Globally Optimal Direction Fields

Felix Knöppel
TU Berlin

Keenan Crane
Caltech

Ulrich Pinkall
TU Berlin

Peter Schröder
Caltech

Abstract

We present a method for constructing smooth n-direction fields (line fields, cross fields, *etc.*) on surfaces that is an order of magnitude faster than state-of-the-art methods, while still producing fields of equal or better quality. Fields produced by the method are globally optimal in the sense that they minimize a simple, well-defined quadratic smoothness energy over all possible configurations of singularities (number, location, and index). The method is fully automatic and can optionally produce fields aligned with a given guidance field such as principal curvature directions. Computationally the smoothest field is found via a sparse eigenvalue problem involving a matrix similar to the cotan-Laplacian. When a guidance field is present, finding the optimal field amounts to solving a single linear system.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: discrete differential geometry, digital geometry processing, direction fields, curvature lines, singularities

Links: DL PDF

1 Introduction

A direction field φ of degree $n \in \mathbb{N}$ associates a collection of n evenly spaced unit tangent vectors to each point of a surface. For instance, $n = 1, 2$ and 4 correspond to direction, line, and cross fields, respectively (Fig. 2). In general such fields must have *singularities*, *i.e.*, isolated points where the field fails to vary smoothly.

At first glance, computing the smoothest n-direction field appears to be a difficult combinatorial optimization problem for two reasons. First, we must determine the optimal number, placement, and indices of singularities. Second, we must identify directions that differ in angle by integer multiples of $2\pi/n$. For a *fixed* configuration of singularities, Crane *et al.* [2010] demonstrate that an optimal solution can be found by solving a pair of sparse linear systems. In many situations, however, it is desirable to place singularities automatically. Historically this task has been formulated in terms of difficult nonconvex optimization problems where little can be said about global optimality (Sec. 1.1). In this paper we describe a simple quadratic smoothness energy that easily admits a global minimum with respect to all possible configurations of singularities. Fig. 1 shows one example.

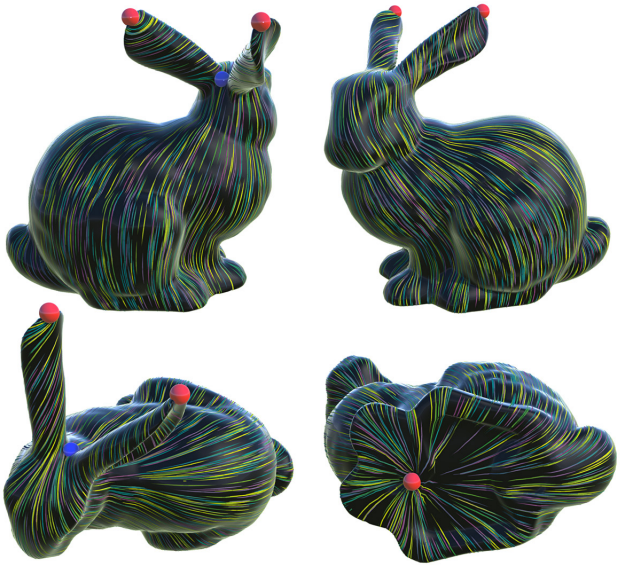

Figure 1: *Smoothest unit vector field on the Stanford bunny over all possible configurations of singularities, computed by solving a single eigenvector problem. Red and blue spheres indicate positive and negative singularities, respectively. (471ms, $|T| = 28k$)*

Our method has two key ingredients. **First**, we represent n-direction fields by storing the n^{th} power of a complex number at each vertex, together with an arbitrary (but fixed) tangent basis direction. Optimizing the smoothness of such a field does not require period jumps or trigonometric functions as in previous methods (Sec. 1.1). **Second**, we measure the smoothness of an n-direction field using the ground state energy of an appropriate *Schrödinger operator*. Unlike many methods, this formulation does not require a nonconvex unit-norm constraint on each vector, and is well-defined even for singular n-direction fields (Sec. 3). In addition, we introduce a continuum of "geometry-aware" smoothness energies that provide a tradeoff between the straightness of field lines and the total number of singularities. Finally, we allow a tradeoff between smoothness and alignment with a guidance field, which in the case of principal curvature alignment leads to a simple, automatic scheme without the need for careful tuning of parameters.

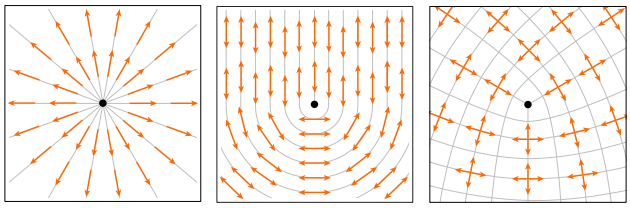

Figure 2: Left to right: *examples of n-direction fields for $n = 1$ (direction), $n = 2$ (line), and $n = 4$ (cross), near singularities of index $+1$, $+\frac{1}{2}$, and $+\frac{1}{4}$, respectively.*

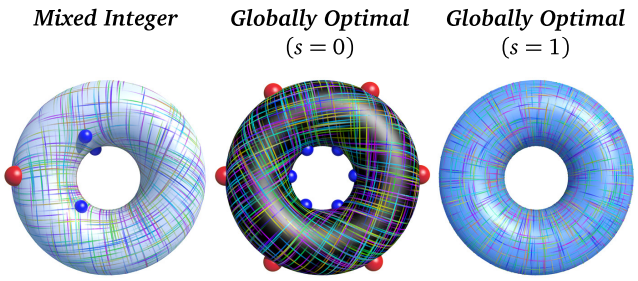

Figure 3: *Smoothness energies involving integer variables are easy to formulate but difficult to minimize.* Left: *solution produced by the mixed-integer method of Bommes* et al. *[2009].* Center: *solution produced by our method with* $s = 0$. Right: *solution produced by our method with* $s = 1$. *Both solutions are global minimizers; the parameter s offers a tradeoff between smoothness and number of singularities.*

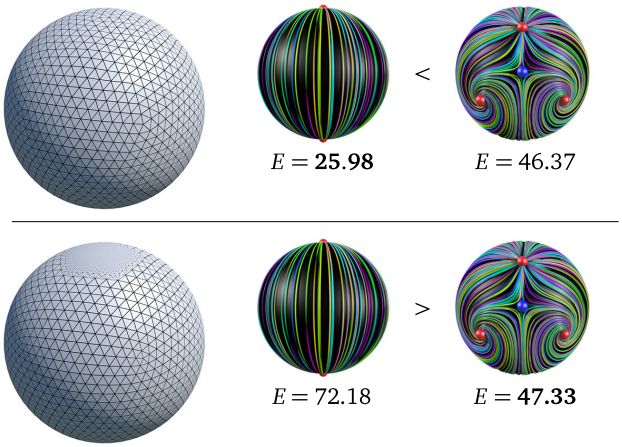

Figure 4: *Dirichlet energy of a* unit *n-direction field is not a reliable measure of quality, since under refinement the energy contributed by a singularity grows without bound.* Top: *uniform tessellation yields the expected result: antipodal singularities have lower energy than the configuration on the right.* Bottom: *refining around the poles increases the energy of the antipodal pair, reversing this relationship. This phenomenon is a consequence of the smooth formulation and is not fixed by,* e.g.*, more accurate area weighting.*

For triangle meshes our algorithm requires the setup of a matrix similar to the cotan-Laplacian, but with complex instead of real variables. To build this system we need a *connection*, *i.e.*, a prescription for mapping tangent vectors between tangent spaces, expressed via unit complex numbers (rotations) associated with edges. Finding an eigenvector corresponding to the smallest eigenvalue of this matrix gives a global minimizer of the smoothness energy. If alignment with a given field is desired, solving a single linear system yields a global minimizer of a weighted smoothness and alignment energy. An outline of these algorithms is given at the end of Sec. 6.

Throughout we will distinguish between *directions* which have unit length, and *vectors* which can have any length.

1.1 Related Work

Computation of smooth n-direction fields on surfaces, in particular the case $n = 4$, is an essential component in applications ranging from nonphotorealistic rendering [Hertzmann and Zorin 2000] to texture synthesis [Lefebvre and Hoppe 2006], parameterization [Ray et al. 2006], and remeshing [Kälberer et al. 2007; Bommes et al. 2009; Nieser et al. 2012], to name a few. In these applications one is often interested in smooth fields which approximate directions of *principal curvature*, *i.e.*, the most extreme directions of bending. While it is relatively straightforward to *formulate* energies which encode the desired effects, they are typically nonconvex and often NP-hard to minimize. Finding a simple *convex* formulation that yields the globally smoothest n-direction fields – in *any* well-defined sense – has so far eluded researchers in this area.

A principal difficulty is the need to identify directions modulo $2\pi/n$. One possible approach is to simply multiply angles by n. For example, Hertzmann and Zorin [2000] work with 4-direction fields and use an energy

$$E(\phi) = -\sum_{ij \in E} \cos\left(4(\phi_i - \phi_j + \theta_{ij})\right),$$

where angles ϕ_i specify directions at vertices and θ_{ij} is the angle between neighboring tangent frames. A variant of this "cosine" energy was later adopted by Kälberer *et al.* [2007]. Owing to nonconvexity, little can be said about optimality of solutions found via local descent; moreover, results depend on initialization.

Another possibility is to use the *representation vectors* $(u_i, v_i) := (\cos(n\phi_i), \sin(n\phi_i))$ themselves as variables [Ray et al. 2006; Palacios and Zhang 2007; Ray et al. 2009]. Now however an additional unit constraint $u_i^2 + v_i^2 = 1$ enters, leading to an NP-hard nonconvex problem [Ling et al. 2009]. In practice the unit constraint is included as a penalty term or relaxed through iterative renormalization. As before, little can be said about global optimality.

A different route was taken by Ray *et al.* [2008] who work with angles directly and use so-called *period jumps* $p_{ij} \in \mathbb{Z}$ to compare angles across a given edge e_{ij}, leading to a smoothness energy

$$E(\phi, p) = \sum_{e_{ij} \in E} \left(\phi_j - \phi_i + \theta_{ij} + \tfrac{2p_{ij}\pi}{n}\right)^2, \qquad (1)$$

which is quadratic in ϕ and p. Since the period jumps are integer variables, the feasible set of solutions is nonconvex. In general, mixed-integer optimization problems are NP-hard and a relaxation must be applied. Ray *et al.* apply a direct rounding procedure after first treating the p_{ij} as real valued, while Bommes *et al.* [2009] use a greedy rounding procedure which tends to produce solutions with lower energy. Again, global optimality of the resulting solutions cannot be asserted.

In all these approaches the user need not provide locations for the singularities (which must be present on surfaces of arbitrary topology). If singularities are prescribed ahead of time, very simple quadratic formulations are possible as demonstrated by Fisher *et al.* [2007] for vector fields and by Crane *et al.* [2010] for n-direction fields.

Our representation is perhaps closest to that of Kass and Witkin [1987] who used squared complex numbers to extract unoriented line fields (*i.e.*, $n = 2$) from images. Palacios and Zhang [2007] propose a similar idea for arbitrary n but use real representation vectors $(R\cos(n\phi), R\sin(n\phi))$ and encode transformations between tangent spaces via matrices. These representations are attractive, but neither work addresses the question of finding optimally smooth fields on surfaces.

2 n-Vector Fields on Surfaces

We begin with a description of our representation of n-vector fields on surfaces in the continuous setting; in Sec. 6 we discuss discretization on triangle meshes. Here and throughout ι denotes the imaginary unit. We use $|\cdot|$ to denote the magnitude of a complex number and "arg" to denote the angle it makes with the real axis. Recall that a vector $z \in \mathbb{C}$ rotated by an angle $\theta \in \mathbb{R}$ can be expressed as $e^{\iota\theta} z$.

