CS 468 (Spring 2013) — Discrete Differential Geometry

Lecture 13 | 13 May 2013 | Tensors and Exterior Calculus

Lecturer: Adrian Butscher
Scribe: Cassidy Saenz

1 Vectors, Dual Vectors and Tensors

1.1 Inner Product Spaces

Definition: Let $\mathcal{V}$ be an abstract vector space of dimension n. We define an inner product on $\mathcal{V}$ as a bilinear, symmetric, and positive definite function:

$$\langle \cdot, \cdot \rangle : \mathcal{V} \times \mathcal{V} \to \mathbb{R}$$

(1) *Bilinear*: The inner product is linear in each vector "slot". This means that:

$$\langle au + w, v \rangle = a\langle u, v \rangle + \langle w, v \rangle$$

$$\langle u, av + w \rangle = a\langle u, v \rangle + \langle u, w \rangle$$

(2) *Symmetric*: $\langle u, v \rangle = \langle v, u \rangle$

(3) *Positive Definite*: $\langle v, v \rangle \geq 0$, and $\langle v, v \rangle = 0 \iff v = \vec{0}$

Later on, we'll consider a more concrete vector space (namely, the tangent space of a surface), but for now we'll discuss in general terms. Once we have an abstract vector space with an inner product, then all the usual constructions follow, which we are familiar with from basic vector geometry:

- The norm of a vector is $\|v\| := \sqrt{\langle v, v \rangle}$.
- Vectors v, w are orthogonal if $\langle v, w \rangle = 0$.
- If $\mathcal{S}$ is a subspace of $\mathcal{V}$ then every vector $v \in \mathcal{V}$ can be uniquely decomposed as $v := v^{\|} + v^{\perp}$ where $v^{\|} \in \mathcal{S}$ and $v^{\perp} \perp \mathcal{S}$.
- The mapping $v \mapsto v^{\|}$ is the orthogonal projection onto $\mathcal{S}$.

1.2 Dual Vectors

Definition: Let $\mathcal{V}$ be an abstract vector space of dimension n. The dual space of the vector space $\mathcal{V}$ consists of all linear functions on $\mathcal{V}$:

$$\mathcal{V}^* := \{\xi : \mathcal{V} \to \mathbb{R} : \xi \text{ is linear}\}$$

This newly defined dual space $\mathcal{V}^*$ is itself a vector space of dimension n. We know $\mathcal{V}^*$ is a vector space simply because if we take ξ and ζ to be two functionals in $\mathcal{V}^*$, then:

$$\begin{matrix} \xi + \zeta \\ a\xi \end{matrix} \in \mathcal{V}^*$$

Furthermore, we can show that $\mathcal{V}^*$ is of dimension n. To show this, we construct a basis for $\mathcal{V}^*$. Since $\mathcal{V}^*$ is built upon the vector space $\mathcal{V}$, it stands to reason that we will build the basis of $\mathcal{V}^*$ from the basis of $\mathcal{V}$. If $\{E_i\}$ is a basis for $\mathcal{V}$, then let's define $\{\omega^i\}$ as a basis for $\mathcal{V}^*$ where:

$$\omega^i(E_s) = \begin{cases} 1 & i = s \\ 0 & \text{otherwise} \end{cases}$$

1.3 The Dual Space of an Inner Product Space

Let $\mathcal{V}$ be a vector space with inner product $\langle \cdot, \cdot \rangle$, and dual space $\mathcal{V}^*$. There are three useful constructions available to us in an inner product space.

(1) *Flat Operator:* Given a vector $v \in \mathcal{V}$, we have $v^\flat \in \mathcal{V}^*$, where we define $v^\flat(w) := \langle v, w \rangle \; \forall \, w \in \mathcal{V}$. In other words, the flat operator takes a vector v from $\mathcal{V}$ to an element $v^\flat$ in the dual space $\mathcal{V}^*$, by means of taking an inner product.

