Day 5: Chern-Weil theory

Arun Debray

July 10, 2020

Today's plan

1. The generalized splitting principle and examples
2. Connections and curvature for principal G-bundles
3. Invariant polynomials and the Chern-Weil homomorphism
4. Examples of the Chern-Weil homomorphism

The generalized splitting principle: setup

- Fix our Lie group G and a maximal torus T of rank n (i.e. $T \cong \mathbb{T}^n$)
- Via the inclusion $i\colon T \hookrightarrow G$, we have models of BG as EG/G and BT as EG/T, so $Bi\colon BT \to BG$ is a fiber bundle with fiber G/T
  - (Note: T is generally not normal in G)
- Let $P \to X$ be a principal G-bundle classified by a map $f_P\colon X \to BG$, and $q\colon Y \to X$ be the pullback of Bi:

$$\begin{array}{ccc} Y & \xrightarrow{g} & BT \\ \downarrow{\scriptstyle q} & & \downarrow{\scriptstyle Bi} \\ X & \xrightarrow{f_P} & BG. \end{array}$$

The generalized splitting principle

- Here's the diagram again:

$$\begin{array}{ccc} Y & \xrightarrow{g} & BT \\ \downarrow{\scriptstyle q} & & \downarrow{\scriptstyle Bi} \\ X & \xrightarrow{f_P} & BG. \end{array}$$

- Theorem, part 1: there is a canonical reduction of structure group of $q^*P \to Y$ to T
- Theorem, part 2: $q^* : H^*(X; \mathbb{Q}) \to H^*(Y; \mathbb{Q})$ is injective

So?

- Suppose c is a characteristic class for principal G-bundles
- Via Bi, it also defines a characteristic class for principal T-bundles
- Since q^* is injective, that characteristic class for the principal T-bundle we obtained over Y determines $c(P)$
- Since $T \cong \mathbb{T}^n$, that principal T-bundle decomposes (in a sense) as a product of n principal $\mathbb{T}$-bundles
- Therefore the characteristic class also factors as a product $\prod_{i=1}^{n}(1+x_i)$, where the x_i are the c_1s of the principal $\mathbb{T}$-bundle summands. The x_i are called the *roots* of P

Example: Chern classes

- The diagonal matrices are a maximal torus in U_n of rank n
- Using associated bundles to pass between principal U_n-bundles and complex vector bundles, this tells us that a complex vector bundle $V \to X$ splits as a sum of line bundles $L_1, \ldots, L_n$ when pulled back to Y
- $c_1(L_i)$ is called the i^{th} *Chern root*, and $c_k(V)$ is the k^{th} symmetric polynomial in the Chern roots
- $H^*(\mathrm{U}_n; \mathbb{Z})$ is free, so we can work over $\mathbb{Z}$

Chern classes and the flag manifold

- Y has a more concrete description in this case
- Namely, the *flag manifold* for $V \to X$
- A *flag* of an inner product space W is a decomposition of W as a sum of one-dimensional, orthogonal subspaces
- The flag manifold $Y \to X$ is a fiber bundle whose fiber at $x \in X$ is the space of flags of V_x
  - (Ok, you need a Hermitian metric to define this, but the isomorphism type of the fiber bundle does not depend on this choice)

Example: Pontrjagin classes, part 1

- $G = \mathrm{SO}_{2n}$: one maximal torus is the diagonal matrices in $\mathrm{U}_n \subset \mathrm{SO}_n$, which has rank n
- Upshot: if V is an oriented rank-$2n$ vector bundle, q^*V splits as a sum of complex line bundles $L_1, \ldots, L_n$, but the symmetric polynomial gets squared:

$$p_i(q^*V) = \sigma_i(c_1(L_1), \ldots, c_1(L_n))^2.$$

- This is because the Pontrjagin classes of V are the Chern classes of $V_{\mathbb{C}}$, and the Chern roots of $V_{\mathbb{C}}$ come in pairs $\pm x_1, \ldots, \pm x_n$
- The Euler class also splits:

$$e(q^*V) = \sigma_n(c_1(L_1), \ldots, c_n(L_n))$$

Example: Pontrjagin classes, part 2

- $G = \mathrm{SO}_{2n+1}$ has a similar story: one maximal torus is the diagonal matrices in $\mathrm{U}_n \subset \mathrm{SO}_{2n+1}$ (so the last diagonal entry is always 1)
- So an oriented rank-$(2n+1)$ real vector bundle V, pulled back to Y, splits as a direct sum of n complex line bundles and a trivial real line bundle
- The Pontrjagin and Euler classes of V admit the same description as in the case SO_{2n}