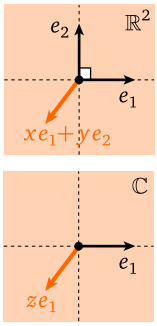

Real Planes and Complex Lines Traditionally, a tangent space T_pM over a point p of a surface M is viewed as a copy of the real Euclidean plane $\mathbb{R}^2$. In this case, any tangent vector can be expressed as a real linear combination of two basis vectors e_1, e_2, *i.e.*, $xe_1 + ye_2$ for some pair of coefficients $x, y \in \mathbb{R}$. A useful alternative is to think of T_pM as a copy of the complex numbers $\mathbb{C}$, in which case any vector can be written as a *complex* multiple of a *single* basis vector e_1. (For instance, any point in $\mathbb{C}$ can be expressed as some complex number z times the real unit vector 1.) In this sense, $\mathbb{C}$ is a *one*-dimensional complex vector space, also known as the *complex line*.

n-Vector Fields When viewed as a collection of complex lines, the tangent bundle TM is referred to as a *complex line bundle*, and a choice of unit basis vector $X_p \in T_pM$ at each point $p \in M$ is called a *basis section*. The *complex structure* J represents a 90-degree rotation in each tangent space. For surfaces in $\mathbb{R}^3$, J is induced by a quarter turn around the unit normal N, *i.e.*, for any tangent vector field Z we let $(JZ)_p := N_p \times Z_p$.

As suggested above, a vector field Z can be expressed as $Z = zX$ for some coefficient function $z : M \to \mathbb{C}$ relative to a basis section X. Likewise, any n-vector field ψ can be expressed as a collection of complex functions

$$\{e^{\iota 2k\pi/n} z, k = 0, \ldots, n-1\}$$

which describes n copies of Z, each rotated by some integer multiple of $2\pi/n$. Alternatively, we can raise any of these functions to the n^{th} power yielding the single complex function

$$u := z^n,$$

as depicted in Fig. 5. This function provides a concise representation of an n-vector field via $\psi = uX$. Moreover if $b = |u|$ and $\phi = \arg(u)$, we can then recover individual vector fields by computing the n^{th} roots

$$\{b^{1/n} e^{\iota(\phi/n + 2k\pi/n)}, k = 0, \ldots, n-1\}.$$

Note that these functions carry meaning *only with respect to the chosen basis section*. In particular, we cannot treat them as ordinary scalar functions, but instead have to consider *parallel transport*, as discussed below. Formally, an n-vector field is a section of the n^{th} order tensor product $L := TM^{\otimes n}$ and the basis section X is naturally identified with a basis section of L.

Parallel Transport The main benefit of this representation is that, at least within a given tangent space T_pM, we can measure the difference between two n-vectors via the simple quadratic expression $|u_1(p) - u_2(p)|^2$. When comparing vectors from different tangent spaces, however, we must be more careful – in particular,

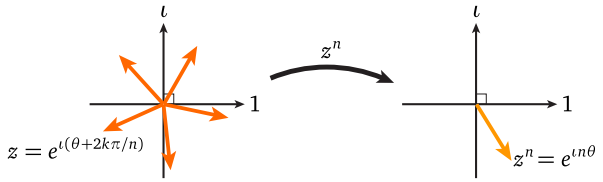

Figure 5: *We represent tangent vectors as complex numbers; directions that differ in angle by some multiple of $2\pi/n$ become indistinguishable when raised to the n^{th} power.*

we must first map both vectors into a common tangent space via *parallel transport*. More explicitly, let X_p and X_q be basis vectors for the two tangent spaces T_pM and T_qM, respectively, and let γ be a geodesic from p to q, as depicted in Fig. 6. If θ_p, θ_q are the angles made by γ with the two basis vectors and $\theta_{pq} := \theta_q - \theta_p$ is the difference between these angles, then the *parallel transport map* $P_{pq} : T_pM \to T_qM$ is given by

$$P_{pq}(z_p X_p) := e^{\iota\theta_{pq}} z_p X_q. \tag{2}$$

In other words, the coefficient z_p with respect to X_p gets mapped to the coefficient $e^{\iota\theta_{pq}} z_p$ with respect to X_q. Geometrically, the parallel transport map translates vectors from one tangent space to another without any in-plane rotation (Levi-Civita). The difference between two vectors is then defined as

$$|e^{\iota\theta_{pq}} z_p - z_q|^2,$$

and more generally we can write the difference between two n-vectors as

$$|(e^{\iota\theta_{pq}} z_p)^n - z_q^n|^2 = |e^{\iota n\theta_{pq}} z_p^n - z_q^n|^2 = |r_{pq} u_p - u_q|^2, \tag{3}$$

where for convenience we define the angle

$$\rho_{pq} := n\theta_{pq} \tag{4}$$

and the corresponding *transport coefficient*

$$r_{pq} := e^{\iota\rho_{pq}}.$$

Importantly, Eq. (3) remains quadratic in z since the coefficient r is constant, depending only on the geometry of M, the choice of basis section X, and the degree n of the field. Note that by introducing the Levi-Civita connection, we also introduce a dependence on the Riemannian metric.

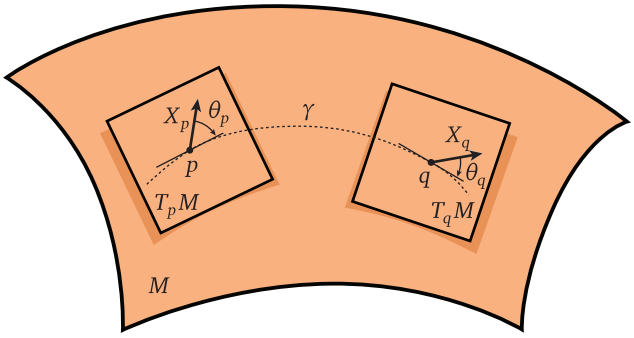

Figure 6: *To transport a vector from one tangent space to another it is sufficient to measure the angle it makes against a geodesic connecting the two tangent spaces.*

3 Smooth n-Direction Fields

We now focus on *n-direction* fields, *i.e.*, unit n-vector fields. The smoothness of a function is often measured via its *Dirichlet energy*; this energy can also be used for n-vector fields $\psi = uX$:

$$E_D(\psi) := \tfrac{1}{2}\int_M |\nabla\psi|^2 dA. \tag{5}$$

Here ∇ denotes the covariant derivative, *i.e.*, the Levi-Civita connection on M. For *n-direction* fields there are two main problems with this measure of smoothness. First, we must enforce the pointwise constraint $|u| = 1$, resulting in hard optimization problems as discussed in Sec. 1.1. Second, the Dirichlet energy of a *singular* n-direction field is not in general well-defined. In particular, consider that any vector field Z can be expressed as the product of a unit vector field $\tilde{Z}$ and a real nonnegative scale factor a, *i.e.*, $Z = a\tilde{Z}$. Since the change in a *unit* vector field is orthogonal to the field itself, we get $\nabla Z = (\nabla a)\tilde{Z} + a\omega J\tilde{Z}$, where ω is the rotation speed of Z. The Dirichlet energy of Z is then

$$\tfrac{1}{2}\int_M |\nabla Z|^2 dA = \tfrac{1}{2}\int_M |\nabla a|^2 + a^2|\omega|^2 dA = \tfrac{1}{2}\langle\langle(\Delta + |\omega|^2)a, a\rangle\rangle$$

where Δ is the positive semidefinite Laplace-Beltrami operator on M and $\langle\langle .,. \rangle\rangle$ denotes the L_2 inner product. For a *direction* field the scale factor is $a \equiv 1$, hence the Dirichlet energy is simply $\int_M |\omega|^2$. The rotation speed ω at a distance r from a singularity is proportional to $1/r$ (as $r \to 0$) since the total angle of rotation $2k\pi$ is divided by the circumference $2\pi r$ for a circle of radius r. *Consequently, the Dirichlet energy of a direction field with singularities is undefined (infinite).*

On a mesh, Dirichlet energy is finite but blows up under refinement. For instance, even if we add area weights to the energy $E(\phi, p)$ (Eq. (1)), refinement around singularities increases energy so much that visibly inferior solutions yield lower energy (Fig. 4). This observation motivates a different treatment of smoothness. In particular, we define the energy $\hat{E}$ of an n-direction field φ as the lowest Dirichlet energy among all n-vector fields $\psi = a\varphi$ parallel to φ. In other words, to evaluate the energy of a fixed unit field φ we look for an optimal scaling $a \geq 0$:

$$\hat{E}(\varphi) := \min_{a\geq 0, \|a\|=1} \int_M |\nabla(a\varphi)|^2 dA, \tag{6}$$

where the constraint $||a|| = 1$ prevents the trivial solution $a \equiv 0$.

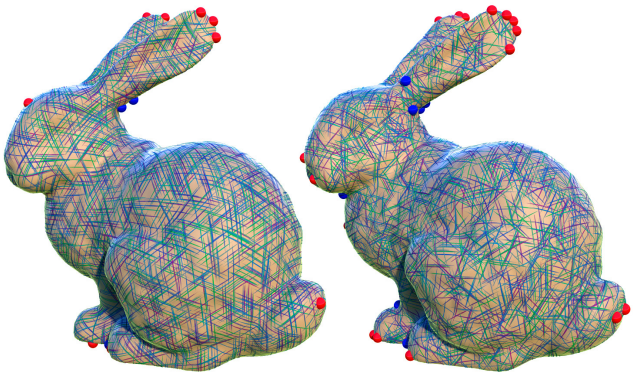

Figure 7: *Our algorithm can produce fields of any degree – here $n = 3$ (left) and $n = 5$ (right) using the holomorphic energy (Sec. 4).*

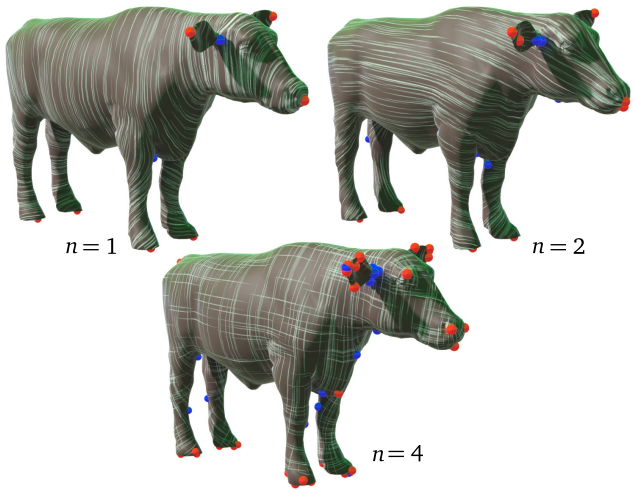

Figure 8: *Smoothest direction fields for varying n, as measured by the holomorphic energy. For higher n singularities arise in groups, as if splitting from those for smaller n. (382ms, $|T| = 25k$)*

At first glance this energy seems ill-defined, since the potential $|\omega|^2$ is unbounded for singular n-direction fields. However, minimizing the quadratic form $\langle\langle(\Delta + |\omega|^2)a, a\rangle\rangle$ over unit-norm functions a is equivalent to solving the eigenvalue problem

$$(\Delta + |\omega|^2)a = \lambda a$$

for the smallest eigenvalue λ and corresponding eigenfunction a. In physics, this equation is known as a *time-independent Schrödinger equation*, where the scalar function $|\omega|^2$ can be interpreted as the potential energy of a particle moving on the surface. The eigenfunction a with smallest eigenvalue is equivalent to the wave function of a free particle with lowest energy, which is always well-defined. From here, the globally optimal n-direction field is obtained by minimizing $\hat{E}(\varphi)$ over all unit fields φ:

$$\min_{|\varphi|=1} \hat{E}(\varphi) = \min_{|\varphi|=1}\left(\min_{a\geq 0, \|a\|=1}\int_M |\nabla(a\varphi)|^2 dA\right).$$

But since the set of rescaled n-direction fields is no different from the set of all n-vector fields, we can instead solve

$$\min_{\|\psi\|=1}\int_M |\nabla\psi|^2 dA,$$

which amounts to a smallest eigenvalue problem $\Delta\psi = \lambda\psi$, and recover the solution to our original problem via $a \leftarrow |\psi|$, $\varphi \leftarrow \psi/a$. In the end, we obtain a smoothest n-direction field (among all possible configurations of singularities) via a surprisingly simple procedure: find the smoothest n-vector field and normalize the resulting vectors. The physical interpretation of $\hat{E}(\varphi)$ also provides some insight into why the globally smoothest n-direction field has a small (though not necessarily minimal) number of singularities: adding too many poles to ω increases the potential energy $|\omega|^2$, hence it also increases the ground-state energy of the corresponding wave function. We emphasize, however, that we never explicitly build the Schrödinger operator $\Delta + |\omega|^2$ or minimize the energy $\hat{E}$ – these objects are introduced simply to analyze the intermediate variational problem. In practice we need only build the Laplacian Δ used in the final eigenvalue problem.