(2) *Sharp Operator:* Given an element $\xi \in \mathcal{V}^*$, there exists some vector $\xi^\sharp \in \mathcal{V}$, such that $\xi(w) = \langle \xi^\sharp, w \rangle \; \forall \, w \in \mathcal{V}$. In other words, the sharp operator takes an element ξ from the dual space $\mathcal{V}^*$ to a vector $\xi^\sharp$ in $\mathcal{V}$, so that taking the inner product against this new vector $\xi^\sharp$ yields the original dual element ξ.

(3) *Induced Inner Product:* If the vector space $\mathcal{V}$ carries an inner product, then $\mathcal{V}^*$ inherits the inner product $\langle \xi, \zeta \rangle_{\mathcal{V}^*} := \langle \xi^\sharp, \zeta^\sharp \rangle \; \forall \, \xi, \zeta \in \mathcal{V}^*$. What this means is we can take the inner product of two functionals ξ and ζ in the dual space, by first applying the sharp operator on each to get their counterparts in $\mathcal{V}$, then taking the inner product.

Note: Sharp and Flat are inverse operations, i.e. $(v^\flat)^\sharp = v$ and $(\xi^\sharp)^\flat = \xi$. Intuitively, this makes sense. The Flat operator says: give me a vector v in the vector space $\mathcal{V}$, and I'll take the inner product; the Sharp operator says: give me an element ξ in the dual space $\mathcal{V}^*$, and I'll give you the vector whose inner product would yield that dual element ξ.

1.4 Basis Representations

All the definitions up until now are perfectly fine basis-free. And practically speaking, there are many cases where you would not want to work in a basis. For example, if working with a very large (possible infinite) dimensional vector space, we don't want to construct a basis, but instead want to work with concepts that exist independent of the basis. Secondly, in differential geometry, in order to prove a geometric concept is independent of parameterization (which we have learned is crucial to nearly every proof of concept), it is very useful to work independent of a basis.

Upper vs. Lower Indices: When we have a basis constructed, we represent a vector as a linear combination of the basis vectors, and therefore need a way to encode the indices. The indices and their position convey a certain meaning depending on context. Let's define a vector in an n-dimensional space with basis $\{E_i\}$:

$$v = \sum_{i=1}^{n} v^i E_i$$

Here, the upper index on v indicates the ith vector component, and similarly the lower index on E indicates the ith basis of our vector space. Now, let's define an element of the n-dimensional dual space with dual basis $\{w_i\}$:

$$\xi = \sum_{i=1}^{n} \xi_i w^i$$

Notice that in the dual space, the upper index on w now indicates the ith basis of the dual space, and the lower index on ξ indicates the ith dual vector component.

Definition: Let $\{E_i\}$ denote a basis for a vector space $\mathcal{V}$, and define $g_{ij} := \langle E_i, E_j \rangle$. Note that $\{E_i\}$ is arbitrary and not necessarily an orthonormal basis, which is why we have these g_{ij} coefficients. Now let g^{ij} be the components of the inverse of the matrix $[g_{ij}]$. With this, we have the following:

(1) The dual basis is $\omega^i := \sum_j g^{ij} E_j^\flat$.

Proof:

$$w^i(E_s) = \sum_j g^{ij} E_j^\flat(E_s) = \sum_j g^{ij} \langle E_j, E_s \rangle = \sum_j g^{ij} g_{js} = \delta_{ij} = \begin{cases} 1 & i = j \\ 0 & \text{otherwise} \end{cases}$$

(2) If $v = \sum_i v^i E_i$ then $v^\flat = \sum_i v_i \omega^i$ where $v_i := \sum_j g_{ij} v^j$. This statement gives some intuition behind the flat operator, and why its name makes sense. Given the vector v, as a linear combination of the basis vectors, we know that $v^\flat$ is the dual space counterpart of v, which can be written as simply a linear combination of the dual basis vectors. Now, what are the components v_i for the linear combination? As the formula shows, we take the original vector components (upper indices), and sum them against the matrix coefficients g_{ij}, to get the dual space components (lower indices). By lowering the indices, we "flatted" the original vector, in the same way we can flat a note in music by lowering the pitch one half step! Yes, it's cute.

