A Stroll through the Garden of Functional Analysis: the Hodge Theorem

Colby Riley

January 14, 2025

Contents

Many fields of mathematics seem to have the unfortunate property that the more beautiful they become, the less useful they are. Like a fish out of water, Functional Analysis reverses this trend. Indeed, the Hodge Decomposition theorem is the very beginning of Hodge theory, and its proof relies heavily on functional tools. In this stroll through mathematics, the goal will be to see as little brute force computation as possible, using whenever possible the deep theorems of functional instead. The proof uses large swaths of what may be contained in a standard introduction to functional analysis, including many of its best theorems (shoutout to Banach-Alaoglu).

While we will assume standard functional analysis, we will briefly introduce and cite more specialized topics when needed. We will use the proof as an opportunity to introduce the key theorems of Sobolev spaces, Elliptic operators, and distributions as relating to the proof. We will also introduce the needed differential geometry. At the end, there is a rather longer section where we can see some of the more immediate applications of the decomposition. This is important to me because it shows very quickly and tangibly that this theorem is not just cool, but useful too!

1 Introduction

Throughout the text, all manifolds will be smooth, closed, orientable and Riemannian, and X will denote an n-dimensional such manifold. The Riemannian metric g is essentially a smoothly varying inner product on TX.

For any point $p \in X$, there is a chart $U \ni p$ and a local coordinate system $\{e_1, ..., e_n\}$ on U such that $\{e_1, ..., e_n\}$ is orthonormal wrt to the metric g, and that the partials of g with respect to the e_k are 0 at p. These are called Riemannian normal coordinates, and they come from exponentiating an orthonormal basis for T_pX. Letting $\Omega^k(X)$ denote the space of smooth k-forms on X, then a basis for $\bigwedge^k(U)$ consists of elements of the form de_I, for I a multi-index of $(1...n)$. For example, if $I = (1, 2, 5)$ then $de_I = de_1 \wedge de_2 \wedge de_5$. Then elements of $\Omega^k(X)$ are locally of the form $\sum_I f_I de_I$ with $f_I : U \to \mathbb{R}^n$ a smooth function.

Now we can come to an explicit definition of the Hodge dual to define an inner product on $\Omega^k(X)$. With respect to the coordinate chart defined above, and for a basis element $de_I \in \bigwedge^k(U)$, we define

$$\star de_I := \pm de_J$$

Where (IJ) concatenated is a permutation of $(1...n)$. The sign is positive if (IJ) is an even permutation and negative otherwise. We then extend by linearity to define $\star$ over all of $\bigwedge^k(U)$ and again extend in the obvious way to define it on $\Omega^k(X)$ locally. It is easy to check that $\star$ does not depend on the proper choice of basis, so that this definition makes sense. For example, in $\mathbb{R}^3$, we have that

$$\star(e_1^2 \; de_1 \wedge de_3) = -e_1^2 de_2$$

The negative sign because (132) is an odd permutation of (123). Note that $\star\star = \pm\mathbb{1}$, with the plus or minus coming from the dimensions of the involved quantities. The specific sign will not matter for the theory, but for computations it is useful to work it out oneself. For properties of the Hodge star see [3].

Now, for $\omega_1, \omega_2 \in \Omega^k$, we can define

$$\langle \omega_1, \omega_2 \rangle = \int_X \omega_1 \wedge \star\omega_2$$

Notice that on the right, $\omega_1 \wedge \star\omega_2$ always has degree n since ω_1 has degree k and ω_2 has degree $n - k$. It is straightforward to verify that this is an inner product because one can verify all properties locally in normal coordinates, where the hodge star is explicit. This integral is always finite precisely because all of our manifolds are closed and orientable.

Definition. *Let $\Omega(X) := \bigoplus_{k=0}^n \Omega^k(X)$, with an inner product structure the direct sums of the inner products on $\Omega^k(X)$. Its norm will be denoted as $||\cdot||$.*

$\Omega(X)$ is NOT a Hilbert space, for the same reason that $C^\infty(\mathbb{R})$ isn't under the standard inner product. For that reason, we will later complete $\Omega(X)$. For now, however, it is desireable to carry on as if it were, and find the adjoint of d, when it is defined. We will call this the formal adjoint.

Lemma 1. *The formal adjoint d^* of d is equal to $\pm \star d\star$, in the sense that $\langle d^*\omega_1, \omega_2\rangle = \langle \omega_1, d\omega_2\rangle$ whenever $\omega_1, \omega_2 \in \Omega(X)$. The sign depends on the dimension.*

Proof. Since this lemma is local in nature, one can take normal coordinates around each point, then do a bit of messy algebra.

□

Remark. Just like d, we have that $(d^*)^2 = 0$

We can finally define the Laplace-Beltrami (or Laplace-de Rham (or just Laplace)) operator:

Definition. *Let $\omega \in \Omega(X)$. Then $\Delta\omega := dd^*\omega + d^*d\omega$. That is, $\Delta := dd^* + d^*d$. We say that ω is a harmonic form, or just harmonic, if $\Delta\omega = 0$. Let $\mathscr{H} \subset \Omega(X)$ be the space of Harmonic forms.*

This definition of course specializes to the regular laplacian over $\mathbb{R}^n$, which can be confirmed using the convenient explicit description of the Hodge $\star$. This finally gives us the vocabulary needed to state the Hodge Decomposition Theorem in its full glory.

