Math 55b: Honors Advanced Calculus and Linear Algebra

Introduction to exterior algebra and differential forms

Rudin does not explicitly develop the algebraic structure that underlies the definition of and basic operations on differential forms. These notes are meant as an introduction to this structure.

Let V be a finite-dimensional real vector space, and $E \subseteq V$ an open subset. A *differential form of order* k, or briefly a k*-form*, on E is a continuous function from E to the k-th *exterior power* $\bigwedge^k V^*$ of the dual space V^*. We begin by defining the exterior powers of an arbitrary vector space W over any field F, and describing them in the finite-dimensional case.

The exterior algebra of a vector space. The *exterior algebra* of W is the F-algebra $\bigwedge^\bullet W$ generated by W and an operation $\wedge : \bigwedge^\bullet W \times \bigwedge^\bullet W \to \bigwedge^\bullet W$ called the *exterior product*, satisfying the following relations:[1]

i) $\wedge$ is associative: $\omega_1 \wedge (\omega_2 \wedge \omega_3) = (\omega_1 \wedge \omega_2) \wedge \omega_3$ for all $\omega_1, \omega_2, \omega_3 \in \bigwedge^\bullet W$;

ii) $\wedge$ is bilinear: $\omega \wedge (a_1\omega_1 + a_2\omega_2) = a_1\omega \wedge \omega_1 + a_2\omega \wedge \omega_2$ and $(a_1\omega_1 + a_2\omega_2) \wedge \omega = a_1\omega_1 \wedge \omega + a_2\omega_2 \wedge \omega$ for all $\omega, \omega_1, \omega_2 \in \bigwedge^\bullet W$;

iii) $\wedge$ is anti-symmetric *on* W: $w \wedge w = 0$ for all $w \in W$. As usual (take $w = w_1 + w_2$) this implies $w_1 \wedge w_2 = -w_2 \wedge w_1$ for all $w_1, w_2 \in W$.

The k-th exterior power of W is then the F-vector subspace $\bigwedge^k W$ of $\bigwedge^\bullet W$ generated by elements of the form $w_1 \wedge w_2 \wedge \cdots \wedge w_k$. Here k may be any nonnegative integer, and $\bigwedge^\bullet W = \oplus_{k=0}^\infty \bigwedge^k W$. If $k = 0$, the empty exterior product is interpreted as $1 \in F$, so F is contained in $\bigwedge^\bullet W$ as $\bigwedge^0 W$.

When spelled out, this gives a rather unwieldy definition of $\bigwedge^\bullet W$ as the quotient $\mathcal{A}/\mathcal{I}$ of two rather large F-vector spaces. Here $\mathcal{A}$ is the space of formal F-linear combinations of elements of the form $w_1 \wedge w_2 \wedge \cdots \wedge w_m$, i.e. $\mathcal{A}$ has one basis element for each choice of $m = 0, 1, 2, \ldots$ and $w_1, \ldots, w_m \in W$. The subspace $\mathcal{I}$ is then the "ideal" generated by all the elements declared to be zero by conditions (i), (ii), (iii); for instance, (i) means that $\mathcal{I}$ contains all linear combinations of elements of the form $\omega \wedge \big[\omega_1 \wedge (\omega_2 \wedge \omega_3) - (\omega_1 \wedge \omega_2) \wedge \omega_3\big] \wedge \omega'$ for $\omega, \omega', \omega_1, \omega_2, \omega_3 \in \mathcal{A}$. Fortunately a much wieldier(?) description is available at least when W is finite-dimensional, which is usually the only case in which $\bigwedge^\bullet W$ is of interest.

Let $(e_1, \ldots, e_n)$, then, be a basis for W. We claim that $\wedge^k W$ has a basis consisting of the $\binom{n}{k}$ products of the form $e_{i_1} \wedge e_{i_2} \wedge \cdots \wedge e_{i_k}$, where $1 \leq i_1 < i_2 < \cdots < i_k \leq n$. We usually abbreviate this to e_I where I is this k-element subset $\{i_1, \ldots, i_k\}$ of $[1, n] \cap \mathbf{Z}$. The full exterior algebra will then be generated by the 2^n such products where I ranges over all subsets of $[1, n] \cap \mathbf{Z}$. For any two subsets I, J, the product $e_I \wedge e_J$ must then be as follows: if I, J have any elements in common, they can be brought together in $e_I \wedge e_J$ by iterating rule (iii), so the product vanishes; otherwise,

[1] Conditions (i) and (ii) alone define the *tensor algebra* $\otimes^* W$ of W; if (iii) is replaced by commutativity, we get the *symmetric algebra* $\mathrm{Sym}(W)$.