(3) If $\xi = \sum_i \xi_i \omega^i$ then $\xi^\sharp = \sum_i \xi^i E_i$ where $\xi^i := \sum_j g^{ij} \xi_j$. A similar intuition follows from this statement, regarding the sharp operator's purpose and name. Instead, here we sum the dual space vector components against the inverse matrix coefficients g^{ij} to raise the indices and recover the original vector components. Hence, sharp!

(4) If $\xi = \sum_i \xi_i \omega^i$ and $\zeta = \sum_i \zeta_i \omega^i$ then $\langle \xi, \zeta \rangle = \sum_{ij} g^{ij} \xi_i \zeta_j$. This is similar to the previous sharp operator case, only now we take two different functionals ξ and ζ in the dual space, and take a bilinear combination against the inverse matrix coefficients g^{ij} to recover the induced inner product carried by the dual space. Recall that the induced inner product first required applying the sharp operator to each functional to recover the original vectors that produced them. This is where the last equation comes from.

Note: If $\{E_i\}$ is orthonormal, then $g_{ij} = \delta_{ij} = \begin{cases} 1 & i = j \\ 0 & \text{otherwise} \end{cases}$, and $v_i = v^i$, and $\xi_i = \xi^i$. These second and third claims are trivial to see, given the first. Plugging in $g_{ij} = I$ leaves the vector components unchanged under the sharp and flat operators.

1.5 Tensors

The subject of tensors is very large, with many applications. We'll do a discussion of the basic facts about tensors, then focus on a couple of key tensors that will be important for our purposes.

Definition: Let $\mathcal{V}$ be a vector space of dimension n. The space of k-covariant and ℓ-contravariant tensors is

$$\mathcal{V}^* \overbrace{\otimes \cdots \otimes}^{k \text{ times}} \mathcal{V}^* \otimes \mathcal{V} \overbrace{\otimes \cdots \otimes}^{\ell \text{ times}} \mathcal{V} := \left\{ \begin{array}{c} f : \mathcal{V} \overbrace{\times \cdots \times}^{k \text{ times}} \mathcal{V} \to \mathcal{V} \overbrace{\times \cdots \times}^{\ell \text{ times}} \mathcal{V} \\ \text{such that } f \text{ is multilinear} \end{array} \right\}$$

Tensors are multilinear functions on $\mathcal{V}$ with multi-vector output. This means they are linear in each of the $(k + \ell)$ input vector "slots":

$$f(\cdots, \underbrace{v_i + a w_i}_{i\text{th slot}}, \cdots) = f(\cdots, v_i, \cdots) + a f(\cdots, w_i, \cdots), \; \forall i \leq (k + \ell)$$

Example: The metric of a surface is an example of a bilinear $(k, 0)$-tensor, where $k = 2$.

Basic facts:
(1) Tensors yield a vector space of dimension $n^{k+\ell}$. A basis can be written in terms of the vector space basis components E_i and the dual space basis components ω^i.

(2) If there is an inner product in the vector space $\mathcal{V}$, then the tensor inherits an inner product from $\mathcal{V}$, and also has its own flat and sharp operators. In particular, it has ℓ flat operators for each vector slot, and k sharp operators for each dual vector slot.

(3) There are operations on tensors called "contractions", where a $\mathcal{V}$ factor is combined with a $\mathcal{V}^*$ factor. This takes a (k,ℓ)-tensor to a $(k-1,\ell-1)$-tensor. We'll see an example later.

2 Symmetric and Alternating Tensors

2.1 Symmetric Bilinear Tensors

Definition: A symmetric (2,0)-tensor is an element $A \in \mathcal{V}^* \otimes \mathcal{V}^*$ such that $A(v,w) = A(w,v)$ for all $v, w \in \mathcal{V}$.

We've already seen an example of this kind of tensor – the metric. Another example we've seen before is the second fundamental form at a point, which takes two vectors as input, and outputs a scalar.

