CSE 274 (WI 2022) Homework 4

Harmonic Functions

Exercise 4.1 — Harmonic functions (10 pts). Let M be a compact and connected Riemannian manifold. A function $u\colon M \to \mathbb{R}$ (*i.e.* a 0-form $u \in \Omega^0(M)$) is said to be **harmonic** if

$$-\Delta u = \delta du = 0. \tag{1}$$

(a) Show that if $\partial M = \varnothing$ then every harmonic function is constant. (You can use the fact that u is constant if and only if $du = 0$ if and only if $\|du\|^2 = 0$.)

(b) Give an example of a non-constant harmonic function when M has boundary. (You can pick your own M.)

Hint *$\langle\langle d\alpha, \beta\rangle\rangle = \langle\langle \alpha, \delta\beta\rangle\rangle$ if M has no boundary.*

Foucault Pendulum

The Foucault pendulums are large and heavy pendulums displayed in several science museums over the world. They swing like a usual pendulum in a vertical plane. After several hours (say a day), this vertical plane will turn with a certain angle. This is the result of the Gauss–Bonnet theorem for parallel transported tangent vector.

The direction of the pendulum is a tangent vector of the earth. The rotation of the globe brings the entire device (the museum building) along the circle of constant latitude counted from the equator), and in the least dissipative manner (adiabatic process) the tangent vector (pendulum's swinging direction) is parallel transported along this circle.

Exercise 4.2 — Foucault pendulum (10 pts). UCSD is located approximately at 32.9° N, 117.2° W. If we set up a Foucault pendulum in UCSD, in how many degrees (clockwise or counterclockwise?) would the swing direction rotate after a full rotation of the earth (*i.e.* in 23 hours 56 minutes 4.091 seconds)?

Connections on Vector Bundles

We have a descent idea of what a **tangent vector field** on a manifold is. A tangent vector field $X \in \Gamma(TM)$ on a manifold M is an assignment of a tangent vector $X_p \in T_pM$ at each point $p \in M$. Since at different points $p, q \in M$ the vector spaces T_pM, T_qM are two different spaces, a priori there isn't a comparison like "$X_q - X_p$" we can make to measure derivative of a tangent vector field. That's why we have **covariant derivative**, also called **connection derivative**, for measuring rate of change of a tangent vector field. On a manifold with metric (Riemannian manifold), there is a canonical Levi-Civita connection for this connection. The discovery of the uniqueness of the Levi-Civita connection (fundamental theorem of Riemannian geometry) basically removes all the ambiguity of what the rate of change of a tangent vector means.

What about a more general vector-valued field that is not necessarily tangent vector field? In the general case there is not a unique connection like Levi-Civita connection.

A **vector bundle** E over M is the union of vector spaces $E = \bigcup_{p\in M} E_p$ based at p. (Of course, an example is $E_p = T_pM$ and $E = TM$.) A **section** $\psi \in \Gamma(E)$ is an assignment of $\psi_p \in E_p$ per point $p \in M$. (For example, a tangent vector field is a section of the tangent bundle $E = TM$.)

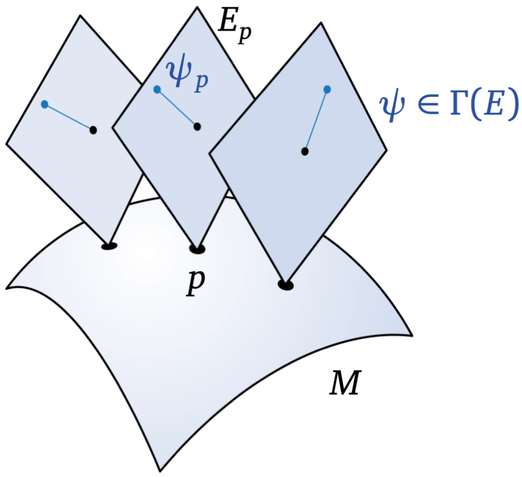

Figure 1 Vector bundle E and a section $\psi \in \Gamma(E)$.

