STUDENT ARTICLE

2

The Poincaré Lemma and de Rham Cohomology

Daniel Litt[†]
Harvard University '10
Cambridge, MA 02138
dalitt@fas.harvard.edu

Abstract

The Poincaré Lemma is a staple of rigorous multivariate calculus—however, proofs provided in early undergraduate education are often overly computational and are rarely illuminating. We provide a conceptual proof of the Lemma, making use of some tools from higher mathematics.

The concepts here should be understandable to those familiar with multivariable calculus, linear algebra, and a minimal amount of group theory. Many of the ideas used in the proof are ubiquitous in mathematics, and the Lemma itself has applications in areas ranging from electrodynamics to calculus on manifolds.

2.1 Introduction

Much of calculus and analysis—the path-independence of line- or surface-integrals on certain domains, Cauchy's Theorem (assuming the relevant functions are C^1) on connected complex regions and the more general residue theorem, and various ideas from physics—depends to a large extent on a powerful result known as the Poincaré Lemma. On the way to the statement and proof of this Lemma, we will introduce the concepts of the exterior power and differential forms, as well as de Rham cohomology.

2.2 Linear Algebra and Calculus Preliminaries

2.2.1 The Exterior Power

We begin by defining some useful objects; on the way, we will digress slightly and remark on their interesting properties. We will begin by defining a vector space called the exterior power, in order to extend the notion of a determinant.

Definition 1 (Alternating Multilinear Form, Exterior Power)**.** Let V be a finite-dimensional vector space over a field F. A n**-linear form** is a map $B : \underbrace{V \times \cdots \times V}_{n} \to W$, where W is an arbitrary vector space over F, that is linear in each term, i.e. such that

$$B(a_1, a_2, \ldots, a_n) + B(a_1', a_2, \ldots, a_n) = B(a_1 + a_1', a_2, \ldots, a_n)$$

[†]Daniel Litt, Harvard '10, is a mathematics concentrator living in Leverett House. Originally from Cleveland, Ohio, he has performed research in computational biology, physics, combinatorics, game theory, and category theory. His other academic interests include epistemology and moral philosophy, as well as poetry and music. He is a founding member of The HCMR and currently serves as Issue Production Director.

and similarly in the other variables, and such that

$$s \cdot B(a_1, a_2, \ldots, a_n) = B(s \cdot a_1, a_2, \ldots, a_n) = B(a_1, s \cdot a_2, \ldots, a_n) = \cdots$$

We say that such forms are **multilinear**.

A multilinear form B is **alternating** if it satisfies

$$B(a_1, \ldots, a_i, a_{i+1}, \ldots, a_n) = -B(a_1, \ldots, a_{i+1}, a_i, \ldots, a_n)$$

for all $1 \leq i < n$.

Then the **n-th exterior power** of V, denoted $\bigwedge^n(V)$, is a vector space equipped with an alternating multilinear map $\wedge : \underbrace{V \times \cdots \times V}_{n} \to \bigwedge^n(V)$ such that any alternating multilinear map

$$f : \underbrace{V \times \cdots \times V}_{n} \to W$$

factors uniquely through $\wedge$, that is, there exists a unique $f' : \bigwedge^n(V) \to W$ such that the diagram

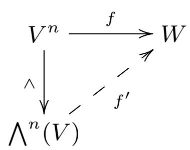

commutes, i.e. $f' \circ \wedge = f$.

It is not immediately clear that such a vector space exists or is unique. For existence, see [DF, p. 446]; the construction is not important for our purposes so we relegate it to a footnote.[1] Uniqueness follows immediately from the fact that the above definition is a **universal property**. However, we provide the following proof to elucidate this notion:

Proposition 2. *The n-th exterior power of a vector space V is unique up to canonical isomorphism.*

Proof. Consider vector spaces $\bigwedge_1^n(V)$ and $\bigwedge_2^n(V)$ with associated maps $\wedge_1$ and $\wedge_2$ satisfying the definition above. As $\wedge_1, \wedge_2$ are both alternating and multilinear, they must factor through one another; that is, we must have that there exist unique $\wedge_1', \wedge_2'$ such that the diagram

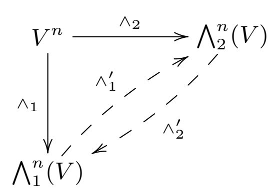

[1] We have that $\bigwedge^0(V) = F$ and $\bigwedge^1(V) = V$. There are three equivalent ways to construct the exterior power for $n \geq 2$. First, the exterior power $\bigwedge^n(V)$ can be viewed as the span of formal strings $v_1 \wedge \cdots \wedge v_n$, where $\wedge$ is a formal symbol satisfying the properties of the wedge product.

Second, for readers familiar with the tensor power, we may, for $n \geq 2$, let $I^2(V) \subset \bigotimes^2(V)$ be the subspace spanned by vectors of the form $v \otimes v$ in $\bigotimes^2(V)$ and for $n > 2$ let $I^n(V) = (V \otimes I^{n-1}(V)) \bigoplus (I^{n-1}(V) \otimes V)$; then $\bigwedge^n(V) \simeq \bigotimes^n(V)/I^n(V)$.

Finally, let $J^2(V) \subset V \otimes V$ be the subspace spanned by vectors of the form $(v \otimes w - w \otimes v)$. Then for $n \geq 2$, $\bigwedge^n(V) \simeq \bigcap_{i=0}^{n-2} \left(\bigotimes^i(V) \otimes J^2(V) \otimes \bigotimes^{n-i-2}(V) \right) \subset \bigotimes^n(V)$.