Fig. 8 visualizes the fields we obtain by varying n on the same shape; Fig. 7 shows more exotic examples.

4 Quadratic Smoothness Energies

The most commonly used smoothness energy is the Dirichlet energy E_D. However, we can obtain a richer family of energies by orthogonally splitting the covariant derivative into a sum of Cauchy-Riemann derivatives, *i.e.*, $\nabla\psi = \bar{\partial}\psi + \partial\psi$ where

$$\bar{\partial}_Z\psi := \tfrac{1}{2}\left(\nabla_Z\psi + J\nabla_{JZ}\psi\right), \quad \partial_Z\psi := \tfrac{1}{2}\left(\nabla_Z\psi - J\nabla_{JZ}\psi\right),$$

and Z is an arbitrary vector field (see [Wirtinger 1927] and [Ahlfors 1966, p. 27]). We call an n-vector field *holomorphic* if $\bar{\partial}_Z\psi = 0$ for all vector fields Z, and *anti-holomorphic* if $\partial_Z\psi = 0$ [Napier and Ramachandran 2011, Ch. 2.4]. These definitions mirror the standard notion that a function $f : \mathbb{C} \to \mathbb{C}$ is holomorphic (*i.e.*, angle- and orientation-preserving) if it satisfies the Cauchy-Riemann equation $0 = \bar{\partial} f = \frac{1}{2}(\frac{\partial}{\partial x} + \imath\frac{\partial}{\partial y})f = 0$.

Using the orthogonal splitting described above, the Dirichlet energy decomposes into holomorphic and anti-holomorphic terms E_H and E_A, respectively:

$$E_D(\psi) = E_H(\psi) + E_A(\psi) := \tfrac{1}{2}\int_M |\bar{\partial}\psi|^2 dA + \tfrac{1}{2}\int_M |\partial\psi|^2 dA.$$

We therefore define our smoothness energy as

$$E_s := (1+s)E_H + (1-s)E_A = E_D - s(E_A - E_H), \tag{7}$$

providing a continuum from anti-holomorphic ($s = -1$) to Dirichlet ($s = 0$) to holomorphic ($s = 1$) energy. Equivalently, the parameter s controls the deviation of E_s from the standard Dirichlet energy by the difference

$$E_A(\psi) - E_H(\psi) = \tfrac{1}{2}\int_M nK|\psi|^2 dA - \tfrac{1}{2}\int_{\partial M} \mathrm{Im}\langle\nabla\psi, \psi\rangle,$$

where K denotes Gaussian curvature (App. C). The potential in the Schrödinger operator now becomes $|\omega|^2 - \frac{s}{2}nK$, biasing singularities towards areas of high ($s < 0$) or low ($s > 0$) Gaussian curvature – this is the sense in which our smoothness energy is "geometry aware." More precisely, direction fields obtained for $s = \pm 1$ depend only on the conformal structure; for all other s, results depend on the metric. As shown in Fig. 9, the parameter s provides a useful tradeoff between number of singularities and "straightness" (*i.e.*, geodesic curvature) of integral curves.

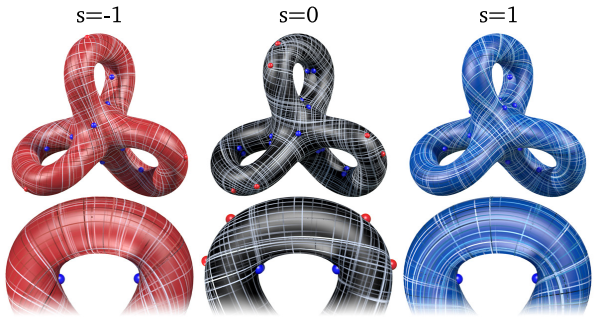

Figure 9: *When seeking a smooth vector field, one often faces a choice: fewer singularities, or straighter field lines? Our method offers a tradeoff between the two, determined by a continuously varying parameter s. As seen above, $s = 1$ (holomorphic energy) typically yields fewer singularities, whereas $s = 0$ (Dirichlet energy) gives straighter lines; $s = -1$ (anti-holomorphic energy) provides a compromise between the two. (405ms, $|T| = 25k$)*

Equation to Solve Let A be the positive semidefinite quadratic form given by $E_s(\psi) = \frac{1}{2}\langle\langle A\psi, \psi\rangle\rangle$. We find a global minimizer ψ of E_s among fields with unit norm by solving the problem

$$A\psi = \lambda\psi \tag{8}$$

for the smallest eigenvalue $\lambda \in \mathbb{R}$.

Discussion The solution ψ to Eq. (8) is unique up to a complex constant with unit norm – this constant determines the global "phase," *i.e.*, it rotates every tangent vector by the same angle. In the case where λ is a repeated eigenvalue we obtain one of many possible solutions; note that all such solutions are global minimizers of our smoothness energy. Alignment constraints (Sec. 5) can be used to specify additional criteria.

A 1-vector field Y on M is anti-holomorphic if and only if the associated real-valued 1-form $X \mapsto \langle Y, X\rangle$ is harmonic, where $\langle\cdot,\cdot\rangle$ denotes the Riemannian metric. Likewise, since vector fields on any open set $U \subset \mathbb{C}$ can be naturally identified with complex functions, we can say that Y is holomorphic whenever it can be expressed as a holomorphic function in each local coordinate chart. These relationships help us make the connection with existing methods from computer graphics. For instance, holomorphic *1-forms* (corresponding to anti-holomorphic 1-vector fields) appear in [Gu and Yau 2003] in the context of parameterization, while the holomorphic energy for *functions* was first used in the context of parameterization by Lévy *et al.* [2002] and Desbrun *et al.* [2002], where it was referred to as the conformal energy. Anti-holomorphic energy was used by Fisher *et al.* [2007] in a DEC [Desbrun et al. 2008] context to find smoothest 1-vector fields as minimizers of the Hodge-Laplacian $\frac{1}{2}\int_M |\nabla \times Z|^2 + |\nabla \cdot Z|^2 = \int_M |\partial Z|^2 = 2E_A(Z)$. The holomorphic energy of 1-vector fields also appeared in the work of Ben-Chen *et al.* [2010, Eq. 8], where the Killing energy of a 1-vector field is defined as $E_K(Z) = 4\int_M \frac{1}{2}|\nabla \cdot Z|^2 + |\bar{\partial} Z|^2$, encoding the fact that a Killing 1-vector field is both holomorphic and divergence free.

5 Alignment Energies

It is often desirable to balance smoothness and alignment with a given field ϕ. Alignment is accomplished via the functional

$$E_l(\psi) = \int_M \mathrm{Re}\left(\langle\phi, \psi\rangle\right) dA = \mathrm{Re}\left(\langle\langle\phi, \psi\rangle\rangle\right).$$

where ϕ is normalized so that $||\phi|| = 1$. We then let

$$E_{s,t}(\psi) := (1-t)E_s(\psi) - tE_l(\psi) \tag{9}$$

where $t \in [0, 1]$ controls the strength of alignment. Just as before we minimize $E_{s,t}$ over all fields ψ with $\|\psi\| = 1$. The pointwise magnitude $|\phi|$ determines the local weighting between alignment and smoothness terms. For example, setting $\phi = 0$ on regions with missing or unreliable information will smoothly interpolate data from regions where $|\phi| > 0$.

Equation to Solve Let A be the quadratic form corresponding to $E_{s,t}$. We find the global minimizer of $E_{s,t}$ by solving

$$(A - \lambda_t I)\widetilde{\psi} = \phi, \tag{10}$$

and normalizing the result, *i.e.*, $\psi \leftarrow \widetilde{\psi}/\|\widetilde{\psi}\|$ (App. A). Letting λ_1 be the smallest eigenvalue of A, the parameter $\lambda_t \in (-\infty, \lambda_1)$ controls the tradeoff between alignment ($\lambda_t \to -\infty$ for $t \to 1$) and smoothness ($\lambda_t \to \lambda_1$ for $t \to 0$), as illustrated in Fig. 10. See Appendix A for further discussion.

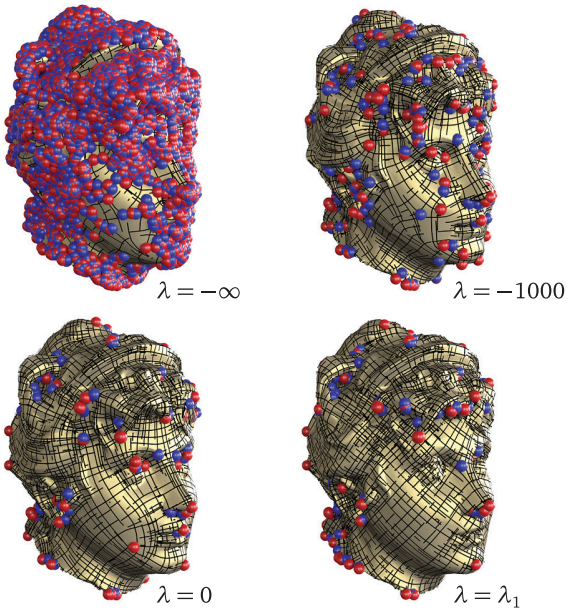

Figure 10: *Adding an alignment term, we can interpolate between a curvature aligned field (upper left) and a smooth field (lower right). Here $s = 0$ and λ_1 is the smallest eigenvalue; we use principal curvature as a guidance field. (498ms, $|T| = 65k$)*

6 Computation on Triangle Meshes

So far we have described our approach in the smooth setting. In this section we collect all details needed for numerical computation on triangle meshes, which are based on a finite element discretization using piecewise linear (PL) basis sections.

Triangle Meshes We assume that our input is an oriented 2-manifold triangle mesh of arbitrary topology with or without boundary. We use V, E, and T to denote the set of vertices v_i, edges e_{ij} and triangles t_{ijk}, respectively. Vertices have coordinates $p_i \in \mathbb{R}^3$, with the rest of the surface defined via linear interpolation.

Vertex Tangent Spaces and Parallel Transport Each vertex v_i carries a unit basis vector X_i (for convenience we use one of the incident edge directions) and a coefficient $u_i \in \mathbb{C}$ representing the associated n-vector with respect to this basis. To measure angles in tangent spaces, we rescale Euclidean angles at each vertex to sum to 2π [Polthier and Schmies 1998; Zhang et al. 2006]. This rescaling effectively "flattens" the vertices, pushing their curvature into the incident triangles $t_{ijk} \ni i$. Let

$$s_i := \frac{2\pi}{\sum_{t_{ijk} \ni i} \alpha_i^{jk}}, \tag{11}$$

where α_i^{jk} denotes the Euclidean angle at v_i in t_{ijk}, opposite e_{jk}. (On the boundary, $s_i := 1$.) For parallel transport from v_i to v_j let

$$\rho_{ij} = n\left(\theta_j(X_j, e_{ij}) - \theta_i(X_i, e_{ij})\right) \tag{12}$$

(*cf.* Eq. (4)), with $\theta_i(X_i, e_{ij})$ denoting the *rescaled Euclidean angle* from X_i to e_{ij} in T_iM (and similarly for θ_j). Each oriented edge of the mesh is assigned the transport coefficient $r_{ij} = e^{\imath\rho_{ij}}$, which is needed for matrix assembly.

