

DIFFERENTIAL FORMS AND INTEGRATION

TERENCE TAO

The concept of integration is of course fundamental in single-variable calculus. Actually, there are *three* concepts of integration which appear in the subject: the *indefinite integral* $\int f$ (also known as the *anti-derivative*), the *unsigned definite integral* $\int_{[a,b]} f(x) dx$ (which one would use to find area under a curve, or the mass of a one-dimensional object of varying density), and the *signed definite integral* $\int_a^b f(x) dx$ (which one would use for instance to compute the work required to move a particle from a to b). For simplicity we shall restrict attention here to functions $f : \mathbf{R} \rightarrow \mathbf{R}$ which are continuous on the entire real line (and similarly, when we come to differential forms, we shall only discuss forms which are continuous on the entire domain). We shall also informally use terminology such as “infinitesimal” in order to avoid having to discuss the (routine) “epsilon-delta” analytical issues that one must resolve in order to make these integration concepts fully rigorous.

These three integration concepts are of course closely related to each other in single-variable calculus; indeed, the fundamental theorem of calculus relates the signed definite integral $\int_a^b f(x) dx$ to any one of the indefinite integrals $F = \int f$ by the formula

$$\int_a^b f(x) dx = F(b) - F(a) \tag{1}$$

while the signed and unsigned integral are related by the simple identity

$$\int_a^b f(x) dx = - \int_b^a f(x) dx = \int_{[a,b]} f(x) dx \tag{2}$$

which is valid whenever $a \leq b$.

When one moves from single-variable calculus to several-variable calculus, though, these three concepts begin to diverge significantly from each other. The indefinite integral generalises to the notion of a *solution to a differential equation*, or of an *integral* of a connection, vector field, or bundle. The unsigned definite integral generalises to the *Lebesgue integral*, or more generally to *integration on a measure space*. Finally, the signed definite integral generalises to the *integration of forms*, which will be our focus here. While these three concepts still have some relation to each other, they are not as interchangeable as they are in the single-variable setting. The integration on forms concept is of fundamental importance in differential topology, geometry, and physics, and also yields one of the most important examples of *cohomology*, namely *de Rham cohomology*, which (roughly speaking) measures precisely the extent to which the fundamental theorem of calculus fails in higher dimensions and on general manifolds.

To motivate the concept, let us informally revisit the one of the basic applications of the signed definite integral from physics, namely to compute the amount of work required to move a one-dimensional particle from point a to point b , in the presence of an external field (e.g. one may move a charged particle in an electric field). At the infinitesimal level, the amount of work required to move a particle from a point $x_i \in \mathbf{R}$ to a nearby point $x_{i+1} \in \mathbf{R}$ is (up to small errors) linearly proportional to the displacement $\Delta x_i := x_{i+1} - x_i$, with the constant of proportionality $f(x_i)$ depending on the initial location x_i of the particle¹, thus the total work required here is approximately $f(x_i)\Delta x_i$. Note that we do not require that x_{i+1} be to the right of x_i , thus the displacement Δx_i (or the infinitesimal work $f(x_i)\Delta x_i$) may well be negative. To return to the non-infinitesimal problem of computing the work $\int_a^b f(x) dx$ required to move from a to b , we arbitrarily select a discrete path $x_0 = a, x_1, x_2, \dots, x_n = b$ from a to b , and approximate the work as

$$\int_a^b f(x) dx \approx \sum_{i=0}^{n-1} f(x_i)\Delta x_i. \quad (3)$$

Again, we do *not* require x_{i+1} to be to the right of x_i (nor do we require b to be to the right of a); it is quite possible for the path to “backtrack” repeatedly, for instance one might have $x_i < x_{i+1} > x_{i+2}$ for some i . However, it turns out in the one-dimensional setting, with $f : \mathbf{R} \rightarrow \mathbf{R}$ assumed to be continuous, that the effect of such backtracking eventually cancels itself out; regardless of what path we choose, the right-hand side of (3) always converges to the left-hand side as long as we assume that the maximum step size $\sup_{0 \leq i \leq n-1} |\Delta x_i|$ of the path converges to zero, *and* the total length $\sum_{i=0}^{n-1} |\Delta x_i|$ of the path (which controls the amount of backtracking involved) stays bounded. In particular, in the case when $a = b$, so that all paths are *closed* (i.e. $x_0 = x_n$), we see that signed definite integral is zero:

