NOTES ON DIFFERENTIAL FORMS. PART 1: FORMS ON $\mathbb{R}^n$

1. What is a form?

Since we're not following the development in Guillemin and Pollack, I'd better write up an alternate approach. In this approach, we're first going to define forms on $\mathbb{R}^n$ via unmotivated formulas, then prove that forms behave nicely, and only **then** go back and interpret forms as a kind of tensor with certain (anti)symmetry properties.

On $\mathbb{R}^n$, we start with the symbols $dx^1, \ldots, dx^n$, which at this point are pretty much meaningless. We define a multiplication operation on these symbols, denoted by a $\wedge$, with the condition

$$dx^i \wedge dx^j = -dx^j \wedge dx^i.$$

Of course, we also want the usual properties of multiplication to also hold. If α, β, γ are arbitrary products of dx^i's, and if c is any constant, then

$$\begin{aligned} (\alpha+\beta)\wedge\gamma &= \alpha\wedge\gamma + \beta\wedge\gamma \\ \alpha\wedge(\beta+\gamma) &= \alpha\wedge\beta + \alpha\wedge\gamma \\ (\alpha\wedge\beta)\wedge\gamma &= \alpha\wedge(\beta\wedge\gamma) \\ (c\alpha)\wedge\beta &= \alpha\wedge(c\beta) = c(\alpha\wedge\beta) \end{aligned} \tag{1}$$

Note that the anti-symmetry implies that $dx^i \wedge dx^i = 0$. Likewise, if $I = \{i_1, \ldots, i_k\}$ is a list of indices where some index gets repeated, then $dx^{i_1} \wedge \cdots \wedge dx^{i_k} = 0$, since we can swap the order of terms (while keeping track of signs) until the same index comes up twice in a row. For instance,

$$dx^1 \wedge dx^2 \wedge dx^1 = -dx^1 \wedge dx^1 \wedge dx^2 = -(dx^1 \wedge dx^1) \wedge dx^2 = 0.$$

- A *0-form* on $\mathbb{R}^n$ is just a function.
- A *1-form* is an expression of the form $\sum_i f_i(x) dx^i$, where $f_i(x)$ is a function and dx^i is one of our meaningless symbols.
- A *2-form* is an expression of the form $\sum_{i,j} f_{ij}(x) dx^i \wedge dx^j$.
- A *k-form* is an expression of the form $\sum_I f_I(x) dx^I$, where I is a subset $\{i_1, \ldots, i_k\}$ of $\{1, 2, \ldots, n\}$ and dx^I is shorthand for $dx^{i_1} \wedge \cdots \wedge dx^{i_k}$.
- If α is a k-form, we say that α has *degree* k.

For instance, on $\mathbb{R}^3$

Date: April 14, 2016.

- 0-forms are functions
- 1-forms look like $Pdx + Qdy + Rdz$, where P, Q and R are functions and we are writing dx, dy, dz for dx^1, dx^2, dx^3.
- 2-forms look like $Pdx \wedge dy + Qdx \wedge dz + Rdy \wedge dz$. Or we could just as well write $Pdx \wedge dy - Qdz \wedge dx + Rdy \wedge dz$.
- 3-forms look like $f dx \wedge dy \wedge dz$.
- There are no (nonzero) forms of degree greater than 3.

When working on $\mathbb{R}^n$, there are exactly $\binom{n}{k}$ linearly independent dx^I's of degree k, and 2^n linearly independent dx^I's in all (where we include $1 = dx^I$ when I is the empty list). If I' is a permutation of I, then $dx^{I'} = \pm dx^I$, and it's silly to include both $f_I dx^I$ and $f_{I'} dx^{I'}$ in our expansion of a k-form. Instead, one usually picks a preferred ordering of $\{i_1, \dots, i_k\}$ (typically $i_1 < i_2 < \cdots < i_k$) and restrict our sum to I's of that sort. When working with 2-forms on $\mathbb{R}^3$, we can use $dx \wedge dz$ or $dz \wedge dz$, but we don't need both.