Theorem 1 (Hodge Decomposition)**.** *Let X be a closed, Riemannian manifold. Then $\Omega(X) = \mathscr{H} \oplus \Delta(\Omega(X))$, and $\mathscr{H}$ is finite dimensional.*

Otherwise put, for any $\omega \in \Omega(X)$, $\omega = \eta + \Delta\sigma$, for $\sigma \in \Omega(X)$ and η harmonic, and η and σ are unique for a given ω. In the context of functional analysis, [2] notes that the theorem is remarkable because it looks very close to the orthogonal decomposition of a bounded self-adjoint operator on a Hilbert space, except that Δ is not bounded and $\Omega(X)$ is not a Hilbert space! The whole idea of the Hodge theorem is correcting these facts: re-defining the domain and codomains of Δ so that it really is a bounded operator on a Hilbert space.

2 The Proof

Our first goal will be to complete $\Omega(X)$ into a Hilbert space. By taking normal coordinates once more, we can define forms that are not necessarily smooth by requiring only that in each coordinate system, $\omega = \sum_I f_I de_I$ for each f_I measureable. Then it still makes sense to integrate over n-forms via the Lebesgue integral, so that the inner product $\langle \cdot, \cdot \rangle$ from the introduction makes sense. Then we can define $L^2(X)$ as the space of forms ω which are measureable functions in each coordinate system, and s.t. $\langle \omega, \omega \rangle < \infty$. It is easy to show that $L^2(X)$ is complete by looking locally; since we are dealing with compact manifolds and only require a finite number of coordinate patches and can bound each one very generously, then use the completeness of $L^2(\mathbb{R}^n)$. This is the same definition as the formal completion of $\Omega(X)$ would give you, so of course $\Omega(X)$ is dense in $L^2(X)$.

We summarize the situation as follows:

Definition. *We define $L^2(X)$ to be the Hilbert space of square-integrable forms. $\Omega(X)$ is dense in $L^2(X)$. for ω a measureable form, $\omega \in L^2(X)$ precisely if*

$$\langle \omega, \omega \rangle = \int_X \omega \wedge \star\omega < \infty$$

This solves the incompleteness of $\Omega(X)$, but does not solve the unboundedness of Δ. For this, we have to introduce Sobolev spaces. This in turn requires us to extend the definition of d and d^* to $L^2(X)$. What we mean is the following:

Definition. *Let $\omega \in L^2$. Then we define $d\omega := \eta$ if the following holds:*

$$\forall \alpha \in \Omega(X), \langle \omega, d^*\alpha \rangle = \langle \eta, \alpha \rangle$$

And we say that η is the weak differential of ω. Likewise, we define $d^*\omega := \eta$ if

$$\forall \alpha \in \Omega(X), \langle \omega, d\alpha \rangle = \langle \eta, \alpha \rangle$$

And we say that η is the weak co-differential of ω.

The definitions are meant to mimic the adjoint-ness of d and d^* on $\Omega(X)$. If a differential (or co-differential) exists, it must be unique: suppose that η_1, η_2 are both co-differentials of ω. Then for all $\alpha \in \Omega(X)$,

$$\langle \eta_1 - \eta_2, \alpha \rangle = \langle \omega - \omega, \alpha \rangle = 0$$

But since $\Omega(X)$ is dense in $L^2(X)$, for any $\sigma \in L^2(X)$, approximate it by $(\sigma_n)_n \subset \Omega(X)$. Then

$$0 = \lim_n \langle \eta_1 - \eta_2, \sigma_n \rangle = \langle \eta_1 - \eta_2, \lim_n \sigma_n \rangle = \langle \eta_1 - \eta_2, \sigma \rangle$$

Thus $\langle \eta_1 - \eta_2, \sigma \rangle = 0$ for any $\sigma \in L^2(X)$, and thus $\eta_1 - \eta_2 = 0$ almost everywhere. The co-differential is the same. This implies in particular that if $\omega \in \Omega(X)$, its weak and co-differentials are exactly $d\omega$ and $d^*\omega$ in the smooth sense, so that our notation makes sense and is consistent.

It is easy to verify the key properties of d and d^*, such as $d^2 = (d^*)^2 = 0$, and that $\langle d\omega_1, \omega_2 \rangle = \langle \omega_1, d^*\omega_2 \rangle$ for these weak differentials (when the relevant weak differentials actually exist), so we will mostly not worry about the difference between a weak differential and a differential. We do have to be careful, though, only to take a differential when inside of $W^{1,2}$.

Definition. *Let $W^{1,2} \subset L^2(X)$ be the subspace of forms which have both weak-differentials and co-differentials. Rather than the subspace topology, give it the following norm: for $\omega \in W^{1,2}$, let*

$$||\omega||_{W^{1,2}} = ||\omega||_{L^2} + ||d\omega||_{L^2} + ||d^*\omega||_{L^2}$$

With this norm it is clearly also a Hilbert space.