Definition 1 — Connection. A **connection** is an operator

$$\nabla_{(\#1)}(\#2)\colon TM \times \Gamma(E) \to \Gamma(E) \tag{2}$$

such that

- $\nabla_{(\#1)}(\#2)$ is **linear** in the first entry; *i.e.* $\nabla_{c_1X_1+c_2X_2}\psi = c_1\nabla_{X_1}\psi + c_2\nabla_{X_2}\psi$;
- $\nabla_{(\#1)}(\#2)$ is **additive** in the second entry; *i.e.* $\nabla_X(\psi_1 + \psi_2) = \nabla_X\psi_1 + \nabla_X\psi_2$;
- Scalar function multiplication obeys the **product rule** for the second entry of $\nabla_{(\#1)}(\#2)$; *i.e.* $\nabla_X(f\psi) = (d_Xf)\psi + f\nabla_X\psi$ for any $f\colon M \to \mathbb{R}$. Here $d_Xf = df(X)$ denotes the standard directional derivative of f.

The last property makes ∇ a bit "non-linear" as scalar (functions) don't just factor out. In fact it is the product rule that makes connection behave more like a derivative; otherwise it is just some pointwise linear operator $E_p \xrightarrow{\text{linear}} E_p$.

Definition 2 — Tensorial. An operator $A\colon \Gamma(E) \to \Gamma(E)$ is said to be **tensorial** if

- It is additive; *i.e.* $A(\psi_1 + \psi_2) = A\psi_1 + A\psi_2$;
- Scalar function multiplication factors out; *i.e.* $A(f\psi) = fA(\psi)$ for all $f\colon M \to \mathbb{R}$ and $\psi \in \Gamma(E)$.

$\nabla_X(\cdot)$ are not tensorial because of the product rule.

Exercise 4.3 — Offsets between connections (5 pts). Let $\nabla, \tilde{\nabla}$ be two connections for the vector bundle E. Define $A_{(\#1)}(\#2)\colon TM \times \Gamma(E) \to \Gamma(E)$ by the difference between the two connections:

$$A_X(\psi) := \tilde{\nabla}_X\psi - \nabla_X\psi. \tag{3}$$

Show that $A_X(\cdot)$ must be tensorial (*i.e.* check that it is additive and scalar function multiplication factors out).

This exercise shows that every connection $\tilde{\nabla}$ on E can always be expressed by $\nabla + A$ using any other connection ∇ and some endomorphism-valued 1-form A.[a]

[a] $A_X\psi$ is linear in the tangent vector X so it is a 1-form. $A_X(\cdot)$ at each point p is a linear map from E_p to E_p and hence an endomorphism.

Ok, so far everything is exciting. We can do calculus. Given a field $\psi \in \Gamma(E)$, and a differential operator ∇, we have a differential $\nabla\psi$. And what is $\nabla\psi$? It is a E-valued 1-form; it is a 1-form because it is to be *linearly* paired with some $X \in T_pM$ so that $\nabla_X\psi$ is some element in E_p. So, E-valued k-form seems important. We will use the notation

$$\Omega^k(M;E) := \text{space of } E\text{-valued } k\text{-forms.} \tag{4}$$

For example, $\psi \in \Gamma(E) = \Omega^0(M;E)$ and $\nabla\psi \in \Omega^1(M;E)$. Neat.

From here we can build the whole exterior calculus for differential forms just that everything is E-valued.

Definition 3 — Connection exterior derivative. Given a connection ∇, extend it to

$$d^\nabla : \Omega^k(M;E) \to \Omega^{k+1}(M;E) \tag{5}$$

so that

- For $\psi \in \Gamma(E) = \Omega^0(M;E)$, we have the base case $d^\nabla\psi = \nabla\psi$;
- Leibniz rule. For a real-valued k-form $\alpha \in \Omega^k(M;\mathbb{R})$ and an E-valued ℓ-form $\Psi \in \Omega^\ell(M;E)$, we have

$$d^\nabla(\alpha \wedge \Psi) = (d\alpha) \wedge \Psi + (-1)^k \alpha \wedge (d^\nabla \Psi). \tag{6}$$

Exercise 4.4 — (5 pts). Apply the Leibniz rule and show that for every $\alpha \in \Omega^k(M;\mathbb{R})$ and $\Psi \in \Omega^\ell(M;E)$

$$d^\nabla d^\nabla(\alpha \wedge \Psi) = \alpha \wedge d^\nabla d^\nabla \Psi. \tag{7}$$

That is, $d^\nabla d^\nabla$ is *tensorial*.