$e_I \wedge e_J = \pm e_{I\cup J}$, where the sign is even or odd according to whether the number of transpositions needed to sort the list $(i_1, \ldots, i_{|I|}, j_1, \ldots, j_{|J|})$ is even or odd. We then extend this bilinearly to determine the exterior product of any linear combinations of the e_I. To verify that this works, we need only check that our multiplication is associative, which by linearity reduces to $e_I \wedge (e_J \wedge e_K) = (e_I \wedge e_J) \wedge e_K$. The only nontrivial case is when I, J, K are pairwise disjoint, when we must show that the $\pm$ signs are consistent; but this is equivalent to a computation we have already done in proving that the sign of a permutation is well-defined.

The appearance of the sign of a permutation both here and in the definition of the determinant is no accident. Let $w_1, \ldots, w_n$ be any vectors in W, which is the same as $\bigwedge^1 W$. What is $w_1 \wedge \cdots \wedge w_n$? This product lives in $\bigwedge^n W$, a space of dimension 1 generated by $e_1 \wedge \cdots \wedge e_n$. So, $w_1 \wedge \cdots \wedge w_n = \delta\, e_1 \wedge \cdots \wedge e_n$ for some $\delta \in F$. What is δ? Let $w_j = \sum_i a_{ij} e_j$. Expand $w_1 \wedge \cdots \wedge w_n$, and note which of the resulting n^n terms survive. These turn out to be just the $n!$ terms of the determinant of the matrix formed by the a_{ij}, each with its correct sign. So, a determinant is essentially an exterior product of the columns (or rows) of a square matrix! (This can be seen more structurally by observing that $\delta(w_1, \ldots, w_n)$ satisfies the axioms the characterize the determinant.) In particular, $w_1 \wedge \cdots \wedge w_n = 0$ if and only if $w_1, \ldots, w_n$ are linearly dependent. This turns out to be true even if $w_1, \ldots, w_n$ are elements of a vector space not necessarily of dimension n. Can you prove it? (There's a big hint a few sections down.)

The algebra of differential forms on E. We now return to the case $F = \mathbf{R}$, $W = V^*$. An identification of V with $\mathbf{R}^n$ yields a natural choice of $(e_1, \ldots, e_n)$, namely e_i is the i-th coordinate function. Recall that a k-form ω on E is by definition a continuous function from E to $\bigwedge^k V^*$. Since $\bigwedge^k V^*$ is a finite-dimensional vector space, this means that the coordinates of ω are continuous functions. A simple example of a k-form is the constant function sending all points of E to the basis element e_I of $\bigwedge^k V^*$, for some k-element subset I of $[1, n] \cap \mathbf{Z}$. We call this k-form $dx_{i_1} \wedge dx_{i_2} \wedge \cdots \wedge dx_{i_k}$, and abbreviate it as before by dx_I. (These are the "basic k-forms" of Rudin, p.257. The d appearing here should be regarded for the time being as an abstract symbol, but see the next section.) The general k-form is then $\sum_I b_I dx_I$, where I ranges over all $\binom{n}{k}$ subsets and each b_I is a continuous function from E to $\mathbf{R}$.

Differential forms inherit an associative algebra structure from $\bigwedge^\bullet V^*$. We again call the resulting multiplication "exterior product" and denote it by $\wedge$, defined as follows:

$$\left(\sum_I b_I dx_I\right) \wedge \left(\sum_J c_J dx_J\right) := \sum_I \sum_J b_I c_J dx_I \wedge dx_J.$$

Here again before $dx_I \wedge dx_J$ is defined to be zero if $I \cap J \neq \emptyset$, and $\pm dx_{I\cup J}$ if I, J are disjoint, with the sign determined as before. Thus for each p, q the exterior product is a bilinear map: $\{p\text{−forms}\} \times \{q\text{−forms}\} \to \{(p+q)\text{−forms}\}$. This map is symmetric unless p, q are both odd, in which case it is alternating: if β, γ are differential forms of order p, q respectively then $\gamma \wedge \beta = (-1)^{pq} \beta \wedge \gamma$. A 0-form is just a continuous function $f : E \to \mathbf{R}$; if $p = 0$ then $f \wedge \gamma$ is just $\sum_J f c_J dx_J$. In this case we usually drop the "$\wedge$" and call the product simply $f\gamma$. Note too that this is consistent with our notations for dx_I; for instance, $b\, dx_1 \wedge dx_2$ really is the exterior product of $b\, dx_1$ with dx_2, as well as the product of the 0-form b with the 2-form $dx_1 \wedge dx_2$.

The operator d and change of variable. Recall that we first defined d as a linear map from $\mathcal{C}^1(E, \mathbf{R})$ to $\mathcal{C}(E, V^*)$ by

$$df := \sum_{i=1}^{n} D_i f \, dx_i.$$

[We use "dx_i" for the basis elements of V^* because we're thinking of V^* as $\bigwedge^1(V)$, and of elements of $\mathcal{C}(E, V^*)$ as 1-forms on E.] The 1-form df is sometimes called the "total differential" of f. As it stands, its definition depends on the choice of coordinates on V. But from the chain rule it easily follows that df is actually coordinate independent. (This could also be surmised from our formula for integrating df over a 1-surface, though of course that formula was proved using the chain rule as well...)