If $\alpha = \sum \alpha_I(x) dx^I$ is a k-form and $\beta = \sum \beta_J(x) dx^J$ is an ℓ-form, then we define

$$\alpha \wedge \beta = \sum_{I,J} \alpha_I(x) \beta_J(x) dx^I \wedge dx^J.$$

Of course, if I and J intersect, then $dx^I \wedge dx^J = 0$. Since going from (I, J) to (J, I) involves $k\ell$ swaps, we have

$$dx^J \wedge dx^I = (-1)^{k\ell} dx^I \wedge dx^J,$$

and likewise $\beta \wedge \alpha = (-1)^{k\ell} \alpha \wedge \beta$. Note that the wedge product of a 0-form (aka function) with a k-form is just ordinary multiplication.

2. Derivatives of forms

If $\alpha = \sum_I \alpha_I dx^I$ is a k-form, then we define the *exterior derivative*

$$d\alpha = \sum_{I,j} \frac{\partial \alpha_I(x)}{\partial x^j} dx^j \wedge dx^I.$$

Note that j is a single index, not a multi-index. For instance, on $\mathbb{R}^2$, if $\alpha = xydx + e^x dy$, then

$$\begin{aligned} d\alpha &= ydx \wedge dx + xdy \wedge dx + e^x dx \wedge dy + 0dy \wedge dy \\ &= (e^x - x) dx \wedge dy. \end{aligned} \tag{2}$$

If f is a 0-form, then we have something even simpler:

$$df(x) = \sum \frac{\partial f(x)}{\partial x^j} dx^j,$$

which should look familiar, if only as an imprecise calculus formula. One of our goals is to make such statements precise and rigorous. Also, remember that x^i is actually a function on

$\mathbb{R}^n$. Since $\partial_j x^i = 1$ if $i = j$ and 0 otherwise, $d(x^i) = dx^i$, which suggests that our formalism isn't totally nuts.

The key properties of the exterior derivative operator d are listed in the following

Theorem 2.1. (1) *If α is a k-form and β is an ℓ-form, then*

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

(2) $d(d\alpha) = 0$. *(We abbreviate this by writing $d^2 = 0$.)*

Proof. For simplicity, we prove this for the case where $\alpha = \alpha_I dx^I$ and $\beta = \beta_J dx^J$ each have only a single term. The general case then follows from linearity.

The first property is essentially the product rule for derivatives.

$$\begin{aligned}
\alpha \wedge \beta &= \alpha_I(x)\beta_J(x) dx^I \wedge dx^J \\
d(\alpha \wedge \beta) &= \sum_j \partial_j(\alpha_I(x)\beta_J(x)) dx^j \wedge dx^I \wedge dx^J \\
&= \sum_j (\partial_j \alpha_I(x))\beta_J(x) dx^j \wedge dx^I \wedge dx^J \\
&\quad + \sum_j \alpha_I(x)\partial_j\beta_J(x) dx^j \wedge dx^I \wedge dx^J \\
&= \sum_j (\partial_j \alpha_I(x)) dx^j \wedge dx^I \wedge \beta_J(x) dx^J \\
&\quad + (-1)^k \sum_j \alpha_I(x) dx^I \wedge \partial_j \beta_J(x) dx^j \wedge dx^J \\
&= (d\alpha) \wedge \beta + (-1)^k \alpha \wedge d\beta. \qquad (3)
\end{aligned}$$

The second property for 0-forms (aka functions) is just "mixed partials are equal":

$$\begin{aligned}
d(df) &= d(\sum_i \partial_i f dx^i) \\
&= \sum_j \sum_i \partial_j \partial_i f dx^j \wedge dx^i \\
&= -\sum_{i,j} \partial_i \partial_j f dx^i \wedge dx^j \\
&= -d(df) = 0, \qquad (4)
\end{aligned}$$

where in the third line we used $\partial_j \partial_i f = \partial_i \partial_j f$ and $dx^i \wedge dx^j = -dx^j \wedge dx^i$. We then use the first property, and the (obvious) fact that $d(dx^I) = 0$, to extend this to k-forms:

$$\begin{aligned}
d(d\alpha) &= d(d\alpha_I \wedge dx^I) \\
&= (d(d\alpha_I)) \wedge dx^I - d\alpha_I \wedge d(dx^I) \\
&= 0 - 0 = 0. \qquad (5)
\end{aligned}$$

where in the second line we used the fact that $d\alpha_I$ is a 1-form, and in the third line used the fact that $d(d\alpha_I)$ is d^2 applied to a function, while $d(dx^I) = 0$. □