Curvature Transporting an n-vector around the boundary of t_{ijk} computes the *holonomy* $\Omega_{ijk} \in (-\pi, \pi]$, which represents the curvature of our line bundle L over t_{ijk}:

$$e^{\imath\Omega_{ijk}} := r_{ij} r_{jk} r_{ki}. \tag{13}$$

For each triangle we compute this value as $\Omega_{ijk} = \arg(r_{ij} r_{jk} r_{ki})$, which is needed for matrix assembly (Sec. 6.1.1). Note that this definition implicitly assumes that curvature is constant in each triangle, and thus proportional to area – conceptually

$$\Omega_{ijk} = \int_{t_{ijk}} nK\,dA = nK_{ijk}|t_{ijk}| \tag{14}$$

where $|t_{ijk}|$ and K_{ijk} are the area and Gaussian curvature in triangle t_{ijk}.

PL Finite Elements For purposes of the finite element method the unit basis vectors at vertices are parallel transported into the incident triangles and attenuated linearly with the standard hat function. This gives PL basis sections Ψ_i supported on the triangles t_{ijk} incident to v_i.

From now on all fields ψ are piecewise linear, given as complex linear combinations of basis sections (App. D.1)

$$\psi = \sum_{v_i \in V} u_i \Psi_i,$$

with u now denoting the length $|V|$ vector of coefficients of ψ.

6.1 Algorithms

Smoothest Field The continuous problem of Eq. (8) turns into the generalized matrix eigenvector problem

$$Au = \lambda M u. \tag{15}$$

Here A is now a $|V| \times |V|$ Hermitian matrix representing the energy E_s with respect to the PL basis sections Ψ_i

$$A_{ij} = \langle\langle A\Psi_i, \Psi_j \rangle\rangle,$$

while M is the Hermitian mass matrix

$$M_{ij} = \langle\langle \Psi_i, \Psi_j \rangle\rangle.$$

Since we care about an eigenvector belonging to the smallest eigenvalue, we use an inverse power iteration to solve Eq. (15) – see Sec. 7 for further details.

Aligned Field The continuous problem of Eq. (10) turns into the matrix problem

$$(A - \lambda_t M)\tilde{u} = Mq, \tag{16}$$

with A and M as above, while q is the coefficient vector for the PL version of the guidance field ϕ. The unit vector u minimizing the discretized version of $E_{s,t}$ is given by $u = \tilde{u}/\|\tilde{u}\|$ for $t = (1 + \|\tilde{u}\|)^{-1}$. While λ_t depends on t monotonically, we have no closed form expression for the relationship. In practice we find that $\lambda_t = 0$ is often a good (starting) value. To solve this problem we use Cholesky factorization followed by back substitution.

6.1.1 Matrix Entries

Because the entries of A and M are given as integrals we can assemble them via a loop over all triangles, *i.e.*, a sum of integrals over t_{ijk}. The entries of these *local* 3×3 matrices can be computed in closed form (see App. D for their derivation). Writing $\langle\langle .,. \rangle\rangle_{ijk}$ the subscript reflects the integration over t_{ijk} only. Assembly into the global matrices through summation is assumed.

Mass Matrix On triangle t_{ijk} the local mass matrix induced by the L_2 inner product is given by (App. D.2)

$$M_{ii} = \langle\langle \Psi_i, \Psi_i \rangle\rangle_{ijk} = \tfrac{1}{6}|t_{ijk}|$$
$$M_{jk} = \langle\langle \Psi_j, \Psi_k \rangle\rangle_{ijk} = \bar{r}_{jk}|t_{ijk}| \frac{6e^{\imath\Omega_{ijk}} - 6 - 6\imath\Omega_{ijk} + 3\Omega_{ijk}^2 + \imath\Omega_{ijk}^3}{3\Omega_{ijk}^4}. \quad (17)$$

The right hand side fraction has a removable singularity in the limit of no curvature ($\Omega_{ijk} \to 0$) where it takes on the value $1/12$ as expected.

Energy Matrix We begin with the Laplacian, *i.e.*, the Dirichlet energy terms (App. D.3)

$$\Delta_{ii} = \langle\langle \nabla\Psi_i, \nabla\Psi_i \rangle\rangle_{ijk} = \frac{1}{4|t_{ijk}|}\left[|p_{jk}|^2 + \Omega_{ijk}^2 \frac{|p_{ij}|^2 + \langle p_{ij}, p_{ik}\rangle + |p_{ki}|^2}{90} \right]$$
$$\Delta_{jk} = \langle\langle \nabla\Psi_j, \nabla\Psi_k \rangle\rangle_{ijk}$$
$$= \frac{\bar{r}_{jk}}{|t_{ijk}|}\left[(|p_{ij}|^2 + |p_{ki}|^2) f_1(\Omega_{ijk}) + \langle p_{ij}, p_{ik}\rangle f_2(\Omega_{ijk}) \right],$$

where $p_{ij} := p_j - p_i$ is the edge vector along e_{ij} and

$$f_1(s) := \tfrac{1}{s^4}\left(3 + \imath s + \tfrac{s^4}{24} - \tfrac{\imath s^5}{60} + (-3 + 2\imath s + \tfrac{s^2}{2})e^{\imath s} \right)$$
$$f_2(s) := \tfrac{1}{s^4}\left(4 + \imath s - \tfrac{\imath s^3}{6} - \tfrac{s^4}{12} + \tfrac{\imath s^5}{30} + (-4 + 3\imath s + s^2)e^{\imath s} \right).$$

As before the singularities in f_1 and f_2 are removable and

$$\langle\langle \nabla\Psi_j, \nabla\Psi_k \rangle\rangle_{ijk} \to -\frac{\bar{r}_{jk}\langle p_{ij}, p_{ik}\rangle}{4|t_{ijk}|} = -\bar{r}_{jk}\tfrac{1}{2}\cot\alpha_i^{jk} \quad \text{for} \quad \Omega_{ijk} \to 0,$$

recovering the standard cotan-Laplacian in the flat setting [MacNeal 1949].

With this, the per triangle energy matrix as a function of the parameter $s \in [-1, 1]$ follows as (App. D.4)

$$A_{ii} = \Delta_{ii} - s\frac{\Omega_{ijk}}{|t_{ijk}|}M_{ii}$$
$$A_{jk} = \Delta_{jk} - s\left(\frac{\Omega_{ijk}}{|t_{ijk}|}M_{jk} - \epsilon_{jk}\frac{\imath\bar{r}_{jk}}{2} \right), \quad (18)$$

where $\epsilon_{jk} = \pm 1$ according to the orientation (positive or negative) of e_{jk} with respect to t_{ijk}.

Boundaries Assembling the energy matrix triangle by triangle we automatically account for the boundary of the mesh (Fig. 11). In the case of the Dirichlet energy ($s = 0$) this amounts to zero Neumann conditions, while for $s \neq 0$ we have the boundary terms due to $E_A - E_H$ (App. C).

Numerical Evaluation The expressions given above for the off-diagonal entries of the mass and Dirichlet matrices have removable singularities as $\Omega_{ijk} \to 0$. To avoid numerical difficulties and ensure efficient, reliable, and accurate (to machine precision) evaluation of these expressions, we employ Chebyshev expansions [Gil et al. 2007]. Due to the equiripple property of Chebyshev polynomials, these can guarantee an error on the order of the least significant bit in double precision over an entire interval. Since the expressions for M_{jk} and Δ_{jk} are only evaluated for $\Omega_{ijk} \in (-\pi, \pi]$ and due to the symmetry resp. anti-symmetry of their real resp. imaginary parts, it is sufficient to find Chebyshev expansions for arguments in $[0, \pi]$. We computed these expansions with the aid of Mathematica and include ready to use code in the .

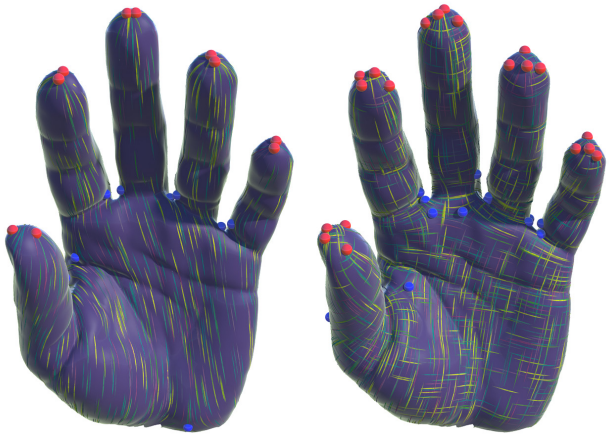

Figure 11: *Assembly of local matrices per triangle into the global matrix automatically accounts for boundary conditions. Here an example using the anti-holomorphic energy and curvature alignment ($\lambda = 0$) to the minimum (left; $n = 2$) and both (right; $n = 4$) principal curvature directions. (976ms, $|T| = 106k$)*

6.1.2 Curvature Alignment

An important practical example for the alignment energy $E_{s,t}$, seeks to compute smoothed versions of curvature lines ($n = 2$) and crosses ($n = 4$). In this case the alignment field

$$\phi = \sum_{v_i \in V} q_i \Psi_i$$

is a PL approximation of the Hopf differential, *i.e.*, the trace-free part of the shape operator. In the discrete setting it is only accessible as a distribution ϕ^δ concentrated along edges [Cohen-Steiner and Morvan 2003]. Pairing it with our basis sections Ψ_i we compute coefficients (App. D.5)

$$\tilde{q}_i = \int_{t_{ijk} \ni i} \phi^\delta \Psi_i dA = -\tfrac{1}{4}\sum_{e \ni i} r_{ie}\beta_e |p_e|,$$

where β_e denotes the dihedral angle at e and the transport coefficients $r_{ie} = e^{\imath 2\theta_i(X_i, e)}$ depend on the rescaled Euclidean angle between X_i and e. The coefficients q of ϕ are the solution of

$$Mq = \tilde{q}.$$

Note that $|q|$ is proportional to the squared difference of principal curvatures $(\kappa_1 - \kappa_2)^2$. In umbilic regions, where principal directions are ill-defined, smoothness is therefore automatically favored over alignment (at least for $t < 1$ – see Eq. (9)).

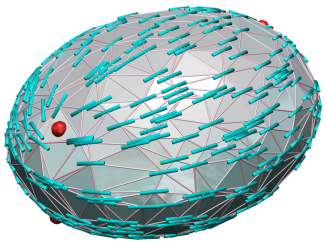

An example of this behavior is demonstrated in the inset image, where q is extremely noisy due to irregular tessellation. Due to the smoothing term we still obtain good results, reliably producing principal curvature directions and the four expected umbilics. In this example we used $-q$ to align with minimum curvature directions; for 4-direction fields we can use q_i^2. Fig. 10 shows an example of the global trade off between smoothness and alignment and how it controls spurious singularities. Several results in Fig. 12 demonstrate further examples of alignment.