In fact, much more is true: d behaves well under any $\mathcal{C}^1$ map! To see this, note first that our definition of d is consistent with our earlier notation: dx_i is in fact the image of the $\mathcal{C}^1$ function x_i under d. Let T be a $\mathcal{C}^1$ map from E to an open set[2] E_1 in some real vector space E_1 of finite dimension m. Define a linear map from $\mathcal{C}(E_1, V_1^*)$ to $\mathcal{C}(E, V^*)$ as follows. Choose a basis for V_1, and let $dy_1, \ldots, dy_m$ be the dual basis for V_1^*. It is then enough to specify the image of $\omega \in \mathcal{C}(E_1, V_1^*)$ of the form $b_i dy_i$, since these generate $\mathcal{C}(E_1, V_1^*)$. We declare that the image is the continuous function from E to V^* taking any $\mathbf{x} \in E$ to

$$\sum_{j=1}^{n} b(T(\mathbf{x}))(D_j t_i)(\mathbf{x}) dx_j,$$

where t_i is the i-th coordinate of T with respect to our basis: $T = (t_1, \ldots, t_m)$. This map is called "pullback by T"; Rudin uses the notation $\omega \mapsto \omega_T$ for this.[3] Again, we have defined this using coordinates, but the chain rule shows that ω_T is the same regardless of the coordinates used for V, V_1. Indeed the chain rule is built into the definition of ω_T, which is made so that change of variables in line integrals [integrals over 1-surface] will work correctly. It will thus not surprise you that if $f : E_1 \to \mathbf{R}$ is any $\mathcal{C}^1$ function then the 1-form $(df)_T$ on E is the same as the total differential of $f \circ T$. The resulting formula looks nicer if we also regard $f \circ T$ as the pullback of f by T, and denote it by f_T. We then have simply $d(f_T) = (df)_T$. This is proved, using the chain rule, on p.263, formula 69.

[2] We must depart from Rudin's notation in p.262: Rudin uses V for an open set containing $T(E)$, but V is already taken.

[3] In the modern literature one normally sees $T^*\omega$, but I refrain from using that notation here not only to minimize divergences from the textbook but also to avoid confusion with linear duality.

Similarly it should be no surprise that pullback behaves well under composition of $\mathcal{C}^1$ maps: if we have such maps $T : E_1 \to E_2$ and $S : E_2 \to E_3$, and ω is any 1-form on E_3, then its pullback ω_{ST} under the composite map can be obtained by pulling back first to E_2, then to E_1: $\omega_{ST} = (\omega_S)_T$.

So far we have only pulled back 0-forms (functions) and 1-forms. But this is enough to define the pullback by T of a differential form ω of any order: ω is a linear combination of terms $b_I dx_{i_1} \wedge \cdots \wedge dx_{i_k}$, and we know how to pull back b_I and each dx_i, so the whole term pulls back to $(b_I)_T (dx_{i_1})_T \wedge \cdots \wedge (dx_{i_k})_T$. It follows automatically that ω_T is independent of coordinate choices, that $(\omega \wedge \lambda)_T = \omega_T \wedge \lambda_T$ for any differential forms ω, λ, and that $\omega_{ST} = (\omega_S)_T$ holds also in this setting. (This is Rudin's formula 71, a.k.a. Thm. 10.23.)

As noted in class, the definition of d generalizes to a linear map from $\mathcal{C}^1(E, \bigwedge^k V^*)$ to $\mathcal{C}(E, \bigwedge^{k+1} V^*)$. We gave the definition (formula 60) in class: if $\omega = \sum_I b_I dx_I$, with each b_I a $\mathcal{C}^1$ function on E, we define

$$d\omega := \sum_I (db_I) \wedge dx_I = \sum_I \sum_{j=1}^{n} D_j b_I \, dx_j \wedge dx_I.$$

Again this definition is coordinate-dependent. But we readily show that if ω, λ are $\mathcal{C}^1$ differential forms of orders k, m respectively then

$$d(\omega \wedge \lambda) = (d\omega) \wedge \lambda + (-1)^k \omega \wedge d\lambda$$

(formula 63). This lets us deduce the coordinate-independence of d on $\mathcal{C}^1(E, \bigwedge^k V^*)$ from its coordinate-independence on $\mathcal{C}^1(E, \mathbf{R})$, since we known that $\wedge$ does not depend on choice of coordinates. It is even true that if ω is a $\mathcal{C}^1$ form then $d(\omega_T) = (d\omega)_T$, though once $k > 0$ this identity involves second partial derivatives of T and thus requires that T be a $\mathcal{C}^2$ map. The fact that mixed partials of a $\mathcal{C}^2$ function commute then enters into the proof in a key step (see formula 70), in the guise of the the identity $d^2 = 0$ (Thm. 10.20b).