Exercise 1: On $\mathbb{R}^3$, there are interesting 1-forms and 2-forms associated with each vector field $v(x) = (v_1(x), v_2(x), v_3(x))$. (Here v_i is a component of the vector v, not a vector in its own right.) Let $\omega_v^1 = v_1dx + v_2dy + v_3dz$, and let $\omega_v^2 = v_1dy\wedge dz + v_2dz\wedge dx + v_3dx\wedge dy$. Let f be a function. Show that (a) $df = \omega^1_{\text{grad f}}$, (b) $d\omega_v^1 = \omega^2_{\text{curl v}}$, and (c) $d\omega_v^2 = (\text{div v})dx\wedge dy\wedge dz$, where grad, curl, and div are the usual gradient, curl, and divergence operations.

Exercise 2: A form ω is called *closed* if $d\omega = 0$, and *exact* if $\omega = d\nu$ for some other form ν. Since $d^2 = 0$, all exact forms are closed. On $\mathbb{R}^n$ it happens that all closed forms of nonzero degree are exact. (This is called the Poincare Lemma). However, on subsets of $\mathbb{R}^n$ the Poincare Lemma does not necessarily hold. On $\mathbb{R}^2$ minus the origin, show that $\omega = (xdy - ydx)/(x^2 + y^2)$ is closed. We will soon see that ω is not exact.

3. Pullbacks

Suppose that $g : X \to Y$ is a smooth map, where X is an open subset of $\mathbb{R}^n$ and Y is an open subset of $\mathbb{R}^m$, and that α is a k-form on Y. We want to define a *pullback form* $g^*\alpha$ on X. Note that, as the name implies, the pullback operation reverses the arrows! While g maps X to Y, and dg maps tangent vectors on X to tangent vectors on Y, g^* maps forms on Y to forms on X.

Theorem 3.1. *There is a unique linear map g^* taking forms on Y to forms on X such that the following properties hold:*

1. *If $f : Y \to \mathbb{R}$ is a function on Y, then $g^*f = f \circ g$.*
2. *If α and β are forms on Y, then $g^*(\alpha \wedge \beta) = (g^*\alpha) \wedge (g^*\beta)$.*
3. *If α is a form on Y, then $g^*(d\alpha) = d(g^*(\alpha))$. (Note that there are really two different d's in this equation. On the left hand side d maps k-forms on Y to $(k+1)$-forms on Y. On the right hand side, d maps k forms on X to $(k+1)$-forms on X.)*

Proof. The pullback of 0-forms is defined by the first property. However, note that on Y, the form dy^i is d of the *function* y^i (where we're using coordinates $\{y^i\}$ on Y and reserving x's for X). This means that $g^*(dy^i)(x) = d(y^i \circ g)(x) = dg^i(x)$, where $g^i(x)$ is the i-th component of $g(x)$. But that gives us our formula in general! If $\alpha = \sum_I \alpha_I(y)dy^I$, then

$$g^*\alpha(x) = \sum_I \alpha_I(g(x))dg^{i_1} \wedge dg^{i_2} \wedge \cdots \wedge dg^{i_k}. \tag{6}$$

Using the formula (6), it's easy to see that $g^*(\alpha \wedge \beta) = g^*(\alpha) \wedge g^*(\beta)$. Checking that $g^*(d\alpha) = d(g^*\alpha)$ in general is left as an exercise in definition-chasing. □

Exercise 3: Do that exercise!

An extremely important special case is where $m = n = k$. The n-form $dy^1 \wedge \cdots \wedge dy^n$ is called the *volume form* on $\mathbb{R}^n$.

