WHAT IS COHOMOLOGY?

ARUN DEBRAY
MARCH 2, 2016

ABSTRACT. Cohomology is a very powerful topological tool, but its level of abstraction can scare away interested students. In this talk, we'll approach it as a generalization of concrete statements from vector calculus, which allows a definition of cohomology which is just as precise, but easier to grasp. This talk should be understandable to students who have taken linear algebra and vector calculus classes.

1. THE THREE STOOGES: DIV, GRAD, AND CURL

You might remember the following theorems from vector calc.

Theorem 1.

- *If f is a smooth function on $\mathbb{R}^3$,* $\operatorname{curl}(\nabla f) = 0$.
- *If $\mathbf{v}$ is a smooth vector field on $\mathbb{R}^3$,* $\operatorname{div}\operatorname{curl}\mathbf{v} = 0$.

These can be proven with a calculation, but what was more interesting is that the converse is also true.

Theorem 2. *Let $\mathbf{v}$ be a smooth vector field on $\mathbb{R}^3$.*

- *If* $\operatorname{curl}\mathbf{v} = 0$*, then* $\mathbf{v} = \nabla f$ *for some function f.*
- *If* $\operatorname{div}\mathbf{v} = 0$*, then* $\mathbf{v} = \operatorname{curl}\mathbf{w}$ *for some vector field $\mathbf{w}$.*

It turns out that Theorem 1 is true in considerably more generality, but Theorem 2 doesn't generalize well.

Example 3. One way we can generalize is to open subsets of $\mathbb{R}^3$. Let $U = \mathbb{R}^3 \setminus \{(0, 0, z)\}$, so we've removed the z-axis, and consider the vector field

$$\mathrm{d}\theta = \left(\frac{-y}{x^2+y^2}, \frac{x}{x^2+y^2}, 0\right).$$

Then,

$$\operatorname{curl}(\mathrm{d}\theta) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial_x & \partial_y & \partial_z \\ \frac{-y}{x^2+y^2} & \frac{x}{x^2+y^2} & 0 \end{vmatrix} = 0\mathbf{i} + 0\mathbf{j} + \frac{(1-1)(\text{stuff})}{(x^2+y^2)^2}\mathbf{k} = 0.$$

However, $\mathrm{d}\theta$ is not a conservative vector field: if S^1 denotes the unit circle in the xy-plane, then $\oint_{S^1} \mathrm{d}\theta = 2\pi$. For conservative vector fields, the integral across any loop vanishes, so this means $\mathrm{d}\theta \neq \nabla f$ for any f.

Though this feels unfortunate, since Theorem 2 was a nice theorem, this actually is the seed of something quite powerful: it tells us that calculus on $\mathbb{R}^3$ is different than calculus on U, and is the start of understanding something called de Rham cohomology.

Next, one might ask, how many vector fields violate Theorem 2?

- This isn't quite the right question to ask: any scalar multiple of $\mathrm{d}\theta$ has the same property, and if I take $\mathrm{d}\theta + \nabla f$, I'll usually get something with the same property.
- There are a lot of these, but it is a fact that any two such vector fields differ by a conservative vector field. In this sense there's "only one" exception.

A big motivation for what we're about to do is to make this intuition precise.

2. Manifolds

The first step in a story is to provide the setting. We'd like to make formal the notion of a "space where you can do calculus."

- We can start with $\mathbb{R}^n$ or open subsets of it, which have a set of coordinates $(x_1, \dots, x_n)$ that allow us to do calculus.
- There may be more than one system of coordinates; for example, we can look at $\mathbb{R}^2$ minus the nonnegative x-axis with polar coordinates, or with rectangular coordinates. In each case, we've restricted the coordinates to only part of $\mathbb{R}^2$. The important fact is that the change-of-coordinates map is smooth; this means that notions of calculus (integrability, differentiability, etc.) are unchanged.
- The final step is to allow coordinates to be local: that is, we can use several different coordinate systems on different parts of a space, where each is only defined "locally," meaning on some patch of the space. Again, we need the transition maps to be smooth, so we can do calculus in different coordinate systems, and we want to have coordinates near any point.

