The Hodge Star Operator

Rich Schwartz

April 22, 2015

1 Basic Definitions

We'll start out by defining the Hodge star operator as a map from $\wedge^k(\boldsymbol{R}^n)$ to $\wedge^{n-k}(\boldsymbol{R}^n)$. Here $\wedge^k(\boldsymbol{R}^n)$ denotes the vector space of alternating k-tensors on $\boldsymbol{R}^n$. Later on, we will extend this definition to alternating tensors on a finite dimensional vector space that is equipped with an inner product.

Let $I = (i_1, ..., i_k)$ be some increasing multi-index of length k. That is $i_1 < i_2 < i_3 < ....$ Let $J = (j_1, ..., j_{n-k})$ be the complementary increasing multi-intex. For instance, if $n = 7$ and $I = (1, 3, 5)$ then $J = (2, 4, 6, 7)$. Let K_0 denote the full multi-index $(1, ..., n)$.

We first define $*$ on the usual basis elements:

$$*(dX_I) = \epsilon(IJ) \ dx_J. \tag{1}$$

Here $\epsilon(IJ)$ is the sign of the permutation which takes IJ to K_0. Sometimes, we will just write $*dx_I$ in place of $*(dx_I)$. In general, we define

$$*\left(\sum a_I \ dx_I\right) = \sum a_I \ (*dx_I). \tag{2}$$

Lemma 1.1 *For any* $\omega \in \wedge^k(\boldsymbol{R}^n)$ *we have* $* * \omega = (-1)^{k(n-k)}\omega$. *That is* $* * \omega = \omega$ *unless k is odd and n is even, in which case* $* * \omega = -\omega$.

Proof: It suffices to check this on a basis element dX_I. Let J be the complementary multi-index as above. Note that $* * dx_I = \pm dX_I$. We just have to get the sign right. When $I = (1, ..., k)$, the sign does work out. So, for any other choice of I, we just have to show that the sign only depends on n and k. That tells us that the sign works out in all cases.

We want to show that

$$\epsilon(IJ \to K_0)\epsilon(JI \to K_0)$$

only depends on n and k. Let $\overline{K}$ denote the multi-index which is the reverse of K. (If K is increasing then $\overline{K}$ is decreasing.) For the second permutation, we can reverse everything and we would get the same answer. That is:

$$\epsilon(JI \to K_0) = \epsilon(\overline{IJ} \to \overline{K}_0).$$

But then

$$\epsilon(\overline{IJ} \to \overline{K}_0) = \epsilon(\overline{I} \to I)\epsilon(\overline{J} \to J)\epsilon(\overline{K}_0 \to K_0)\epsilon(IJ \to K_0).$$

Putting everything together, we get

$$\epsilon(IJ \to K_0)\epsilon(JI \to K_0) =$$

$$(\epsilon(IJ \to K_0))^2\epsilon(\overline{I} \to I)\epsilon(\overline{J} \to J)\epsilon(\overline{K}_0 \to K_0) =$$

$$\epsilon(\overline{I} \to I)\epsilon(\overline{J} \to J)\epsilon(\overline{K}_0 \to K_0).$$

Since I and J are increasing, the two quantities $\epsilon(\overline{I} \to I)$ and $\epsilon(\overline{J} \to J)$ only depend on the length of I and J respectively and not on their specific terms. The quantity $\epsilon(K_0 \to \overline{K}_0)$ is the same in all cases. ♠

2   Rotational Symmetry

Let $SO(n)$ denote the group of orientation preserving orthogonal matrices. Given any $M \in SO(n)$ we have an induced map $M^* : \wedge^\ell(\boldsymbol{R}^n) \to \wedge^\ell(\boldsymbol{R}^n)$. The purpose of this section is to prove that

$$*M^*(\omega) = M^*(*\omega), \tag{3}$$

for all $\omega \in \wedge^k(\boldsymbol{R}^n)$. This equation expresses the fact that the Hodge star operator is rotationally symmetric. In other words, the star operator commutes with the action of $SO(n)$.

Say that an element $M \in SO(n)$ is *good* if Equation 3 holds for M. For instance, the identity element is obviously good.

Lemma 2.1 *Suppose that A and B are good. Then AB is also good.*

Proof: The only thing to note is that $(AB)^* = B^*A^*$. So, now we compute

$$(AB)^*(*\omega) = B^*A^*(*\omega) = B^*(*A^*(\omega)) = *B^*A^*(\omega) = *(AB)^*(\omega).$$

That's it. ♠

Before stating the next result, we observe that it concerns an orthogonal matrix whose determinant is -1. So, the next result does not contradict the main thing we are trying to prove.