We leave checking that these constructions satisfy the definition of the exterior power (and are thus isomorphic) as an exercise; the reader may look at the given reference [DF] for the solution.

commutes. Note however that $\wedge_1, \wedge_2$ must be surjective, so $\wedge_1', \wedge_2'$ must be mutually inverse by the commutativity of the diagram above. But then $\bigwedge_1^n(V) \simeq \bigwedge_2^n(V)$ as desired, and we have uniqueness. □

In accordance with convention, for $v_1, \ldots, v_n \in V$, we define

$$v_1 \wedge \cdots \wedge v_n := \wedge(v_1, \ldots, v_n).$$

This is referred to as the **wedge product** of $v_1, \ldots, v_n$. By the fact that $\wedge$ is multilinear and alternating, note that

1. $v \wedge w = -w \wedge v$ (this immediately implies that $v \wedge v = 0$),
2. $a \cdot (v \wedge w) = (a \cdot v) \wedge w = v \wedge (a \cdot w)$,
3. $v \wedge w + v \wedge w' = v \wedge (w + w')$,

with the appropriate generalizations to higher-order exterior powers. We now calculate, for an n-dimensional vector space V, the dimension of $\bigwedge^s(V)$.

Proposition 3. *We have that*

$$\dim \bigwedge^s(V) = \binom{n}{s}.$$

In particular, given a basis $\{e_1, \ldots, e_n\}$ *of* V*, the vectors*

$$e_{i_1} \wedge \cdots \wedge e_{i_s} \text{ for } 1 \leq i_1 < i_2 < \cdots < i_s \leq n$$

form a basis of $\bigwedge^s(V)$*.*

Proof. We begin with the second claim. We first show that

$$\bigwedge^s(V) = \text{span}(e_{i_1} \wedge \cdots \wedge e_{i_s} \mid 1 \leq i_1 < i_2 < \cdots < i_s \leq n).$$

As $\{e_1, \ldots, e_n\}$ is a basis of V, we may write any $v_1 \wedge \cdots \wedge v_s \in \bigwedge^s(V)$ as

$$\left(\sum_{i=1}^n a_i^1 \cdot e_i\right) \wedge \cdots \wedge \left(\sum_{i=1}^n a_i^s \cdot e_i\right)$$

for some $(a_i^j) \in F$. We may distribute the wedge product over this summation by multilinearity—(3) above—and rearrange the terms appropriately by (1) above, so that we have a linear combination of vectors in the desired form.

To see that these vectors are linearly independent, we produce linear maps $B_{i_1,\ldots,i_k} : \bigwedge^s(V) \to F$ such that $B_{i_1,\ldots,i_s}(e_{i_1} \wedge \cdots \wedge e_{i_s}) = 1$ and for $\{i_1', \ldots, i_s'\} \neq \{i_1, \ldots, i_s\}$, $B_{i_1,\ldots,i_s}(e_{i_1'} \wedge \cdots \wedge e_{i_s'}) = 0$. This is sufficient because if $e_{i_1} \wedge \cdots \wedge e_{i_s}$ were to equal

$$\sum_{\{j_1,\ldots,j_s\} \neq \{i_1,\ldots i_s\}} a_{j_1,\ldots,j_s} \cdot e_{j_1} \wedge \cdots \wedge e_{j_s}$$

with some $a_{j_1,\ldots,j_s} \in F$ nonzero, we would have $B_{j_1,\ldots,j_s}(e_{i_1} \wedge \cdots \wedge e_{i_s}) = a_{j_1,\ldots,j_s}$, a contradiction.

Given an ordered n-tuple of basis elements $(e_{i_1}, \ldots, e_{i_s}) \in V^n$, with no two indices equal, let $\sigma_{i_1,\ldots,i_s}$ be the unique permutation that orders the indices of the basis elements. Consider the map $\Sigma_{i_1,\ldots,i_s} : \underbrace{V \times \cdots \times V}_{s} \to F,\ 1 \leq i_1 < \cdots < i_s \leq n$ defined by

$$\Sigma_{i_1,\ldots,i_s}(e_{j_1}, \ldots, e_{j_s}) = \begin{cases} 0, & \text{if } \{i_1, \ldots, i_s\} \neq \{j_1, \ldots, j_s\} \\ \text{sign}(\sigma_{j_1,\ldots,j_s}) & \text{if } \{i_1, \ldots, i_s\} = \{j_1, \ldots, j_s\} \end{cases}$$

and extended by imposing multilinearity. It is easy to check that this map is multilinear and alternating, so it must factor uniquely through $\wedge$; that the resulting map $\bigwedge^s(V) \to F$ is $B_{i_1,\ldots,i_s}$ is also easy to check.

But the number of sets of strictly ordered indices $\{i_1, \ldots, i_s\}$ is $\binom{n}{s}$, as claimed, which completes the proof. □

Corollary 4. *Let V be an n-dimensional vector space. Then* $\dim \bigwedge^n(V) = 1$.

Proof. By Proposition 2, we have $\dim \bigwedge^n(V) = \binom{n}{n} = 1$. □

We now extend the standard notion of the determinant. Given an endomorphism $T : V \to V$, we define $\bigwedge^s(T) : \bigwedge^s(V) \to \bigwedge^s(V)$ to be the map

$$v_1 \wedge \cdots \wedge v_s \mapsto T(v_1) \wedge \cdots \wedge T(v_s).$$

This map is linear as T is linear and as $\wedge$ is multilinear.

Definition 5 (Determinant)**.** The **determinant** $\det(T)$ of an endomorphism T of an n-dimensional vector space V is the map

$$\det(T) := \bigwedge^n(T).$$

In particular, note that $\bigwedge^n(T)$ is a map on a one-dimensional vector space (by Corollary 1), and is thus simply multiplication by a scalar. We claim that, having chosen a basis for $\bigwedge^n(V)$, this scalar is exactly the standard notion of the determinant; proving this is an exercise in algebra, which we recommend the reader pursue. Furthermore, this definition allows one to prove easily that the determinant of T is nonzero if and only if T is invertible; the proof follows below.