Exercise 4: Let g is a smooth map from $\mathbb{R}^n$ to $\mathbb{R}^n$, and let ω be the volume form on $\mathbb{R}^n$. Show that $g^*\omega$, evaluated at a point x, is $\det(dg_x)$ times the volume form evaluated at x.

Exercise 5: An important property of pullbacks is that they are *natural.* If $g : U \to V$ and $h : V \to W$, where U, V, and W are open subsets of Euclidean spaces of various dimensions, then $h \circ g$ maps $U \to W$. Show that $(h \circ g)^* = g^* \circ h^*$.

Exercise 6: Let $U = (0, \infty) \times (0, 2\pi)$, and let V be $\mathbb{R}^2$ minus the non-negative x axis. We'll use coordinates (r, θ) for U and (x, y) for V. Let $g(r, \theta) = (r\cos(\theta), r\sin(\theta))$, and let $h = g^{-1}$. On V, let $\alpha = e^{-(x^2+y^2)}dx \wedge dy$.
(a) Compute $g^*(x)$, $g^*(y)$, $g^*(dx)$, $g^*(dy)$, $g^*(dx \wedge dy)$ and $g^*\alpha$ (preferably in that order).
(b) Now compute $h^*(r)$, $h^*(\theta)$, $h^*(dr)$ and $h^*(d\theta)$.

The upshot of this exercise is that pullbacks are something that you have been doing for a long time! Every time you do a change of coordinates in calculus, you're actually doing a pullback.

4. INTEGRATION

Let α be an n-form on $\mathbb{R}^n$, and suppose that α is compactly supported. (Being compactly supported is overkill, but we're assuming it to guarantee integrability and to allow manipulations like Fubini's Theorem. Later on we'll soften the assumption using partitions of unity.) Then there is only one multi-index that contributes, namely $I = \{1, 2, \ldots, n\}$, and $\alpha(x) = \alpha_I(x)dx^1 \wedge \cdots \wedge dx^n$. We define

$$\int_{\mathbb{R}^n} \alpha := \int_{\mathbb{R}^n} \alpha_I(x)|dx^1 \cdots dx^n|. \tag{7}$$

The left hand side is the integral of a form that involves wedges of dx^i's. The right hand side is an ordinary Riemann integral, in which $|dx^1 \cdots dx^n|$ is the usual volume measure (sometimes written dV or d^nx). Note that the order of the variables in the wedge product, x^1 through x^n, is implicitly using the standard orientation of $\mathbb{R}^n$. Likewise, we can define the integral of α over any open subset U of $\mathbb{R}^n$, as long as α restricted to U is compactly supported.

We have to be a little careful with the left-hand-side of (7) when $n = 0$. In this case, $\mathbb{R}^n$ is a single point (with positive orientation), and α is just a number. We take $\int \alpha$ to be that number.

Exercise 7: Suppose g is an orientation-preserving diffeomorphism from an open subset U of $\mathbb{R}^n$ to another open subset V (either or both of which may be all of $\mathbb{R}^n$). Let α be a compactly supported n-form on V. Show that

$$\int_U g^*\alpha = \int_V \alpha.$$

How would this change if g were orientation-reversing? [Hint: use the change-of-variables formula for multi-dimensional integrals. Where does the Jacobian come in?]

Now we see what's so great about differential forms! The way they transform under change-of-coordinates is perfect for defining integrals. Unfortunately, our development so far only allows us to integrate n-forms over open subsets of $\mathbb{R}^n$. More generally, we'd like to integrate k-forms over k-dimensional objects. But this requires an additional level of abstraction, where we define forms on manifolds.

Finally, we consider how to integrate something that isn't compactly supported. If α is not compactly supported, we pick a partition of unity $\{\rho_i\}$ such that each ρ_i is compactly supported, and define $\displaystyle\int \alpha = \sum \int \rho_i \alpha$. Having this sum be independent of the choice of partition-of-unity is a question of absolute convergence. If $\displaystyle\int_{\mathbb{R}^n} |\alpha_I(x)| dx^1 \cdots dx^n$ converges as a Riemann integral, then everything goes through. (The proof isn't hard, and is a good exercise in understanding the definitions.)