Definition 4 — Curvature of a connection. The **curvature** of a connection ∇ is the endomorphism-valued 2-form F^∇ defined by

$$F^\nabla \wedge \Psi = d^\nabla d^\nabla \Psi. \tag{8}$$

As a matter of fact, for a 2-dimensional surface M and for its tangent bundle TM with Levi-Civita connection ∇, the curvature tensor F^∇ is $F^\nabla = -\mathbb{i}K\sigma$ where $\mathbb{i}$ is the $90°$ rotation operator, K the Gaussian curvature and σ the area 2-form. (Can you see it using some Stokes theorem and some version of the Gauss–Bonnet theorem?)

Another fact. When there is a shift in connection $\tilde{\nabla} = \nabla + A$, the connection exterior derivative is changed into $d^{\tilde{\nabla}} = d^\nabla + A\wedge$, and the curvature is modified by (this is actually easy to check)

$$F^{\tilde{\nabla}} = F^\nabla + d^\nabla A + A \wedge A. \tag{9}$$

In particular, for complex line bundle with $\nabla = d + \mathbb{i}\alpha$, the curvature is $F^\nabla = \mathbb{i}d\alpha$.

If you are interested in science fiction here are some stories. In the **Yang–Mills theory** in the standard model of particle physics, the main energy is given by the L^2-norm of the term $d^\nabla A + A \wedge A$. You can also check that we always have $d^\nabla F = 0$, which is called the **2nd Bianchi identity**. (Hint, start with $d^\nabla d^\nabla d^\nabla \Psi = d^\nabla(F^\nabla \wedge \Psi) = F^\nabla \wedge d^\nabla \Psi$ and apply Leibniz rule.) If E is the tangent bundle of the 4D spacetime, and ∇ is the Levi-Civita connection, then $d^\nabla F = 0$ says that there is some quantity that is conserved (it is saying that a 2 form is closed after all). By equating that quantity (known as the Einstein tensor) with the energy and matter of the universe (which we also believe is conserved) then we obtain Einstein's field equation of general relativity. In summary, the state of the art physics is characterized in terms of some curvatures of some connections of some bundles over the world.

Now, in the end of this short quarter, you have learned exterior calculus and the most general bundle-valued exterior calculus. In the case where the connection for the bundle is not *flat*, we have some nonzero $d^\nabla \circ d^\nabla$ and that is the curvature. In the implementation part, we are once again reprising the Laplace matrix. But this time, we are building it using some d^∇ with curvature. What do we expect for the Dirichlet$^\nabla$ energy minimizing field?

Implementation Part (40 pts)

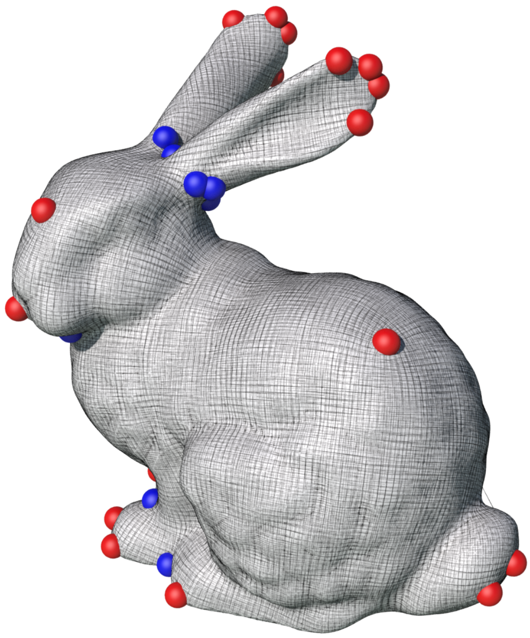

Figure 2 The ground state cross field on the bunny.

