grad in terms of exterior calculus. More interesting is the case of the Laplacian of a vector field. Let $v:\mathbb{R}^3\to\mathbb{R}^3$ be a vector field then the standard vector Laplacian reads as

$$\Delta v = (\Delta v_1, \Delta v_2, \Delta v_3)^\mathsf{T},$$

i.e., each components of $v = (v_1, v_2, v_3)$ is treated as a function. Now let $\alpha = v^\flat$ and show

that
$$(\Delta v)^\flat = (*d*d - d*d*)\alpha.$$

2.0.8 Stokes' Theorem

The central theorem and workhorse of exterior calculus is **Stokes' theorem**. Let ω be a $m-1$-form and M a compact m-manifold with boundary ∂M. Then

$$\int_M d\omega = \int_{\partial M} \omega.$$

Or in words, integrating the differential of a form over a region is the same as integrating the form itself over the boundary of the region. If we think of integration as a bi-linear pairing $\langle .,. \rangle$ between a domain and a differential form this can also be expressed as

$$\langle M, d\omega \rangle = \langle \partial M, \omega \rangle$$

In other words the topological operator ∂ is the adjoint of the differential operator d under this pairing. This way of thinking about it will become important later on when we *define* the discrete exterior derivative using Stokes' theorem and the boundary operator.

Stokes' theorem unifies (and extends to higher dimension) a number of theorems from (vector) calculus. For example, the **fundamental theorem of calculus** is the 1D version

$$\int_{[a,b]} df = \int_{\partial[a,b]} f = f(b) - f(a)$$

since $\partial[a,b] = b - a$. Another example is the **divergence theorem**, which states that the divergence of a vector field $v: \mathbb{R}^3 \to \mathbb{R}^3$ over a region $C \subset \mathbb{R}^3$ is equal to the net in-/out-flow across the boundary

$$\int_C \nabla \cdot v \, dV = \int_{\partial C} v \cdot n \, dS$$

where dV is the standard volume form, n the outward pointing normal of ∂C, and dS the area form on ∂C. For $\alpha = v^\flat$ Stokes' tells us that

$$\int_C d*\alpha = \int_{\partial C} *\alpha$$

where we used $*1 = dV$ and $*(v^\flat)(X,Y) = \det(v,X,Y)$. Going down by one order we get **Green's theorem** for 2D domains $S \subset \mathbb{R}^2$

$$\int_{\partial S} L\,dx + M\,dy = \int_S \left(\frac{\partial M}{\partial x} - \frac{\partial L}{\partial y}\right) dx \wedge dy$$

resp. the **Kelvin-Stokes theorem** for surfaces $\Sigma \subset \mathbb{R}^3$

$$\int_\Sigma \nabla \times v \cdot d\Sigma = \int_{\partial\Sigma} v \cdot d\ell$$

where $d\Sigma$ is the area form on Σ which $d\ell$ is the length element along the curve $\partial\Sigma$.

Exercise 2.12. For α a 1-form and X, Y tangent vectors prove that $*\alpha(X, Y) = \det(\alpha^\sharp, X, Y)$.

Exercise 2.13. Green's theorem tells us that the area of a region $S \subset \mathbb{R}^2$, $A = \int_S dx \wedge dy$ can be found by a contour integral around the boundary ∂S so long as we manage to find a 1-form, $\alpha = a\,dx + b\,dy$ such that $d\alpha = dx \wedge dy$. Calculate a (simple!) such α.

Proof of Stokes' Theorem

We will show Stokes' theorem for k-dimensional cubes $[0,1]^k$ only. Decomposing a domain into cubes and making sure that the various limits involved are all well defined under suitable smoothness assumptions is a tedium we will allow ourselves the luxury to skip here. So we want to show

$$\int_{[0,1]^k} d\omega = \int_{\partial[0,1]^k} \omega$$

for any $(k-1)$-form ω. Since the expression is linear in ω we restrict ourselves to $\omega = f^i\, du^1 \wedge \ldots \wedge \widehat{du^i} \wedge \ldots \wedge du^k$, *i.e.*, a pure wedge containing all basis differentials except for the i^{th}. To describe integration over the $(k-1)$-dimensional walls of the k-dimensional unit cube we define the inclusion functions $\iota_{i,\alpha}(u^1, \ldots, u^{k-1}) = (u^1, \ldots, u^{i-1}, \alpha, u^i, \ldots, u^{k-1})$ for $\alpha = \{0, 1\}$. Hence

$$\int_{\partial[0,1]^k} \omega = \sum_{j,\alpha} (-1)^{j+\alpha} \int_{[0,1]^{k-1}} \iota_{j,\alpha}^* \omega.$$

To understand this latter expression consider what ι_j^* does. It maps from the tangent space of $[0,1]^{k-1}$ to the tangent space of $[0,1]^k$ except for leaving out the j^{th} direction. For ω which has the i^{th} slot missing this implies

$$\iota_{j,\alpha}^* \omega = (f^i \circ \iota_{j,\alpha}) \det\left(\frac{\partial(u^l \circ \iota_{j,\alpha})}{\partial u^m}\right).$$

