

Extrinsically Smooth Direction Fields

Zhiyang Huang Tao Ju
Washington University in St. Louis

Abstract

We consider the problem of finding a unit vector field (i.e., a direction field) over a domain that balances two competing objectives, smoothness and conformity to the shape of the domain. Common examples of this problem are finding normal directions along a curve and tangent directions over a surface. In a recent work, Jakob et al. observed that minimizing *extrinsic variation* of a tangent direction field on a surface achieves both objectives without the need for parameter-tuning or the use of additional constraints. Inspired by their empirical observations, we analyze the relation between extrinsic smoothness, intrinsic smoothness, and shape conformity in a continuous and general setting. Our analysis not only explains their observations but also suggest that an extrinsically smooth normal field along a curve can strike a similar balance between smoothness and shape-awareness. Our second contribution is offering extension of, justification for and improvement over the optimization framework of Jakob et al. In our experiments, we demonstrate the suitability of extrinsically smooth field in a variety of applications and compared with existing solutions.

Keywords:

direction field, surface tangents, curve normals

1. Introduction

Unit vector fields (which we call *direction fields* in this work) in \mathbb{R}^3 have various applications in computer graphics. For example, a sequence of normal directions along a 1-dimensional spatial curve is useful for sweep-based modeling [1, 2] and visualization of tubular structures [3]. A tangent direction field over a 2-dimensional surface is useful for texture synthesis [4, 5, 6] and non-photorealistic rendering [7]. In these and many other problem settings, the direction field is restricted to some *plane* at each location of the domain. In the two examples above, the plane is either normal to the curve or tangent to the surface.

We consider the general problem of finding a direction field g over some manifold $M \subset \mathbb{R}^3$ such that g is orthogonal to a given direction field f over M (i.e., g is restricted to the plane orthogonal to f). Since this is an under-constrained problem, regularization is needed to ensure that g is “good”. While the criteria of “goodness” vary with applications, a typical criteria is being smooth or having low variation. For some applications, an additional requirement is that g should conform to the shape of the domain M . For example, a tangent direction field that aligns to surface features is desirable for depicting the surface with strokes [7] or textures [8].

Most existing methods formulate these two criteria independently. In particular, the variation of g is usually measured *intrinsically*, as the amount of twisting of g around the input field f . However, intrinsic variation does not capture the variation of f itself, and the latter is intimately related to the shape of M in the case when f is

derived from M (e.g., being its tangents or normals). To align g with the shape features of M , these methods need to couple intrinsic energy with additional shape-related constraints, such as user sketches, feature lines, and the curvature field (see brief review in the next section). However, parameter-tuning is required to achieve the desirable balance between these two terms, smoothness and constraints. Moreover, obtaining the shape-related constraints often requires user interaction or computation of high-order differential quantities, such as curvatures.

In a recent work, Jakob et al. [9] proposed an elegant solution to the problem of designing a tangent direction field (and more generally, a symmetric n -vector field) over a surface. In their discrete implementation, they found that the two criteria, smoothness and shape conformity, can be achieved simultaneously by minimizing the *extrinsic* variation of g in the embedded space \mathbb{R}^3 . The resulting field naturally follows prominent feature lines of the surface and varies smoothly away from these features. There is no need for parameter-tuning, user interaction, or computing high-order quantities.

The authors of [9] did not offer any theoretical explanation to their observed behaviors. In particular, it was not clear why minimizing extrinsic variation leads to alignment with shape features, or how it balances between smoothness and alignment. Answers to these questions would shed light into the role of extrinsic variation in other direction field problems (e.g., normal directions on curves).

To provide answers to these questions, we analyze the extrinsic variation in the continuous and general setting –

not just for tangent fields on surfaces, but for any direction field g orthogonal to some input field f (i.e., the general problem we mentioned earlier). Our key findings are:

1. The extrinsic variation of g can be decomposed into a twisting term, which is the commonly studied intrinsic variation, and a bending term, which measures how close g is to the *direction of greatest variance of f* .
2. The bending term is scaled by the amount of anisotropy in the variation of f : the greater difference exists in the variation of f in different directions, the greater the role that bending is playing in the extrinsic variation of f .

The analysis allows us to conclude that an extrinsically smooth tangent field over a surface would tend to align with the principal curvature direction with a smaller absolute curvature value, if the local principal curvatures have very different magnitudes, and maximize intrinsic smoothness otherwise. This explains the observations made in [9]. The analysis also shows that an extrinsically smooth normal field along a curve strikes a similar balance between aligning to the binormals and being intrinsically smooth.

Minimizing the extrinsic variation is a non-linear optimization problem. As our second contribution, we extend, justify, and improve the discrete optimization framework of Jakob et al. [9]. First, we extend their discretization of extrinsic variation to simplices at other dimensions, such as edges (1D) and tetrahedra (3D), so that we can handle more general direction field problems. Second, we justify a key step in their iterative Gauss-Seidel step for optimizing the field. Third, and perhaps more importantly, we offer a simpler and parameter-free alternative to their hierarchical approach for initializing the field. Our approach only needs to solve a single eigenvalue problem on the original surface.

Our experimental results show the superiority of extrinsically smooth direction fields over results of existing, intrinsic methods on both curves and surfaces. To demonstrate the generality of our analysis and algorithms, we explored two other direction field problems. In the first, we generate a direction field orthogonal to a non-tangent field along a spatial curve, with the application of camera orientation. In the second, we compute a direction field within a tetrahedral volume to create a ribbon visualization of an input vector field.

Contributions. We make the following theoretical and practical contributions:

1. We analyze the extrinsic variation in general direction fields and give a geometric form of the bending term. The analysis allows us to explain observations made in a recent work [9] for extrinsically smooth tangent fields on surfaces and suggest a similar property of extrinsically smooth normal fields along a curve (Section 3).

2. We extend, justify, and improve the discrete optimization framework of Jakob et al. [9] for computing extrinsically smooth direction fields in general (Section 4).
3. We experimentally demonstrate the suitability of extrinsic smoothness in a variety of direction field problems (Section 5).

2. Related work

Most of the previous work on direction fields are concerned with either normal vectors along a curve or tangent vectors over a surface. We briefly review literature on these two topics, with an eye towards the balancing of the two goals (smoothness and shape conformity).

Normal fields on a curve. Natural choices of normals over a smooth curve include the principal normal (derivative of unit tangent) and the binormal (cross-product of tangent and unit principal normal). Together with the tangent, they constitute the Frenet frame. While reflecting local curve shape, the Frenet frame is undefined at C^1 connections or where the curvature is zero, and it can twist significantly near low-curvature points.