6.1.3 Index Computation

Given an n-direction field ψ we want to label each triangle t by an integer $\text{index}_t\,\psi$ (an index of ± 1 indicating the presence of a singularity). Let ψ be given by complex numbers u_i of norm one for each vertex. Then for each edge e_{ij} we define the *rotation angle* of ψ as the unique number $\omega_{ij} \in (-\pi, \pi)$ such that $u_j = e^{\iota\omega_{ij}} r_{ij} u_i$, and define

$$\text{index}_t\,\psi := \tfrac{1}{2\pi}(\omega_{ij} + \omega_{jk} + \omega_{ki} + \Omega_{ijk}) \in \{-1, 0, 1\}. \qquad (19)$$

See App. B for more details on this definition and a proof of the associated discrete Poincaré-Hopf theorem. Singularities of index $+1$ and -1 are plotted as red and blue spheres, respectively, in all images.

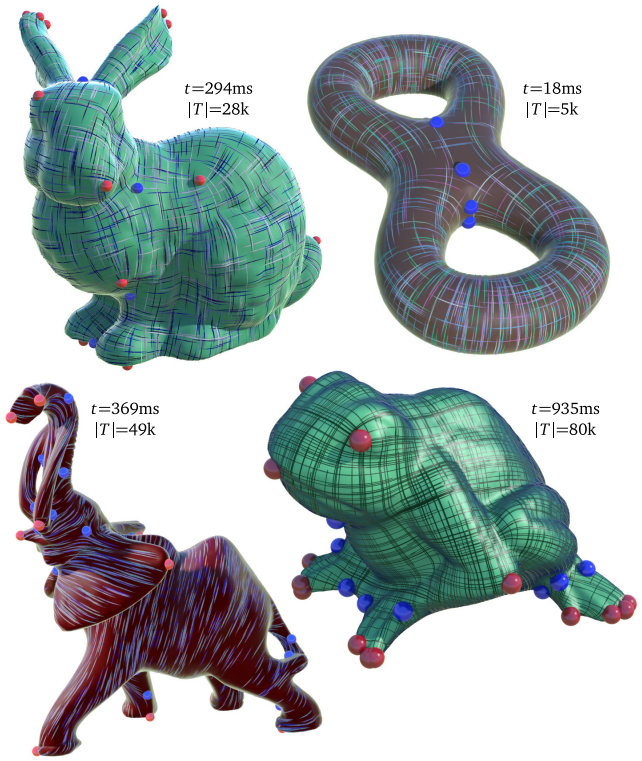

Figure 12: *A gallery of examples. The bunny with a 4-direction field aligned ($\lambda = 0$) with (q_i^2) using $s = 1$. The two-holed torus aligning with (q_i^2) using $s = -1$. The frog shows a smoothest ($s = 1$) 4-direction field. The Elephant uses $s = -1$ and alignment ($\lambda = 0$) with the minimum curvature direction $-q$.*

Algorithm 1 Setup

Input: $(\{V, E, T\}, s \in (-1, 1), n \in \mathbb{N})$
Output: (M, A)
Compute s_i at each vertex. ▷ Eq. (11)
Pick arbitrary unit vector X_i at each vertex. ▷ Sec. 2
Compute transport at edges $r_{ij} \leftarrow e^{\iota\rho_{ij}}$. ▷ Eq. (12)
Compute curvatures $\Omega_{ijk} \leftarrow \arg(r_{ij} r_{jk} r_{ki})$. ▷ Eq. (13)
Assemble M. ▷ Eq. (17)
Assemble A. ▷ Eq. (18)
$A \leftarrow A + \varepsilon M$ ▷ Sec. 7

Algorithm 2 Smoothest Field (Sec. 6.1)

Input: (M, A)
Output: u
$LL^T \leftarrow \text{Cholesky}(A)$
$u \leftarrow \text{UniformRand}(-1, 1) \in \mathbb{C}^{|V|}$
for $i = 1$ to nPowerIterations **do**
$x \leftarrow \text{BackSubstitute}(LL^T, Mu)$ ▷ Eq. (15)
$u \leftarrow x / \sqrt{x^T M x}$
end for

Algorithm 3 Aligned Field (Sec. 6.1)

Input: $(q, \lambda_t \in (-\infty, \lambda_0))$
Output: u
$LL^T \leftarrow \text{Cholesky}(A - \lambda_t M)$
$x \leftarrow \text{BackSubstitute}(LL^T, Mq)$ ▷ Eq. (16)
$u \leftarrow x / \sqrt{x^T M x}$

7 Evaluation

To evaluate our algorithm, we compared it to the state-of-the-art *mixed integer* method of Bommes *et al.* [2009]. We implemented both algorithms in a common C++ framework, using CoMISo [Bommes et al. 2012] for the mixed integer problem, and CHOLMOD [Chen et al. 2009] for linear systems required in our method. To avoid factorization issues in case A is rank deficient we add a small multiple ($\varepsilon = 10^{-8}$) of M to A prior to factorization; this shift does not change the eigenvectors. All timings were taken on the same 2.4GHz Intel Core 2 Duo machine. For our eigenvector problem (Eq. (15)) we factorized A and applied a fixed number of power iterations (20 for all examples shown in this paper). Note that since we need only the *smallest* eigenvalue, we do not require a sophisticated eigensolver like ARPACK, making our method particularly easy to implement. To minimize alignment energy we performed factorization and a single back-substitution. On a mesh of over 350k triangles, total time for our method was 11s for smoothing, 7s for alignment; the mixed integer method required more than three minutes for smoothing and just under three minutes for alignment. Fig. 13 gives more detailed timing information; overall speedup was approx. 19x on average.

To assess the quality of our results we compared fields produced by *both* methods using the energy described in Eq. (1), *i.e.*, the same energy minimized by Bommes *et al.*Note that our method produces smoother fields, even though we do not explicitly minimize this energy – in these examples we minimize our Dirichlet energy ($s = 0$), which is most similar to Eq. (1). For curvature-aligned fields, we do not attempt a direct comparison of energies due to the large number of parameters involved in the method of Bommes *et al.*; we report only the number of singularities S and the time t.

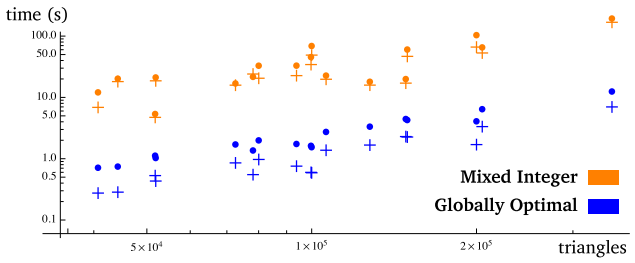

Figure 13: *Log-log plot of computation time for a variety of common models. Dots and crosses correspond to smoothest and curvature-aligned fields, respectively. Overall, we observe an average speedup of 19.3x relative to the method of Bommes* et al.

Fig. 14 shows a rounded cube – here we achieve a lower energy than the mixed integer method both when optimizing smoothness (Dirichlet) and alignment. Even with no guidance field present (top row), our method places the expected 8 singularities at cube corners. A final set of examples is shown in Fig. 15. On these more complex shapes the two methods find similar configurations of singularities, yet we achieve our results in *significantly* shorter time.

Robustness Inspired by [Bommes et al. 2009, Fig. 8], Fig. 16 shows the results of our curvature-aligned smoothing algorithm applied to an extremely poor triangulation (left) of an object with sharp features (here we use $s = 0$). Our method places singularities in the expected locations; adding noise *(middle)* or regularization *(right)* only slightly perturbs singularity locations in otherwise flat regions.

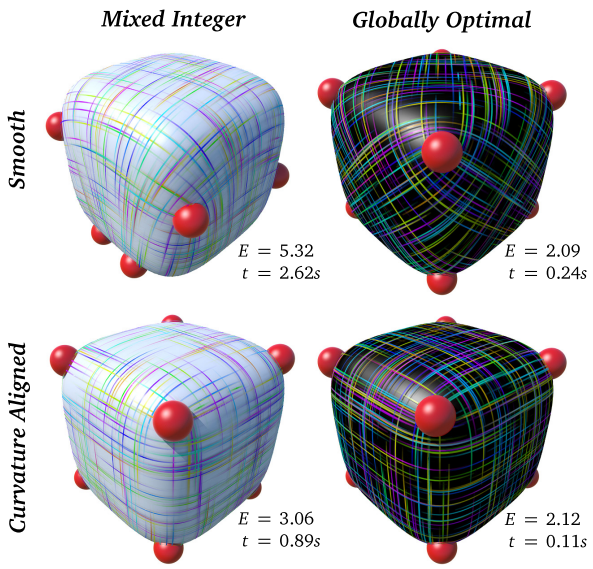

Figure 14: *A comparison of results on the rounded cube. Smoothest fields (top row) and curvature aligned fields (bottom row) by method. Optimization of Eq. (1) (left column) versus our Dirichlet energy (right column). In each case the energy according to Eq. (1) is reported.*

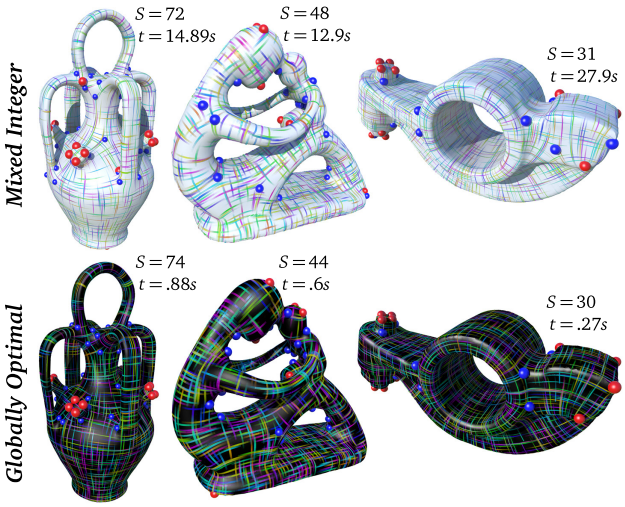

Figure 15: *Comparison of results using optimization of Eq. (1) (top) with minimization of our Dirichlet energy (bottom). Note in particular the difference in timings.*

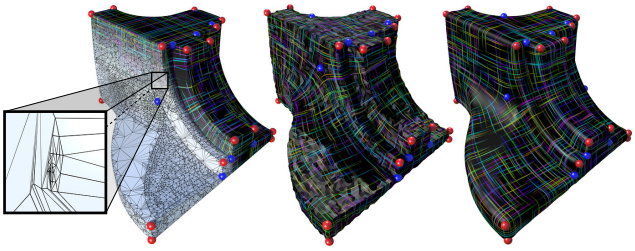

Figure 16: *Our method is robust to ill-conditioned, noisy, and over-regularized meshes.*

8 Conclusion

We have presented a principled formulation of n-direction fields on surfaces that naturally leads to efficient numerical algorithms. Although we have focused primarily on automatic singularity placement, our formulation provides a useful theory that can be applied in a variety of applications, especially in light of well-established connections with existing work (Sec. 4). These connections open up a number of interesting avenues for future investigation. For example, in some applications it is still desirable to edit fields manually in order to accommodate other (*e.g.* æsthetic) criteria. Editing could easily be achieved in our framework via a dualized version of the *trivial connections* algorithm of Crane *et al.* [2010]. One might also use holomorphic and/or anti-holomorphic n-vector fields as the basis for surface parameterization, creating a direct path to quad meshing (*e.g.*, [Kälberer et al. 2007]). Different weighting schemes for the magnitude of the field ϕ may also permit the incorporation of other interesting criteria via the alignment term E_l.

Acknowledgments This research was supported by a Google PhD Fellowship, the Hausdorff Research Institute for Mathematics, BMBF Research Project GEOMEC, SFB / Transregio 109 "Discretization in Geometry and Dynamics," and the TU München Institute for Advanced Study, funded by the German Excellence Initiative. Meshes provided by the Stanford Computer Graphics Laboratory and the AIM@SHAPE Shape Repository.

References

AHLFORS, L. V. 1966. *Complex Analysis*, 2nd Ed. McGraw-Hill.