$$\int_a^a f(x) dx = 0. \quad (4)$$

In the language of forms, this is asserting that any one-dimensional form $f(x)dx$ on the real line \mathbf{R} is automatically *closed*. (The fundamental theorem of calculus then asserts that such forms are also automatically *exact*.) The concept of a closed form corresponds to that of a *conservative force* in physics (and an exact form corresponds to the concept of having a *potential function*).

From this informal definition of the signed definite integral it is obvious that we have the concatenation formula

$$\int_a^c f(x) dx = \int_a^b f(x) dx + \int_b^c f(x) dx \quad (5)$$

regardless of the relative position of the real numbers a, b, c . In particular (setting $a = c$ and using (4)) we conclude that

$$\int_a^b f(x) dx = - \int_b^a f(x) dx.$$

¹In analogy with the Riemann integral, we could use $f(x_i^*)$ here instead of $f(x_i)$, where x_i^* is some point intermediate between x_i and x_{i+1} . But as long as we assume f to be continuous, this technical distinction will be irrelevant.

Thus if we reverse a path from a to b to form a path from b to a , the sign of the integral changes. This is in contrast to the *unsigned definite integral* $\int_{[a,b]} f(x) dx$, since the set $[a, b]$ of numbers between a and b is exactly the same as the set of numbers between b and a . Thus we see that paths are not quite the same as sets; they carry an *orientation* which can be reversed, whereas sets do not.

Now we move from one dimensional integration to higher-dimensional integration (i.e. from single-variable calculus to several-variable calculus). It turns out that there will be *two* dimensions which will be relevant: the dimension n of the ambient space² \mathbf{R}^n , and the dimension k of the path, oriented surface, or oriented manifold S that one will be integrating over.

Let us begin with the case $n \geq 1$ and $k = 1$. Here, we will be integrating over a continuously differentiable path (or *oriented rectifiable curve*³) γ in \mathbf{R}^n starting at some point $a \in \mathbf{R}^n$ and ending at point $b \in \mathbf{R}^n$ (which may or may not be equal to a , depending on whether the path is closed or open); from a physical point of view, we are still computing the work required to move from a to b , but are now moving in several dimensions instead of one. In the one-dimensional case, we did not need to specify exactly which path we used to get from a to b (because all backtracking cancelled itself out); however, in higher dimensions, the exact choice of the path γ becomes important. Formally, a path from a to b can be described (or more precisely, *parameterised*) as a continuously differentiable function $\gamma : [0, 1] \rightarrow \mathbf{R}^n$ from the standard unit interval $[0, 1]$ to \mathbf{R}^n such that $\gamma(0) = a$ and $\gamma(1) = b$. For instance, the line segment from a to b can be parameterised as $\gamma(t) := (1-t)a + tb$. This segment also has many other parameterisations, e.g. $\tilde{\gamma}(t) := (1-t^2)a + t^2b$; it will turn out though (similarly to the one-dimensional case) that the exact choice of parameterisation does not ultimately influence the integral. On the other hand, the reverse line segment $(-\gamma)(t) := ta + (1-t)b$ from b to a is a genuinely different path; the integral on $-\gamma$ will turn out to be the negative of the integral on γ .