As a remedy, a normal field with zero rotation (i.e., intrinsic variation) over an open curve can be defined by the parallel-transport of any initial normal vector from one end of the curve [10]. In fact, parallel-transport creates a *frame field*, which gives rise to a one-parameter family of normal fields that all have zero twist. To handle closed curves and additional constraints (e.g., fixed frames at certain points), one may define a frame field that minimizes the total rotation measured in angular space [11] or quaternion space [12]. Various numerical techniques have been proposed for rotation-minimizing frames (see the in-depth review in [11]), and these frames have been widely used in applications such as axial deformation [13, 14], sweep-based modeling [1, 2], visualization of tubular structures [3], and physical simulation of rods [15]. However, framing methods in general do not suggest a unique normal field. More importantly, without additional constraints, normal fields derived from rotation-minimizing frames do not capture the shape of the curve (see comparisons in Section 5).

Tangent fields on a surface. The principal curvature fields are natural choices of tangent fields over a smooth surface. However, the principal curvature direction is undefined at umbilic points, and can be highly unstable in near-umbilic regions. In addition, a principal curvature direction does not differentiate between its forward and backward orientations.

To achieve shape alignment while maintain smoothness of the field, the majority of existing field creation methods solve a constrained optimization problem. Many methods [16, 7, 4, 6, 17, 8, 18, 19, 20, 21, 22] utilize sparse directional constraints, which are either computed automatically from local curvature (e.g, feature lines [8, 20]

and feature regions [7, 6]) or provided interactively by users. Other methods [23, 24] utilize the curvature field as a guidance field and vary the guidance strength based on the curvature values. The recent work [25] follows a similar principle to produce line directions (called *regularized curvature lines*) whose alignment with the curvature field is moderated by the so-called *curvature anisotropy*, or the difference between the two principal curvature values. Some of these above-mentioned methods also require topological constraints (e.g., location and index of singularities) to be explicitly specified [18, 21, 22]. Note that many of these methods can handle the more general, symmetric n -vector fields. With a few exceptions [4, 8, 16, 23], majority of these methods use *intrinsic variation* to measure smoothness of the field. However, parameter-tuning is generally needed for this class of methods, either for balancing between smoothness and alignment to constraints in the energy formulation or for creating the constraints.

To the best of our knowledge, tangent direction fields that minimize *extrinsic* variation were first considered by Gonzalez-Davila and Vanhecke [26]. In the special case of surfaces of revolution, they studied the uniqueness of such fields given constraints on the boundary of the surface. However, the relation between extrinsically smooth fields and the shape of the underlying surface was not mentioned. Without realizing this relation, some researchers even couple extrinsic variation with explicit curvature direction constraints to create shape-aligned fields [23]. The relation is observed only recently, by Jakob et al. [9], for creating n -vector fields on a triangular mesh. They measured extrinsic variation as a weighted sum of vector differences over all edges of the mesh, and they proposed a simple yet effective iterative optimization method. Inspired by their empirical observation, this paper seeks to explore the theoretical relation between extrinsic variation and shape, and do so in a more general context. We also enrich and improve their computational framework for other direction field problems.

3. Theoretical analysis

The goal of section is to understand how minimizing extrinsic variation leads to direction fields that balance between intrinsic smoothness and conformity to shape. We consider a smooth submanifold M of \mathbb{R}^3 that can be a curve (1-dimensional), surface (2-dimensional) or volume (3-dimensional), and a differentiable unit vector field f over M . While our analysis applies to arbitrary choices of f , we note that in many applications f is *derived* from M (e.g., being tangent or normal to M). In the latter case, the variation of f naturally captures the shape of M .

We define an *extrinsically smooth* direction field g for the pair $\{M, f\}$ as a vector field over M that minimizes the Dirichlet energy,

$$\int_M |\nabla g|^2 \quad (1)$$

subject to constraints $f^T g = 0, g^T g = 1$. Here, ∇g denotes the derivative of the vector field g , and $|\cdot|^2$ is the squared Frobenius norm. To study the energy over domains of different dimensionality in a consistent manner, ∇g is written as a 3×3 matrix whose rows are gradients of the X, Y, and Z coordinates of g over the tangent space T_M . More precisely, let $\{v_1, \dots, v_k\}$ be any set of mutually orthogonal unit tangent vectors in T_M , where k equals the dimension of T_M (and M). Then $\nabla g = \sum_{i=1}^k g_{v_i} v_i^T$ where g_{v_i} is the derivative of g in the direction of v_i . The dot product of ∇g and any tangent direction $v \in T_M$ gives the derivative of g in the direction of v , or g_v .

The two constraints, orthogonality to f and unity, are critical in our formulation to prevent trivial minimizers of Equation 1. If we drop orthogonality, any constant unit vector field g would have zero energy. In we drop unity, the zero vector field g would have zero energy. These two constraints avoid the need for boundary conditions that are usually required for Dirichlet-type problems.

In the following, we first investigate the integrand $|\nabla g|^2$, which measures the amount of *local* extrinsic variation (Section 3.1). We shall arrive at a decomposition of $|\nabla g|^2$ into a *twisting* term, which measures intrinsic variation, and a *bending* term, which is related to the variation of f itself. Since intrinsic variation has been extensively studied in the past, we focus on the bending term and give a simple and geometric expression of bending. Next, we will apply our expression to understand our two motivating problem settings, normal directions along a curve and tangent directions on a surface (Section 3.2).

3.1. Local extrinsic variation

The derivative of g along any tangent direction $v \in T_M$ (here T_M refers to the tangent space of M), or g_v , can be decomposed into two orthogonal vectors: one orthogonal to f , as $(I - f f^T)g_v$ (I is the identity matrix), and another parallel to f , as $f f^T g_v$ (see Figure 1 (a)). As a result, we can decompose $|\nabla g|^2$ as the sum of two terms:

$$|\nabla g|^2 = |(I - f f^T)\nabla g|^2 + |f f^T \nabla g|^2 \quad (2)$$

The first term, the variation of g within the plane orthogonal to f , describes how much g rotates around f . We call it the *twisting term*. Since the twist is oblivious of




Figure 1: Illustrating the decomposition of the derivative g_v (a) and the eigenvectors of matrix J (b).

the change in f itself, it captures the intrinsic variation of g with respect to f . The second term, the variation of g parallel to f , describes how much g bends towards or away from f . We call it the *bending term*. This decomposition has been known, for example, when M is a surface and f is its normal field [27].

We are primarily interested in how $|\nabla g|^2$ relates to the variation of f , as the latter is often tied to the shape of M . As a result, we focus on understanding the bending term, $|f^T \nabla g|^2$. Since $f^T g = 0$, by differentiation we have $f^T \nabla g = -g^T \nabla f$. This observation allows us to re-write bending as a quadratic form,

$$|f^T \nabla g|^2 = |g^T \nabla f|^2 = g^T J g, \quad (3)$$

where $J = \nabla f (\nabla f)^T$ is a positive semi-definite matrix. In other words, the bending is the inner product of g with itself in the new metric J that encodes the variation of f .