The implementation is the best way to understand the **complex line bundle** field design in geometry processing.

The goal is to

- (**40 pts**) Compute the "smoothest cross field" as shown in the figure above by the smallest eigenvalue problem of a complex Laplacian.
- (**Bonus +10 pts**) Find and visualize all the singularities of the cross field as shown in the figure.

Representation of the cross field

In the template file `hw4.hipnc`, the final cross field is generated by copying the cross geometry (or any quarter-turn symmetric geometry) to all the points of the mesh with a certain orientation defined pointwise. This orientation has been setup in the pointwrangle node "`setup_orientation`."

Indeed the cross orientation on each 2D tangent plane has only one rotation-about-normal degree of freedom. As shown there, a reference direction "`vector basis`" is given by the direction of a half-edge `i@hedge` sourced from this point. (The point attribute `i@hedge` has been set up for you; it is a pointer from point to vertex as an arbitrary half-edge sourced from this point.) With respect to this base direction "`basis`", the cross field's orientation is the rotation by angle `@phase/4` about the normal from this direction. The division by 4 in `@phase/4` makes it so that as `@phase` traverses from 0 to 2π, the actual rotation traverses from 0 to $\pi/2$ (quarter turn). Indeed the cross field returns to the original state after only a quarter turn.

Now go to the pointwrangle node "`select_eigenfunction`" where

$$\texttt{@phase} = \texttt{atan2}(\mathrm{Im}(\psi), \mathrm{Re}(\psi)) = \arg(\psi), \quad \textit{i.e. } \psi = re^{\mathring{\imath}(\texttt{@phase})} \tag{10}$$

for some $\mathbb{C}$ value ψ (with magnitude r). That is, the unknown `@phase` is now encoded in the complex-valued function ψ defined on points. Again, as `@phase` traverses from 0 to 2π, the cross field rotates by a quarter turn and returns to the original state, and simultaneously the complex variable ψ also traverses through all complex phases and returns to the original state.

Smallest eigenvalue problem of the connection Laplacian

To measure how the neighboring cross fields are parallel to each other, we take their difference: let ij be a half-edge

$$(d^\nabla \psi)_{ij} = e^{-\mathring{\imath}\alpha_{ij}}\psi_j - \psi_i \tag{11}$$

where the angle-valued 1-form α (the vertex attribute `f@A` in the Houdini file) is necessary to account for the fact that the reference directions (`i@hedge`) are also chosen incoherently on each point.

The crucial part is to implement a correct α otherwise the resulting cross field would just be a random mess.

Suppose a correct α is given, then minimizing

$$\sum_{(ij)\in\text{hedges}} \begin{pmatrix}\text{hedge}\\ \text{weight}\end{pmatrix}_{ij} \left|(d^\nabla\psi)_{ij}\right|^2 \tag{12}$$

is accomplished by finding the eigenfunction corresponding to the smallest eigenvalue of the eigenvalue problem

$$\underbrace{\overline{d^\nabla}^\intercal \star_1 d^\nabla}_{L}\, \psi = \lambda \star_0 \psi. \tag{13}$$

Here the conjugate-transpose $\overline{(\cdot)}^\intercal$ is called the Hermitian transpose. (In `scipy.sparse`, while `transpose()` gives the transpose of the matrix, `getH()` gives the Hermitian transpose of the matrix.)

Suppose you have α, you can modify the familiar Laplacian into this connection Laplacian (by modifying d to d^∇).

The connection α (`f@A`)

In the context of vector field (rather than cross field) the standard connection is the **Levi-Civita connection** given by the following.