BEN-CHEN, M., BUTSCHER, A., SOLOMON, J., AND GUIBAS, L. 2010. On Discrete Killing Vector Fields and Patterns on Surfaces. *Comp. Graph. Forum 29*, 5, 1701–1711.

BOMMES, D., ZIMMER, H., AND KOBBELT, L. 2009. Mixed-Integer Quadrangulation. *ACM Trans. Graph. 28*, 3.

BOMMES, D., ZIMMER, H., AND KOBBELT, L. 2012. Practical Mixed-Integer Optimization for Geometry Processing. In *Proc. 7th Int. Conf. Curves & Surfaces*, 193–206. Project page: http://www.graphics.rwth-aachen.de/software/comiso.

CHEN, Y., DAVIS, T. A., HAGER, W. W., AND RAJAMANICKAM, S. 2009. Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate. *ACM Trans. Math. Softw. 35*, 3, 22:1–22:14.

COHEN-STEINER, D., AND MORVAN, J.-M. 2003. Restricted Delaunay Triangulations and Normal Cycle. In *Proc. Symp. Comp. Geom.*, 312–321.

CRANE, K., DESBRUN, M., AND SCHRÖDER, P. 2010. Trivial Connections on Discrete Surfaces. *Comp. Graph. Forum 29*, 5, 1525–1533.

DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic Parameterizations of Surface Meshes. *Comp. Graph. Forum 21*, 3, 209–218.

DESBRUN, M., KANSO, E., AND TONG, Y. 2008. Discrete Differential Forms for Computational Modeling. In *Discrete Differential Geometry*, A. I. Bobenko, P. Schröder, J. M. Sullivan, and G. M. Ziegler, Eds., Vol. 38 of *Oberwolfach Seminars*. Birkhäuser Verlag, 287–324.

FISHER, M., SCHRÖDER, P., DESBRUN, M., AND HOPPE, H. 2007. Design of Tangent Vector Fields. *ACM Trans. Graph. 26*, 3.

GIL, A., SEGURA, J., AND TEMME, N. M. 2007. *Numerical Methods for Special Functions*. SIAM.

GU, X., AND YAU, S.-T. 2003. Global Conformal Surface Parameterization. In *Proc. Symp. Geom. Proc.*, 127–137.

HERTZMANN, A., AND ZORIN, D. 2000. Illustrating Smooth Surfaces. In *Proc. ACM/SIGGRAPH Conf.*, 517–526.

KÄLBERER, F., NIESER, M., AND POLTHIER, K. 2007. QuadCover - Surface Parameterization using Branched Coverings. *Comp. Graph. Forum 26*, 3, 375–384.

KASS, M., AND WITKIN, A. 1987. Analyzing Oriented Patterns. *Comp. Vis., Graph., Im. Proc. 37*, 3, 362–385.

LEFEBVRE, S., AND HOPPE, H. 2006. Appearance-space Texture Synthesis. *ACM Trans. Graph. 25*, 3, 541–548.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. *ACM Trans. Graph. 21*, 3, 362–371.

LING, C., NIE, J., QI, L., AND YE, Y. 2009. Biquadratic Optimization Over Unit Spheres and Semidefinite Programming Relaxations. *SIAM J. on Opt. 20*, 3, 1286–1310.

MACNEAL, R. 1949. *The Solution of Partial Differential Equations by means of Electrical Networks*. PhD thesis, Caltech.

MULLEN, P., TONG, Y., ALLIEZ, P., AND DESBRUN, M. 2008. Spectral Conformal Parameterization. *Comp. Graph. Forum 27*, 5, 1487–1494.

NAPIER, T., AND RAMACHANDRAN, M. 2011. *An Introduction to Riemann Surfaces*. Birkhäuser.

NIESER, M., PALACIOS, J., POLTHIER, K., AND ZHANG, E. 2012. Hexagonal Global Parameterization of Arbitrary Surfaces. *IEEE Trans. Vis. Comp. Graph. 18*, 6, 865–878.

PALACIOS, J., AND ZHANG, E. 2007. Rotational Symmetry Field Design on Surfaces. *ACM Trans. Graph. 26*, 3.

PINKALL, U., AND POLTHIER, K. 1993. Computing Discrete Minimal Surfaces and Their Conjugates. *Experiment. Math. 2*, 1, 15–36.

POLTHIER, K., AND SCHMIES, M. 1998. Straightest Geodesics on Polyhedral Surfaces. In *Mathematical Visualization: Algorithms, Applications and Numerics*, H.-C. Hege and K. Polthier, Eds. Springer Verlag, 135–152.

RAY, N., LI, W. C., LÉVY, B., SHEFFER, A., AND ALLIEZ, P. 2006. Periodic Global Parameterization. *ACM Trans. Graph. 25*, 4, 1460–1485.

RAY, N., VALLET, B., LI, W. C., AND LÉVY, B. 2008. N-Symmetry Direction Field Design. *ACM Trans. Graph. 27*, 2, 10:1–10:13.

RAY, N., VALLET, B., ALONSO, L., AND LÉVY, B. 2009. Geometry Aware Direction Field Processing. *ACM Trans. Graph. 29*, 1.

WIRTINGER, W. 1927. Zur formalen Theorie der Funktionen von mehr komplexen Veränderlichen. *Math. Ann. 97*, 1, 357–375.

ZHANG, E., MISCHAIKOW, K., AND TURK, G. 2006. Vector Field Design on Surfaces. *ACM Trans. Graph. 25*, 4, 1294–1326.

A Minimum of Alignment Energy

In Section 5 we claimed that the normalized solution $\psi = \tilde{\psi}/\|\tilde{\psi}\|$ of

$$(A - \lambda(t)I)\tilde{\psi} = \phi$$

for $\lambda(t) \in (-\infty, \lambda_1)$ yields a global minimum of $E_{s,t}(\psi)$ over all $\|\psi\| = 1$. Here we prove this fact and show the relationship between λ and t. For simplicity we assume throughout that $\|\phi\| = 1$. Furthermore, we will only use the real part of the product $\langle\langle .,. \rangle\rangle$ in this section.

At a constrained minimum the constraint must be fulfilled and the gradients of energy and constraint must be parallel, *i.e.*, there exists a Lagrange multiplier $\lambda/2$ such that

$$A\psi - \tfrac{t}{1-t}\phi = \lambda(t)\psi. \tag{20}$$

Substituting $\psi = \tilde{\psi}/\|\tilde{\psi}\|$ we find it to solve this equation for $t = (1 + \|\tilde{\psi}\|)^{-1}$.

Notice now how for $t \to 0$, ψ approaches the eigenvector belonging to the smallest eigenvalue λ_1 (we are working with a positive semidefinite Hermitian operator), since we are looking for a minimum of the energy. To see the behavior for $t \to 1$ consider the inner product of Eq. (20) with ψ

$$\langle\langle A\psi, \psi\rangle\rangle = \lambda(t) + \tfrac{t}{1-t}\langle\langle \phi, \psi\rangle\rangle.$$

Together with

$$\langle\langle A\psi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle,$$

which we will show in a moment, we get

$$\lambda(t) \leq \langle\langle A\phi, \phi\rangle\rangle - \tfrac{t}{1-t}\langle\langle \phi, \psi\rangle\rangle.$$

As $t \to 1$ the scalar product $\langle\langle \phi, \psi\rangle\rangle \to 1$ and so $\lambda(t) \to -\infty$.

To see that

$$\langle\langle A\psi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle$$

we consider that for a minimizer ψ

$$\langle\langle A\psi, \psi\rangle\rangle - \tfrac{t}{1-t}\langle\langle \phi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle - \tfrac{t}{1-t}\|\phi\|^2.$$

With $\|\phi\| = 1$ we get

$$\langle\langle A\psi, \psi\rangle\rangle \leq \langle\langle A\phi, \phi\rangle\rangle - \tfrac{t}{1-t}(1 - \langle\langle \phi, \psi\rangle\rangle)$$

and since $\langle\langle \phi, \psi\rangle\rangle \leq 1$ the desired bound.

Together we see that as $t \to 0$, $\lambda(t) \to \lambda_1$ and for $t \to 1$, $\lambda(t) \to -\infty$.

Finally we show that there exists a strictly monotone analytic function $t(\lambda) : (-\infty, \lambda_1) \to (0, 1)$.

Since the resolvent

$$R(\lambda) = (A - \lambda I)^{-1}$$

is a bounded self-adjoint operator which is an analytic function of λ so is

$$\psi(\lambda) = \frac{R(\lambda)\phi}{\|R(\lambda)\phi\|}$$

and $t(\lambda) = (1 + \|R(\lambda)\phi\|^{-1})$.

Now assume by contradiction that $t(\lambda)$ is not strictly monotone. Then there exists a $\lambda_0 \in (-\infty, \lambda_1)$ for which $t'(\lambda_0) = 0$. Taking the derivative at λ_0 of Eq. (20) we find

$$(A - \lambda_0 I)\psi'(\lambda_0) - \psi(\lambda_0) = 0.$$

Taking the inner product of this equation with $\psi'(\lambda_0)$ the second term vanishes (since $\|\psi\| = 1$) leading to $\langle\langle (A - \lambda_0 I)\psi'(\lambda_0), \psi'(\lambda_0)\rangle\rangle = 0$, which contradicts the positive definiteness of $A - \lambda_0 I$.

Together these results show that solving

$$(A - \lambda(t)I)\tilde{\psi} = \phi$$

for a choice of $\lambda \in (-\infty, \lambda_1)$ finds a global minimizer of $E_{s,t}(\psi)$ for $t(\lambda)$, a strictly monotone function with range $(0, 1)$. For $\lambda(t) \to \lambda_1$, $t \to 0$ (*i.e.*, no alignment) while for $\lambda(t) \to -\infty$, $t \to 1$ (*i.e.*, no smoothing).

B Poincaré-Hopf

In this Section we assign to every n-direction field ψ an index $\text{index}_t\,\psi \in \{-1, 0, 1\}$ on each triangle t. This is how we locate the singularities and label them as positive or negative. Under some smoothness assumption on the surface we will prove a discrete version of the Poincaré-Hopf theorem. A proof in the smooth case can be found in [Ray et al. 2008].

Recall that we describe the parallel transport of n-vectors by transport coefficients r_{ij} defined for each oriented edge e_{ij}. They can be thought of as an angle valued 1-form, *i.e.*, they satisfy $r_{ji} = r_{ij}^{-1}$. For each face t_{ijk} there is a unique real number $\Omega_{ijk} \in (-\pi, \pi)$ such that

$$r_{ij}r_{jk}r_{ki} = e^{\imath\Omega_{ijk}}.$$

We call Ω the curvature 2-form of the transport.

For a face $t = t_{ijk}$ we call the total curvature pushed into the triangle from its three vertices the *geometric curvature*

$$\sigma_t := s_i\alpha_i^{jk} + s_j\alpha_j^{ki} + s_k\alpha_k^{ij} - \pi$$

of t. A triangle mesh is called n-smooth if for each face we have

$$|\sigma_t| < \tfrac{\pi}{n}.$$

On an n-smooth mesh we have $\Omega_t = n\sigma_t$ and therefore a Gauß-Bonnet type theorem:

$$\sum_{t \in T} \Omega_t = 2n\pi\chi.$$

Here χ is the Euler characteristic of the mesh.