Proposition 6. *A linear map $T : V \to V$ is invertible if and only if* $\det(T) \neq 0$.

Proof. Note that $\bigwedge^n(\text{id}_V) = \text{id}_{\bigwedge^n(V)}$ and that, given two endomorphisms $T, S : V \to V$, $\bigwedge^n(T \circ S) = \bigwedge^n(T) \circ \bigwedge^n(S)$; that is, $\bigwedge$ respects identity and composition.[2]

To see necessity, note that we have

$$\text{id}_{\bigwedge^n(V)} = \bigwedge^n(\text{id}_V) = \bigwedge^n(T \circ T^{-1}) = \bigwedge^n(T) \circ \bigwedge^n(T^{-1}).$$

But then $\bigwedge^n(T)$ is non-zero, as it is invertible ($\bigwedge^n(T^{-1})$ is its inverse).

To see sufficiency, we show the contrapositive; that is, for non-invertible T, $\det(T) = 0$. Assume that $\dim T(V) < n$. Let $m = \dim T(V)$, and let $\{e_1, \ldots, e_m\}$ be a basis for $T(V)$. But then given any $v_1 \wedge \cdots \wedge v_n \in \bigwedge^n(V)$, we have, distributing, that $\bigwedge^n(T)(v_1 \wedge \cdots \wedge v_n) = \sum a_{i_1,\ldots,i_n} \cdot e_{i_1} \wedge \cdots \wedge e_{i_n}$; as $m < n$, we have by the pigeonhole principle that each term contains a repeated index. But then by (1) above, the determinant is zero as claimed. □

[2] In fact, $\bigwedge^n(-)$ is a functor.

2.2.2 Homotopies

The motivation here is to classify maps and domains by the existence of continuous transformations between them; we give some definitions that will be useful later. In particular, we wish to characterize the types of domains on which the the Poincaré Lemma will hold.

Definition 7 (Homotopy)**.** Two continuous maps $g_0, g_1 : U \to V$ with $U \subset \mathbb{R}^m, V \subset \mathbb{R}^n$ are said to be **homotopic** if there exists a continuous map $G : [0,1] \times U \to V$ such that for all $x \in U, g_0(x) = G(0,x)$ and $g_1(x) = G(1,x)$.

Intuitively, the notion behind this definition is that $G(t, -)$ interpolates continuously between g_0, g_1. We may use this idea to characterize certain types of domains, which may, speaking imprecisely, be continuously squished to a point.

Definition 8 (Contractibility)**.** We say a domain $U \in \mathbb{R}^m$ is **contractible** if, for some point $c \in U$, the constant map $x \mapsto c$ is homotopic to the identity on U.

Note that all convex and star-shaped domains are contractible.

2.2.3 The Change of Variables Formula

We now begin the calculus preliminaries to the Poincaré Lemma. A well-known theorem from single-variable calculus states that for integrable f defined and continuously differentiable g on $[a,b]$ with integrable derivative, and with f defined on $g([a,b])$, we have

$$\int_{g(a)}^{g(b)} f(x)\, dx = \int_a^b f \circ g(x) \cdot g'(x)\, dx. \tag{2.1}$$

This is proved easily using the chain rule and the fundamental theorem of calculus.

The theorem (2.1) may be generalized to multivariate functions as follows:

Proposition 9 (The Change of Variables Theorem)**.** *Consider open $U, V \subset \mathbb{R}^n$ with $g : U \to V$ an injective differentiable function with continuous partials and an invertible Jacobian for all $x \in U$. Then given continuous f with compact, connected support in $g(U)$, we have*

$$\int_{g(U)} f(x) = \int_U f \circ g(x) \cdot \left| \det(dg|_x) \right|$$

Proof. See [Sp, p. 66-72]. □

While this fact initially seems quite counterintuitive, careful thought gives good reason for the above formula. Consider a small rectangular prism in U with volume v; as long as it is non-degenerate, the vectors parallel to its sides form a basis for $\mathbb{R}^n$. For some x in this prism, we may approximate g at x as $g \approx T + dg|_x$, for some translation T. Then the action of g on this prism is (approximately) to transform the basis vectors parallel to its sides by $dg|_x$, inducing a new parallelepiped, which is non-degenerate if and only if $dg|_x$ is invertible. It is well-known, from computational geometry, that the volume of this new parallelepiped is $\left| \det dg|_x \right| \cdot v$. Considering the definition of the integral from Riemann sums, we have a geometrical motivation for this formula—the volume of each box in the summation is dilated by a factor of $\left| \det dg|_x \right|$.

2.3 Differential Forms

2.3.1 Motivation

The Change of Variables Theorem has an odd implication—that is, that integration is not coordinate-independent. In particular, diffeomorphic distortions of the coordinate system (that is, continuously differentiable and invertible maps, whose inverse is also continuously differentiable) change the integrals of maps, even though no information is added or lost. This is undesirable because there is no obvious reason why any particular coordinate system is "better" than any other.

Much of mathematics seeks to escape from this type of arbitrary choice—an analogous motivation gives the dot product. The dot product gives a coordinate-free definition of length and angle; similarly, we would like to define a concept of the integral that is invariant under as many diffeomorphisms as possible.

Intuitively, the idea is to construct a class of objects that contain information about how they behave in any given coordinate system. In particular, we wish them to have some notion of "infinitesimal" area at any given point, which can be transformed by diffeomorphisms—we wish to formalize Leibniz's notion of infinitesimals. (The notation we will use will reflect this intention.) The goal is to have such objects encode the change of variables theorem as closely as possible.