As in the one-dimensional case, we will need to approximate the continuous path γ by a discrete path

$$x_0 = \gamma(0) = a, x_1 = \gamma(t_1), x_2 = \gamma(t_2), \dots, x_n = \gamma(1) = b.$$

Again, we allow some backtracking: t_{i+1} is not necessarily larger than t_i . The displacement $\Delta x_i := x_{i+1} - x_i \in \mathbf{R}^n$ from x_i to x_{i+1} is now a *vector* rather than a scalar. (Indeed, one should think of Δx_i as an infinitesimal *tangent vector* to the ambient space \mathbf{R}^n at the point x_i .) In the one-dimensional case, we converted the scalar displacement Δx_i into a new number $f(x_i)\Delta x_i$, which was linearly related to the original displacement by a proportionality constant $f(x_i)$ depending on the position x_i . In higher dimensions, the analogue of a “proportionality constant” of

²We will start with integration on Euclidean spaces \mathbf{R}^n for simplicity, although the true power of the integration on forms concept is only apparent when we integrate on more general spaces, such as abstract n -dimensional manifolds.

³Some authors distinguish between a path and an oriented curve by requiring that paths to have a designated parameterisation $\gamma : [0, 1] \rightarrow \mathbf{R}^n$, whereas curves do not. This distinction will be irrelevant for our discussion and so we shall use the terms interchangeably. It is possible to integrate on more general curves (e.g. the (unrectifiable) Koch snowflake curve, which has infinite length), but we do not discuss this in order to avoid some technicalities.

a linear relationship is a *linear transformation*. Thus, for each x_i we shall need a linear transformation $\omega_{x_i} : \mathbf{R}^n \rightarrow \mathbf{R}$ that takes an (infinitesimal) displacement $\Delta x_i \in \mathbf{R}^n$ as input and returns an (infinitesimal) scalar $\omega_{x_i}(\Delta x_i) \in \mathbf{R}$ as output, representing the infinitesimal “work” required to move from x_i to x_{i+1} . (In other words, ω_{x_i} is a *linear functional* on the space of tangent vectors at x_i , and is thus a *cotangent vector* at x_i .) In analogy to (3), the net work $\int_\gamma \omega$ required to move from a to b along the path γ is approximated by

$$\int_\gamma \omega \approx \sum_{i=0}^{n-1} \omega_{x_i}(\Delta x_i). \quad (6)$$

If ω_{x_i} depends continuously on x_i , then (as in the one-dimensional case) one can show that the right-hand side of (6) is convergent if the maximum step size $\sup_{0 \leq i \leq n-1} |\Delta x_i|$ of the path converges to zero, *and* the total length $\sum_{i=0}^{n-1} |\Delta x_i|$ of the path stays bounded. The object ω , which continuously assigns⁴ a cotangent vector to each point in \mathbf{R}^n , is called a *1-form*, and (6) leads to a recipe to integrate any 1-form ω on a path γ (or, to shift the emphasis slightly, to integrate the path γ against the 1-form ω). Indeed, it is useful to think of this integration as a *binary* operation (similar in some ways to the dot product) which takes the curve γ and the form ω as inputs, and returns a scalar $\int_\gamma \omega$ as output. There is in fact a “duality” between curves and forms; compare for instance the identity

$$\int_\gamma (\omega_1 + \omega_2) = \int_\gamma \omega_1 + \int_\gamma \omega_2$$

(which expresses (part of) the fundamental fact that integration on forms is a linear operation) with the identity

$$\int_{\gamma_1 + \gamma_2} \omega = \int_{\gamma_1} \omega + \int_{\gamma_2} \omega$$

(which generalises (5)) whenever⁵ the initial point of γ_2 is the final point of γ_1 , where $\gamma_1 + \gamma_2$ is the *concatenation* of γ_1 and γ_2 . This duality is best understood using the abstract (and much more general) formalism of *homology* and *cohomology*.

Because \mathbf{R}^n is a Euclidean vector space, it comes with a dot product $(x, y) \mapsto x \cdot y$, which can be used to describe 1-forms in terms of vector fields (or equivalently, to identify cotangent vectors and tangent vectors): specifically, for every 1-form ω there is a unique vector field $F : \mathbf{R}^n \rightarrow \mathbf{R}^n$ such that $\omega_x(v) := F(x) \cdot v$ for all $x, v \in \mathbf{R}^n$. With this representation, the integral $\int_\gamma \omega$ is often written as $\int_\gamma F(x) \cdot dx$. However, we shall avoid this notation because it gives the misleading impression that Euclidean structures such as the dot product are an essential aspect of the integration on differential forms concept, which can lead to confusion when one generalises this concept to more general manifolds on which the natural analogue of the dot product (namely, a *Riemannian metric*) might be unavailable.