Let now ψ be an n-direction field given by complex numbers u_i of norm one for each vertex. Then for each edge e_{ij} we define the *rotation angle* of ψ as the unique number $\omega_{ij} \in (-\pi, \pi)$ such that

$$u_j = e^{\imath\omega_{ij}} r_{ij} u_i.$$

Then for each face $t = t_{ijk}$ we have $\omega_{ij} + \omega_{jk} + \omega_{ki} + \Omega_{ijk} \in (-4\pi, 4\pi)$ and $e^{\imath(\omega_{ij}+\omega_{jk}+\omega_{ki}+\Omega_{ijk})} = 1$. This means that for every face we have an integer

$$\text{index}_t\,\psi := \tfrac{1}{2\pi}(\omega_{ij} + \omega_{jk} + \omega_{ki} + \Omega_{ijk}) \in \{-1, 0, 1\}.$$

In the language of Discrete Exterior Calculus this could be expressed as

$$\text{index}\,\psi = \tfrac{1}{2\pi}(d\omega + \Omega)$$

If we sum this equation over all faces, the rotation angles cancel and we are left with the

Discrete Poincaré-Hopf theorem: On an n-smooth closed triangle mesh the index sum of every n-direction field equals $2n\chi$ where χ is the Euler characteristic of the mesh.

C Difference of Anti-Holomorphic and Holomorphic Energy

Here we show that for any smooth n-vector field ψ we have $E_A(\psi) - E_H(\psi) = \frac{1}{2}\int_M nK|\psi|^2 dA - \frac{1}{2}\int_{\partial M} \mathrm{Im}\langle \nabla\psi, \psi\rangle$. First we spell out all the requirements, each of which is fulfilled in our setting, of the theorem and a few facts that we will use during the derivation.

Let M be an oriented surface with Riemannian metric. Denote by J the complex structure on M. Suppose we have a complex line bundle L over M, *i.e.*, L is a 2-dimensional real vector bundle with a complex structure J. Our ψ are sections of such a line bundle. Assume that L comes with a complex connection ∇, meaning ∇ is compatible with J: $\nabla_X(J\psi) = J\nabla_X\psi$ for all sections ψ of L and all vector fields X. Then a section ψ is called holomorphic resp. anti-holomorphic if

$$0 = \bar\partial_X\psi := \tfrac{1}{2}(\nabla_X\psi + J\nabla_{JX}\psi)$$
$$0 = \partial_X\psi := \tfrac{1}{2}(\nabla_X\psi - J\nabla_{JX}\psi).$$

If ψ is a section of L, at each point $p \in M$ the linear map $\nabla\psi$ from the 2-dimensional real vector space T_pM into the 2-dimensional real vector space L_p can be decomposed into a complex linear part $(\partial\psi)_p$ and an anti-linear part $(\bar\partial\psi)_p$ and we can write

$$\nabla = \partial + \bar\partial.$$

Now suppose in addition that each L_P comes with a Hermitian scalar product $\langle .,. \rangle_p$ that is invariant under J and that ∇ is a metric connection, *i.e.*,

$$X\langle\psi,\varphi\rangle = \langle\nabla_X\psi,\varphi\rangle + \langle\psi,\nabla_X\varphi\rangle.$$

Replacing here X by some Lie bracket $[X,Y]$ we obtain

$$\begin{aligned}
&\langle\nabla_{[X,Y]}\psi,\varphi\rangle + \langle\psi,\nabla_{[X,Y]}\varphi\rangle = [X,Y]\langle\psi,\varphi\rangle\\
&\quad = X(\langle\nabla_Y\psi,\varphi\rangle + \langle\phi,\nabla_Y\varphi\rangle) - Y(\langle\nabla_X\psi,\varphi\rangle + \langle\psi,\nabla_X\varphi\rangle)\\
&\quad = \langle\nabla_X\nabla_Y\psi,\varphi\rangle + \langle\psi,\nabla_X\nabla_Y\varphi\rangle - \langle\nabla_Y\nabla_X\psi,\varphi\rangle - \langle\psi,\nabla_Y\nabla_X\varphi\rangle
\end{aligned}$$

Thus the curvature tensor R of L defined by

$$R(X,Y)\psi = \nabla_X\nabla_Y\psi - \nabla_Y\nabla_X\psi - \nabla_{[X,Y]}\psi$$

satisfies

$$\langle R(X,Y)\psi,\varphi\rangle + \langle\psi,R(X,Y)\varphi\rangle = 0.$$

In particular this implies that $\langle R(X,Y)\psi,\psi\rangle$ is always imaginary. Thus there is a real-valued 2-form Ω on M such that

$$R(X,Y)\psi = \Omega(X,Y)J\psi.$$

In case L is the tangent bundle TM, the Gaussian curvature K of M at a point p is defined in terms of a unit vector $X \in T_pM$ as

$$K = \langle R(X,JX)JX,X\rangle = -\Omega(X,JX).$$

Thus, if σ denotes the volume form, the curvature 2-form of the tangent bundle is $\Omega = -K\sigma$. We are mostly interested in the case $L = TM^{\otimes n}$, in which case we have $\Omega = -nK\,dA$ and therefore

$$R(X,Y)\psi = -nK\,dA(X,Y)J\psi.$$

The Dirichlet energy of a section ψ of L is defined as

$$E_D(\psi) = \frac{1}{2}\int_M |\nabla\psi|^2 dA,$$

where we view $\nabla\psi$ as a 1-form with values in L and the squared norm of such a 1-form ω at a point $p \in M$ is defined as

$$|\omega_p|^2 = |\omega(X)|^2 + |\omega(Y)|^2$$

where $\{X,Y\}$ form an orthonormal basis for T_pM.

Similarly, we define the holomorphic resp. anti-holomorphic energy of ψ as

$$E_H(\psi) = \frac{1}{2}\int_M |\bar\partial\psi|^2 dA \qquad E_A(\psi) = \frac{1}{2}\int_M |\partial\psi|^2 dA.$$

Theorem: The holomorphic and anti-holomorphic energies of a section ψ are related as

$$E_A(\psi) - E_H(\psi) = \frac{1}{2}\int_M nK|\psi|^2 dA - \frac{1}{2}\int_{\partial M} \mathrm{Im}\langle\nabla\psi,\psi\rangle.$$

Proof: Because of

$$\mathrm{Re}\langle\nabla\psi,\psi\rangle = \tfrac{1}{2}d\langle\psi,\psi\rangle$$

we can define a real-valued 1-form η on M via

$$\langle\nabla\psi,\psi\rangle = \tfrac{1}{2}d\langle\psi,\psi\rangle + \imath\eta(X).$$

Then for a locally defined unit vector field X we have

$$\begin{aligned}
\imath d\eta(X,JX) &= X\langle\nabla_{JX}\psi,\psi\rangle - (JX)\langle\nabla_X\psi,\psi\rangle - \langle\nabla_{[X,JX]}\psi,\psi\rangle\\
&= \langle R(X,JX)\psi,\psi\rangle + \imath(\langle J\nabla_{JX}\psi,\nabla_X\psi\rangle + \langle\nabla_X\psi,J\nabla_{JX}\psi\rangle)\\
&= \imath(nK|\psi|^2 + |\bar\partial\psi|^2 - |\partial\psi|^2)dA(X,JX).
\end{aligned}$$

Thus

$$d\eta = (nK|\psi|^2 + |\bar\partial\psi|^2 - |\partial\psi|^2)dA$$

and our claim follows by Stokes theorem.

D Integrals

In this section we give a derivation of all the integrals needed in the finite element discretization of the smooth theory, and begin with the construction of the basis sections.

D.1 PL Basis Sections

The PL basis sections Ψ_i are defined by extending the basis X_i from each vertex into the incident triangles through parallel transport along radii, giving us a unit basis section Φ_i supported on all incident triangles. Φ_i is then linearly attenuated with the standard PL hat function at v_i to give us Ψ_i. On a single incident triangle, say t_{ijk}, this amounts to

$$\Psi_i = b_i\Phi_i,$$

where b_i is the standard barycentric coordinate function for v_i in t_{ijk}.

While this procedure, together with the fact that the curvature is defined everywhere, uniquely defines the PL basis sections Ψ_i we do not have a closed form expression for them. Yet we can work out all required integrals in closed form due of our earlier assumption that the curvature is constant in each triangle (Eq. (14)). Specifically, the constancy of curvature over t_{ijk} gives

us the holonomy angle $\Omega_{t'}$ (Eq. (13)) along the boundary of any sub-triangle $t' \subset t_{ijk}$ as a fraction of the area:

$$\Omega_{t'} = \int_{t'} nK\,dA = nK_{ijk}|t'|,$$

where $|t'|$ denotes the area of t' (see Eq. (14)).

This brings us to the starting position.

D.2 Mass Matrix

Consider the Hermitian product (complex anti-linear resp. linear in the left resp. right factor) of basis sections restricted to the triangle t_{ijk}

$$\langle\langle \Psi_j, \Psi_k \rangle\rangle_{ijk} := \int_{t_{ijk}} \langle \Psi_j, \Psi_k \rangle dA$$

These determine the L_2 metric on the space of PL sections over a triangle and define the local mass matrix.

Let v be inside t_{ijk} with barycentric coordinates b_i, b_j, b_k

$$v = b_i v_i + b_j v_j + b_k v_k,$$

and consider the integrand of the mass matrix as a function of v

$$\langle \Psi_j(v), \Psi_k(v) \rangle = b_j b_k \langle \Phi_j(v), \Phi_k(v) \rangle. \tag{21}$$

We know that Φ_j is parallel along e_{jk}, implying

$$\Phi_j(v_k) = r_{jk} \Phi_k(v_k)$$

(*cf.* Eq. (2); Fig. 17). Moreover Φ_k is parallel along the ray from

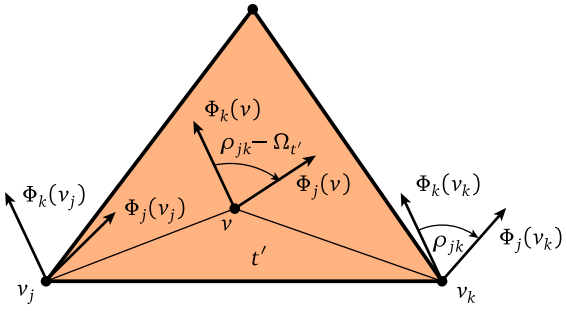

Figure 17: *The angle between two basis sections at some point v in t_{ijk} can be deduced from the holonomy $\Omega_{t'}$ of the sub-triangle $t' = \{v_k, v, v_j\}$ and the known transport along e_{jk}.*

v_k to v and Φ_j along the ray from v to v_j. Since parallel transport around $t' = \{v_k, v, v_j\}$, *i.e.*, from v_k to v on to v_j and finally back to v_k recovers the holonomy angle of t' we get

$$\langle \Phi_j(v), \Phi_k(v) \rangle r_{jk} = e^{\imath \Omega_{t'}} = e^{\imath \Omega_{ijk} b_i}. \tag{22}$$

(using Eqs. (13) and (14). Putting Eqs. (21) and (22) together and integrating symbolically we get

$$\langle\langle \Psi_j, \Psi_k \rangle\rangle_{ijk} = \bar{r}_{jk} |t_{ijk}| \frac{6e^{\imath\Omega_{ijk}} - 6 - 6\imath\Omega_{ijk} + 3\Omega_{ijk}^2 + \imath\Omega_{ijk}^3}{3\Omega_{ijk}^4}.$$

For the product of a basis section with itself the curvature does not play a role and one obtains $\|\Psi_i\|^2_{ijk} = |t_{ijk}|/6$.

This completely determines the metric on the space of basis n-vector fields.

D.3 Dirichlet Energy

To compute the Dirichlet energy we need the covariant derivatives of Φ_j and Φ_k. To this end we will employ a particular (linear) parameterization $f(x,y)$ of the embedding p_{ijk} of t_{ijk} with $f(0,0) = p_i$, $f(1,0) = p_j$ and $f(0,1) = p_k$, and will treat Φ and Ψ as defined on the image of f in this section. Denote by ∂_x and ∂_y the tangent coordinate frame corresponding to the coordinates x, y (Fig. 18).