The 1 in $W^{1,2}$ stands for the 1st differentials existing, and the 2 comes from the use of L^2. We can also define $W^{k,2}$ inductively, by taking more differentials:

Definition. *For $k > 1$, let $W^{k,2} \subset W^{1,2}$ be the subspace of forms ω such that $d\omega, d^*\omega \in W^{k-1,2}$. Rather than the subspace topology, give it the following norm: for $\omega \in W^{k,2}$, let*

$$||\omega||_{W^{k,2}} = ||\omega||_{L^2} + ||d\omega||_{W^{k-1,2}} + ||d^*\omega||_{W^{k-1,2}}$$

With this norm it is clearly also a Hilbert space. By convention, $W^{0,2} := L^2(X)$.

Dealing with so many different topologies can be confiusing: if we write $W^{k,2}$, we mean with its Hilbert space topology. If we want to refer to a space via the subspace topology, for example $W^{1,2}$ under the L^2 norm, we will write $(W^{1,2}, L^2)$ or whatever else is relevant.

We can consider differentials as operators $d, d^* : W^{k,2} \to W^{k-1,2}$. Almost by definition,

$$||d\omega||_{W^{k-1,2}}, ||d^*\omega||_{W^{k-1},2} \leq ||\omega||_{W^{k,2}}$$

So they are bounded operators. In particular, $\Delta : W^{2,2} \to L^2(X)$ is a bounded operator, meaning among other things that there is an adjoint $\Delta^* : L^2(X) \to W^{2,2}$. This immediately gives us the following corollary:

Corollary (Our main Corollary)**.**

$$L^2(X) = \ker \Delta^* \oplus (\ker \Delta^*)^\perp = \ker \Delta^* \oplus \overline{\text{Image}\, \Delta}$$

Just by defining the "correct" domains of Δ, functional analysis has spat out an already very promising result! The remainder of the proof will consist of strengthening this corollary: we need to get rid of the adjoint and show that the image is already closed. We also need to replace the $L^2(X)$ with $\Omega(X)$, and of course show that $\ker \Delta$ is finite dimensional, since that is part of the theorem. In the spirit of our stroll, we will prove the last goal first.

2.1 Sobolev Embeddings: $\ker \Delta$

The main results concerning Sobolev spaces that we will need are the following two embedding theorems:

Theorem 2 (Sobolev embedding 1, [4])**.** *The embedding $i : W^{k,2} \hookrightarrow W^{k-1,2}$ is a compact operator.*

Theorem 3 (Sobolev embedding 2, [4])**.**

$$\cap_{k=1}^{\infty} W^{k,2} = \Omega(X)$$

The important thing about the second embedding theorem is that if we want to show that a form is smooth, it suffices to show that it is weakly differentiable infinitely many times, not just differentiable. For more information about Sobolev spaces and the embedding theorems, see [4].

Lemma 2. *For $\omega \in W^{2,2}$, $\Delta\omega = 0$ iff $d\omega = d^*\omega = 0$. That is, ω and is both closed and coclosed.*

Proof. If $\Delta\omega = 0$, then

$$0 = \langle \Delta\omega, \omega \rangle = \langle dd^*\omega, \omega \rangle + \langle d^*d\omega, \omega \rangle$$
$$= \langle d^*\omega, d^*\omega \rangle + \langle d\omega, d\omega \rangle = ||d^*\omega||^2 + ||d\omega||^2$$

If $\Delta\omega = 0$, then $\langle \Delta\omega, \omega \rangle = 0$ and thus both norms are actually equal to 0, so that $d^*\omega = d\omega = 0$. And if $d\omega = d^*\omega = 0$, then of course $\Delta\omega = d(d^*\omega) + d^*(d\omega) = d0 + d^*0 = 0$. □

Lemma 3. *If $\Delta\omega \in \Omega(X)$, then $\omega \in \Omega(X)$. In particular,* $\ker\Delta \subset \Omega(X)$ *so* $\ker\Delta = \mathscr{H}$.

Proof. This lemma is a proof by abuse of Sobolev embedding 2. If $\Delta\omega \in \Omega(X)$, then

$$\Omega(X) \ni d\Delta\omega = ddd^*\omega + dd^*d\omega = dd^*d\omega$$

Since $d^2 = 0$. Since it is also evident that $d^*d^*d\omega = 0 \in \Omega(X)$, we have that both $d(d^*d\omega)$ and $d^*(d^*d\omega)$ are in $\Omega(X)$ (and in particular in $W^{k,2}$ for every k). Thus $d^*d\omega$ is in $W^{k,2}$ for every k, and by the Sobolev embedding, $d^*d\omega \in \Omega(X)$. Since $dd\omega = 0 \in \Omega(X)$ as well, for the same reason again $d\omega \in \Omega(X)$. The analogous proof shows that $d^*\omega \in \Omega(X)$. One more application of Sobolev embedding 2 gives that $\omega \in \Omega(X)$.

□

Theorem 4. *$\mathscr{H}$ is finite dimensional.*

Proof. Since kernels are closed, $\ker\Delta$ is a closed vector subspace of $W^{2,2}$ under the subspace topology. The goal is to show that the unit ball in $\ker\Delta$ is compact: this implies for any TVS that the space is finite dimensional. Let B be the closed unit ball of $\ker\Delta$ in the subspace topology.