To make the relation between bending and the variation of f more precise, consider the eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ of J and the corresponding eigenvectors e_1, e_2, e_3 that form an orthonormal basis. Since $f^T f = 1$, by differentiation we have $f^T \nabla f = 0$, which gives $J f = 0$. Hence f is an eigenvector of J with eigenvalue 0. Let $e_3 = f$ and $\lambda_3 = 0$. Since g, e_1, e_2 all lie in the plane orthogonal to f , we can express g in the bases of e_1, e_2 as $(g^T e_1)e_1 + (g^T e_2)e_2$, which leads to

$$g^T J g = \lambda_1 (g^T e_1)^2 + \lambda_2 (g^T e_2)^2 \quad (4)$$

$$= \lambda_2 + (\lambda_1 - \lambda_2) \cos^2 \theta \quad (5)$$

where θ is the angle between g and e_1 (see Figure 1 (b)).

It is now clear that the bending term is proportional to the angular agreement between g and $\pm e_1$, or equivalently the angular deviation from g to $\pm e_2$. The closer g is to $\pm e_1$, the greater the bending. The angle measure is scaled by the multiplier $(\lambda_1 - \lambda_2)$. What remains unclear is the geometric meaning of e_1 and $\lambda_1 - \lambda_2$, which we reveal next.

The following lemma shows that e_1 is in fact the *direction of the greatest variation* of f :

Lemma 3.1. *The derivative of f along any unit tangent vector $v \in T_M$, denoted as f_v , achieves maximal magnitude when f_v is a scalar multiple of e_1 .*

Proof. Consider unit tangent vector v^* such that f_{v^*} achieves the maximal magnitude. Since $f_v = (\nabla f)v$, we have

$$v^* = \arg \max_{v \in T_M, |v|=1} ((\nabla f)v)^T (\nabla f)v. \quad (6)$$

Note that we can drop the tangent requirement without changing the choice of v^* : since $(\nabla f)v = 0$ for any vector v orthogonal to the tangent space T_M , the magnitude of $(\nabla f)v$ will not change after replacing v by its component in T_M . Hence we have

$$v^* = \arg \max_{|v|=1} ((\nabla f)v)^T (\nabla f)v. \quad (7)$$

By Rayleigh quotient theorem, v^* is the eigenvector of $K = (\nabla f)^T \nabla f$ with the largest eigenvalue. On the other hand, there is a one-to-one correspondence between eigenvectors of J and K : for any eigenvector u of K with eigenvalue δ , $(\nabla f)u$ is an eigenvector of J with the eigenvalue δ . Conversely, for any eigenvector e of J with eigenvalue λ , $(\nabla f)^T e$ is an eigenvector of K with the eigenvalue λ . Therefore we conclude that e_1 is a scalar multiple of $(\nabla f)v^* = f_{v^*}$. \square

The proof of the lemma also shows that λ_1 , the largest eigenvalue of J , is the *magnitude of the greatest variation* of f . That is, $\lambda_1 = \max |f_v|$ for all unit tangent vectors $v \in T_M$, and the maximum is achieved when f_v is a scalar multiple of e_1 . The second eigenvalue, λ_2 , is the magnitude of f_v in the orthogonal direction e_2 . Hence the difference $(\lambda_1 - \lambda_2)$ captures the non-uniformity, or *anisotropy*, in the variation of f on its orthogonal plane.

To summarize, the local extrinsic variation of g is the sum of a twisting term, which measures the intrinsic variation of g with respect to f , and a bending term, which is related to the variation of f . In particular, bending is proportional to the agreement between g and the direction of the greatest variation of f and is scaled by the anisotropy of variation of f .

3.2. Case studies

From the above analysis of local extrinsic variation, we see that minimizing total extrinsic variation (i.e., Equation 1) has two effects: it maximizes local intrinsic smoothness on one hand and, on the other hand, encourages alignment to the direction orthogonal to the greatest variation of f (i.e., e_2). The balance between these two effects depends locally on the anisotropy of variation of f (i.e., $\lambda_1 - \lambda_2$). In regions where f varies nearly uniformly in all directions orthogonal to f , the twisting term dominates, and the extrinsically smooth field g will be intrinsically smooth. In regions where f varies mostly in one direction (i.e., e_1), the bending term dominates, and g will align with e_2 .

In problem settings where f is derived from the domain M , the variation of f , and hence the extrinsically smooth field g , is intimately related to the shape of M . We make such relation precise in the following two specific problem settings.

Normal directions on a curve. Let M be an oriented 1-dimensional curve (open or closed) and f identify with the unit tangent vector field t along M . By the Frenet–Serret formulas, ∇f has the form

$$\nabla f = \kappa n t^T, \quad (8)$$

where κ is the curvature and n is the unit principal normal. The first two eigenvectors of matrix J are $e_1 = n, e_2 = b$ with eigenvalues $\lambda_1 = \kappa^2, \lambda_2 = 0$, where $b = t \times n$ is the unit binormal (see Figure 2 (a)). Applying these settings




Figure 2: Illustrating the eigenvectors e_1, e_2, e_3 of J when f is the tangent direction of a curve (a) or the normal direction on a surface (b).

to Equation 5, the bending term in extrinsic variation becomes

$$g^T Jg = \kappa^2 \cos^2 \theta, \quad (9)$$

where θ is the angle between g and the principal normal n .

It becomes clear that the curvature, κ , plays the balancing role between the smoothness and shape-alignment in an extrinsically smooth normal field. Specifically, minimizing extrinsic variation will either encourage g to align with the binormal b , in parts of the curve where the curvature κ is large, or to minimize rotation in the normal plane of the curve, in parts where κ is small. In the extreme case where $\kappa = 0$ (i.e., the curve is straight), the bending term vanishes, and any intrinsically smooth direction field (in this case, constant normal directions) is also extrinsically smooth.

Extrinsic smoothness over a curve also has a physical explanation. Suppose the curve M is the centerline of an elastic, inextensible rod with an anisotropy bending response on its cross-sections such that the response is zero in the direction $f \times g$ (i.e., bending the rod in direction orthogonal to g incurs no energy). The Dirichlet energy in Equation 1 is precisely the elastic energy of the rod [15] (assuming that both bending and twisting modulus of the material are 1). This observation justifies the use of the extrinsic smooth direction field in applications such as creating least-bending ribbons from a curve (see Section 5).