Example (sort of): Stiefel-Whitney classes

- O_n isn't connected, so we can't use the theorem
- Nonetheless, enough of the structure persists with $\mathbb{Z}/2$ coefficients and the subgroup $O_1^n \subset O_n$ to prove something via similar methods
- One can prove q^* is an injection on mod 2 cohomology and q^*P admits a canonical reduction of structure group to a principal O_1^n-bundle
- Upshot: a rank-n real vector bundle V, after pullback to Y, splits as a direct sum of n real line bundles $L_1, \ldots, L_n$, and

$$w_k(V) = \sigma_i(w_1(L_1), \ldots, w_1(L_n))$$

Invariant polynomials

- G is a Lie group, $\mathfrak{g} := T_e G$ is its Lie algebra
- G acts on $\mathfrak{g}$ by the *adjoint representation*, differentiating the action of G on itself by conjugation
  - Example: $G = \mathrm{GL}_n(\mathbb{R})$, $\mathfrak{g}$ is all $n \times n$ matrices, and the adjoint representation is matrix conjugation
  - If G is a subgroup of $\mathrm{GL}_n(\mathbb{R})$, $\mathfrak{g}$ is a subspace of all $n \times n$ matrices, and the adjoint representation is also conjugation

Invariant polynomials

- The vector space of polynomial functions on $\mathfrak{g}$ is $\mathrm{Sym}^\bullet(\mathfrak{g}^*)$
  - That is, take the dual vector space $\mathfrak{g}^*$, and take symmetric functions on it
  - Why does this get to be called polynomials? There is an identification $\mathrm{Sym}^\bullet((\mathbb{R}^n)^*) \cong \mathbb{R}[x_1, \ldots, x_n]$, but it uses the basis, so using $\mathrm{Sym}^\bullet(\mathfrak{g}^*)$ is basis-independent
- The *invariant polynomials* on $\mathfrak{g}$ are the subspace of $\mathrm{Sym}^\bullet(\mathfrak{g}^*)$ fixed by the G-action. This subspace is denoted $I^\bullet(G)$
  - Warning: $I^\bullet(G)$ is graded strangely: a degree-q polynomial carries grading $2q$

Examples of invariant polynomials

- G connected and abelian (e.g. $\mathbb{R}^n$, $\mathbb{T}^n$): conjugation is trivial, so all polynomials are invariant
- $G = \mathrm{U}_2$ (2×2 unitary complex matrices), with Lie algebra $\mathfrak{u}_2$ (2×2 complex matrices satisfying $A = -A^*$)
  - Both $\mathrm{tr}(X)$ (degree 2) and $\mathrm{tr}(X^2)$ are U_2-invariant polynomials on $\mathfrak{u}_2$
  - In fact, this is everything! $I^\bullet(\mathrm{U}_2) \cong \mathbb{R}[\mathrm{tr}(X), \mathrm{tr}(X^2)]$
  - The same is true for $\mathrm{GL}_2(\mathbb{C})$
- $G = \mathrm{SU}_2$: $\mathfrak{su}_2$ consists of traceless matrices, so $I^\bullet(\mathrm{SU}_2) = \mathbb{R}[\mathrm{tr}(X^2)]$

Examples of invariant polynomials

- $G = \mathrm{SO}_{2n}$: Lie algebra $\mathfrak{so}_{2n}$ is skew-symmetric real $2n \times 2n$ matrices
  - We get a bunch of invariant polynomials f_{4i} of degree $4i$ via

$$\det(t - X) = \sum_{i=0}^{n} f_{4i}(X) t^{2i}$$

    because the determinant is invariant by conjugation by SO_{2n} (similarly for U_n)
  - Also a new one:
    - $X \in \mathfrak{so}_{2n}$ iff, as a map $X : \mathbb{R}^n \to (\mathbb{R}^n)^*$, $X^* = -X$
    - In $\mathrm{Hom}(\mathbb{R}^n, (\mathbb{R}^n)^*) = (\mathbb{R}^n)^* \otimes (\mathbb{R}^n)^*$, X is alternating, so identified with $\omega_X \in \Lambda^2(\mathbb{R}^n)^*$
    - $\mathrm{Det}(\mathbb{R}^n) := \Lambda^n(\mathbb{R}^n)^*$ is one-dimensional, so