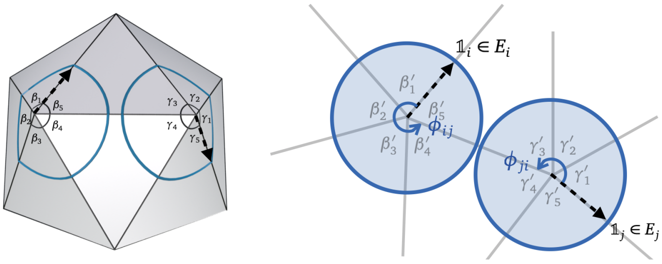

Let i, j be neighboring points connected by the half-edge ij as well as ji. Using the reference direction (labeled as $\mathbb{1}_i$ and $\mathbb{1}_j$) we can assign the angle ϕ_{ij} and ϕ_{ji} of the shared half edges ij and ji respectively. This angle ϕ_{ij}, ϕ_{ji} is the partial sum of the *scaled* interior angles (labeled as β'_k and γ'_ℓ) where $\beta'_k = \frac{2\pi}{\sum_\ell \beta_\ell}\beta_k$ and similarly for all other vertices.

Now insisting that the shared half-edge vectors $e^{\mathbb{i}\phi_{ij}}$ (in reference to $\mathbb{1}_i$) and $e^{\mathbb{i}\phi_{ji}}$ (in reference to $\mathbb{1}_j$) are parallel vectors but with opposite signs, we know that the Levi-Civita connection $\alpha_{ij}^{\mathrm{LC}}$ must satisfy

$$e^{\mathbb{i}\alpha_{ij}^{\mathrm{LC}}} e^{\mathbb{i}\phi_{ij}} = -e^{\mathbb{i}\phi_{ji}} \tag{14}$$

Since $-1 = e^{\mathbb{i}\pi}$ the above equation simply states

$$\alpha_{ij}^{\mathrm{LC}} = \phi_{ji} - \phi_{ij} - \pi \mod 2\pi. \tag{15}$$

How about the connection α for the cross field? Since a rotation $e^{\mathbb{i}\alpha}\psi$ by angle α to the complex number ψ amounts to only rotating $\alpha/4$ of the actual physical cross field in the tangent plane, we would need to set

$$\alpha_{ij} = 4\alpha_{ij}^{\mathrm{LC}}. \tag{16}$$

Input: A closed triangle mesh. A reference half edge per point.

1: For each point i, assign ϕ_{ij} as the angle of the half-edge ij relative to the reference half edge at i.
2: Compute the Levi-Civita connection $\alpha_{ij}^{\mathrm{LC}}$ for the tangent bundle.
3: Compute the connection for cross fields $\alpha_{ij} = 4\alpha_{ij}^{\mathrm{LC}}$.
4: Build the covariant derivative matrix d^∇ so that $(d^\nabla\psi)_{ij} = e^{-\mathbb{i}\alpha_{ij}}\psi_j - \psi_i$.
5: Build $\star_0, \star_1$ by the area weights and cotan weights as usual, and build $L = \overline{d^\nabla}^\intercal \star_1 d^\nabla$.
6: Solve the smallest eigenvalue problem $L\psi = \lambda \star_0 \psi$.
7: Let `@phase` be the complex phase of ψ. Then set the cross field direction as the one rotated from the reference half edge direction by angle `@phase/4`.

By the way...

the cross field is a quarter-turn symmetric object. You can replace the "4" by any n in the entire setup then you will obtain an n-direction field. An n-direction symmetric object will return to its original state after a ${}^{2\pi}/_{n}$ turn. For $n = 1$ we have vector field. For $n = 2$ we have bidirection field. $n = 3, 5, 6$ also gives interesting patterns.

A singularity of an n direction field is that, as we walk around that singularity, the field has turned $m \cdot {}^{2\pi}/_{n}$ for some integer m analogous to the Poincaré index for vector fields. The total sum of indices m will be $n\chi(M) = (2 - 2g)n$.

In fact, this hairyball formula does not prevent us from considering $n = {}^{1}/_{2}$, since the Euler characteristic is always even. The corresponding n-direction field would be objects that returns to original state only after 4π turn! (Quaternions are like that.) The n-direction fields with $n = {}^{1}/_{2}$ are called the *spinor fields*. This is what it means in quantum physics that some elementary particles are spin-${}^{1}/_{2}$; though in that context M is a 4D spacetime.