Tangent directions on a surface. Now let M be a smooth 2-dimensional surface (open or closed) and f be the unit normal vector field n over M . By basic differential geometry, ∇f is a symmetric matrix

$$\nabla f = (t_1, t_2) \cdot \begin{pmatrix} \kappa_1 & 0 \\ 0 & \kappa_2 \end{pmatrix} \cdot (t_1, t_2)^T, \quad (10)$$

where κ_1, κ_2 are the two principal curvature values (we ask that $|\kappa_1| \geq |\kappa_2|$), and $t_1, t_2 \in T_M$ are the corresponding

principal curvature directions. The first two eigenvectors of matrix J are $e_1 = t_1, e_2 = t_2$ with eigenvalues $\lambda_1 = \kappa_1^2, \lambda_2 = \kappa_2^2$ (see Figure 2 (b)). The bending term of the extrinsic variation has the form

$$g^T Jg = \kappa_2^2 + (\kappa_1^2 - \kappa_2^2) \cos^2 \theta, \quad (11)$$

where θ is the angle between g and t_1 .

The difference, $(\kappa_1^2 - \kappa_2^2)$, is the key that balances intrinsic smoothness with alignment to the surface shape. A larger difference means that the surface has a dominant bending direction (t_1), and minimizing the extrinsic variation encourages g to align with the orthogonal direction (t_2). On the other hand, a small difference implies that the surface locally lacks a single dominant direction. That is, the surface is nearly umbilic, nearly flat, or having a near-zero mean curvature. In these locations, minimizing the extrinsic variation encourages g to be intrinsically smooth.

We make a special note of the direction t_2 , which is the principal curvature direction with a smaller *absolute* curvature value. This is different from standard definition of the *minimum curvature direction*, which is associated with the smaller *signed* curvature value. This difference is important when it comes to describing features of M : while t_2 is identical with the minimum curvature direction along ridges of M , where they generally follow the ridge line, the two directions are orthogonal along valleys of M , where only t_2 follows the valley line (see Figure 3 (b)). The minimum curvature direction along a valley is usually transverse to the valley line because the curvature has a smaller, negative value in that direction. Since an extrinsically smooth field aligns with t_2 , it follows strong (in terms of the difference $\kappa_1^2 - \kappa_2^2$) ridge and valley features of the surface. This explains the observation made in [9].

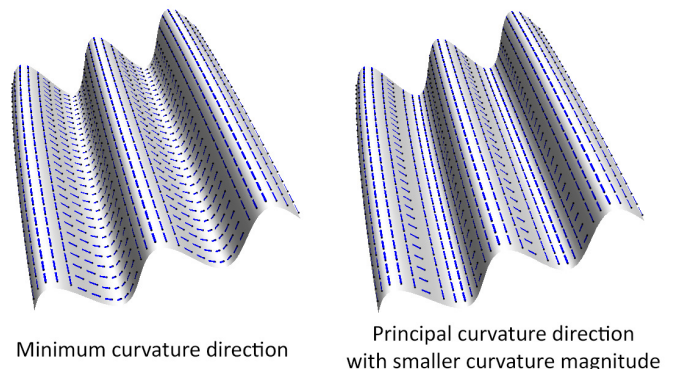


Figure 3: Comparing minimum curvature direction (a) and the principal curvature direction with a smaller absolute curvature value (t_2) (b). Note that while the latter is aligned with both ridges and valleys, the former is transverse to the valleys.

We make a final remark on singularities. As pointed out in [24], the twisting term of Equation 2, also known as the intrinsic energy, approaches infinity in the vicinity of a singularity of g on a surface (i.e., where g is not continuous). Since the Dirichlet energy is composed of the

twisting term and the bending term, it also blows up near singularities. Unfortunately, on many non-trivial surfaces, any tangent vector field has to contain singularities due to the Poincare-Hopf theorem. For these surfaces, the minimizer of the (infinite) Dirichlet energy becomes undefined. Nonetheless, we can define a discrete approximation of the Dirichlet energy over a simplicial domain that is finite by construction (see Section 4). Since the discrete energy is high around a singularity, albeit finite, minimizing the discrete energy tends to create a field with low counts of singularities. This observation was made in [9] as well as in our own experiments (see Section 5).

4. Discrete optimization

We consider the problem of computing extrinsically smooth direction fields in a discrete setting. We assume the input domain M is represented as a simplicial complex (e.g., a polyline, a triangular mesh, or a tetrahedral volume) and the direction field f is sampled at each vertex of M . Our output is a set of directions, one at each vertex, that encode the extrinsically smooth field g .

Our method is built on top of the discrete optimization method of Jakob et al. [9], which is simple but effective for the problem of finding tangent vectors on a surface. We enrich their method in three aspects. First, we provide justifications for their discrete energy formulation and extend it to domains in dimensions other than 2 (Section 4.1). Second, we provide justification for a key step in their iterative optimization (Section 4.2). Third, and more importantly, we offer an alternative approach to initialize the optimization that is simpler than theirs and requires no parameters (Section 4.3).

4.1. Discrete energy

To measure extrinsic variation of tangent directions over a surface, Jakob et al. [9] uses a particularly simple approximation as the weighted sum of squared differences over all edges of the surface. However, no rigorous justification is provided for their formulation or the choice of weights. Here we show that their edge-based formulation, with the proper choice of weights, is in fact the exact continuous Dirichlet energy (Equation 1) under the assumption of a coordinate-wise piecewise linear (PL) interpolation of g . This analysis also allows us to extend their formulation to other domains, such as curves and volumes.

The discretization of Dirichlet energy over M depends on the choice of interpolation model. The typical PL model linearly interpolates the values at vertices over each simplex. To apply this model to a vector field g while keeping it unit-length and orthogonal to f everywhere on M , one would need to first linearly interpolate f over each simple σ and then interpolate g while ensuring its unity and orthogonality to the interpolated f at each location in σ . However, it is difficult to obtain a closed-form integral of the Dirichlet energy of g in this way (although

this is possible for intrinsic energy [24, 28]). Instead, we shall drop the two constraints, unity and orthogonality, when interpolating g . That is, we perform *coordinate-wise* linear interpolation of g over each simplex. Such interpolation yields a *constant* $|\nabla g|^2$ for any point in a simplex, which makes integration trivial.

Given directions g_i at each vertex i , the Dirichlet energy defined by the coordinate-wise PL interpolation over M has been well-studied [29] and can be written as a sum over all pairs of edge-adjacent vertices i, j ,

$$\int_M |\nabla g|^2 = C_1 \sum_{\{i,j\} \in M} \omega_{ij} (g_i - g_j)^2. \quad (12)$$

Here, C_1 is a global constant and ω_{ij} are the well-known *discrete harmonic weights*. Equation 12 holds for M at any dimension (1,2,3). The form of ω_{ij} at each dimension can be easily derived or found in the literature as:

- M is a curve: $\omega_{ij} = 1/l_{ij}$ where $l_{i,j}$ is the length of the edge $\{i, j\}$.
- M is a surface: $\omega_{ij} = \cot \alpha + \cot \beta$ where α, β are the angles opposite to edge $\{i, j\}$ in the two triangles sharing this edge [29]. This is one of the weights suggested in [9].
- M is a volume: $\omega_{ij} = \sum_{\sigma} l_{\sigma} \cot \alpha_{\sigma}$ where the summation is over all tetrahedra σ sharing the edge $\{i, j\}$, l_{σ} is the length of the edge in σ opposite to $\{i, j\}$, and α_{σ} is the dihedral angle in σ along that opposite edge [30].