$$\frac{\omega_X^n}{n!} = P_n(X)\mathrm{vol},$$

      where vol is the volume form induced by the inner product
    - P_n is a degree-n invariant polynomial, and $P_n(X^2) = \det(X)$, so $I^\bullet(\mathrm{SO}_{2n})$ is not a polynomial ring

Connections on principal G-bundles: background

- We'll need a more abstract perspective on connections today
- This perspective is useful in other places in geometry
- Refresher: If $E \to M$ is a vector bundle, $\Omega^k_M(E) := \Gamma(\Lambda^k T^*M \otimes E)$
  - If V is a vector space, $\Omega^k_M(V) := \Omega^k_M(\underline{V})$. $\Omega^k_M := \Omega^k_M(\mathbb{R})$

Connections on principal G-bundles

- Let $\pi : P \to M$ be a principal G-bundle over a smooth manifold M
- We know what it means for a tangent vector $v \in T_pP$ to be *vertical*: $d\pi|_p(v) = 0$
- This defines a short exact sequence of vector bundles

$$0 \to \ker(\pi_*) \to TP \to \pi^*TM \to 0$$

- A *connection* is a G-invariant splitting of this sequence
  - i.e., a choice of what it means for tangent vectors to P to be horizontal

The connection 1-form

- There is an identification $\ker(\pi_*) \cong \underline{\mathfrak{g}}$, and a section $\pi^* TM \to TP$ is equivalent to a section $TP \to \ker(\pi^*)$, i.e. an element of $\Gamma_P(T^*P \otimes \ker(\pi_*)) = \Omega^1_P(\mathfrak{g})$
- This element is denoted Θ and called the *connection* 1-*form*; the horizontal vectors are $\ker(\Theta)$

The curvature 2-form

- Tensor the de Rham differential $\mathrm{d}\colon \Lambda^k T^*P \to \Lambda^{k+1} T^*P$ with $\mathfrak{g}$ to obtain $\mathrm{d}\colon \Omega^k_P(\mathfrak{g}) \to \Omega^{k+1}_P(\mathfrak{g})$
- The wedge product has the form $\Omega^k_P(\mathfrak{g}) \times \Omega^\ell_P(\mathfrak{g}) \to \Omega^{k+\ell}_P(\mathfrak{g} \otimes \mathfrak{g})$
- The Lie bracket $[-,-]\colon \mathfrak{g} \otimes \mathfrak{g} \to \mathfrak{g}$ induces $\Omega^k_P(\mathfrak{g} \otimes \mathfrak{g}) \to \Omega^k_P(\mathfrak{g})$
- Therefore we may define the *curvature* of Θ to be

$$\Omega := \mathrm{d}\Theta + [\Theta \wedge \Theta] \in \Omega^2_P(\mathfrak{g})$$

Descending the curvature to M

- Under the G-action on P, Ω transforms in the adjoint representation
- Therefore it descends to M, but valued in the *adjoint bundle* $\mathfrak{g}_P \to M$, the associated bundle $P \times_G \mathfrak{g} \to M$: $\Omega \in \Omega^2_M(\mathfrak{g}_P)$
- If $G = \mathrm{GL}_n(\mathbb{R})$, so that P is associated to a rank-n vector bundle $V \to M$, $\mathfrak{g}_P = \operatorname{End} V$ (idea: $\mathfrak{g} = \operatorname{End}(\mathbb{R}^n)$), and this recovers the description of a connection from Monday

The Chern-Weil homomorphism

- Fix G and a principal G-bundle $P \to M$ (M is a smooth manifold)
- The Chern-Weil homomorphism is a map $I^\bullet(G) \to \Omega^\bullet(M)$
- $f \mapsto \omega_f := f(\Omega^{\wedge(|f|)})$
- Typechecking the degree: Ω is degree 2, and $|f|$ is in $I^\bullet(G)$ (so twice the polynomial degree), so this is a degree-preserving map