Definition. A *manifold* is a space with a system of local coordinates such that

(1) every point is in some local coordinate system, and
(2) if a point is in two different local coordinate systems, their transition map is smooth.

This is intentionally a little vague; it can be made precise, but that takes a bit longer and doesn't add much more in the way of intuition.

For example, $\mathbb{R}^n$, $\mathbb{R}^2 \setminus 0$ with rectangular and different kinds of polar coordinates, spheres, and the space of lines in $\mathbb{R}^2$.

3. Differential Forms and the Exterior Derivative

To generalize functions and vector fields, we need to pin down exactly what it means to say "a thing that can be integrated." On $\mathbb{R}^n$, there are n particularly simple things to integrate against: $\mathrm{d}x_1, \dots, \mathrm{d}x_n$, which are just "integrating in the x_i-direction." We can also take combinations of these, e.g. $\mathrm{d}x_1 \mathrm{d}x_2$ means integrating on the $x_1 x_2$-plane, and we can integrate functions with them. Note that $\mathrm{d}x_1\,\mathrm{d}x_2 \neq \mathrm{d}x_2\,\mathrm{d}x_1$; the first one specifies the standard orientation, and the second specifies the opposite orientation.

This suggests that "things which can be integrated" are of the form $f(x)\,\mathrm{d}x_{i_1}\,\mathrm{d}x_{i_2} \cdots \mathrm{d}x_{i_k}$, where f is a smooth function, and indeed, this defines a *differential k-form* on $\mathbb{R}^n$. (We'll also allow finite sums of these.) A manifold might have more than one notion of coordinates, and so to define a differential k-form on a manifold, we simply ask that in each coordinate system, it look like a differential k-form on $\mathbb{R}^n$. In this more general setting, $f(x)\,\mathrm{d}x_{i_1} \cdots \mathrm{d}x_{i_k}$ is generally notated $f(x)\,\mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}$; the why is beyond the scope of this talk, but it might help you remember that the order of forms matters: you can switch the order of two things, but you have to add a minus sign.

Differential k-forms are a vector space $\Omega^k(M)$. We allow $k = 0$: a 0-form is just a smooth function.

We thought of this as something that can be integrated, and indeed, a k-form can be integrated on something that's k-dimensional. For example, if I have a curve γ and a 1-form $f\,\mathrm{d}x + g\,\mathrm{d}y$, I can certainly integrate $\int_\gamma f\,\mathrm{d}x + g\,\mathrm{d}y$ just as in multivariable calculus.

We can also differentiate forms, using an operator called the *exterior derivative*, or d. This is defined according to the following rules.

- If $f \in \Omega^0(M)$, then in any coordinate system $(x_1, \dots, x_n)$,
$$\mathrm{d}f = \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}\,\mathrm{d}x_i.$$
- More generally, in a coordinate system $(x_1, \dots, x_n)$, the exterior derivative of $\omega = f(x)\,\mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}$ is
$$\mathrm{d}\omega = \mathrm{d}f \wedge \mathrm{d}x_{i_1} \wedge \cdots \wedge \mathrm{d}x_{i_k}.$$
That is, expand $\mathrm{d}f$ according to the first rule, and then
- cancel out any $\mathrm{d}x_i$ that appears two or more times in a monomial.

Thus, the derivative of a k-form is a $(k+1)$-form.

For example,

$$d(x^2\,dy + yz\,dz) = 2x\,dx \wedge dy + z\,dy \wedge dz + y\,dz \wedge dz = 2x\,dx \wedge dy + z\,dy \wedge dz.$$

Lemma 4. $d \circ d = 0$ *on any manifold.*

This is something you can work out purely algebraically. The intuition is that the boundary of a boundary is zero, or that this should generalize Theorem 1. In fact, that theorem is a special case.