4.2. Optimization

Minimizing the right-hand side of Equation 12, under the constraints that each g_i is a unit vector orthogonal to an input vector f_i , is a non-linear optimization problem. A standard solution strategy is to use iterative local updates. Jakob et al. [9] uses Gauss-Seidel iterations, which we found effective in our experiments as well. Starting with an initial assignment of g_i at all vertices (to be discussed in next subsection), in each iteration, this method visits the vertices in a pre-defined order and updates the g_i at vertex i using vectors at all vertices sharing an edge with i (i.e., the 1-ring neighbors of i , denoted as N_i). The update involves first taking the weighted average of those vectors, then projecting onto the plane orthogonal to f_i , and finally normalizing to have unit length:

$$\begin{aligned} g_i^* &\leftarrow \sum_{j \in N_i} \omega_{ij} g_j \\ g_i' &\leftarrow (I - f_i f_i^T) g_i^* \\ g_i &\leftarrow g_i' / |g_i'| \end{aligned} \quad (13)$$

Although this update step is simple and intuitive, one may ask if it is *optimal* in the sense that updated vector g_i minimizes the discrete energy (Equation 12) locally (i.e., assuming g_j at all other vertices $j \neq i$ are fixed). We give an affirmative answer by the following lemma:

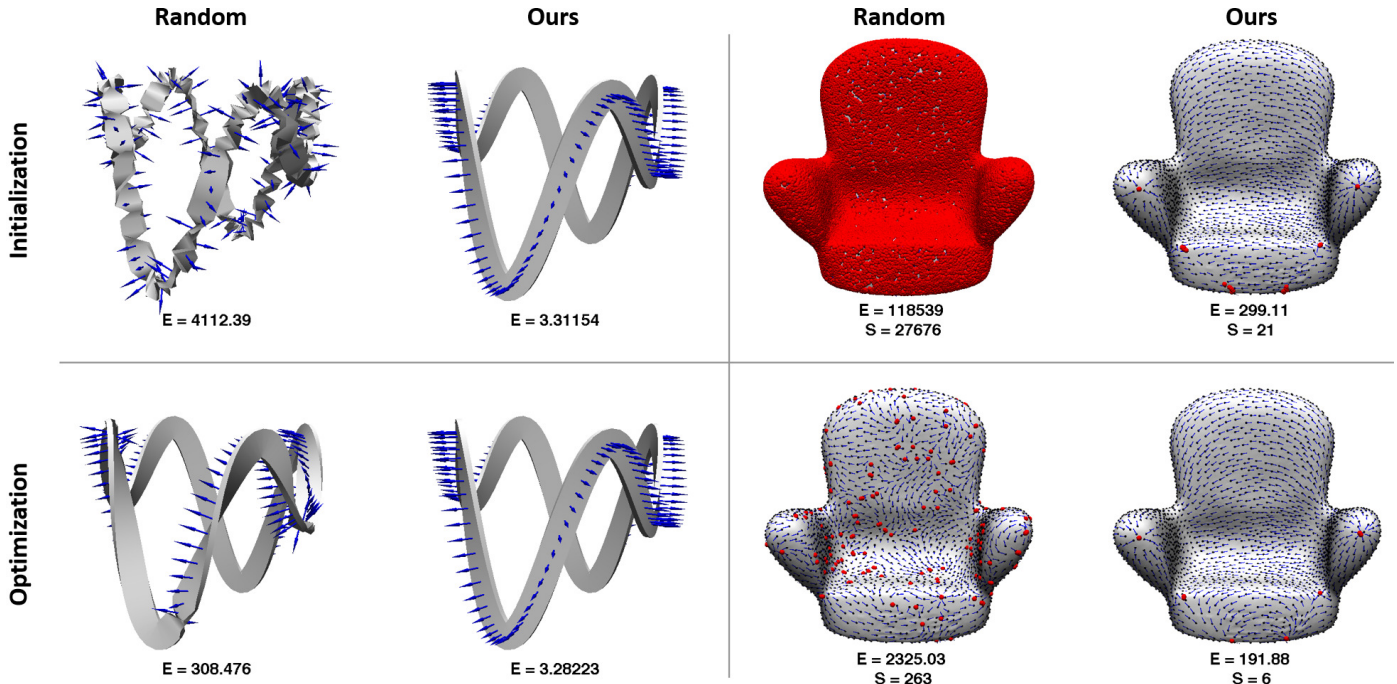


Figure 4: Comparing random assignment with our eigenvector approach for initializing the normal direction field on a curve (left) and the tangent direction field on a chair surface (right). We show both the initial fields (top row) and after iterative optimization (bottom row). The curves are visualized as ribbons whose cross-sections are elongated in the direction orthogonal to the normals, and the normals are drawn as blue arrows. On a surface, the tangent directions are drawn as blue lines and the red dots are singularities. The Dirichlet energy (E) for each field is shown, and so is the number of singularities (S) on the surface. Note that the initialization using our eigenvector approach yields better optimization results.

Lemma 4.1. *The vector g_i obtained by Equation 13 is the solution to the following constrained minimization problem:*

$$\arg \min_v D(v), \text{ s.t. } v^T v = 1, v^T f_i = 0 \quad (14)$$

where

$$D(v) = \sum_{j \in N_i} \omega_{ij} (v - g_j)^2$$

Proof. If we drop the constraint that v needs to be a unit vector, $D(v)$ is a quadratic function over the plane orthogonal to f whose minimum is achieved precisely at g'_i in Equation 13. In addition, the gradient of $D(v)$ over the plane orthogonal to f is $v - g'_i$. Bringing back the unity constraint, the solution to Equation 14 is the unit vector v where the gradient, $v - g'_i$, is parallel to v . There are two such unit vectors, $g'_i/|g'_i|$ and $-g'_i/|g'_i|$, but the former is closer to g'_i and hence is the true minimizer. \square

A direct implication of Lemma 4.1 is that the iterative optimization of [9] always converges to a local minimum, since the energy decreases monotonically at each iteration. The optimization can thus be considered as a block coordinate descent of the discrete Dirichlet energy.

4.3. Initialization

The iterative optimization is sensitive to the initial assignment of g_i . For example, a random assignment easily

leads to sub-optimal results that have high energy and contain numerous singularities (see Figure 4). To address this limitation, Jakob et al. [9] employs the standard hierarchical optimization method [4, 5], which involves creating a sequence of meshes with decreasing levels of complexity, optimizing the field at coarser levels and propagating the result to the next finer level as the initial assignment. While delivering impressive results, the hierarchical method requires additional parameters (e.g., the number of meshes, amount of simplification between meshes at different levels, etc.), and adds complication to the optimization framework (e.g., need to generate multiple meshes and maintain their relations). Moreover, the hierarchical method was only designed for surface meshes.