⁴More precisely, one can think of ω as a *section* of the *cotangent bundle*.

⁵One can remove the requirement that γ_2 begins where γ_1 leaves off by generalising the notion of an integral to cover not just integration on paths, but also integration on *formal sums* or *differences* of paths. This makes the duality between curves and forms more symmetric.

Note that to any continuously differentiable function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ one can assign a 1-form, namely the *derivative* df of f , defined as the unique 1-form such that one has the Taylor approximation

$$f(x + v) \approx f(x) + df_x(v)$$

for all infinitesimal v , or more rigorously that $|f(x + v) - f(x) - df_x(v)|/|v| \rightarrow 0$ as $v \rightarrow 0$. Using the Euclidean structure, one can express $df_x = \nabla f(x) \cdot dx$, where ∇f is the *gradient* of f ; but note that the derivative df can be defined without any appeal to Euclidean structure. The fundamental theorem of calculus (1) now generalises as

$$\int_{\gamma} df = f(b) - f(a) \tag{7}$$

whenever γ is any oriented curve from a point a to a point b . In particular, if γ is closed, then $\int_{\gamma} df = 0$. A 1-form which can be expressed in the form df for some continuously differentiable function is called *exact*; the fundamental theorems of calculus assert that a 1-form is exact if and only if its integral against every closed curve vanishes. This turns out to be a general fact, valid for all manifolds. There is the slightly weaker notion of a 1-form being *closed*, which means that its integral against every *sufficiently small* closed curve around any given point vanishes. Clearly every exact 1-form is closed. Is the converse true? In a Euclidean space (or more generally, any other *simply connected* manifold), then the answer is yes (basically because one can decompose any large closed curve into small closed curves on such manifolds), but it is not true for general domains; in modern terminology, this demonstrates that the *de Rham cohomology* of such domains can be non-trivial.

Now we turn to integration on k -dimensional sets with $k > 1$; for simplicity we discuss the two-dimensional case $k = 2$, i.e. integration of forms on (oriented) surfaces in \mathbf{R}^n , as this already illustrates many features of the general case. Physically, such integrals arise when computing a *flux* of some field (e.g. a magnetic field) across a surface; a more intuitive example⁶ would arise when computing the net amount of force exerted by a wind blowing on a sail. We parameterised one-dimensional oriented curves as continuously differentiable functions $\gamma : [0, 1] \rightarrow \mathbf{R}^n$ on the standard (oriented) unit interval $[0, 1]$; it is thus natural to parameterise two-dimensional oriented surfaces as continuously differentiable functions $\phi : [0, 1]^2 \rightarrow \mathbf{R}^n$ on the standard (oriented) unit square $[0, 1]^2$ (we will be vague here about what “oriented” means). This will not quite cover all possible surfaces one wishes to integrate over, but it turns out that one can cut up more general surfaces into pieces which can be parameterised using “nice” domains such as $[0, 1]^2$.

In the one-dimensional case, we cut up the oriented interval $[0, 1]$ into infinitesimal oriented intervals from t_i to $t_{i+1} = t_i + \Delta t$, thus leading to infinitesimal curves from $x_i = \gamma(t_i)$ to $x_{i+1} = \gamma(t_{i+1}) = x_i + \Delta x_i$. Note from Taylor expansion that Δx_i and Δt are related by the approximation $\Delta x_i \approx \gamma'(t_i)\Delta t_i$. In the two-dimensional case,

⁶Actually, this example is misleading for two reasons. Firstly, net force is a vector quantity rather than a scalar quantity; secondly, the sail is an unoriented surface rather than an oriented surface. A more accurate example would be the net amount of light falling on *one side* of a sail, where any light falling on the opposite side counts *negatively* towards that net amount.