Consider the embedding $i : W^{2,2} \hookrightarrow L^2(X)$. By Sobolev embedding 1, this is a compact mapping, and the restriction $j : \ker\Delta \hookrightarrow L^2(X)$ is compact as well. Thus $j(B)$ is compact. But for $\omega \in \ker\Delta$,

$$||\omega||_{\ker\Delta} = ||\omega||_{W^{2,2}} = ||\omega||_{L^2} + ||d\omega||_{W^{1,2}} + ||d^*\omega||_{W^{1,2}}$$

$$= ||\omega||_{L^2}$$

Thus $\ker\Delta$ is also closed under the L^2 topology, and in fact the closed unit balls are exactly the same! thus, the Sobolev embedding theorem is just stating that B is compact in the $W^{2,2}$ topology, which means that $\ker\Delta$ is finite dimensional. The stated result is then an application of the previous lemma. □

The nice thing about the above proof is that we did not require any theory of elliptic operators. For the next two parts, we will however need that theory.

2.2 Elliptic Regularity: Image Δ

In order to continue any further, we must use the deeper theory of elliptic operators. This is to be expected, since otherwise there would be nothing special about Δ, and we would be able to prove a Hodge theorem for almost any old operator. There are many wonderful intrinsic definitions of elliptic operators. Here, we will do just enough as is necessary to convince the wanting reader that Δ is elliptic, following the proof given by Warner in [7].

First, we introduce them over $\mathbb{R}^n$. We start with a function $u : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^m$, such that $u(x, \zeta)$ is smooth in x and polynomial in $\zeta = (\zeta_i)_i$ (by which I mean each coordinate function is polynomial in ζ). Then there is an associated operator L which is obtained by formally replacing each ζ_i with ∂_i. This is called a differential operator, and takes as input a vector field $f : \mathbb{R}^n \to \mathbb{R}^m$ and outputs another such vector field.

The *degree* of the operator is the highest degree of u as a polynomial in ζ. Since it is a polynomial in ζ, we can consider a matrix consisting of just the top degree terms, discarding any term with degree lower than the highest. The resulting homogenous polynomial is called the *principle symbol*, often denoted as $\sigma(L)$. For example, $\sigma(\partial_1 + \partial_2\partial_3 + \sin(x)\partial_1^2) = \zeta_2\zeta_3 + \sin(x)\zeta_1^2$.

Definition. *Let L be a differential operator. Then L is elliptic at the point x_0 if $\sigma(L)(x_0) : \mathbb{R}^n \to \mathbb{R}^m$ is invertible at every point $0 \neq \zeta \in \mathbb{R}^n$.*

An intuition behind elliptic operators is that they are differential operators which are "almost" invertible. This will turn out to be rigorously true in some well defined sense later as we will see in the applications section. Next we have to extend our definition of ellipticity to manifolds. To do this we will use the characterization:

Lemma 4 ([7])**.** *a degree d differential operator L acting on vector fields $\mathbb{R}^n \to \mathbb{R}^m$ is elliptic at x_0 iff for every smooth $\phi : \mathbb{R}^n \to \mathbb{R}$ and $u : \mathbb{R}^n \to \mathbb{R}^m$, with in addition $u(x_0) \neq 0$, $\phi(x_0) = 0$, $d\phi(x_0) \neq 0$.*

$$L(\phi^d u)(x_0) \neq 0$$

This finally lets us talk about differential operators on $\Omega(X)$. Whenever we take a local coordinate chart of X, the elements of $\Omega(X)$ are vector fields. We say, then, that an operator $\Omega(X) \to \Omega(X)$ is differential if it, locally, takes the form of a differential operator on $\mathbb{R}^n \to \mathbb{R}^m$. Similarly, an operator is elliptic if it is locally elliptic. This would normally require taking coordinate systems, which is why lemma 4 is so nice. The reason why the Laplace-Beltrami operator is elliptic ultimately boils down to its symmetry for a completely algebraic reason.

Lemma 5. *Let A, B, C be finite dimensional inner product spaces. Then if*

$$A \xrightarrow{T} B \xrightarrow{V} C$$

Is an exact sequence, then $TT^* + V^*V$ is an automorphism of B.

Proof. Since B is finite-dimensional, we only have to prove injectivity. Let $0 \neq b \in B$. Then

$$\langle (TT^* + V^*V)b, b\rangle = \langle T^*b, T^*b\rangle + \langle Vb, Vb\rangle = ||T^*b||^2 + ||Vb||^2$$

If $Vb \neq 0$, we are done, since the norm will be positive. Otherwise, $Vb = 0$. In that case, then by exactness, $b = Ta$ for some $a \in T$. Thus,

$$\langle T^*b, a\rangle = \langle Ta, Ta\rangle = ||Ta||^2 = ||b||^2 > 0 \implies T^*b \neq 0$$

Thus either way the first inner product is non-zero, and thus it must be that $(TT^* + V^*V)b \neq 0$. □

This local automorphism is what we want to show "happens" pointwise on X.