We propose a simple and parameter-free initialization method that leads to pleasing optimization results in practice. We minimize a modified energy that replaces the unity constraint on g in our original problem with the unity constraint on the *integral of the squared length of g over M* . This change allows us to find the global optimal solution by solving a single eigenvector problem. The resulting g is normalized before serving as the initial assignment for iterative optimization. This strategy was originally proposed in [24] for computing intrinsically smooth fields and later adapted to compute smooth complex functions [31], both over surface meshes. Here we further adapt the strategy to initialize extrinsically smooth fields over 1-, 2-, or 3-dimensional domains.

Specifically, our initialization step solves the following continuous problem:

$$\arg \min_g \int_M |\nabla g|^2, \quad \text{s.t. } f^T g = 0, \int_M g^T g = 1 \quad (15)$$

Note that g is no longer a “direction” field as it has varying magnitude. In our discrete setting, dropping the unity constraint allows us to express each g_i at a vertex in a *linear* form, $a_i \alpha_i + b_i \beta_i$, where α_i, β_i is a fixed pair of orthonormal basis in the plane orthogonal to f_i , and a_i, b_i become the unknown scalars to be solved. Denote $a = \{a_1, \dots, a_m\}^T$ and $b = \{b_1, \dots, b_m\}^T$ where m is the number of vertices of M .

Using coordinate-wise PL interpolation, the field g can be expressed using a set of linear bases Φ_i (resp. Ψ_i), each being a hat-function that evaluates to be α_i (resp. β_i) at vertex i and attenuates (in magnitude) linearly in each adjacent triangle of vertex i . Denote $\Phi = \{\Phi_1, \dots, \Phi_m\}$ and $\Psi = \{\Psi_1, \dots, \Psi_m\}$. We can write the interpolated g as

$$g = (\Phi \ \Psi) \begin{pmatrix} a \\ b \end{pmatrix} \quad (16)$$

Using this representation, we can write the Dirichlet energy in a quadratic form:

$$\int_M |\nabla g|^2 = C_2 \begin{pmatrix} a^T & b^T \end{pmatrix} \begin{pmatrix} A_{\alpha,\alpha} & A_{\alpha,\beta} \\ A_{\beta,\alpha} & A_{\beta,\beta} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \quad (17)$$

where C_2 is a constant and each $A_{\square,\Delta}$ is an $n \times n$ matrix (\square, Δ denotes either α or β). The (i, j) -th entry of $A_{\square,\Delta}$ is non-zero only if $i = j$ or $j \in N_i$ (i.e., $\{i, j\}$ is an edge of M). In the former case, the entry is $\square_i^T \Delta_i \sum_{j \in N_i} \omega_{ij}$, where ω_{ij} is the discrete harmonic weight. In the latter case, the entry is $-\square_i^T \Delta_j \omega_{ij}$. Note that these matrix entries on a 2-dimensional M are similar to those used in [31], where the dot products $\square_i^T \Delta_j$ are replaced by entries of some rotation matrix defined on the edge $\{i, j\}$.

By deriving integrals of products of barycentric coordinates, the integral of squared norm of g can also be written in a quadratic form:

$$\int_M g^T g = C_3 \begin{pmatrix} a^T & b^T \end{pmatrix} \begin{pmatrix} M_{\alpha,\alpha} & M_{\alpha,\beta} \\ M_{\beta,\alpha} & M_{\beta,\beta} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \quad (18)$$

where C_3 is a constant and each (i, j) -th entry of the $n \times n$ matrix $M_{\square,\Delta}$ (\square, Δ denotes either α or β) is non-zero only if $i = j$ or $j \in N_i$. In the former case, the entry is $\square_i^T \Delta_i \sum_{\sigma} V_{\sigma}$, where the summation is taken over all highest-dimensional simplices σ (i.e., edge, triangle and tetrahedron for a 1-, 2- and 3- dimensional M) sharing vertex i , and V_{σ} is the length, area or volume of σ . In the latter case, the entry is $\square_i^T \Delta_j \sum_{\sigma} V_{\sigma}/2$, where the summation is now taken over all simplices σ that share the edge $\{i, j\}$.

The solution (a, b) to the minimization problem of Equation 15 is thus the eigenvector corresponding to the

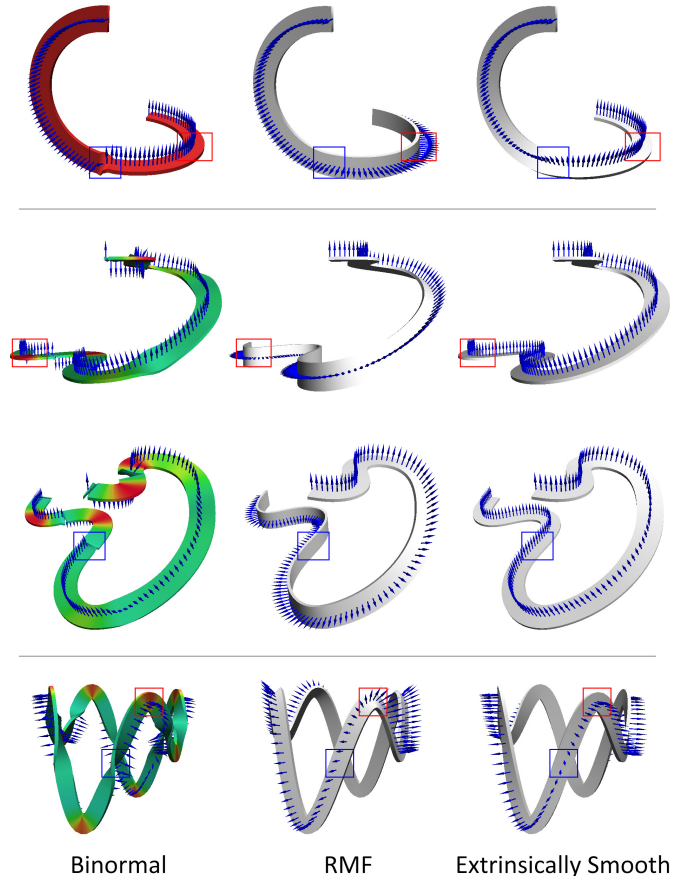


Figure 5: Comparing the binormal (left), the Rotation Minimizing Frames [11] (middle), and extrinsically smooth normal fields (right) on two open curves (second one shown from two views) and a closed curve. The normal vectors are shown in blue arrows, and the curves are shown as ribbons. The ribbons in the left column are colored by the squared curvature from green (low) to red (high). Note how the extrinsically smooth field aligns with the binormal where the curvature is high (e.g., see red outlines) and minimizes twist where the curvature is low (e.g., see blue outlines).

smallest eigenvalue in the following generalized eigenvalue problem:

$$\begin{pmatrix} A_{\alpha,\alpha} & A_{\alpha,\beta} \\ A_{\beta,\alpha} & A_{\beta,\beta} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \lambda \begin{pmatrix} M_{\alpha,\alpha} & M_{\alpha,\beta} \\ M_{\beta,\alpha} & M_{\beta,\beta} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \quad (19)$$

In our experiments, we have observed that the initial fields obtained using our eigenvector approach are already very close to being extrinsically smooth with few singularities (see examples in Figure 4). In most of our tests, iterative optimization only results in minor improvement of the Dirichlet energy upon this initial field.