It is interesting to note that, as a byproduct of this discussion, we will provide a formal, mathematical motivation for the div, grad, and curl operators, which are usually motivated only physically. We will also provide a generalization of these operators, and justify the physical intuition that they are connected to one another through more than just notation.

2.3.2 Definitions

Let U be a domain in $\mathbb{R}^n$.

Definition 10 (Differential Forms)**.** A **differential k-form** on U is a continuous, infinitely differentiable map $\omega : U \to \bigwedge^k(\mathbb{R}^{n*})$, where $\mathbb{R}^{n*}$ is the dual of $\mathbb{R}^n$ as a vector space. The set of all differential k-forms on U is denoted $\Omega^k(U)$. A k-form ω is also said to be of **degree** k, denoted $\deg \omega = k$.

In particular, we may let $x_1, \ldots, x_n$ be a basis for $\mathbb{R}^n$, and let $dx_i \in \mathbb{R}^{n*}$ be the unique linear map $\mathbb{R}^n \to \mathbb{R}$ that satisfies

$$dx_i(x_j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j. \end{cases}$$

Then in this basis, we may write any differential k-form ω as

$$\omega(x) = \sum_{1 \leq i_1 < \ldots < i_k \leq n} f_{i_1,\ldots,i_k}(x) \cdot dx_{i_1} \wedge \cdots \wedge dx_{i_k},$$

for some $f_{i_1,\ldots,i_k}$. Intuitively, we may consider $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$ to be an oriented, k-dimensional volume element; this is the notation of "infinitesimal volume" we were looking for above. Note that this notion conforms geometrically with the properties of the exterior power: if one extends one dimension of a parallelepiped or formally sums two parallelepipeds, the volume changes linearly, and the orientation alters when one transposes two neighboring edges.

Differential forms admit a natural multiplication map $\wedge : \Omega^k(U) \times \Omega^l(U) \to \Omega^{k+l}(U)$, which is induced by the wedge product. However, this multiplication is *neither* antisymmetric nor symmetric; in particular, for $\omega \in \Omega^k(U), \alpha \in \Omega^l(U)$, we have

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

by reordering the dx_i.

Let V be a domain in $\mathbb{R}^k$; let g be a continuous, differentiable map $V \to U$.

Definition 11. The **pullback** of a k-form $\omega \in \Omega^k(U)$ through g, denoted $g^*(\omega) \in \Omega^k(V)$, is, with ω written as above, the map

$$g^*(\omega)(x) = \sum_{1 \leq i_1 < \ldots < i_k \leq n} f_{i_1,\ldots,i_k} \circ g(x) \cdot \bigwedge^k (dg|_x^*)(dx_{i_1} \wedge \cdots \wedge dx_{i_k})$$

where $dg|_x^*$ is the adjoint of the linear map $dg|_x$.

We claim that the pullback is the desired, (almost) coordinate-free transform we were looking for earlier. To see this, we must first define the integral of a differential form. In fact, for $U \subset \mathbb{R}^n$, we need only worry about n-forms, i.e. $\omega \in \Omega^n(U)$. Writing $\omega(x) = f(x) \cdot dx_1 \wedge \cdots \wedge dx_n$, we define

$$\int_U \omega := \int_U f(x)$$

where the integral on the right is the standard integral on real-valued functions. (Note the illustrative abuse of notation; if we write the term on the left out, we have

$$\int_U f \cdot dx_1 \wedge \cdots \wedge dx_n.$$

Omitting the wedges gives the standard notation from multivariate calculus.)

It is important to note that the sign of the integral is non-canonical; we have chosen an orientation of $\mathbb{R}^n$ by choosing an ordering of its basis vectors.

2.3.3 The Change of Variables Formula Revisited

We claim that the value of the integral is invariant under diffeomorphism, up to a sign. Let U, V be domains in $\mathbb{R}^n$, with $\omega \in \Omega^n(U)$. Consider a diffeomorphism $g : V \to U$. Then we claim

$$\int_U \omega = \pm \int_V g^*(\omega).$$

But writing out the term on the right according to our definitions, and writing $\omega(x)$ as $f(x) \cdot dx_1 \wedge \ldots \wedge dx_n$ gives us exactly

$$\begin{aligned}
\int_V g^*(\omega) &= \int_V f \circ g \cdot \bigwedge^k (dg|_x^*)(dx_{i_1} \wedge \cdots \wedge dx_{i_k}) \\
&= \int_V f \circ g \cdot \det(dg|_x^*) \cdot dx_{i_1} \wedge \cdots \wedge dx_{i_k} \\
&= \int_V f \circ g \cdot \det(dg|_x^*) \\
&= \int_V f \circ g \cdot \det(dg|_x),
\end{aligned}$$

where we use the fact that $\det(dg|_x) = \det(dg|_x^*)$. But this is exactly the change of variables formula without the absolute value sign. As g is a diffeomorphism, $dg|_x$ is always invertible and thus has nonzero derivative; continuity implies then that $\det(dg|_x)$ is everywhere-positive or everywhere-negative. So

$$\int_U \omega = \int_V f \circ g \cdot \left|\det(dg|_x)\right| = \pm \int_V f \circ g \cdot \det(dg|_x) = \pm \int_V g^*(\omega),$$

as claimed. We say that a diffeomorphism g is orientation-preserving if $\det(dg|_x)$ is everywhere-positive; in this case, we have strict equality above.

In fact, we may use this notion to redefine the notions of the line integral, the surface integral, and so on; in general, we may take the k-integral of a k-form ω on $U \subset \mathbb{R}^n$. Consider a domain $V \in \mathbb{R}^k$. Then the k-integral over a curve $g : V \to U$ is just the integral of $g^*(\omega)$ as defined above. Note that this is the integral of a k-form in $\mathbb{R}^k$, and is thus well-defined. This definition immediately gives invariance of the integral under appropriate re-parameterization, as before.