we will cut up the oriented unit square $[0, 1]^2$ into infinitesimal oriented squares⁷, a typical one of which may have corners $(t_1, t_2), (t_1 + \Delta t, t_2), (t_1, t_2 + \Delta t), (t_1 + \Delta t, t_2 + \Delta t)$. The surface described by ϕ can then be partitioned into (oriented) regions with corners $x := \phi(t_1, t_2), \phi(t_1 + \Delta t, t_2), \phi(t_1, t_2 + \Delta t), \phi(t_1 + \Delta t, t_2 + \Delta t)$. Using Taylor expansion in several variables, we see that this region is approximately an (oriented) parallelogram in \mathbf{R}^n with corners $x, x + \Delta_1 x, x + \Delta_2 x, x + \Delta_1 x + \Delta_2 x$, where $\Delta_1 x, \Delta_2 x \in \mathbf{R}^n$ are the infinitesimal vectors

$$\Delta_1 x := \frac{\partial \phi}{\partial t_1}(t_1, t_2) \Delta t; \quad \Delta_2 x := \frac{\partial \phi}{\partial t_2}(t_1, t_2) \Delta t.$$

Let us refer to this object as the infinitesimal parallelogram with *dimensions* $\Delta_1 x \wedge \Delta_2 x$ with base point x ; at this point, the symbol \wedge is a meaningless placeholder. In order to integrate in a manner analogous with integration on curves, we now need some sort of functional ω_x at this base point which should take the above infinitesimal parallelogram and return an infinitesimal number $\omega_x(\Delta_1 x \wedge \Delta_2 x)$, which physically should represent the amount of “flux” passing through this parallelogram.

In the one-dimensional case, the map $\Delta x \mapsto \omega_x(\Delta x)$ was required to be linear; or in other words, we required the axioms

$$\omega_x(c\Delta x) = c\omega_x(\Delta x); \quad \omega_x(\Delta x + \widetilde{\Delta x}) = \omega_x(\Delta x) + \omega_x(\widetilde{\Delta x})$$

for any $c \in \mathbf{R}$ and $\Delta x, \widetilde{\Delta x} \in \mathbf{R}^n$. Note that these axioms are intuitively consistent with the interpretation of $\omega_x(\Delta x)$ as the total amount of work required or flux experienced along the oriented interval from x to $x + \Delta x$. Similarly, we will require that the map $(\Delta_1 x, \Delta_2 x) \mapsto \omega_x(\Delta_1 x \wedge \Delta_2 x)$ be *bilinear*, thus we have the axioms

$$\begin{aligned} \omega_x(c\Delta x_1 \wedge \Delta x_2) &= c\omega_x(\Delta x_1 \wedge \Delta x_2) \\ \omega_x((\Delta x_1 + \widetilde{\Delta x}_1) \wedge \Delta x_2) &= \omega_x(\Delta x_1 \wedge \Delta x_2) + \omega_x(\widetilde{\Delta x}_1 \wedge \Delta x_2) \\ \omega_x(\Delta x_1 \wedge c\Delta x_2) &= c\omega_x(\Delta x_1 \wedge \Delta x_2) \\ \omega_x(\Delta x_1 \wedge (\Delta x_2 + \widetilde{\Delta x}_2)) &= \omega_x(\Delta x_1 \wedge \Delta x_2) + \omega_x(\Delta x_1 \wedge \widetilde{\Delta x}_2) \end{aligned}$$

for all $c \in \mathbf{R}$ and $\Delta x_1, \Delta x_2, \widetilde{\Delta x}_1, \widetilde{\Delta x}_2$. These axioms are also physically intuitive, though it may require a little more effort to see this than in the one-dimensional case. There is one additional important axiom we require, namely that

$$\omega_x(\Delta x \wedge \Delta x) = 0 \tag{8}$$

for all $\Delta x \in \mathbf{R}^n$. This reflects the geometrically obvious fact that when $\Delta_1 x = \Delta_2 x = \Delta x$, the parallelogram with dimensions $\Delta x \wedge \Delta x$ is degenerate and should thus experience zero net flux. Any continuous assignment $\omega : x \mapsto \omega_x$ that obeys the above axioms is called⁸ a *2-form*.