5. Experimental results

We demonstrate the extrinsically smooth direction fields computed by our method for a number of problems and compare with existing methods.

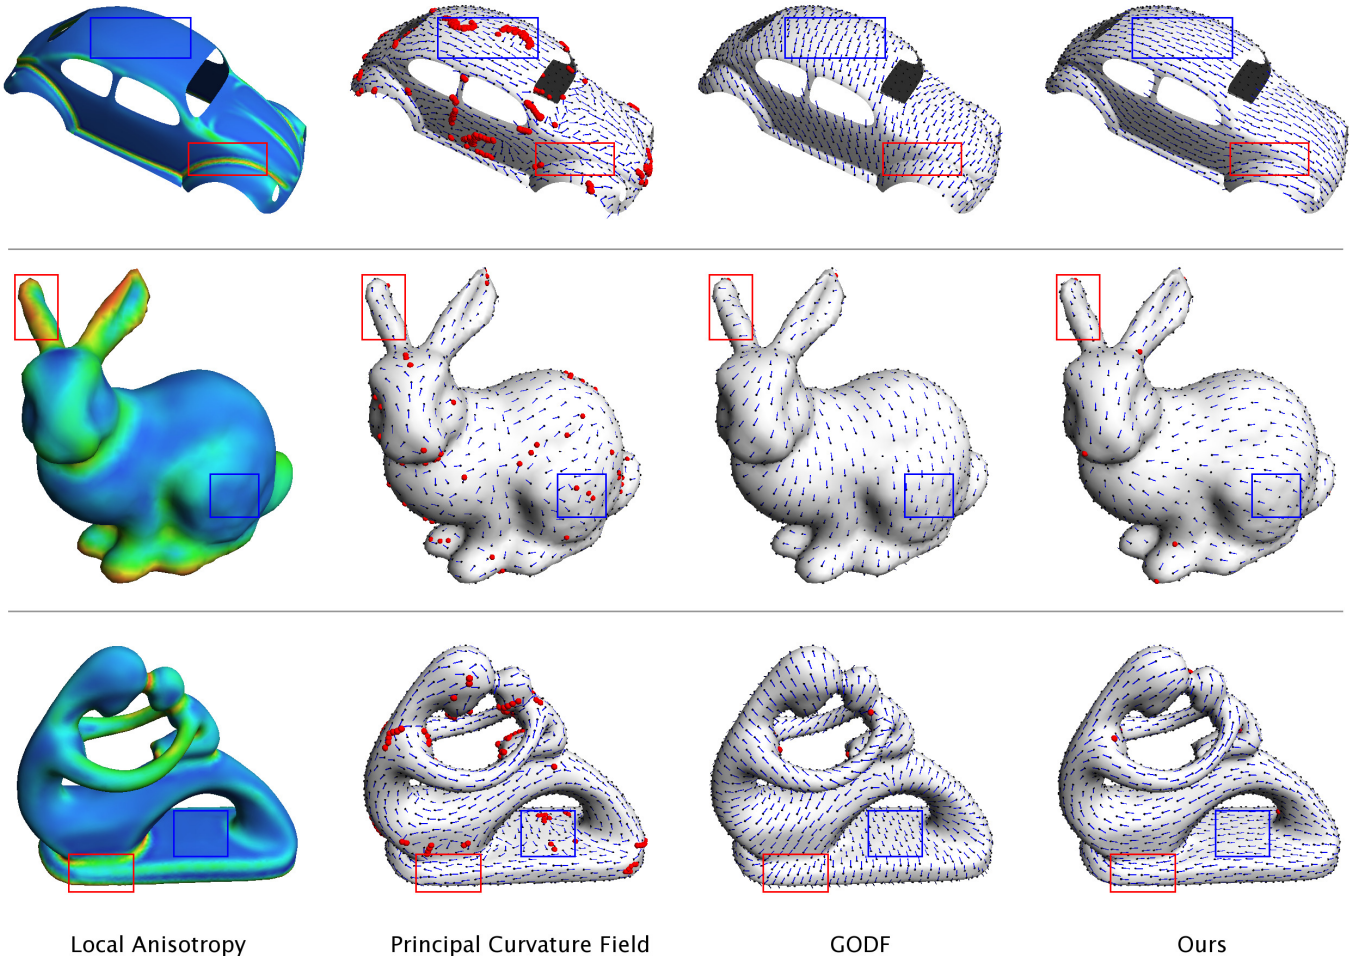


Figure 6: Comparing the principal curvature fields t_2 (middle-left), the Global Optimal Direction Fields [24] (middle-right), and extrinsically smooth tangent fields (right) on three surfaces (the Beetle has boundaries). The normal vectors are shown in blue arrows and singularities in red. The left column colors each surface by local anisotropy, $(\kappa_1^2 - \kappa_2^2)$, from blue (low) to red (high). Note how the extrinsically smooth field aligns with t_2 where the anisotropy is high (e.g., see red outlines) and minimizes twist where the anisotropy is low (e.g., see blue outlines).

Normal directions on a curve. We compare extrinsically smooth normal directions with the binormal along the curve and the Rotation Minimizing Frames (RMF) [11] that minimize intrinsic variation. RMF produces a sequence of aligned normal *frames* along the curve, and any sequence of corresponding vectors in these frames are equally smooth in the intrinsic sense. For comparison purposes, we pick the sequence of vectors that matches our extrinsically smooth direction at one end point of an open curve or an arbitrary point on a closed curve. The results on three example curves are shown in Figure 5.

Observe that the binormal captures the local shape of the curve. For example, the binormal along a near-planar curve segment is normal to the underlying plane. However, the binormal is undefined where the curve is only C^1 continuous (e.g., the first example) or where the curvature is zero (e.g., the second example), and it can vary significantly where the curvature is low (e.g., the third example). On the other hand, although the direction fields derived from RMF appear very smooth, they do not reveal

the shape of the curve: they are insensitive of whether the curve is locally planar or not. The extrinsically smooth field seems to combine the best of these two fields: by aligning to the binormal where the curvature is high, it better captures the local shape of the curve than RMF (see regions outlined in red); by maximizing intrinsic smoothness where the curvature is low, it avoids the instability and twists of the binormal (see regions outlined in blue).