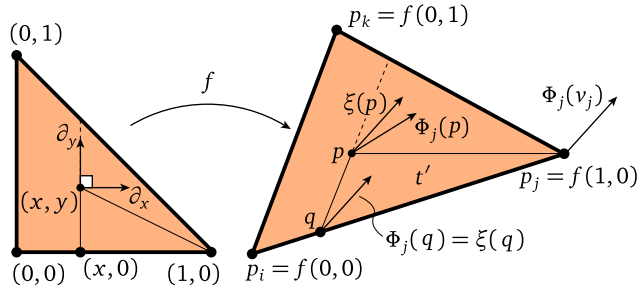

Figure 18: *Using the parameterization f we compute the covariant derivative of a basis n-vector field Φ_j using the local section ξ.*

To compute $\nabla_{\partial_y} \Phi_j$ at some point p in the interior of the triangle, let ξ be a section on the triangle t_{ijk} that agrees with Φ_j for $y = 0$ and is parallel along the lines $\{x = \text{const}\}$, giving $\nabla_{\partial_y} \xi = 0$. For $p = f(x,y)$ let $q := f(x,0)$, then the holonomy around t_{pqj} is $\Omega_{pqj} = (1-x)y\Omega_{ijk}$ and consequently

$$e^{\imath(1-x)y\Omega_{ijk}} \xi(p) = \Phi_j(p).$$

This gives us

$$\nabla_{\partial_y} \Phi_j = (\partial_y e^{\imath(1-x)y\Omega_{ijk}})\xi = \imath(1-x)\Omega_{ijk}\Phi_j.$$

The derivative $\nabla_{\partial_x} \Phi_k$ follows immediately through interchange of x and y. Taking account of the orientation, we obtain

$$\nabla_{\partial_x} \Phi_k = -\imath(1-y)\Omega_{ijk}\Phi_k.$$

Since by construction Φ_j is parallel along rays from v_j and Φ_k along rays from v_k we get a linear relationship between the covariant derivative of Φ_j with respect to ∂_x and ∂_y (and similarly for Φ_k)

$$0 = (1-x)\nabla_{\partial_x}\Phi_j - y\nabla_{\partial_y}\Phi_j \quad 0 = x\nabla_{\partial_x}\Phi_k + (y-1)\nabla_{\partial_y}\Phi_k$$

which give us all the covariant derivatives for Φ_j and Φ_k. For $\Psi_j = x\Phi_j$ and $\Psi_k = y\Phi_k$ this results in

$$\nabla_{\partial_x}\Psi_j = (1 + \imath\Omega_{ijk}xy)\Phi_j, \qquad \nabla_{\partial_y}\Psi_j = \imath\Omega x(1-x)\Phi_j,$$
$$\nabla_{\partial_x}\Psi_k = -\imath\Omega_{ijk}y(1-y)\Phi_k, \qquad \nabla_{\partial_y}\Psi_k = (1 - \imath\Omega_{ijk}xy)\Phi_k.$$

To simplify the computation of the integrals we now switch from the basis $\{\partial_x, \partial_y\}$ to an orthogonal basis $\{E_1, E_2\}$. Letting

$$g_{11} = |p_j - p_i|^2, \quad g_{12} = \langle p_j - p_i, p_k - p_i \rangle = g_{21}, \quad g_{22} = |p_k - p_i|^2,$$

the orthogonal basis follows as

$$E_1 = \frac{1}{\sqrt{g_{11}}}\partial_x, \quad E_2 = \frac{1}{2|t_{ijk}|\sqrt{g_{11}}}(g_{11}\partial_y - g_{12}\partial_x).$$

With respect to this basis we get

$$\begin{aligned}
\nabla_{E_1}\Psi_j &= \tfrac{1}{\sqrt{g_{11}}}(1+\imath\Omega_{ijk}xy)\Phi_j,\\
\nabla_{E_2}\Psi_j &= \tfrac{1}{2|t_{ijk}|\sqrt{g_{11}}}(-g_{12}+\imath\Omega_{ijk}(g_{11}x(1-x)-g_{12}xy))\Phi_j,\\
\nabla_{E_1}\Psi_k &= \tfrac{1}{\sqrt{g_{11}}}(-\imath\Omega_{ijk}y(1-y))\Phi_k,\\
\nabla_{E_2}\Psi_k &= \tfrac{1}{2|t_{ijk}|\sqrt{g_{11}}}(g_{11}+\imath\Omega_{ijk}(g_{12}y(1-y)-g_{11}xy))\Phi_k.
\end{aligned}$$

With this we have all the necessary components to compute the Dirichlet products and integrate them to yield

$$\begin{aligned}
\langle\!\langle\nabla\Psi_j,\nabla\Psi_j\rangle\!\rangle_{ijk} =& \tfrac{1}{4|t_{ijk}|}\left[g_{22}+\Omega_{ijk}^2\tfrac{3g_{11}-3g_{12}+g_{22}}{90}\right],\\
\langle\!\langle\nabla\Psi_j,\nabla\Psi_k\rangle\!\rangle_{ijk} =& \tfrac{\bar r_{jk}}{|t_{ijk}|\Omega_{ijk}^4}[(3g_{11}+4g_{12}+3g_{22})\\
&+\imath\Omega_{ijk}(g_{11}+g_{12}+g_{22})-\imath\Omega_{ijk}^3\tfrac{g_{12}}{6}\\
&+\Omega_{ijk}^4\tfrac{g_{11}-2g_{12}+g_{22}}{24}-\imath\Omega_{ijk}^5\tfrac{g_{11}-2g_{12}+g_{22}}{60}\\
&+(-(3g_{11}+4g_{12}+3g_{22})\\
&+\imath\Omega_{ijk}(2g_{11}+3g_{12}+2g_{22})\\
&+\Omega_{ijk}^2\tfrac{g_{11}+2g_{12}+g_{22}}{2})e^{\imath\Omega_{ijk}}].
\end{aligned}$$

D.4 Boundary Terms

The proof in App. C applies to a single triangle t and we only need to compute the boundary term, since the $nK\|\Psi\|_{ijk}^2$ term is the curvature weighted mass matrix (App. D.2).

We begin by observing that the 1-form $\langle\nabla\Psi_j,\Psi_k\rangle$ is nonzero only on e_{jk}. Moreover, since $\langle\nabla\Psi_i,\Psi_i\rangle = d|\Psi_i|^2$, it is real for matching indices giving us

$$\int_{\partial t}\operatorname{Im}\langle\nabla\Psi,\Psi\rangle = \sum_{e_{ij}\in\partial t}\operatorname{Im}\int_{e_{ij}}\left(\bar\alpha_i\alpha_j\langle\nabla\Psi_i,\Psi_j\rangle+\bar\alpha_j\alpha_i\langle\nabla\Psi_j,\Psi_i\rangle\right).$$

Note also that $d\langle\Psi_i,\Psi_j\rangle = \langle\nabla\Psi_i,\Psi_j\rangle+\overline{\langle\nabla\Psi_j,\Psi_i\rangle}$. Since $\Psi_i(v_j) = 0 = \Psi_j(v_i)$ Stokes' theorem yields

$$0=\int_{\partial e_{ij}}\langle\Psi_i,\Psi_j\rangle=\int_{e_{ij}}\langle\nabla\Psi_i,\Psi_j\rangle+\int_{e_{ij}}\overline{\langle\nabla\Psi_j,\Psi_i\rangle},$$

which we use to simplify the boundary edge sum

$$\int_{\partial t}\operatorname{Im}\langle\nabla\Psi,\Psi\rangle=\sum_{e_{ij}\in\partial t}2\operatorname{Im}\left(\bar\alpha_i\alpha_j\int_{e_{ij}}\langle\nabla\Psi_i,\Psi_j\rangle\right).$$

We now turn to an individual edge term. Parameterize e_{jk} with a constant speed $\gamma\colon[0,1]\to M$, *i.e.*, $|\gamma'| = |p_{jk}|$. Due to $db_j(\gamma') = -1$ for b_j the barycentric coordinate function of v_j, and the parallelity of Φ_j along e_{jk}, we get

$$\nabla_{\gamma'}\Psi_j = db_j(\gamma')\Phi_j + b_j\nabla_{\gamma'}\Phi_j = -\Phi_j,$$

and hence

$$\int_{e_{jk}}\langle\nabla\Psi_j,\Psi_k\rangle=\int_0^1\langle\nabla_{\gamma'}\Psi_j,\Psi_k\rangle dt=-\int_0^1 b_k\langle\Phi_j,\Phi_k\rangle dt=-\tfrac{\bar r_{jk}}{2}.$$

This finally yields

$$\int_{\partial t}\operatorname{Im}\langle\nabla\Psi,\Psi\rangle=-\sum_{e_{jk}\in\partial t}\operatorname{Im}\left(\bar r_{jk}\bar\alpha_j\alpha_k\right)$$

Notice that for an edge e_{jk} with two incident triangles (not on the boundary) the corresponding terms from each triangle cancel out, leaving us with only a sum over the boundary of the triangle mesh.

In the flat case ($\Omega_{ijk} = 0$) these boundary terms are simply the areas of the range of ψ seen as a complex function, *i.e.*, $2E_H(\psi) = E_D(\psi) - A(\psi(M))$ (compare with [Mullen et al. 2008, Eq. 2], resp. [Pinkall and Polthier 1993, Eq. 2.1]).

D.5 Curvature Alignment

To perform alignment with principal curvature directions we need to find a 2-vector corresponding to these directions (2-vector since principal curvature directions are indistinguishable under rotations by π). In the smooth setting, the principal curvature directions are the eigendirections of the shape operator S, given by the derivative of the Gauss map

$$dN = df\circ S.$$

The *Hopf differential* Q is the trace-free part of the shape operator

$$S = H\cdot I + Q,$$

where $H = (\kappa_1+\kappa_2)/2$ is the mean curvature and I the identity. Q has the principal curvature directions as eigendirections with eigenvalues $(\kappa_1-\kappa_2)/2$ and $(\kappa_2-\kappa_1)/2$. Thus the information provided by Q contains the (unoriented) direction of maximal curvature together with a positive "intensity." Thus Q can be viewed as a 2-vector field.

Using a local complex coordinate z we can write any tangent vector as $a\frac{\partial}{\partial x}$ where x is the real part of z and $a\in\mathbb{C}$. Q is anti-linear as an endomorphism of the tangent space (it anti-commutes with J) and therefore we have

$$Q(a\tfrac{\partial}{\partial x}) = q\bar a\tfrac{\partial}{\partial x}.$$

So with respect to a complex coordinate Q is described by a complex function q.

In the discrete setting the shape operator survives in a distributional sense concentrated along edges [Cohen-Steiner and Morvan 2003]. Consider an edge e and its two incident triangles and map them isometrically to the plane with e along the x-axis, then

$$S_e = \delta_y\beta_e\tfrac{\partial}{\partial y}dy = \delta_y\tfrac{\beta_e}{2}\tfrac{\partial}{\partial x}dz-\delta_y\tfrac{\beta_e}{2}\tfrac{\partial}{\partial x}d\bar z.$$

Here δ_y is the delta distribution along across e, β_e the dihedral angle at e, $dz = dx+\imath dy$, and $d\bar z = dx-\imath dy$. Hence

$$q_e = -\delta_y\tfrac{\beta_e}{2}.$$

Since q exists only as a distribution, we treat it as a functional on the space of smooth sections. In particular we can pair it with each of our PL 2-vector basis sections

$$\tilde q_i := q(\Psi_i) = \sum_{e\ni i}q_e(\Psi_i) = -\tfrac14\sum_{e\ni i}r_{ie}\beta_e|p_e|.$$

where the transport coefficient $r_{ie} = e^{\imath 2\theta_i(X_i,e)}$ depends on the rescaled Euclidean angle between X_i and e. Hence the coefficients q of the PL Hopf differential solve the matrix problem

$$Mq = \tilde q.$$