5. Differential forms on manifolds

An n-manifold is a (Hausdorff) space that locally looks like $\mathbb{R}^n$. We defined abstract smooth n-manifolds via structures on the coordinate charts. If $\psi : U \to X$ is a parametrization of a neighborhood of $p \in X$, where U is an open set in $\mathbb{R}^n$, then we associate functions on X near p with functions on U near $\psi^{-1}(p)$. We associate tangent vectors in X with velocities of paths in U, or with derivations of functions on U. Likewise, we associated differential forms on X that are supported in the coordinate neighborhood with differential forms on U.

All of this has to be done "mod identifications". If $\psi_{1,2} : U_{1,2} \to X$ are parametrizations of the same neighborhood of X, then p is associated with both $\psi_1^{-1}(p) \in U_1$ and $\psi_2^{-1}(p) \in U_2$. More generally, if we have an atlas of parametrizations $\psi_i : U_i \to X$, and if $g_{ij} = \psi_j^{-1} \circ \psi_i$ is the transition function from the ψ_i coordinates to the ψ_j coordinates on their overlap, then we constructed X as an abstract manifold as

$$X = \coprod U_i / \sim, \qquad x \in U_i \sim g_{ij}(x) \in U_j. \tag{8}$$

We had a similar construction for tangent vectors, and we can do the same for differential forms.

Let $\Omega^k(U)$ denote the set of k-forms on a subset $U \in \mathbb{R}^n$, and let V be a coordinate neighborhood of p in X. We define

$$\Omega^k(V) = \coprod \Omega^k(U_1) / \sim, \qquad \alpha \in \Omega^k(U_j) \sim g_{ij}^*(\alpha) \in \Omega^k(U_i). \tag{9}$$

Note the direction of the arrows. g_{ij} maps U_i to U_j, so the pullback g_{ij}^* maps forms on U_j to forms on U_i. Having defined forms on neighborhoods, we stitch things together in the usual way. A form on X is a collection of forms on the coordinate neighborhoods of X that agree on their overlaps.

Let ν denote a form on V, as represented by a form α on U_j. We then write $\alpha = \psi_j^*(\nu)$. As with the polar-cartesian exercise above, writing a form in a particular set of coordinates is technically pulling it back to the Euclidean space where those coordinates live. Note that $\psi_i = \psi_j \circ g_{ij}$, and that $\psi_i^* = g_{ij}^* \circ \psi_j^*$, since the realization of ν in U_i is (by equation (9)) the pullback, by g_{ij}, of the realization of ν in U_j.

This also tells us how to do calculus with forms on manifolds. If μ and ν are forms on X, then

- The wedge product $\mu \wedge \nu$ is the form whose realization on U_i is $\psi_i^*(\mu) \wedge \psi_i^*(\nu)$. In other words, $\psi_i^*(\mu \wedge \nu) = \psi_i^*\mu \wedge \psi_i^*\nu$.
- The exterior derivative $d\mu$ is the form whose realization on U_i is $d(\psi_i^*(\mu))$. In other words, $\psi_i^*(d\mu) = d(\psi_i^*\mu)$.

Exercise 8: Show that $\mu \wedge \nu$ and $d\mu$ are well-defined.

Now suppose that we have a map $f : X \to Y$ of manifolds and that α is a form on Y. The pullback $f^*(\alpha)$ is defined via coordinate patches. If $\phi : U \subset \mathbb{R}^n \to X$ and $\psi : V \subset \mathbb{R}^m \to Y$ are parametrizations of X and Y, then there is a map $h : U \to V$ such that $\psi(h(x)) = f(\phi(x))$. We define $f^*(\alpha)$ to be the form of X whose realization in U is $h^* \circ (\psi^*\alpha)$. In other words,

$$\phi^*(f^*\alpha) = h^*(\psi^*\alpha). \tag{10}$$

An important special case is where X is a submanifold of Y and f is the inclusion map. Then f^* is the restriction of α to X. When working with manifolds in $\mathbb{R}^N$, we often write down formulas for k-forms on $\mathbb{R}^N$, and then say "consider this form on X". E.g., one might say "consider the 1-form $xdy - ydx$ on the unit circle in $\mathbb{R}^2$". Strictly speaking, this really should be "consider the pullback to $S^1 \subset \mathbb{R}^2$ by inclusion of the 1-form $xdy - ydx$ on $\mathbb{R}^2$," but (almost) nobody is *that* pedantic!