⁷One could also use infinitesimal oriented rectangles, parallelograms, triangles, etc.; this leads to an equivalent concept of the integral.

⁸There are several other equivalent definitions of a 2-form. For instance, as hinted at earlier, 1-forms can be viewed as sections of the cotangent bundle $T^*\mathbf{R}^n$, and similarly 2-forms are sections of the *exterior power* $\wedge^2 T^*\mathbf{R}^n$ of that bundle. Similarly, expressions such as $v \wedge w$, where $v, w \in T_x \mathbf{R}^n$ are tangent vectors at a point x , can be given meaning by using abstract algebra to construct the exterior power $\wedge^2 T_x \mathbf{R}^n$, at which point $(v, w) \mapsto v \wedge w$ can be viewed as a bilinear anti-symmetric map from $T_x \mathbf{R}^n \times T_x \mathbf{R}^n$ to $\wedge^2 T_x \mathbf{R}^n$ (indeed it is the *universal* map with this properties). One can also construct forms using the machinery of *tensors*.

By applying (8) with $\Delta x := \Delta_1 x + \Delta_2 x$ and then using several of the above axioms, we arrive at the fundamental *anti-symmetry property*

$$\omega_x(\Delta x_1 \wedge \Delta x_2) = -\omega_x(\Delta x_2 \wedge \Delta x_1). \quad (9)$$

Thus swapping the first and second vectors of a parallelogram causes a reversal in the flux across that parallelogram; the latter parallelogram should then be considered to have the reverse orientation to the former.

If ω is a 2-form and $\phi : [0, 1]^2 \rightarrow \mathbf{R}^n$ is a continuously differentiable function, we can now define the integral $\int_{\phi} \omega$ of ω against ϕ (or more precisely, the image of the oriented square $[0, 1]^2$ under ϕ) by the approximation

$$\int_{\phi} \omega \approx \sum_i \omega_{x_i}(\Delta x_{1,i} \wedge \Delta x_{2,i}) \quad (10)$$

where the image of ϕ is (approximately) partitioned into parallelograms of dimensions $\Delta x_{1,i} \wedge \Delta x_{2,i}$ based at points x_i . We do not need to decide what order these parallelograms should be arranged in, because addition is both commutative⁹ and associative. One can show that the right-hand side of (10) converges to a unique limit as one makes the partition of parallelograms “increasingly fine”, though we will not make this precise here.

We have thus shown how to integrate 2-forms against oriented 2-dimensional surfaces. More generally, one can define the concept of a k -form¹⁰ on an n -dimensional manifold (such as \mathbf{R}^n) for any $0 \leq k \leq n$ and integrate this against an oriented k -dimensional surface in that manifold. For instance, a 0-form on a manifold X is the same thing as a scalar function $f : X \rightarrow \mathbf{R}$, whose integral on a positively oriented point x (which is 0-dimensional) is $f(x)$, and on a negatively oriented point x is $-f(x)$. By convention, if $k \neq k'$, the integral of a k -dimensional form on a k' -dimensional surface is understood to be zero. We refer to 0-forms, 1-forms, 2-forms, etc. (and formal sums and differences thereof) collectively as *differential forms*.

Scalar functions enjoy three fundamental operations: addition $(f, g) \mapsto f + g$, pointwise product $(f, g) \mapsto fg$, and differentiation $f \mapsto df$, although the latter is only obviously well-defined when f is continuously differentiable. These operations obey various relationships, for instance the product *distributes* over addition

$$f(g + h) = fg + fh$$

and differentiation is a *derivation* with respect to the product:

$$d(fg) = (df)g + f(dg).$$

It turns out that one can generalise all three of these operations to differential forms: one can add or take the *wedge product* of two forms ω, η to obtain new forms

⁹For some other notions of an integral, such as that of an integral of a connection with a non-abelian structure group, one loses commutativity, and so one can only integrate along one-dimensional curves.