What is the value of the determinant? Only when $j = i$ can all basis 1-forms line up with their corresponding dual vector. So for $j = i$ the determinant is 1, while in all other cases it is zero. Consequently

$$\int_{\partial[0,1]^k} \omega = (-1)^{i+1} \int_{[0,1]^{k-1}} f(u^1, \ldots, u^{i-1}, 1, u^i, \ldots, u^{k-1})\, du^1 \wedge \ldots \wedge du^{k-1}$$

$$+(-1)^i \int_{[0,1]^{k-1}} f(u^1,\ldots,u^{i-1},0,u^i,\ldots,u^{k-1})\, du^1 \wedge \ldots \wedge du^{k-1}.$$

Now for the left side of Stokes' theorem. The integrand reads

$$d(f^i\, du^1 \wedge \ldots \wedge \widehat{du^i} \wedge \ldots \wedge du^k) = (-1)^{i-1} \frac{\partial f^i}{\partial u^i}\, du^1 \wedge \ldots \wedge du^k$$

leading to

$$\begin{aligned}
\int_{[0,1]^k} d\omega =& (-1)^{i-1} \int_{[0,1]^k} \frac{\partial f^i}{\partial u^i}\, du^1 \wedge \ldots \wedge du^k \\
=& (-1)^{i-1} \int_{[0,1]^{k-1}} \left(\int_0^1 \frac{\partial f^i}{\partial u^i}\, du^i \right) du^1 \wedge \ldots \wedge \widehat{du^i} \wedge \ldots \wedge du^k \\
=& (-1)^{i-1} \int_{[0,1]^{k-1}} f(u^1,\ldots,u^{i-1},1,u^{i+1},\ldots,u^k)\, du^1 \wedge \ldots \wedge \widehat{du^i} \wedge \ldots \wedge du^k \\
-& (-1)^{i-1} \int_{[0,1]^{k-1}} f(u^1,\ldots,u^{i-1},0,u^{i+1},\ldots,u^k)\, du^1 \wedge \ldots \wedge \widehat{du^i} \wedge \ldots \wedge du^k
\end{aligned}$$

where we turned only the integration along u^i into a standard integral.

Evidently this matches the right hand side calculation above.

2.0.9 The Co-Differential

Now that we have Stokes' theorem we can talk about the **co-differential**, $\delta\colon \Omega^k(M;\mathbb{R}) \to \Omega^{k-1}(M;\mathbb{R})$. It is defined as

$$\delta := (-1)^k *^{-1} d *$$

and one easily checks that it indeed goes from k- to $k-1$-forms. As the name suggests it is something of an adjoint for the exterior derivative d with respect to the inner product on forms. Let $\alpha \in \Omega^{k-1}(M;\mathbb{R})$, $\beta \in \Omega^k(M;\mathbb{R})$, and $\mu \in \Omega^m(M;\mathbb{R})$ the volume form used in the definition of the Hodge star. Then

$$(\langle d\alpha, \beta\rangle - \langle \alpha, \delta\beta\rangle)\,\mu = d(\alpha \wedge *\beta),$$

i.e., d and δ are adjoints of one another up to a differential. For example if we integrate both sides over a manifold M without boundary, $\partial M = \emptyset$, then they are indeed adjoints of one another due to Stokes' theorem

$$\langle\!\langle d\alpha, \beta \rangle\!\rangle = \langle\!\langle \alpha, \delta\beta \rangle\!\rangle$$

where the double angle brackets indicate the L^2 inner product.

Exercise 2.14. Check that the definition of δ verifies the "adjoint up to a differential" property claimed above.

2.0.10 Hodge Laplacian

Now that we have the co-differential we can define the **Hodge Laplacian** for k-forms as

$$\Delta = d\,\delta + \delta\,d\,.$$

Notice that for 0-forms f this amounts to $\Delta f = \delta\,d\,f$ since $\delta f = 0$. Similarly for m-forms μ we get $\Delta\mu = d\,\delta\mu$ since $d\,\mu = 0$.

Exercise 2.15. Check that the Hodge Laplacian Δf corresponds to the earlier definition *up to sign*. The Hodge Laplacian is positive semi-definite while the standard Laplacian of vector calculus $\Delta = \sum_{i=1}^{m} \frac{\partial^2}{\partial (x^i)^2}$ is negative semi-definite.

Exercise 2.16. Show that for a vector field $v : \mathbb{R}^3 \to \mathbb{R}^3$, $(\Delta v^\flat)^\sharp$ is the usual vector Laplacian from calculus *up to sign*. Conclude as an immediate corollary that the 2-form Hodge Laplacian in $\mathbb{R}^3$ also amounts to the standard vector Laplacian (up to sign).