- If $f \in \Omega^0(\mathbb{R}^3)$, so f is a function, then $df = \sum \frac{\partial f}{\partial x_i}\,dx_i$. We can think of a vector field $\mathbf{v} = (v_x, v_y, v_z)$ as $\mathbf{v} = v_x\,dx + v_y\,dy + v_z\,dz \in \Omega^1(\mathbb{R}^3)$, so $d : \Omega^0(\mathbb{R}^3) \to \Omega^1(\mathbb{R}^3)$ is *gradient*.
- The same identification of vector fields and one-forms means that if I have $\mathbf{v} = v_x\,dx + v_y\,dy + v_z\,dz$, then (after some calculation)
$$d\mathbf{v} = \left(\frac{\partial v_z}{\partial y} - \frac{\partial v_y}{\partial z}\right) dy \wedge dz + \left(\frac{\partial v_x}{\partial z} - \frac{\partial v_z}{\partial x}\right) dx \wedge dz + \left(\frac{\partial v_x}{\partial y} - \frac{\partial v_y}{\partial x}\right) dx \wedge dy.$$
Now, we have to make a different identification of a vector field (w_x, w_y, w_z) to the 2-form $w_x\,dy \wedge dz + w_y\,dx \wedge dz + w_z\,dx \wedge dy$, but when we do this, $d : \Omega^1(\mathbb{R}^3) \to \Omega^2(\mathbb{R}^3)$ is *curl*.
- Finally, suppose $\mathbf{v} = (v_x, v_y, v_z)$ is a vector field regarded as a 2-form. Then,
$$d\mathbf{v} = \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z}\right) dx \wedge dy \wedge dz.$$
We can think of a 3-form on $\mathbb{R}^3$ as just a function, by thinking of $f\,dx \wedge dy \wedge dz$ as just f, and when we do, we see that $d : \Omega^2(\mathbb{R}^3) \to \Omega^3(\mathbb{R}^3)$ is *divergence*.

Now, we can better understand the question we asked at the beginning: for a given manifold M, to what degree does Lemma 4 have a converse? In other words, if $d\omega = 0$, is it true that $\omega = d\alpha$ for some α of one degree lower?

Definition. Let ω be a differential form.

- ω is called *closed* if $d\omega = 0$.
- ω is called *exact* if $\omega = d\alpha$ for some form α.

In fewer words: is every closed form exact? We know it's true on $\mathbb{R}^3$.

4. de Rham Cohomology

We have a bunch of maps $d : \Omega^k(M) \to \Omega^{k+1}(M)$; we can put them together into a diagram, called the *de Rham complex*: the objects are vector spaces and the arrows are linear maps.

$$\Omega^0(M) \xrightarrow{d^0} \Omega^1(M) \xrightarrow{d^1} \Omega^2(M) \xrightarrow{d^2} \cdots$$

Here, since the exterior differential is technically a different map at each degree, I've given each one a different name.

The key property of this diagram is that $d^{k+1} \circ d^k = 0$, so $\operatorname{Im}(d^k) \subset \ker(d^{k+1})$.[1] The difference in their sizes is exactly what we were looking for! Hence, we can define the k^{th} *de Rham cohomology* to be the vector space $H^k(M) = \ker(d^{k+1}) / \operatorname{Im}(d^k)$.[2]

You might wonder why one defines a whole vector space instead of just its dimension, but this useful for the following very important reason: given a smooth map $f : M \to N$ between manifolds, one can *pull back* forms on N to forms on M: given a form ω on N, we can define a form on M by "first doing f, then ω;" after defining this formally one can check that closed forms are sent to closed forms and exact forms are sent to exact forms, so this pullback induces a linear map $f^* : H^k(N) \to H^k(M)$ for each k. Thus, we have an invariant not just for manifolds, but also for smooth maps — properties of f^* can be used to learn a lot about f.[3]

So let's break this definition down into smaller parts. An $a \in H^k(M)$ is a coset of objects in $\ker(d^{k+1})$, so a bunch of closed k-forms whose differences are exact. Counting these objects, or their dimension, is what we were getting at earlier: if any two closed-but-not-exact forms differ by an exact form, it's easy to obtain one from the other, so they're really not that different.