Lemma 2.2 *Let $k < n$ be some index. Let M be the orthogonal matrix defined by the rules that $M(e_j) = e_j$ if $j \neq k, k+1$ and $M(e_k) = e_{k+1}$ and $M(e_{k+1}) = e_k$. Then $M^** = - * M^*$.*

Proof: It suffices to prove this on a basis. Let I be an increasing multi-index. We will consider the equation on $\omega = dx_I$. There are 4 cases to consider.

- Suppose that I contains both k and $k+1$. Then $M^*(\omega) = -\omega$ and $*M^*(\omega) = - * \omega$. On the other hand $*\omega = \pm dx_J$, where J does not involve either k or $k+1$. Hence $M^*(*\omega) = *\omega$.

- Suppose that I contains neither k nor $k+1$. This case is similar to the previous case.

- Suppose that I contains k but not $k+1$. Then $M^*(\omega) = dx_{I'}$, where I' is obtained from I by swapping out k and putting in $k+1$. We have $*\omega = \epsilon(IJ)dx_J$ and $M^*(*\omega) = \epsilon(IJ)dx_{J'}$, where J' is obtained from k by swapping out $k+1$ for k. On the other hand, we have $*M^*(\omega) = \epsilon(I'J')dX_{J'}$. To finish the proof in this case, we just have to show that $\epsilon(IJ) = -\epsilon(I'J')$. But the two permutations just differ by composition with the transposition which swaps k and $k+1$. So, the equality holds.

- Suppose that I contains $k+1$ but not k. This case has a similar proof to the previous case.

This takes care of all the cases. ♠

Corollary 2.3 *Suppose that M is good and P is any permutation matrix. Then PMP^{-1} is also good.*

Proof: Call an orthogonal matrix *anti-good* if it has the transformation law given in Lemma 2.2. The same argument as in Lemma 2.1 shows that the product of two anti-good matrices is good, and that the product of a good and an anti-good matrix is anti-good.

Any permutation matrix is the product of finitely many matrices covered by Lemma 2.2. For this reason, it suffices to prove the result when P is such a matrix. In this case $P = P^{-1}$, and P is anti-good. But then MP^{-1} is anti-good, and $P(MP^{-1})$ is good. ♠

Lemma 2.4 *Let M be the element of $SO(n)$ which has the following action:*

- $M(e_j) = e_j$ *for $j = 3, 4, 5, ...$*

- $M(e_1) = e_1 \cos(\theta) + e_2 \sin(\theta),$

- $M(e_2) = -e_1 \sin(\theta) + e_2 \cos(\theta).$

In other words, M rotates by θ in the e_1, e_2 plane and fixes the perpendicular directions. Then M is good.

Proof: It suffices to check this on a basis. We have $*dx_I = \epsilon(IJ)dx_J$. For notational convenience, we'll consider the case when $\epsilon(IJ) = 1$. The other case has the same kind of proof. So, in short, we are considering the case when $*dx_I = dx_J$.

Note that the restriction of M to the e_1, e_2 plane is an orientation-preserving rotation. For this reason $M^*(dx_I) = dx_I$ when I contains both 1 and 2, and also when I contains neither 1 nor 2. Likewise $M^*(dx_J) = dx_J$. So, in either of these two cases, $*M^*(dx_I) = M^* * (dx_I) = dx_J$.

Consider the case when I contains 1 but not 2. Then $dx_I = dx_1 \wedge dx_{I'}$. Here I' is obtained from I by omitting 1. Similiarly, we have the equations

$$*(dx_1 \wedge dx_{I'}) = dx_2 \wedge dx_{J'}, \qquad *(dx_2 \wedge dx_{I'}) = -dx_1 \wedge dx_{J'}.$$

Here J' is obtained from J by omitting 2. The sign change in the second calculation comes from the fact that $\epsilon(IJ) = -\epsilon(\hat{I}\hat{J})$, where $\hat{I}$ and $\hat{J}$ are obtained from I and J by swapping 1 and 2.