It is worth pausing here to examine what we have achieved. A careful reader might say that we have achieved *nothing*, at least insofar as pursuit of truth is concerned. We have just redefined some terms: the integral, and coordinate transformations, to be precise. We have replaced them with ideas that conform to normative notions we have about how the objects in question should behave. In some sense, this evaluation would be true from a purely epistemological view, but it would miss the pedagogical point. By restricting our attention to objects which are coordinate-free, we can examine the coordinate-independent properties of the objects they correspond to with much greater ease. The value of this labor will become clear as we develop this machinery in the next few sections.

2.3.4 The Exterior Derivative

We wish to define an operator on differential forms that is similar to the derivative; in particular, it should satisfy some analogue of the product rule, and in some sense be invariant under coordinate transformations. Taking our cue from the antisymmetry of the wedge product, we want to find a collection of operators (d^k) that satisfies

- d^k is a linear map $\Omega^k(U) \to \Omega^{k+1}(U)$.
- Given two differential forms $\omega \in \Omega^k(U), \alpha \in \Omega^l(U)$, $d^{k+l}(\omega\wedge\alpha) = d^k(\omega)\wedge\alpha+(-1)^k\omega\wedge d^l(\alpha)$. This is analogous to the product rule. This condition and the prior condition make d a **derivation of degree 1**.
- For f a 0-form, i.e. a function $U \to \mathbb{R}, U \subset \mathbb{R}^n$, d coincides with the derivative in the following sense: $d^0(f) = \sum_{i=1}^n \frac{\partial f}{\partial x_i} \cdot dx_i$. That is, in matrix form, $d^0(f)(x)$ is exactly $df|_x$ (albeit in a different space, which, having chosen the bases $\{x_1, \dots, x_n\}, \{dx_1, \dots, dx_n\}$, is non-canonically isomorphic to the usual space).
- $d^{k+1} \circ d^k = 0$ for all k.

In general, we omit the superscript and the parentheses; i.e. $d^k(\omega)$ is written $d\omega$, and we write d^k as simply d for all k; the last condition above would then be written $d \circ d = 0$.

Proposition 12. *The map d is uniquely defined by the above conditions.*

Proof. Consider the function $\chi_i : \mathbb{R}^n \to \mathbb{R}$ given by $(x_1, x_2, \dots, x_n) \mapsto x_i$; this function coincides with dx_i as defined above, but we use dx_i from here on to denote the constant differential form $x \mapsto 1 \cdot dx_i$, by (confusing) convention, just as we might use the constant c to denote the map $x \mapsto c$. Note that, by the third condition above, $d\chi_i = dx_i$. So by the fourth condition above, $d(dx_i) = 0$. We may now proceed to define d^k inductively, through the second condition above. In particular, it is clear that $\Omega^k(U)$ is spanned by the set of differential forms with a single term, e.g. $\omega = f(x) \cdot dx_{i_1} \wedge \dots \wedge dx_{i_k}$. But then

$$\begin{aligned} d^k(f(x) \cdot dx_{i_1} \wedge \dots \wedge dx_{i_{k-1}} \wedge dx_{i_k}) &= d^{k-1}(f(x) \cdot dx_{i_1} \wedge \dots \wedge dx_{i_{k-1}}) \wedge dx_{i_k} \\ &\quad + (-1)^{k-1} f(x) \cdot dx_{i_1} \wedge \dots \wedge dx_{i_{k-1}} \wedge d(dx_{i_k}) \\ &= d^{k-1}(f(x) \cdot dx_{i_1} \wedge \dots \wedge dx_{i_{k-1}}) \wedge dx_{i_k} \end{aligned}$$

and we may extend d linearly to linear combinations of single-term forms. While this proves uniqueness, it is not immediately clear that d is well-defined, as we must check that property (2) holds for all k, l, rather than just for $l = 1$.

To show that d is well-defined, we give an explicit construction that satisfies the inductive construction given above. In particular, for single-term forms $\omega(x) = f(x) \cdot dx_{i_1} \wedge \dots \wedge dx_{i_k}$, we let

$$d\omega = \sum_{j=1}^n \frac{\partial f}{\partial x_j} \cdot dx_j \wedge dx_{i_1} \wedge \dots \wedge dx_{i_k}.$$

The interested reader can check that this construction satisfies the first three conditions above; we check the fourth. The proof is inductive. We have that for 0-forms, e.g. $f(x)$, that

$$\begin{aligned} d \circ d(f) &= d\left(\sum_{i=1}^{n} \frac{\partial f}{\partial x_i} \cdot dx_i\right) \\ &= \sum_{j=1}^{n}\sum_{i=1}^{n} \frac{\partial^2 f}{\partial x_j \partial x_i} \cdot dx_j \wedge dx_i \\ &= \sum_{i<j} \left(\frac{\partial^2 f}{\partial x_i \partial x_j} dx_i \wedge dx_j + \frac{\partial^2 f}{\partial x_j \partial x_i} dx_j \wedge dx_i\right) \\ &= \sum_{i<j} \left(\frac{\partial^2 f}{\partial x_i \partial x_j} dx_i \wedge dx_j - \frac{\partial^2 f}{\partial x_j \partial x_i} dx_i \wedge dx_j\right) \\ &= 0, \end{aligned}$$

where the last step uses the equality of mixed partials. We have already noted that $d(dx_i) = 0$ above. Assume that for $i < p$, we have for $\omega \in \Omega^i(U), d \circ d(\omega) = 0$. Then for $\alpha \in \Omega^p(U)$ with one term, we may write $\alpha = \beta \wedge \gamma$ for $\beta \in \Omega^1(U), \gamma \in \Omega^{p-1}(U)$. We have

$$\begin{aligned} d \circ d(\alpha) &= d \circ d(\beta \wedge \gamma) \\ &= d(d\beta \wedge \gamma - \beta \wedge d\gamma) \\ &= d(d\beta \wedge \gamma) - d(\beta \wedge d\gamma) \\ &= d(d\beta) \wedge \gamma + d\beta \wedge d\gamma - d\beta \wedge d\gamma + \beta \wedge d(d\gamma) \\ &= 0 \end{aligned}$$

by the induction hypothesis; differential forms with more than one term satisfy the same claim by linearity. This completes the proof. □