Theorem 5. $\Delta : \Omega(X) \to \Omega(X)$ *is elliptic.*

Proof. We will not show that it is a differential operator, but it is not hard. The idea is that d and d^* manifest in coordinate charts as taking partial derivatives, which is exactly what makes an operator differential over $\mathbb{R}^n$. It is degree 2 — each d or d^* naturally counts one degree. As for showing that it is elliptic, we use lemma 4. Let $x \in X, \phi, u$ satisfy the conditions of the lemma. Then we want to show that

$$\Delta(\phi^2 u)(x) \neq 0$$

Remembering that ϕ is a scalar function, we calculate

$$d \star d \star (\phi^2 u)(x) = d \star [(d\phi^2) \star u \pm \phi^2 d \star u](x) = (d \star (d\phi^2) \star u)(x)$$

the last equality because $\phi^2(x) = 0$

$$= (d \star 2\phi(d\phi) \star u)(x) = (2(d\phi) \star (d\phi) \star u)(x) = 2(d\phi) \star (d\phi)u(x)$$

Similarly, $\star d \star d(\phi^2 u)(x) = 2 \star (d\phi) \star (d\phi)u(x)$. In this calculation we have suppressed the wedges $\wedge$ for simplicity. Thus

$$\Delta(\phi^2 u)(x) = 2[\pm(d\phi) \star (d\phi) \pm \star(d\phi) \star (d\phi)]u(x)$$

Then we have in the spirit of lemma 5 the exact sequence

$$\bigwedge^{k-1}(T_x^* X) \xrightarrow{\wedge d\phi} \bigwedge^{k}(T_x^* X) \xrightarrow{\wedge d\phi} \bigwedge^{k+1}(T_x^* X)$$

formed by wedging with $d\phi \in \bigwedge^1(T_x^* X)$. It is easy to see that the adjoint of $d\phi$ is $\pm \star (d\phi)\star$ by the construction of the hodge star. Thus the lemma gives us that since $0 \neq u(x) \in \bigwedge^k(T_x^* X)$, so too is $0 \neq [d\phi d\phi^* + d\phi^* d\phi]u(x) = [\pm(d\phi) \star (d\phi) \pm \star(d\phi) \star (d\phi)]u(x)$ □

Remark. We used $\pm$ here because the signs do work out properly and it is not nearly the main focus of the proof, but if one is worried they can go through and track them more carefully using the hodge star.

This proof of ellipticity, while not the most enlightening, showcases heavily the algebraic nature of Δ through lemma 5. Elliptic complexes generalize these properties, and more can be learned about them in [6]. Elliptic operators have very strong regularity conditions, including the one below, which we will not prove:

Theorem 6 (Elliptic Regularity, [7])**.** *For L an order m elliptic operator, there is some $C > 0$ s.t.*

$$||\omega||_{W^{k+m,2}} \leq C(||L\omega||_{W^{k,2}} + ||\omega||_{W^{k,2}})$$

In particular, if $\omega, L\omega \in W^{k,2}$, then $\omega \in W^{k+m}$. This is valid also for $k = 0$.

This theorem is just one manifestation of elliptic regularity; there are many more versions of it which are just beyond our garden walls. Other than just being independently interesting, it also allows us to conclude the section!

Theorem 7. Image Δ *is closed*

Proof. Let $\Delta\omega_k \to \eta$ in L^2, and assume that $(\omega_k)_k$ is bounded in $W^{2,2}$. Then by Elliptic regularity,

$$||\omega_n - \omega_m||_{W^{2,2}} \leq C(||\Delta\omega_n - \Delta\omega_m||_{L^2} + ||\omega_n - \omega_m||_{L^2})$$

Since $\Delta\omega_k \to \eta$, $(\Delta\omega_k)_k$ is Cauchy in L^2. Now, since $(\omega_k)_k$ is bounded in $W^{2,2}$, there is a weakly-convergent subsequence by Banach-Alaoglu, and under the Sobolev embedding, that subsequence converges strongly in L^2. Thus the RHS of the inequality goes to 0 under the subsequence, so the subsequence $(\omega_{j_k})_j$ is Cauchy in $W^{2,2}$; let ω be its limit. Then it must be that $\Delta\omega = \eta$ by continuity.

Now assume for the sake of contradiction that $(\omega_k)_k$ is not a bounded sequence. We can assume WLOG that $\omega_k \in (\ker\Delta)^\perp$. After passing to subsequence if necessary, we can assume that $||\omega_k||_{W^{2,2}} \to \infty$. Then let $\beta_k = \omega_k/||\omega_k||_{W^{2,2}}$. Now, β_k is bounded in $W^{2,2}$, and $||\Delta\beta_k||_{L^2} = ||\Delta\omega_k||/||\omega_k||_{W^{2,2}} \to 0$, since $||\Delta\omega_k|| \to ||\eta||$ is bounded. Thus again by the first part of the proof, $(\beta_k)_k$ is cauchy (after potentially passing to subsequence). Let it approach β. Then by continuity $\Delta\beta = 0$, so $\beta \in \ker\Delta$. But $(\beta_k)_k \subset (\ker\Delta)^\perp$, which is closed! Thus $\beta \in \ker\Delta \cap (\ker\Delta)^\perp \implies \beta = 0$. But this contradicts the fact that $||\beta_k||_{W^{2,2}} = 1$ for all k, so $(\omega_k)_k$ was bounded all along. □

2.3 Distributions: $\ker\Delta^*$

We know that on $\Omega(X)$, Δ is formally self-adjoint, since it is of the form $dd^* + d^*d$ (it is quite literally symmetric!). So it makes sense that Δ is equal to Δ^* whenever possible, and particularly that their kernels are the same. In this way, symmetry on $\Omega(X)$ extends the furthest that it makes sense. To formalize this, we introduce the following lemma:

Lemma 6. *If $\omega, \eta \in W^{2,2}$ then*

$$\langle\Delta\omega, \eta\rangle_{L^2} = \langle\omega, \Delta\eta\rangle_{L^2}$$

Proof. We suppress the subscript L^2 for now. $\Omega(X)$ is dense in $W^{2,2}$, so we can approximate η with smooth η_n in $W^{2,2}$ and thus also in L^2. $d\eta, d^*\eta$ are both smooth as well, so by the definition of weak-differentiability,

$$\begin{aligned}\langle\Delta\omega, \eta_n\rangle = \langle(dd^* + d^*d)\omega, \eta_n\rangle = \langle(dd^* + d^*d)\omega, \eta_n\rangle \\ = \langle\omega, (dd^* + d^*d)\eta_n\rangle = \langle\omega, \Delta\eta_n\rangle\end{aligned}$$

So

$$\begin{aligned}|\langle\Delta\omega, \eta\rangle - \langle\omega, \Delta\eta\rangle| &\leq |\langle\Delta\omega, \eta\rangle - \langle\Delta\omega, \eta_n\rangle| + |\langle\Delta\omega, \eta_n\rangle - \langle\omega, \Delta\eta\rangle| \\ &= |\langle\Delta\omega, \eta - \eta_n\rangle| + |\langle\omega, \Delta\eta_n\rangle - \langle\omega, \Delta\eta\rangle| \\ &= |\langle\Delta\omega, \eta - \eta_n\rangle| + |\langle\omega, \Delta(\eta_n - \eta)\rangle|\end{aligned}$$

By continuity of Δ and the inner product, as $\eta_n \to \eta$, this draws to 0. Thus $\langle\Delta\omega, \eta\rangle = \langle\Delta\omega, \eta\rangle$ on $W^{2,2}$. □

Lemma 7. $\ker\Delta^* \cap W^{2,2} = \ker\Delta \cap W^{2,2} = \ker\Delta$.

Proof. If $\omega \in \ker\Delta^* \cap W^{2,2}$, then by the previous lemma, for any $\eta \in W^{2,2}$,

$$\langle\eta, \Delta\omega\rangle = \langle\Delta\eta, \omega\rangle = \langle\eta, \Delta^*\omega\rangle = 0$$

Since $W^{2,2}$ is dense in L^2, it must be that $\Delta\omega = 0$ in L^2. The other inclusion is exactly the same way, and the final equality is because $\ker\Delta \subset \Omega(X)$. □

The only remaining obstruction is the possibility that there is some non-differentiable $\eta \in L^2$ or $W^{1,2}$ with $\Delta^*\eta = 0$. This should be surprising because of elliptic regularity. Unfortunately, the form of elliptic regularity we currently have isn't quite enough to extend elliptic regularity from Δ to Δ^*. What we need is a form of elliptic regularity for distributions. Fortunately, it extends essentially without change. That is, if $\Delta\eta = \omega$ in a distributional sense, and ω is smooth, then η can be lifted from a distribution to also being a smooth form. In particular, if $\Delta\eta = 0$ as a distribution, then $\eta \in \Omega(X)$ (see [7] once more).

Theorem 8. $\ker \Delta^* \subset \Omega(X)$.

Proof. If $\eta \in \ker \Delta^*$, then for all $\omega \in W^{2,2}(X)$, $\langle \Delta^*\eta, \omega\rangle = 0 \implies \langle \eta, \Delta\omega\rangle = 0$ for all $\omega \in W^{2,2}$. This is exactly the statement that as a distribution, $\Delta\eta = 0$ (Δ of a distribution is defined by the relation $(\Delta\eta)(\omega) = \langle \eta, \Delta\omega\rangle$).

Then by elliptic regularity, η lifts to an actually smooth solution $\Delta\eta = 0$ which agrees with η as a distribution; but this means that $\eta \in L^2(X)$ is actually in $\Omega(X)$. □

This, combined with lemma 7, finishes the proof that $\ker \Delta = \ker \Delta^*$.

Remark. One can think about the elliptic regularity that we applied here as an extension of the ones before. We start by defining, for $k > 0$, $W^{-k,2} := (W^{k,2})^*$. Then, as in the proof of the above, we define $\Delta : W^{-k,2} \to W^{-k-2,2}$ by

$$(\Delta\eta)(\omega) = \eta(\Delta\omega)$$

By Riesz-representation, we can think of η as an inner product against some element of $W^{k,2}$, so we can likewise define it as

$$(\Delta\eta)(\omega) = \langle \eta, \Delta\omega\rangle$$

Just as we did in the proof of the theorem. By defining $W^{-k,2}$ in this way, it is a wonderful fact that elliptic regularity extends without change to allow k negative. This expression of elliptic regularity is covered in [1].