We set $C = \cos(\theta)$ and $S = \sin(\theta)$. An easy computation shows that

$$M^*(dx_1) = Cdx_1 - Sdx_2, \qquad M^*(dx_2) = Sdx_1 + Cdx_2.$$

These calculations tell us that

$$*M^*(dx_I) = *\big((Cdx_1 - Sdx_2) \wedge dx_{I'}\big) =$$

$$*\big(Cdx_1 \wedge dx_{I'}\big) - *\big(Sdx_2 \wedge dx_{I'}\big) = Cdx_2 \wedge dx_{J'} + Sdx_1 \wedge dx_{J'}.$$

Similarly

$$M^*(*dx_I) = M^*(dx_2 \wedge dx_{J'}) =$$

$$(Sdx_1 + Cdx_2) \wedge dx_{J'} = Sdx_1 \wedge dx_{J'} + Cdx_2 \wedge dx_{J'}.$$

The two expressions agree.

There is one more case, when I contains 2 but not 1. This case is similar to the last case, and actually follows from the last case and Lemma 1.1. ♠

Lemma 2.5 *Let $i < j$ be two indices. Let M be the element of $SO(n)$ which has the following action:*

- *$M(e_k) = e_k$ if $k \neq i, j$.*

- *$M(e_i) = e_i \cos(\theta) + e_j \sin(\theta)$,*

- *$M(e_j) = -e_i \sin(\theta) + e_j \cos(\theta)$.*

In other words, M rotates by θ in the e_i, e_j plane and fixes the perpendicular directions. Then M is good.

Proof: Let P be any permutation matrix which maps 1 to i and 2 to j. Let M_{12} be the matrix from Lemma 2.4. Then $M = PM_{12}P^{-1}$. By the preceding lemmas, M is good. ♠

Say that a matrix from Lemma 2.5 is a *basic rotation*.

Lemma 2.6 *Any element of $SO(n)$ is the finite product of basic rotations.*

Proof: The proof goes by induction on $n \geq 2$. The case $n = 2$ is obvious. Let $w_1, ..., w_n$ be any positively oriented orthonormal basis. Applying basic rotations, we can map w_n to e_n. By induction, we can map $w_1, ..., w_{n-1}$ to $e_1, ..., e_n$ using basic rotations which fix e_n. ♠

Lemma 2.5 says that the basic rotations are all good. Lemma 2.1 says that finite products of basic rotations are good. So, by the last result, all elements of $SO(n)$ are good. This completes the proof of the rotational symmetry of the Hodge star operator.

3 A Consequence of the Symmetry

Suppose that $w_1, ..., w_n$ is some other positively oriented orthonormal basis of $\boldsymbol{R}^n$. We define
$$dw_I = dw_{i_1} \wedge ... \wedge dw_{I_k}.$$
Here $dw_1, ..., dw_n$ is the basis of 1-tensors dual to our new orthonormal basis. Here is a consequence of the rotational symmetry of the star operator:

Lemma 3.1 $*dw_I = \epsilon(IJ)dw_J$, where J is the increasing multi-index complementary to I.

Proof: For notational convenience, we'll prove this when $\epsilon(IJ) = 1$. The other case is similar.

There is an element $M \in SO(n)$ such that $M(w_i) = e_i$ for all i. Note that $M^*(dx_i) = dw_i$. Hence $M^*(dx_I) = dw_I$ and $M^*(dx_J) = dw_J$. Now we compute

$$*(dw_I) = *M^*(dx_I) = M^* * (dx_I) = M^*(dx_J) = dw_J.$$

That's it. ♠

Lemma 3.1 says that we could have made the basic definition for the star operator with respect to any positively oriented orthonormal basis and we could have gotten the same answer.

4 General Definition

Suppose now that V is an n-dimensional real vector space and Q is an inner product on V. Suppose also that V is oriented.

There is an isometry $g : \mathbf{R}^n \to V$ which carries the standard basis on $\mathbf{R}^n$ to a positively oriented orthonormal basis on V, with respect to Q. Any two such isometries differ by composition with an element of $SO(n)$.

Given any $\omega \in \wedge^k(V)$, we define

$$*\omega = (g^{-1})^*(*g^*(\omega)).$$

That is, we do the following:

- Pull ω back to $\mathbf{R}^n$ using g^*. Call this new form η.

- Take $*\eta$.

- Pull $*\eta$ back to V using $(g^{-1})^*$. This last form is $*\omega$.

The fact that $*$ commutes with elements of $SO(n)$ guarantees that this definition is independent of the choice of g. In fact, we can equally well define $*$ in the following way: Choose a positively oriented orthonormal basis $v_1, ..., v_n$ for V and define

$$*(dv_I) = \epsilon(IJ)dv_J.$$

The rotational symmetry guarantees that we would get the same answer with respect to any positively oriented orthonormal basis.

The general definition allows us to extend the definition of the star operator from inner product spaces to oriented Riemannian manifolds.