Calculation gives that the exterior derivative commutes with the pullback, e.g. $g^*(d\omega) = d(g^*(\omega))$. That is, in some sense, the exterior derivative "flows" with changes of coordinates; for 0-forms, this is just the chain rule.

For the rest of this subsection, we will restrict our attention to $\mathbb{R}^3$. Note that, from Proposition 2, we have that the dimensions of $\bigwedge^0(\mathbb{R}^3), \bigwedge^1(\mathbb{R}^3), \bigwedge^2(\mathbb{R}^3)$, and $\bigwedge^3(\mathbb{R}^3)$ are $1, 3, 3$, and 1, respectively. In particular, we can identify $\bigwedge^0(\mathbb{R}^3)$ and $\bigwedge^3(\mathbb{R}^3)$ with $\mathbb{R}$ (actually the former is identified as such canonically), and $\bigwedge^1(\mathbb{R}^3), \bigwedge^2(\mathbb{R}^3)$ with $\mathbb{R}^3$. Then d gives maps $\mathbb{R} \to \mathbb{R}^3$, etc., and, by precomposition, maps $\nabla : C^\infty(\mathbb{R}^3) \to (\mathbb{R}^3 \to \mathbb{R}^3), \nabla\times : (\mathbb{R}^3 \to \mathbb{R}^3) \to (\mathbb{R}^3 \to \mathbb{R}^3), \nabla\cdot : (\mathbb{R}^3 \to \mathbb{R}^3) \to C^\infty(\mathbb{R}^3)$. Easy computation gives that these maps correspond, respectively, to the gradient, curl, and divergence. In fact, the first of these three computations follows immediately from the third bullet in the definition of the exterior derivative.

It is important to note that the identifications above are non-canonical; in the most general case, we define a canonical isomorphism called the **Hodge dual**, denoted

$$* : \bigwedge^k(V) \xrightarrow{\sim} \bigwedge^{n-k}(V),$$

where n is the dimension of V. Using this isomorphism, we may extend the ideas of gradient, curl, and divergence given above to vector spaces with arbitrary finite dimension.

2.3.5 The Interior Product and the Lie Derivative

Consider an element $w \in \bigwedge^k(V)$, where V is a finite-dimensional vector space over some field F. Then w can be viewed as an alternating, multilinear mapping $w : V^{*k} \to F$, where V^* is the vector space dual to V (here we take advantage of the canonical isomorphism $V \simeq V^{**}$). Given $a_1, \ldots, a_n \in V^*$, we may define, for each $\alpha \in V^*$, a function ι_α such that $\iota_\alpha(w)(a_1, \ldots, a_n) = w(\alpha, a_1, \ldots, a_n)$.

In particular, we may uniquely define ι_α as follows [Wa, p. 61]:

- $\iota_\alpha : \bigwedge^k(V) \to \bigwedge^{k-1}(V)$,
- For $v \in \bigwedge^1(V), \iota_\alpha(v) = \alpha(v)$,
- ι_α is a derivation of degree -1, i.e. $\iota_\alpha(a \wedge b) = \iota_\alpha(a) \wedge b + (-1)^{\deg a} a \wedge \iota_\alpha(b)$.

The proof that these properties uniquely define ι_α is analogous to the proof for d above and is left to the reader.

We now restrict our attention to $V = (\mathbb{R}^n)^*$. Consider a vector field $\xi : U \to \mathbb{R}^n$, where U is a domain in $\mathbb{R}^n$. Then for a differential form ω on U, we may let ι_ξ act on U point-wise, e.g. $\iota_\xi(\omega)(x) = \iota_{\xi(x)}(\omega(x))$. We define the **Lie Derivative** of a form ω with respect to a vector field ξ by

$$\mathrm{Lie}_\xi(\omega) := d \circ \iota_\xi(\omega) + \iota_\xi \circ d(\omega).$$

In some sense, this operator takes the derivative of a form with respect to a (possibly time-dependent) vector field. This intuition is clear for constant vector field; computation, which we omit, gives that for a constant vector field $\vec{x}$,

$$\mathrm{Lie}_{\vec{x}}(f \cdot dx_1 \wedge \cdots \wedge dx_n) = \frac{\partial f}{\partial x} \cdot dx_1 \wedge \cdots \wedge dx_n.$$

This fact will be useful later, and to remind ourselves of it, we will denote a constant vector field with respect to some coordinate x_i as $\frac{\partial}{\partial x_i}$.

2.4 Chain Complexes

Above, we noted that the exterior derivative satisfies $d \circ d = 0$. This fact suggests a more general structure, which we abstract as follows:

Definition 13 (Chain Complex)**.** A **chain complex** is a sequence of Abelian groups (or algebraic objects with Abelian structure, e.g. modules, vector spaces) $A_{-1}, A_0, A_1, A_2, \ldots$ with connecting homomorphisms $d^k : A_k \to A_{k-1}$ such that for all k, $d^k \circ d^{k+1} = 0$. We denote all of this data as $(A_\bullet, d_\bullet)$.