Nice properties of the Chern-Weil map

- Thus far, we've only defined ω_f as an element of Ω_P^*
- To show it descends to M, argue that contraction $\iota_\zeta \omega_f$ with any vertical vector field ζ gives zero
  - Proof is a computation, using that $\iota_\zeta \Omega = 0$
- ω_f is closed: compute $\mathrm{d}\omega_f$
  - Then follows from the *Bianchi identity* $\mathrm{d}\Omega + [\Theta \wedge \Omega] = 0$
- de Rham cohomology class does not depend on the choice of connection
- Naturality: given $\phi : N \to M$, $\phi^* \omega_f(P) = \omega_f(\phi^* P)$ (so long as you use the pullback connection on $\phi^* P$)

Universal targets

- Naturality suggests maybe we could do this once, on a moduli space $B_\nabla G$ for principal bundles with connection
- Problem: the functor sending a manifold to the stack of principal G-bundles with connection is not representable. That is, no such space $B_\nabla G$ exists
  - Note: it does exist as a stack, in a sense; occasionally this is useful (Freed-Hopkins)
- Nonetheless, if $\pi_0(G)$ is finite, there is a model for $EG \to BG$ in which both are "Hilbert manifolds," i.e. possibly infinite-dimensional manifolds locally modeled on separable Hilbert spaces
  - This is enough structure to run the Chern-Weil argument!

When is the Chern-Weil homomorphism an isomorphism?

- Theorem: if G is compact, $I^* \bullet (BG) \to H^*(BG; \mathbb{R})$ is an isomorphism
- Proof begins with the case G is connected and abelian: easier since the adjoint representation is trivial
- Then relax abelian assumption: choose a maximal torus $T \subset G$, use the Serre spectral sequence to pass information from T to G
- Finally, relaxing connectedness is another Serre spectral sequence argument (!), using the short exact sequence $1 \to G_0 \to G \to \pi_0(G) \to 1$

$G = \mathbb{T}$

- We know what to expect at the end: $H^*(B\mathbb{T};\mathbb{R}) = H^*(\mathbb{CP}^\infty;\mathbb{R}) = \mathbb{R}[c_1]$, $|c_1| = 2$
- The Lie algebra of $\mathbb{T}$ is $i\mathbb{R}$, with trivial Lie bracket because $\mathbb{T}$ is abelian
- Thus $I^\bullet(\mathbb{T}) = \mathbb{R}[x]$, with $|x| = 2$
- The Chern-Weil homomorphism sends $P, \Theta \mapsto [\Omega] \in H^2_{\mathrm{dR}}(M)$, which is (a nonzero multiple of) $c_1(P)$

$G = \mathrm{SU}_2$

- $H^*(B\mathrm{SU}_2;\mathbb{R}) \cong \mathbb{R}[c_2]$, $|c_2| = 4$
- $I^\bullet(\mathrm{SU}_2) = \mathbb{R}[s]$, where $s(A) := \mathrm{tr}(A^2)$
- The Chern-Weil map is an isomorphism, as before
- Both of these can be proven by studying transgression in the Serre spectral sequence for $G \to EG \to BG$

$G = \mathrm{SL}_2(\mathbb{C})$

- G is noncompact; maximal compact is SU_2, so $H^*(BG;\mathbb{R}) \cong H^*(B\mathrm{SU}_2;\mathbb{R}) \cong \mathbb{R}[c_2]$
- $I^\bullet(\mathrm{SL}_2(\mathbb{C})) \cong \mathbb{R}[s]$, where $s(A) := \mathrm{tr}(A^2)$
- Even though G is noncompact, the Chern-Weil map is an isomorphism

Nonexample: $G = \mathrm{SL}_2(\mathbb{R})$

- G is noncompact, with maximal compact SO_2, so $H^*(BG;\mathbb{R}) \cong H^*(B\mathrm{SO}_2;\mathbb{R}) \cong \mathbb{R}[c_1]$, $|c_1| = 2$
- However, $I^\bullet(\mathrm{SL}_2(\mathbb{R})) \cong \mathbb{R}[s]$, with $s(A) := \mathrm{tr}(A^2)$, meaning s has degree 4
- The Chern-Weil map sends s to a multiple of c_1^2