¹⁰One can also define k -forms for $k > n$, but it turns out that the multilinearity and antisymmetry axioms for such forms will force them to vanish, basically because any k vectors in \mathbf{R}^n are necessarily linearly dependent.

$\omega + \eta$ and $\omega \wedge \eta$; and, if a k -form ω is continuously differentiable, one can also form the derivative $d\omega$, which is a $k+1$ -form. The exact construction of these operations requires a little bit of algebra and is omitted here. However, we remark that these operations obey similar laws to their scalar counterparts, except that there are some sign changes which are ultimately due to the anti-symmetry (9). For instance, if ω is a k -form and η is an l -form, the commutative law for multiplication becomes

$$\omega \wedge \eta = (-1)^{kl} \eta \wedge \omega,$$

and the derivation rule for differentiation becomes

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

A fundamentally important, though initially rather unintuitive¹¹ rule, is that the differentiation operator d is nilpotent:

$$d(d\omega) = 0. \tag{11}$$

The fundamental theorem of calculus generalises to *Stokes' theorem*

$$\int_S d\omega = \int_{\partial S} \omega \tag{12}$$

for any oriented manifold S and form ω , where ∂S is the oriented boundary of S (which we will not define here). Indeed one can view this theorem (which generalises (1), (7)) as a definition of the derivative operation $\omega \mapsto d\omega$; thus differentiation is the *adjoint* of the boundary operation. (Thus, for instance, the identity (11) is dual to the geometric observation that the boundary ∂S of an oriented manifold itself has no boundary: $\partial(\partial S) = \emptyset$.) As a particular case of Stokes' theorem, we see that $\int_S d\omega = 0$ whenever S is a *closed* manifold, i.e. one with no boundary. This observation lets one extend the notions of closed and exact forms to general differential forms, which (together with (11)) allows one to fully set up *de Rham cohomology*.

We have already seen that 0-forms can be identified with scalar functions, and in Euclidean spaces 1-forms can be identified with vector fields. In the special (but very physical) case of three-dimensional Euclidean space \mathbf{R}^3 , 2-forms can *also* be identified with vector fields via the famous *right-hand rule*¹², and 3-forms can be identified with scalar functions by a variant of this rule. (This is an example of *Hodge duality*.) In this case, the differentiation operation $\omega \mapsto d\omega$ is identifiable to the *gradient* operation $f \mapsto \nabla f$ when ω is a 0-form, to the *curl* operation $X \mapsto \nabla \times X$ when ω is a 1-form, and the *divergence* operation $X \mapsto \nabla \cdot X$ when ω is a 2-form. Thus, for instance, the rule (11) implies that $\nabla \times \nabla f = 0$ and $\nabla \cdot (\nabla \times X) = 0$ for any suitably smooth scalar function f and vector field X , while Stokes' theorem (12), with this interpretation, becomes the Stokes' theorems for integrals of curves and surfaces in three dimensions that may be familiar to you from several variable calculus.

¹¹It may help to view $d\omega$ as really being a “wedge product” $d \wedge \omega$ of the differentiation operation with ω , in which case (11) is a formal consequence of (8) and the associativity of the wedge product.

¹²This is an entirely arbitrary convention; one could just as easily have used the left-hand rule to provide this identification, and apart from some harmless sign changes here and there, one gets essentially the same theory as a consequence.