Properties:
(1) In a basis we have $A = \sum_{ij} A_{ij}\, \omega^i \otimes \omega^j$ with $A_{ij} = A_{ji}$. Since we are working in the space of $\mathcal{V}^* \otimes \mathcal{V}^*$, the basis representation is now a linear combination of the tensor product of the two bases, where the tensor product is defined as:

$$\omega^i \otimes \omega^j(v,w) = \omega^i(v)\omega^j(w)$$

(2) We define an associated self-adjoint (1,1)-tensor $S \in \mathcal{V}^* \otimes \mathcal{V}$ with the formula $A(v,w) := \langle S(v), w\rangle$. This is the shape operator we associated with the second fundamental form in previous lectures. More formally speaking, we applied a sharp operator to one of the dual vector slots of the second fundamental form to recover its counterpart in the original vector space, leaving us with one dual space "slot" and one vector space slot.

(3) In a basis we have $S = \sum_{ij} S_i^j \omega^i \otimes E_j$ where $S_i^j = \sum_k g^{kj} A_{ik}$. Now that we're in the space of $\mathcal{V}^* \otimes \mathcal{V}$, the basis representation uses the tensor product of one vector space basis and one dual space basis (notice the upper vs. lower indices).

(4) If v and w are defined on a basis, meaning $v = \sum_s v^s E_s$ and $w = \sum_t w^t E_t$, then $\langle v, w\rangle = [v]^\top [g][w]$ – recall we defined $g_{ij} := \langle E_i, E_j\rangle$. Also, we can expand the formula for A from (1) to get the following forms for the second fundamental form and the shape operator:

$$A(v,w) = \sum_{ij} A_{ij}\omega^i(v)\omega^j(w) = \sum_{ij} A_{ij} \underbrace{\omega^i(E_s)v^s}_{\delta_s^i} \underbrace{\omega^j(E_t)w^t}_{\delta_t^j} = \sum_{st} v^s w^t A_{st} = [v]^\top [A][w]$$

$$S = [g]^{-1}[A]$$

(5) As mentioned in the previous section, a contraction takes one upper index (dual space) component, and sums with one lower index (vector space) component to take a (k,ℓ)-tensor to a $(k-1,\ell-1)$-tensor. The contraction of A = the trace of $S = \sum_{ij} g^{ij} A_{ij}$ = the mean curvature!

2.2 Alternating Tensors

Definition: A k-form is an element $\sigma \in \mathcal{V}^* \otimes \cdots \otimes \mathcal{V}^*$ that is antisymmetric in every pair of vector slots. So, for all $v, w \in \mathcal{V}$ and pairs of slots in σ we have:

$$\sigma(\ldots v \ldots w \ldots) = -\sigma(\ldots w \ldots v \ldots) \qquad \text{"Alternating } (k,0)\text{-tensor"}$$

Note: If $\dim \mathcal{V} = 2$ then only $k = 0, 1, 2$ are non-trivial tensors. Also:

$$\mathsf{Alt}^0(\mathcal{V}) = \mathbb{R} \quad \text{and} \quad \mathsf{Alt}^1(\mathcal{V}) = \mathcal{V}^* \quad \text{and} \quad \mathsf{Alt}^2(\mathcal{V}) \cong \mathbb{R}$$

To prove this, let's consider $\sigma(u, v, w)$ to be an alternating 3-form ($k = 3$) in 2-space, i.e. $\mathcal{V} \in \mathbb{R}^2$. Since we are in 2-space, we know that u, v, and w are not linearly independent vectors. Without loss of generality, let's consider w to be a linear combination of u and v, then:

$$\sigma(u, v, w) = \sigma(u, v, au + bv) = a\sigma(u, v, u) + b\sigma(u, v, v)$$

The last step comes from the fact that tensors are multilinear functions. Since alternating tensors are antisymmetric in every pair of slots, we know that:

$$\sigma(u, v, u') = -\sigma(u', v, u)$$

$$\therefore \sigma(u, v, u) = 0$$

Where $u = u'$, and are only different to illustrate the switch made between the 1st and 3rd slots. The same can be shown for the 2nd and 3rd slots, yielding $\sigma(u, v, v) = 0$. This gives us a trivial equation for the 3-form:

$$\sigma(u, v, w) = a\sigma(u, v, u) + b\sigma(u, v, v) = a(0) + b(0) = 0$$

Note: The wedge product of two dual vectors, $\omega^i \wedge \omega^j$ is the basis of the alternating 2-tensor in 2-space. Let $v = \sum_i v^i E_i$ and $w = \sum_i w^i E_i$. Then we define the wedge product via:

$$\omega^1 \wedge \omega^2(v, w) := \det([v\, w])$$

Duality: If $\mathcal{V}$ is an inner product space:

(1) We get a unique alternating 2-form that relates to the metric. The area form $dA \in \mathsf{Alt}^2(\mathcal{V})$:

$$dA(v, w) := \begin{bmatrix} \text{Signed area of} \\ \text{parallelogon} \\ \text{spanned by } v \wedge w \end{bmatrix}$$

(2) There is a special operator called the Hodge-star operator $*$:

$$\omega \wedge *\tau := \langle \omega, \tau \rangle \, dA$$

This will be covered in more depth with practical applications in the next lecture. For now, just know that it exists in the space of alternating 2-tensors, when dealing with an inner product space. Additionally, know these helpful formulas:

$$*dA = 1$$

$$*1 = dA$$

$$\text{and if} \, \omega \in \mathsf{Alt}^1(\mathcal{V}), then * \omega(v) = \omega(R_{\pi/2}(v))$$

3 Tensor Bundles on a Surface

Up until now, we've been dealing with an abstract vector space and linear algebra, with no explicit connection to a surface or constructions applied to some geometric object. Now we will be a little more concrete and consider more than one vector space, and more specifically, the many vector spaces that are the tangent planes to a surface.

Definition: Let S be a surface and let $\mathcal{V}_p := T_pS$. The bundle of (k,ℓ)-tensors over S attaches the vector space $\mathcal{V}_p^{(k,\ell)}$ at each $p \in S$, where:

$$V_p^{(k,\ell)} := \mathcal{V}_p^* \otimes \cdots \otimes \mathcal{V}_p^* \otimes \mathcal{V}_p \otimes \cdots \otimes \mathcal{V}_p$$

Definition: A section of this bundle is the assignment $p \mapsto \sigma_p \in \mathcal{V}_p^{(k,\ell)}$. In other words, to get a section of a bundle, you make a specific choice of a tensor at every point p.

Examples:
(1) $k = 0, \ell = 0$ — sections are functions on S, because for each point on the surface, we output a scalar.

(2) $k = 0, \ell = 1$ — sections are vector fields on S, because for each point on the surface, we output a vector.

(3) $k = 1, \ell = 0$ — sections are one-forms on S, which means we output a dual vector (or functional), at every point.

(4) $k = 2, \ell = 0$ and symmetric — sections are symmetric bilinear forms at each point, such as the metric and the second fundamental form.

(5) $k = 2, \ell = 0$ and antisymmetric — sections are two-forms on S, such as the area form.

4 Covariant Differentiation in a Tensor Bundle

We are familiar with the covariant derivative on vector fields, and it extends naturally to tensor bundles, which are built upon vector fields.

A formula: Choose a basis and suppose

$$\sigma := \sum_{ijkl} \sigma_{ij}^{k\ell} \omega^i \otimes \omega^j \otimes E_k \otimes E_\ell$$

is a (k,ℓ)-tensor. Then:

$$\nabla\sigma := \sum_{ijkls} \nabla_s \sigma_{ij}^{k\ell} [\omega^i \otimes \omega^j \otimes E_k \otimes E_\ell] \otimes \omega_s$$

is also a tensor, where

$$\nabla_s \sigma_{ij}^{k\ell} := \frac{\partial \sigma_{ij}^{k\ell}}{\partial x^s} - \Gamma_{is}^t \sigma_{tj}^{k\ell} - \Gamma_{js}^t \sigma_{it}^{k\ell} + \Gamma_{ts}^i \sigma_{ij}^{t\ell} + \Gamma_{ts}^\ell \sigma_{ij}^{kt}$$

and the Γ_{jk}^i terms are the Christoffel symbols, or "correction factors" which account for the metric changing along the surface.

This new tensor formed by the covariant derivative is a $(k+1,\ell)$-tensor. The additional input vector comes from the choice of direction of differentiation when performing the covariant derivative. This direction also gives us the extra ω_s term in the basis representation.