Suppose that M is a smooth oriented manifold with a Riemannian metric. Then we have a smooth choice of inner product Q_p on each tangent space $T_p(M)$. This allows us to define the Hodge star operator in each tangent space. So if $\omega \in \Omega^k(M)$ is some smooth k form, then $*\omega$ is an $n - k$ form on M.

Lemma 4.1 *$*\omega$ is smooth. That is, $*\omega \in \Omega^{n-k}(M)$.*

Proof: In the neighborhood of any point, M has a smoothly varying and positively oriented orthonormal basis. We just take a smoothly varying positively oriented basis, coming from a coordinate chart, and apply the Gram-Schmidt process in each tangent plane. When ω is defined locally with respect to this basis, we can see that $*\omega$ is smooth. ♠

5   A Variant of Stokes' Theorem

Suppose that M is an n dimensional oriented manifold-with-boundary contained in $\boldsymbol{R}^n$. In this situation, M has a canonical Riemannian metric, coming from the dot product on $\boldsymbol{R}^n$. The same goes for ∂M. A good example to consider is when M is a closed ball and ∂M is the sphere bounding the ball.

Let V be a vector field on M, say $V = (V_1, ..., V_n)$. We have the usual associated 1-form $\omega = \sum V_i dx_i$. Note that $*\omega$ is an $(n-1)$ form on M and $d * \omega$ is an n-form on M. Stokes' theorem, applied to $*\omega$, tells us that

$$\int_M d(*\omega) = \int_{\partial M} *\omega. \tag{4}$$

We're going to re-interpret each half of this equation.

The Left Side: A direct calculation shows that

$$d(*\omega) = \sum \partial V_i / \partial x_i \; dx_1 \wedge ... \wedge dx_n = \mathrm{div}(V) \; dx_1 \wedge ... \wedge dx_n.$$

So, the left hand side of Equation 4 equals

$$\int_M \mathrm{div}(V) \; dx_1...dx_n,$$

the usual integral of the divergence of a vector field.

The Right Side: Now let's consider the right hand side of Equation 4. Consider the form $*\omega$ at a point p of ∂M. We can find an oriented orthonormal basis for $\boldsymbol{R}^n$ at p, say $w_1, ..., w_n$, so that

- $w_1, ..., w_{n-1}$ is an oriented orthonormal basis for $T_p(\partial M)$.

- $\nu = (-1)^{n-1} w_n$ is the normal vector that is compatible with Stokes' theorem.

We can write
$$\omega = \sum b_i w_i.$$

Note that the restriction of $*dw_i$ to ∂M at p is 0 unless $i = n$. Therefore the restriction of $*\omega$ to ∂M at p equals $b_n * dw_n$. That is

$$*\omega|_{\partial M} = b_n (*dw_n).$$

But

$$b_n = \omega(w_n) = (-1)^{n-1}\omega(\nu) = (-1)^{n-1}V \cdot \nu.$$

Finally,

$$*dw_n = (-1)^{n-1}dw_1 \wedge ... \wedge dw_{n-1}.$$

Putting these three equations together, we get

$$*\omega|_{\partial M} = V \cdot \nu \; dw_1 \wedge ... \wedge dw_{n-1}.$$

Our theory of integrating functions on manifolds tells us that the right hand side of Equation 4 is

$$\int_{\partial M} V \cdot \nu.$$

The Interpretation: Putting everything together, we have

$$\int_M \mathrm{div}(V) = \int_{\partial M} V \cdot \nu. \tag{5}$$

On the left hand side, we are integrating with the usual volume measure on Euclidean space, and on the right hand side we are integrating a function on an oriented manifold according to the theory explained in the class and in a previous handout. This is a classical n-dimensional generalization of the usual low dimensional version of Stokes' theorem which involves the divergence.

6 Another Variant

Again consider the case when M is an n-dimensional manifold-with-boundary contained in $\boldsymbol{R}^n$. Let f be a smooth function on M. Note that df is the 1-form corresponding to the gradient ∇f. We know from the previous section that

$$\int_M \mathrm{div}(\nabla f) = \int_{\partial M} \nabla f \cdot \nu.$$

A direct calculation shows that

$$\mathrm{div}(\nabla f) = \Delta f := \sum \partial^2 f \partial x_i^2.$$

Here Δf is the Laplacian of f.

This gives us another variant of Stokes' Theorem:

$$\int_M \Delta f = \int_{\partial M} \nabla f \cdot \nu. \tag{6}$$

In physical terms, this result says that the total flux of f through ∂M equals the integral of the Laplacian of f on the interior of M.