2.4 Putting it all together

From our main corollary, we know that

$$L^2(X) = \ker \Delta^* \oplus \overline{\mathrm{Image}\,\Delta}$$

We from sections 2.2 and 2.3, we can rewrite this as

$$= \ker \Delta \oplus \mathrm{Image}\,\Delta$$

Now, we proved in lemma 3 that $\ker \Delta \subset \Omega(X)$. So we can take the intersection with $\Omega(X)$ for

$$\Omega(X) = \ker \Delta \oplus (\mathrm{Image}\,\Delta \cap \Omega(X))$$

$$= \mathscr{H} \oplus (\mathrm{Image}\,\Delta \cap \Omega(X))$$

Clearly $\Delta(\Omega(X)) \subset \mathrm{Image}\,\Delta \cap \Omega(X)$. For the other inclusion, if $\omega = \omega_1 + \Delta\omega_2$ for $\omega_1 \in \ker \Delta$, then $\Delta\omega_2 = \omega - \omega_1 \in \Omega(X)$. But then by lemma 2, $\omega_2 \in \Omega(X)$. Thus $\mathrm{Image}\,\Delta \cap \Omega(X) \subset \Delta(\Omega(X))$ as well.

Thus, can in fact write that $\Omega(X) = \mathscr{H} \oplus \Delta(\Omega(X))$, proving Hodge decomposition.

3 Applications

Having finally finished our stroll through the garden, it is time to head inside and enjoy a nice Charcuterie board of applications.

Theorem 9 (The Hodge Theorem)**.** *Every de Rham cohomology class is represented by a unique harmonic form. In particular, the de Rham cohomology is finitely generated.*

Proof. Let $[\alpha] \in H^k(X)$ be represented by a closed form α. Then by the Hodge theorem, $\alpha = \gamma + \Delta\beta$ for γ harmonic. Since α is closed,

$$0 = d\alpha = d\gamma + d\Delta\beta = d(dd^*\beta) + dd^*d\beta = dd^*d\beta$$

Where we use the fact that harmonic forms are closed and that $d^2 = 0$. But then,

$$||d^* d\beta||^2 = \langle d^* d\beta, d^* d\beta \rangle = \langle dd^* d\beta, d\beta \rangle = \langle 0, d\beta \rangle = 0$$

Thus actually $d^* d\beta = 0$. Thus we can conclude that

$$\alpha - \gamma = d(d^* \beta)$$

So α and γ belong to the same cohomology class. For uniqueness: assume that $\eta - \gamma = d\omega$ with η, γ both harmonic. Then

$$\begin{aligned} ||\eta - \gamma||^2 &= \langle \eta - \gamma, d\omega \rangle = \langle d^* \eta - d^* \gamma, \omega \rangle \\ &= \langle 0, \omega \rangle = 0 \end{aligned}$$

Where we use that harmonic forms are coclosed. □

Remark. Another method to prove finite generation of cohomology groups of smooth, closed manifolds is to show that they are finite CW complexes via Morse Theory, then apply cellular homology.

Theorem 10 (Poincare Duality made easy)**.** $H^k(M) \cong H^{n-k}(M)$

Proof. Since Cohomology classes are represented by unique harmonic forms, we can define the isomorphism by, for each element $x \in H^k(M)$, picking the unique harmonic $[\gamma] = x$, and mapping $[\gamma] \mapsto [\star\gamma]$. It is easy to see that this is linear, and that $\star\gamma$ is harmonic as well. Since $\star$ is invertible, this is an isomorphism. □

It is of course a classic part of harmonic analysis that Harmonic functions obey the maximum principle, and thus are constant on compact manifolds. Another, funnier proof of this is the following:

Corollary. *A harmonic function on a closed, connected orientable riemannian manifold is constant.*

Proof. We know that every constant function is harmonic, so the dimension of the space of harmonic functions is at least one dimensional containing the constant functions. However, harmonic functions are just smooth harmonic 0-forms, and the 0th de Rham cohomology of a connected manifold is 1 dimensional, so by the de Rham theorem, the space of harmonic functions is exactly 1 dimensional and so consists entirely of constant functions. □

Remark. By Poincare duality, this also shows us that the harmonic n-forms on an n-dimensional such manifold are exactly the constant multiples of the volume form! In the same way, one can conclude that the harmonic functions on a potentially disconnected closed manifold are always locally constant.

Also interesting from a functional analysis perspective is

Theorem 11 ([5])**.** *If* $L : W^{k+m,2} \to W^{k,2}$ *is an elliptic operator, then it is fredholm*

This formalizes the idea that elliptic operators are almost invertible from before; they are invertible exactly up to compact operators. In fact, we have actually shown this in the case of Δ! For indeed, we proved that $\ker \Delta$ is finite, that $\ker \Delta^*$ is finite (since it is equal to $\ker \Delta$), and that the image Image Δ is closed, thus that it is Fredholm. In fact, we showed that Δ is index 0. In that spirit:

Corollary (a Fredholm Alternative)**.** *For any closed Riemannian manifold X, for each k, exactly one of the following holds:*

1. *For any $\eta \in \Omega^k(X)$, there exists an $\omega \in \Omega^k(X)$ such that $\Delta\omega = \eta$.*
2. *There exists $0 \neq \omega \in \Omega^k(X)$ such that $\Delta\omega = 0$.*

Proof. The Hodge decomposition states that $\Omega(X) = \mathscr{H} \oplus \Delta(\Omega(X))$. This clearly restricts to $\Omega^k(X) = \mathscr{H}^k \oplus \Delta(\Omega^k(X))$ (where $\mathscr{H}^k$ is the space of harmonic k-forms), since if ω is a k-form, then $\Delta\omega$ is also a k-form. If $\dim \mathscr{H}^k = 0$, then $\Omega(X) = \Delta(\Omega(X))$ and so the first alternative holds, and if $\dim \mathscr{H}^k \neq 0$, the second holds. □

Last but not least, we can now discuss how Harmonic forms behave under different sorts of topological manipulations.