5 Exterior Calculus

5.1 Exterior Differentiation

Definition: The exterior derivative is the operator $d : \mathsf{Alt}^k(S) \to \mathsf{Alt}^{k+1}(S)$ defined as follows:

(1) Choose a basis.

(2) If $f \in \mathsf{Alt}^0(S)$ we define df geometrically by $df(V) := V(f)$ or

$$df = \sum_i \frac{\partial f}{\partial x^i}\omega^i \qquad \text{Thus } (df)^\sharp = \nabla f$$

(3) If $\omega = \sum_i a_i\omega^i \in \mathsf{Alt}^1(S)$, then $d\omega = \left(\frac{\partial a^1}{\partial x^2} - \frac{\partial a^2}{\partial x^1}\right)\omega^1 \wedge \omega^2$.

(4) If $\omega = a\,\omega^1 \wedge \omega^2 \in \mathsf{Alt}^2(S)$, then $d\omega = 0$.

Basic Facts:
(1) $dd\omega = 0$ for all $\omega \in \mathsf{Alt}^k(S)$ and all k.

(2) $d\omega = \mathsf{Antisym}(\nabla\omega)$, and as such it does not involve the Christoffel symbols like the covariant derivative. It can be thought of as the metric-independent part of the covariant derivative.

5.2 The Co-differential Operator

Definition: The co-differential is the L^2-adjoint of d. It is therefore an operator $\delta : \mathsf{Alt}^{k+1}(S) \to \mathsf{Alt}^k(S)$ that satisfies:

$$\int_S \langle d\omega, \tau\rangle\, dA = \int_S \langle \omega, \delta\tau\rangle\, dA$$

It is given by $\delta := - * d *$, where $*$ signifies the Hodge star operator defined previously.

Some concepts:
(1) If f is a function, then $(df)^\sharp = \nabla f$.

(2) If X is a vector field, then $\delta X^\flat = div(X)$.

(3) If X is a vector field, then $dX^\flat = curl(X)\, dA$.

(4) If f is a function, then $\big(\delta(f\,dA)\big)^\flat = R_{\pi/2}\big(\nabla f\big)$.

6 Stokes' Theorem

Intuition: Generalization of the Fundamental Theorem of Calculus.

> Suppose that c is a $(k+1)$-dimensional submanifold of S with k-dimensional boundary ∂c. Let ω be a k-form on S. Then:
>
> $$\int_c d\omega = \int_{\partial c} \omega$$

This says that if you integrate the differential of a form over an entire surface, it is the same as integrating the form over just the boundary of the surface.

Interpretation: The divergence theorem:

$$\int_S div(X) dA = \int_{\partial S} \langle N_{\partial S}, X \rangle \, d\ell$$

Where X is the vector field on the surface, and $N_{\partial S}$ is the vector field along the boundary of the surface, which is both tangent to the surface and normal to the boundary.

7 The Hodge Theorem

Theorem: $\mathsf{Alt}^1(S) = d\mathsf{Alt}^0(S) \oplus \delta\mathsf{Alt}^2(S) \oplus \mathcal{H}^1$ where $\mathcal{H}^1$ is the set of harmonic one-forms:

$$\begin{aligned} h \in \mathcal{H}^1 \quad &\Leftrightarrow \quad dh = 0 \text{ and } \delta h = 0 \\ &\Leftrightarrow \quad \underbrace{(d\delta + \delta d)}_{\text{“Hodge Laplacian”}} h = 0 \end{aligned}$$

Corollary: Every vector field X on S can be decomposed into a “gradient” part, a “divergence-free” part, and a “harmonic part.”

$$X = \nabla\phi + curl\big(\nabla\psi\big) + h^\sharp \quad \text{with } h \in \mathcal{H}^1$$

Another deep mathematical result:

Theorem: $\mathsf{dim}(\mathcal{H}^1) = 2g(S)$. This is a topological invariant. This also implies that harmonic one-forms only exist on a surface when its genus $g \neq 0$, otherwise the dimension of the vector space would be 0.