In a **cochain complex**, the connecting homomorphisms proceed in the opposite direction; i.e. $d^k : A_k \to A_{k+1}$ and $d^k \circ d^{k-1} = 0$. In this case, we denote the entire collection of Abelian groups and connecting homomorphisms as $(A^\bullet, d^\bullet)$.

Note that the chains and cochains are identical, but with different indexing; the terminology stems from convention. The notion of the pullback suggests the following:

Definition 14 (Map of Complexes)**.** A **map of complexes** $\psi^\bullet : (A^\bullet, d^\bullet) \to (B^\bullet, e^\bullet)$, in the case of a cochain, is a collection of maps $\psi^k : A_k \to B_k$ such that $e^k \circ \psi^k = \psi^{k+1} \circ d^k$, i.e. the diagram in Figure 2.1 commutes. The case of chains is analogous.

It should be clear by now that differential forms on some domain $U \subset \mathbb{R}^n$ form a complex

$$\cdots \longrightarrow 0 \longrightarrow \Omega^0(U) \xrightarrow{d} \Omega^1(U) \xrightarrow{d} \cdots \xrightarrow{d} \Omega^n(U) \longrightarrow 0 \longrightarrow \cdots.$$

and, from the fact that pullbacks commute with the exterior derivative, that pullbacks are maps of complexes. We call this complex the **de Rham complex** and denote the de Rham complex on U

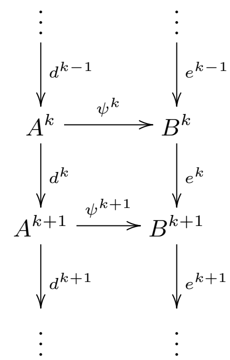

Figure 2.1: A Map of Complexes.

as $(\Omega^\bullet(U), d^\bullet)$. As in the de Rham derivative, we often ignore the superscripts on the connecting homomorphisms, e.g. $d \circ d = 0$, in arbitrary chains or cochains. Also, when we are discussing more than one complex, it is common to use the same symbol for their respective connecting homomorphisms, e.g. $d \circ \psi^k = \psi^{k+1} \circ d$.

Definition 15 (Closed, Exact)**.** Given an element $\omega \in A_k$, we say that ω is **closed** if $d\omega = 0$. We say that ω is **exact** if there exists α such that $d\alpha = \omega$.

Note that, as $d \circ d = 0$, we have that im $d^k \subset$ ker d^{k+1}; that is, all exact elements of a chain or cochain are closed. It is natural to ask when closed elements are exact—in the de Rham complex, the Poincaré Lemma addresses this question to a large extent. Pursuing it in general, we define:

Definition 16 (Homology, Cohomology)**.** The **k-th homology group** of a chain $(A_\bullet, d_\bullet)$ is

$$H_k(A_\bullet) := \frac{\ker d^k}{\text{im } d^{k+1}}.$$

Analogously, in a co-chain $(B^\bullet, d^\bullet)$, the **k-th cohomology group of** $B^\bullet$ is

$$H^k(B^\bullet) := \frac{\ker d^k}{\text{im } d^{k-1}}.$$

Intuitively, this group characterizes those closed forms that are not exact; i.e. elements in any given coset are identical up to an exact form.

Consider two co-chains $A^\bullet, B^\bullet$ and a map of complexes $\phi^\bullet : A^\bullet \to B^\bullet$. We claim that $\phi^\bullet$ induces well-defined maps $H^k(\phi^\bullet) : H^k(A^\bullet) \to H^k(B^\bullet)$. (An analogous claim holds for chains.)

Proof. Consider an element $[a] \in H^k(A^\bullet)$. We claim that the map $H^k(\phi^\bullet) : [a] \to [\phi(a)]$ is well-defined. By definition, any element $a' \in [a]$ differs from a by an exact element ω; as it is exact, we may write $\omega = d\alpha$ for some α. Then $[\phi(a')] = [\phi(a + d\alpha)] = [\phi(a) + \phi(d\alpha)] = [\phi(a) + d(\phi(\alpha))] = [\phi(a)]$, where we use the fact that maps of complexes commute with the complexes' connecting maps. So the map of cohomologies (resp. homologies) is well-defined. □

It is natural to ask when two maps of complexes induce the same map between cohomologies. To this end, we consider the following definition:

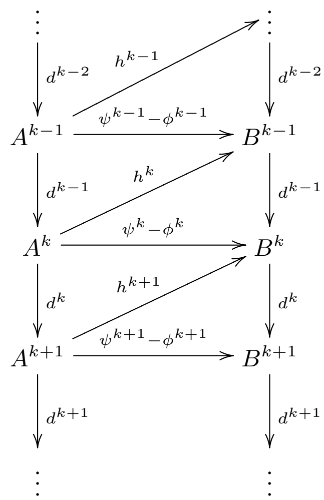

Figure 2.2: A Homotopy of Cochain Complexes.

Definition 17 (Homotopy)**.** In a (justifiable, as we shall see) homonym, we say that two maps $\psi^\bullet, \phi^\bullet : A^\bullet \to B^\bullet$ of complexes are **homotopic** through the **homotopy** (h^k) if there exists a sequence of maps $h^k : A^k \to B^{k-1}$ (in a co-chain, with analogous indexing for chains) such that

$$\psi^k - \phi^k = d \circ h^k + h^{k+1} \circ d.$$

That is, in the diagram in Figure 2.2, the parallelograms commute with the horizontal arrows.