6. Integration on oriented manifolds

Let X be an oriented k-manifold, and let ν be a k-form on X whose support is a compact subset of a single coordinate chart $V = \psi_i(U_i)$, where U_i is an open subset of $\mathbb{R}^k$. Since X is oriented, we can require that ψ_i be orientation-preserving. We then define

$$\int_X \nu = \int_{U_i} \psi_i^*\nu. \tag{11}$$

Exercise 9: Show that this definition does not depend on the choice of coordinates. That is, if $\psi_{1,2} : U_{1,2} \to V$ are two sets of coordinates for V, both orientation-preserving, that $\int_{U_1} \psi_1^*\nu = \int_{U_2} \psi_2^*\nu$.

If a form is not supported in a single coordinate chart, we pick an open cover of X consisting of coordinate neighborhoods, pick a partition-of-unity subordinate to that cover, and define

$$\int_X \nu = \sum \int_X \rho_i \nu.$$

We need a little bit of notation to specify when this makes sense. If $\alpha = \alpha_I(x) dx^1 \wedge \cdots \wedge dx^k$ is a k-form on $\mathbb{R}^k$, let $|\alpha| = |\alpha_I(x)| dx^1 \wedge \cdots \wedge dx^k$. We say that ν is *absolutely integrable* if each $|\psi_i^*(\rho_i \nu)|$ is integrable over U_i, and if the sum of those integrals converges. It's not hard to show that being absolutely integrable with respect to one set of coordinates and partition of unity implies absolute integrability with respect to arbitrary coordinates and partitions of unity. Those are the conditions under which $\int_X \nu$ unambiguously makes sense.

When X is compact and ν is smooth, absolute integrability is automatic. In practice, we rarely have to worry about integrability when doing differential topology.

The upshot is that **k-forms are meant to be integrated on k-manifolds**. Sometimes these are stand-alone abstract k-manifolds, sometimes they are k-dimensional submanifolds of larger manifolds, and sometimes they are concrete k-manifolds embedded in $\mathbb{R}^N$.

Finally, a technical point. If X is 0-dimensional, then we can't construct orientation-preserving maps from $\mathbb{R}^0$ to the connected components of X. Instead, we just take $\int_X \alpha = \sum_{x \in X} \pm\alpha(x)$, where the sign is the orientation of the point x. This follows the general principle that reversing the orientation of a manifold should flip the sign of integrals over that manifold.

Exercise 10: Let $X = S^1 \subset \mathbb{R}^2$ be the unit circle, oriented as the boundary of the unit disk. Compute $\int_X (x dy - y dx)$ by explicitly pulling this back to $\mathbb{R}$ with an orientation-preserving chart and integrating over $\mathbb{R}$. (Which is how you learned to do line integrals way back in calculus.) [Note: don't worry about using multiple charts and partitions of unity. Just use a single chart for the unit circle minus a point.]

Exercise 11: Now do the same thing one dimension up. Let $Y = S^2 \subset \mathbb{R}^3$ be the unit sphere, oriented as the boundary of the unit ball. Compute $\int_X (x dy \wedge dz + y dz \wedge dx + z dx \wedge dy)$ by explicitly pulling this back to a subset of $\mathbb{R}^2$ with an orientation-preserving chart and integrating over that subset of $\mathbb{R}^2$. As with the previous exercise, you can use a single coordinate patch that leaves out a set of measure zero, which doesn't contribute to the integral. Strictly speaking this does *not* follow the rules listed above, but I'll show you how to clean it up in class.