---

[1] In your linear algebra class, the *image* $\operatorname{Im}(f)$ may have been called the column space, and the *kernel* $\ker(f)$ may have been called the null space. There's no difference.

[2] If you haven't seen quotient spaces before, think of this as the orthogonal complement to $\operatorname{Im}(d^k)$ in the space $\ker(d^{k+1})$.

[3] If you like categories, this pullback defines a functor from the category of smooth manifolds to the category of real vector spaces for each k.

Example 5. Let's compute this for $\mathbb{R}^3$. First, if $k > 3$, then any k-form has some x_i appearing twice in it, so $\Omega^k(\mathbb{R}^3) = 0$. This makes sense: we've run out of things to integrate.

The de Rham complex is

$$\Omega^0(\mathbb{R}^3) \xrightarrow{\nabla} \Omega^1(\mathbb{R}^3) \xrightarrow{\text{curl}} \Omega^2(\mathbb{R}^3) \xrightarrow{\text{div}} \Omega^3(\mathbb{R}^3).$$

By Theorem 2, this sequence is *exact*: at each object, the image of the inward arrow is equal to the kernel of the outward arrow. Thus, whenever there are two arrows, the cohomology is $\ker(\mathrm{d}^i)/\ker(\mathrm{d}^i) = 0$, a single point (which is a zero-dimensional vector space). Thus, $H^1(\mathbb{R}^3) = H^2(\mathbb{R}^3) = 0$. $H^3(\mathbb{R}^3) = 0$ too, since $\Omega^4(\mathbb{R}^3) = 0$, so everything has to be in the kernel, but a function is the divergence of its antiderivative, so everything is in the image.[4]

However, $H^0(\mathbb{R}^3)$ is nonzero: there's nothing in the image of d^{-1}, and $\ker(\mathrm{d}^0)$ consists of the constant functions. In particular, $H^0(\mathbb{R}^3) = \ker(\mathrm{d}^0)/0 \cong \mathbb{R}$ is a one-dimensional vector space.

5. Conclusion

This is a small step into a big world of algebraic topology; there is much to say here. Check out Bott and Tu's *Differential Forms in Algebraic Topology* or Madsen and Tornehave's *From Calculus to Cohomology* for the next steps in this journey. I'll end by mentioning a few things that cohomology can be useful for.

- If two manifolds are "the same," their cohomology is also the same. This turns out to also be true for a more flexible notion of sameness called *homotopy*. One use of cohomology is to understand a given manifold, e.g. by showing that it isn't homotopic to things we already know about.
- One very valuable principle is that cohomology classifies "obstructions:" if a certain cohomology group vanishes, then we have some nice thing that's true. For example, if U is a submanifold of $\mathbb{R}^3$, then the existence of a nontrivial class in $H^1(U)$ is equivalent to there being a vector field whose curl is zero, but that isn't conservative. The cohomology $H^{\dim(M)}(M)$ classifies whether a manifold is orientable.
- I mentioned it before, but the fact that a function on manifolds induces a linear map on cohomology is extremely useful: it's used to calculate cohomology, to understand fixed points of functions, to determine when two functions are "the same," and more.

[4]This is not completely true as stated: not all functions on $\mathbb{R}^3$ have antiderivatives. It suffices for f to be zero outside of some bounded set, and in fact these are really the functions you get when you identify $\Omega^3(\mathbb{R}^3)$ with smooth functions. This is a subtle point, so I'm not going to dwell on it.