Proposition 18. *If two maps of complexes $\psi^\bullet, \phi^\bullet$ are homotopic through some homotopy (h^k), then $H^k(\psi^\bullet) = H^k(\phi^\bullet)$.*

Proof. Note that, by linearity $H^k(\psi^\bullet) - H^k(\phi^\bullet) = H^k(\psi^\bullet - \phi^\bullet) = H^k(d \circ h^k + h^{k+1} \circ d) = H^k(d \circ h^k) + H^k(h^{k+1} \circ d)$. We claim that $H^k(d \circ h^k) = H^k(h^{k+1} \circ d) = 0$. To see that $H^k(h^{k+1} \circ d) = 0$, note that for $[a] \in H^k(A^\bullet)$, we have that $a \in \ker(d)$, so $H^k(h^{k+1} \circ d)([a]) = [h^{k+1} \circ d(a)] = 0$. Furthermore, $H^k(d \circ h^k)([a]) = [d(h^k(a))]$. But $d(h^k(a)) \in \operatorname{im}(d)$, so $[d(h^k(a))] = [0]$. But then $H^k(\psi^\bullet) - H^k(\phi^\bullet) = 0$, so $H^k(\psi^\bullet) = H^k(\phi^\bullet)$ as claimed. □

2.5 The Poincaré Lemma

We finally are able to state and prove the Poincaré Lemma. We wish to characterize situations in which closed forms are also exact.

Theorem 19 (The Poincaré Lemma)**.** *Let U be a contractible domain in $\mathbb{R}^n$, and let k be a positive integer. Then for $\omega \in \Omega^k(U)$ such that $d\omega = 0$, there exists $\alpha \in \Omega^{k-1}(U)$ such that $\omega = d\alpha$. In other words all closed differential k-forms on contractible domains are exact.*

Proof. We first prove a general lemma—that is, that the pullbacks through homotopic maps are homotopic as maps of complexes, as is suggested by the terminology.

Lemma 20. *Let V and W be domains, $V \subset \mathbb{R}^n, W \subset \mathbb{R}^m$. Consider maps $g_0, g_1 : V \to W$ that are homotopic, i.e. there is a map $G : I \times V \to W$, where $I = [0,1]$ such that $G(0,x) = g_0(x), G(1,x) = g_1(x)$. Then the maps of complexes $g_0^*, g_1^* : \Omega^k(W) \to \Omega^k(V)$ are homotopic.*

Proof of Lemma 20. Let $G_t : V \to W$ be the map $x \mapsto G(t,x)$. For $\omega \in \Omega^k(W)$, define

$$h^k(\omega)(x) = \int_{t=0}^{t=1} \iota_{\frac{\partial}{\partial t}}(G_t^*(\omega))(x).$$

We claim that this is the desired homotopy of complexes. In particular, we have that

$$\begin{aligned}
(d^{k-1} \circ h^k + h^{k+1} \circ d^k)(\omega) &= d\left(\int_{t=0}^{t=1} \iota_{\frac{\partial}{\partial t}}(G_t^*(\omega))\right) + \int_{t=0}^{t=1} \iota_{\frac{\partial}{\partial t}}(G_t^*(d\omega)) \\
&= \int_{t=0}^{t=1} (d \circ \iota_{\frac{\partial}{\partial t}} + \iota_{\frac{\partial}{\partial t}} \circ d)(G_t^*(\omega)) \\
&= \int_{t=0}^{t=1} \text{Lie}_{\frac{\partial}{\partial t}}(G_t^*(\omega)) \\
&= \int_{t=0}^{t=1} \frac{\partial}{\partial t} G_t^*(\omega) \\
&= G_1^*(\omega) - G_0^*(\omega) \\
&= g_1^*(\omega) - g_0^*(\omega),
\end{aligned}$$

as desired. In the above manipulations, we use the commutation of the differential with the integral and the pullback as well as the fundamental theorem of calculus. □

Corollary 21. *The pullbacks through homotopic maps act identically on the cohomologies, that is, $H^k(g_0^*) = H^k(g_1^*)$. In particular, on a contractible domain, $H^k(id^*) = H^k(c^*)$, where c is the constant map.*

But $H^k(c^*) = 0$. So we have from the corollary that, on contractible domains, $H^k(\text{id}^*) = 0$. But then $H^k(\Omega^k(U)) = 0$, i.e. $\text{im}\, d = \ker d$. And this is precisely what we wanted to prove. □

2.6 Conclusion

It is valuable to consider what, if anything, we have accomplished beyond the Lemma itself. In particular, the ideas here seem somewhat far-removed from those where we started—in the realm of coordinate-invariant objects. What does the Poincaré Lemma tell us? What have we gained by introducing such strange, if elegant, mathematical tools?

To begin with, many more standard proofs of the Lemma are heavily calculational [Sp]; the referenced method proves the theorem only on star-shaped domains, and at the cost of massive amounts of counterintuitive index-juggling. The methods here slightly weaken the hypothesis on the domain and achieve a much cleaner solution.

But more importantly, the tools we have developed have varied applications. One of the better-known such applications occurs in electrodynamics, where Maxwell's equations tell us that, under magneto-static conditions,

$$\nabla \times \vec{\mathbf{E}} = 0,$$

where $\vec{\mathbf{E}}$ denotes the electric field. The Poincaré Lemma implies immediately that there exists a scalar function V such that

$$\vec{\mathbf{E}} = -\nabla V,$$

that is, the electric potential.

2.7 Acknowledgment

The vast majority of the information presented here was demonstrated to the author in some form by Professor Dennis Gaitsgory of Harvard University, and the methods herein owe much to his presentation.

References

[DF] David S. Dummit and Richard M. Foote: *Abstract Algebra*, 3rd Edition. New York: John Wiley and Sons, 2004.

[Sp] Michael Spivak: *Calculus on Manifolds*. New York: Perseus Books Publishing, 1965.

[Wa] Frank W. Warner: *Foundations of Differentiable Manifolds and Lie Groups.* New York: Springer Science+Business Media, 1983.