Lemma 8. *Let X, Y be two closed, orientable Riemannian manifolds, and equip $X \times Y$ with the product metric. Let π_X be the projection onto X and π_Y onto Y. Then for $\alpha \in \Omega(X)$, $\beta \in \Omega(Y)$,*

$$\star(\pi_X^*(\alpha) \wedge \pi_Y^*(\beta)) = \pi_X^*(\star\alpha) \wedge \pi_Y^*(\star\beta)$$

That is to say, the star operator works well with product manifolds. The proof for this is not too hard and can be found on stack exchange. This, along with the Hodge theorem, gives us a new type of Künneth formula:

Theorem 12 (a Künneth formula)**.** *Let X, Y be two closed, orientable, Riemannian manifolds and let $X \times Y$ have the product metric wrt X and Y. Then there is an isomorphism given by*

$$\mathscr{H}(X) \otimes \mathscr{H}(Y) \to \mathscr{H}(X \times Y)$$

$$\alpha \otimes \beta \mapsto \pi_X^*(\alpha) \wedge \pi_Y^*(\beta)$$

In other words, the harmonic forms on $\mathscr{H}(X \times Y)$ are exactly the products of harmonic forms on X and Y separately.

Proof. First, we show the map is well-defined. Now, if α, β are harmonic in X and Y, then we want to show that $d[\pi_X^*(\alpha) \wedge \pi_Y^*(\beta)] = d^*[\pi_X^*(\alpha) \wedge \pi_Y^*(\beta)] = 0$ according to lemma 2. Indeed,

$$d[\pi_X^*(\alpha) \wedge \pi_Y^*(\beta)] = [d\pi_X^*(\alpha)] \wedge \pi_Y^*(\beta) \pm \pi_X^*(\alpha) \wedge [d\pi_Y^*(\beta)]$$

$$= [\pi_X^*(d\alpha)] \wedge \pi_Y^*(\beta) \pm \pi_X^*(\alpha) \wedge [\pi_Y^*(d\beta)] = 0 \wedge \pi_Y^*(\beta) \pm \pi_X^*(\alpha) \wedge 0 = 0$$

The last equality since $d\alpha, d\beta = 0$ since they themselves are harmonic. Next, by the previous lemma,

$$d^*[\pi_X^*(\alpha) \wedge \pi_Y^*(\beta)] = \pm \star d \star [\pi_X^*(\alpha) \wedge \pi_Y^*(\beta)]$$

$$= \pm \star d[\pi_X^*(\star\alpha) \wedge \pi_Y^*(\star\beta)] = \pm \star [\pi_X^*(d \star \alpha) \wedge \pi_Y^*(\star\beta) \pm \pi_X^*(\star\alpha) \wedge \pi_Y^*(d \star \beta)]$$

$$= \pm[\pi_X^*(\star d \star \alpha) \wedge \pi_Y^*(\star \star \beta) \pm \pi_X^*(\star \star \alpha) \wedge \pi_Y^*(\star d \star \beta)] = \pm\pi_X^*(d^*\alpha) \wedge \pi_Y^*(\beta) \pm \pi_X^*(\alpha) \wedge \pi_Y^*(d^*\beta) = 0$$

So indeed the product actually is harmonic. Next, the map is clearly linear. For injectivity, if $\pi_X^*(\alpha) \wedge \pi_Y^*(\beta) = 0$, then since $X \times Y$ carries the product metric, $\pi_X^*(\alpha)$ and $\pi_Y^*(\beta)$ are disjoint to each-other so either $\alpha = 0$ or $\beta = 0$ (more formally, at any point we can pick coordinates $(e_i)_1^n$ so that $\pi_X^*(\alpha)$ uses only the first $\dim X$ coordinates and $\pi_Y^*(\beta)$ only the others). Either way, $\alpha \otimes \beta = 0$ and so we are done. Lastly, surjectivity follows since by the Hodge theorem and the traditional Künneth formula on manifolds and the Hodge theorem, the dimensions of the two vector spaces are the same. ☐

References

[1] Jim Davis. Elliptic regularity and hodge theory. `https://jfdmath.sitehost.iu.edu/notes/Elliptic_Regularity.pdf`. [Online; accessed 6-January-2025].

[2] Daniel Estéves. The hodge theorem. `https://destevez.net/wp-content/uploads/2018/09/hodge.pdf`. [Online; accessed 6-January-2025].

[3] Robert C. Gunning. *An Introduction to Analysis.* Princeton University Press, 2018.

[4] Giovanni Leoni. *A first course in Sobolev spaces*, volume 105 of *Graduate Studies in Mathematics.* American Mathematical Society, 2009.

[5] Richard Melrose. Differential analysis lecture notes for 18.155 and 156. `https://math.mit.edu/~rbm/18-155-F17/Chapter6.pdf`, 2017. [Online; accessed 6-January-2025].

[6] Liviu Nicolaescu. Notes on the atiyah-singer index theorem. 2005.

[7] Frank W. Warner. *Foundations of Differentiable Manifolds and Lie Groups.* Graduate Studies in Mathematics. Springer Science & Business Media, 1983.