Just as the signed definite integral is connected to the unsigned definite integral in one dimension via (2), there is a connection between integration of differential forms and the Lebesgue (or Riemann) integral. On the Euclidean space \mathbf{R}^n one has the n standard co-ordinate functions $x_1, x_2, \dots, x_n : \mathbf{R}^n \rightarrow \mathbf{R}$. Their derivatives dx_1, \dots, dx_n are then 1-forms on \mathbf{R}^n . Taking their wedge product one obtains an n -form $dx_1 \wedge \dots \wedge dx_n$. We can multiply this with any (continuous) scalar function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ to obtain another n -form $f dx_1 \wedge \dots \wedge dx_n$. If Ω is any open bounded domain in \mathbf{R}^n , we then have the identity

$$\int_{\Omega} f(x) dx_1 \wedge \dots \wedge dx_n = \int_{\Omega} f(x) dx$$

where on the left we have an integral of a differential form (with Ω viewed as a positively oriented n -dimensional manifold), and on the right we have the Riemann or Lebesgue integral of f on Ω . If we give Ω the negative orientation, we have to reverse the sign of the left-hand side. This correspondence generalises (2).

There is one last operation on forms which is worth pointing out. Suppose we have a continuously differentiable map $\Phi : X \rightarrow Y$ from one manifold to another (we allow X and Y to have different dimensions). Then of course every point x in X *pushes forward* to a point $\Phi(x)$ in Y . Similarly, if we let $v \in T_x X$ be an infinitesimal tangent vector to X based at x , then this tangent vector also *pushes forward* to a tangent vector $\Phi_* v \in T_{\Phi(x)}(Y)$ based at $\Phi(x)$; informally speaking, $\Phi_* v$ can be defined by requiring the infinitesimal approximation $\Phi(x+v) = \Phi(x) + \Phi_* v$. One can write $\Phi_* v = D\Phi(x)(v)$, where $D\Phi : T_x X \rightarrow T_{\Phi(x)} Y$ is the *derivative* of the several-variable map Φ at x . Finally, any k -dimensional oriented manifold S in X also pushes forward to a k -dimensional oriented manifold $\Phi(S)$ in Y , although in some cases (e.g. if the image of Φ has dimension less than k) this pushed-forward manifold may be degenerate.

We have seen that integration is a duality pairing between manifolds and forms. Since manifolds push forward under Φ from X to Y , we thus expect forms to *pull-back* from Y to X . Indeed, given any k -form ω on Y , we can define the *pull-back* $\Phi^* \omega$ as the unique k -form on X such that we have the *change of variables formula*

$$\int_{\Phi(S)} \omega = \int_S \Phi^*(\omega).$$

In the case of 0-forms (i.e. scalar functions), the pull-back $\Phi^* f : X \rightarrow \mathbf{R}$ of a scalar function $f : Y \rightarrow \mathbf{R}$ is given explicitly by $\Phi^* f(x) = f(\Phi(x))$, while the pull-back of a 1-form ω is given explicitly by the formula

$$(\Phi^* \omega)_x(v) = \omega_{\Phi(x)}(\Phi_* v).$$

Similarly for other differential forms. The pull-back operation enjoys several nice properties, for instance it respects the wedge product,

$$\Phi^*(\omega \wedge \eta) = (\Phi^* \omega) \wedge (\Phi^* \eta),$$

and the derivative,

$$d(\Phi^* \omega) = \Phi^*(d\omega).$$

By using these properties, one can recover rather painlessly the change-of-variables formulae in several-variable calculus. Moreover, the whole theory carries effortlessly over from Euclidean spaces to other manifolds. It is because of this that the theory

of differential forms and integration is an indispensable tool in the modern study of manifolds, especially in differential topology.

Update, Apr 7, 2018: Definition of closed 1-form corrected. Thanks to Michael Drumheller and Nicola Nesa for corrections.

Update, Aug 25, 2018: Further typo corrections (thanks to Mark Muldoon for pointing them out).

Update, Aug 15, 2020: Another typo correction (thanks to Michael Zochowski for pointing it out).

Update, Aug 19, 2020: Another typo correction (thanks to Charles Jin for pointing it out).

Update, Dec 18, 2020: Another typo correction (thanks to Ho Boon Suan for pointing it out).

DEPARTMENT OF MATHEMATICS, UCLA, LOS ANGELES CA 90095-1555

E-mail address: `tao@math.ucla.edu`