One application of the extrinsically smooth normal field is generating ribbon visualizations. In Figure 5, the cross-sections of the ribbons are elongated in the orthogonal direction to the computed normals. Compared with ribbons generated from the binormal and RMF, ribbons generated from extrinsically smooth fields reveal the curve shape while minimizing twists. Physically speaking, such ribbons also have the minimal bending energy (see discussion in Section 3.2).

Tangent directions on a surface. We compare extrinsically smooth tangent fields with the principal curvature field t_2 (associated with the principal curvature of smaller magni-

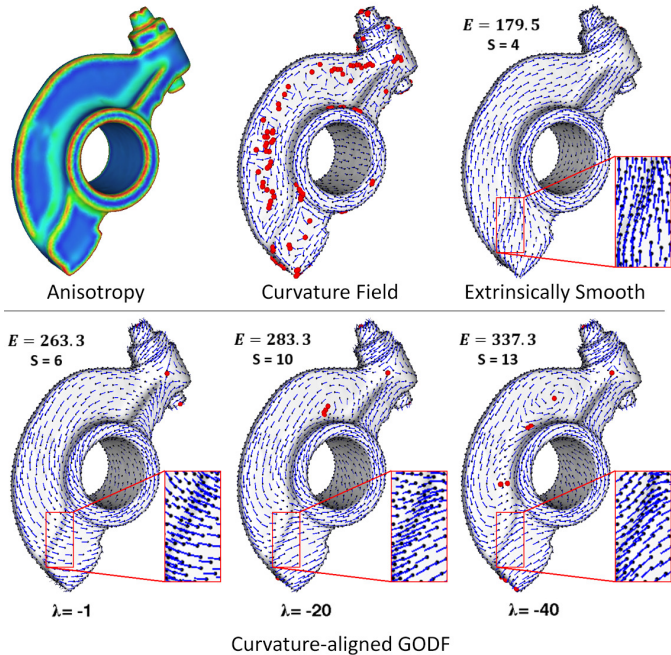


Figure 7: Rockarm: plot of local anisotropy ($\kappa_1^2 - \kappa_2^2$) (top-left), the principal curvature field t_2 (top-middle), the extrinsically smooth tangent field (top-right), and the curvature-aligned GODF [24] (bottom) with increasing alignment strength λ . The Dirichlet energy (E) and number of singularities (S) for each field are shown. Note that a large magnitude of λ allows GODF to align to strong shape features (e.g., the ridge in the red outline) at the cost of creating unwanted twisting and singularities.

tude) and the Global Optimal Direction Fields (GODF) [24] that minimize intrinsic variation. We orient the undirected t_2 using a minimum-spanning tree over the edge graph of the mesh [21]. The results on three examples are shown in Figure 6.

Although the principal curvature direction t_2 follows prominent ridges and valleys (see Figure 3), it is inherently unstable in regions where the anisotropy measure ($\kappa_1^2 - \kappa_2^2$) is low (cf. the plot of anisotropy in the left column). On the other hand, while GODF is smooth and has few singularities, in general it does not follow any features of the surface. The extrinsically smooth field combines the best of both: by aligning to t_2 where anisotropy is high, it better conforms to strong surface features than GODF (see regions outlined in red); by maximizing intrinsic smoothness where the anisotropy is low, it avoids the instability of t_2 (see regions outlined in blue).

The GODF method [24] can optionally produce a curvature-aligned and smooth 1-vector field using a curvature guidance field (we fed the method with the oriented t_2 field). The alignment strength is controlled by a parameter λ , which needs to be manually tuned. In some cases, like the one shown in Figure 7, it is challenging to find a suitable parameter: while a large magnitude of λ is needed to align the field with strong features (including the ridge in the red outline), the same parameter value also creates unwanted twisting and singularities caused by alignment

to less obvious features. In comparison, the extrinsically smooth field achieves better feature alignment with less twisting and fewer singularities, and there is no need for parameter-tuning.

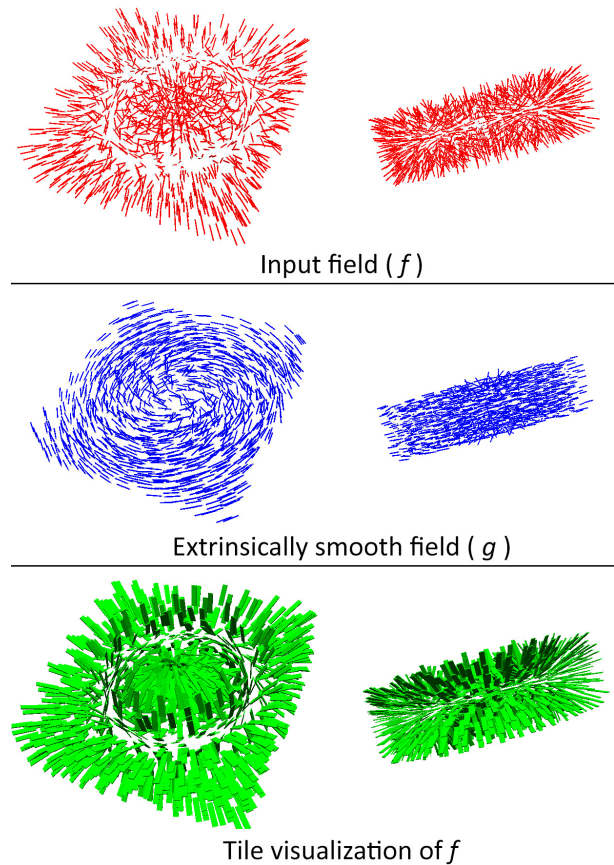


Figure 8: Input field f within a 3-dimensional volume (top), the extrinsically smooth direction field g (middle), and the tile-based visualization (bottom) of f where each tile is created by directions f and g . Note that, compared to the line segments (top), the tiles (bottom) significantly enhance the perception of f .

Other direction field problems. We explore two other problems with practical implications.

First, consider the problem where M is a curve and f is an arbitrary smooth vector field along M . One may consider M as the pre-planned path of a virtual camera, and f being the “LookAt” vector of the camera that aims at a static or moving object as the camera travels along M (Figure 9 far-left). In this interpretation, the sought orthogonal direction field, g , serves as the “Up” vector of the moving camera. It is well known that the standard practice of projecting a global Z vector onto the orthogonal plane of f results in camera flips when f is near-parallel to Z (Figure 9 middle-left). Another approach is using the cross-product of successive LookAt vectors (which in fact approximates the second eigenvector e_2 of matrix J), but the result can be unstable when successive LookAt vectors are near co-linear (Figure 9 middle). One may also minimize the “rolling” of camera around the LookAt vector by the parallel-transport of an initial Up vector along the




Figure 9: Different methods for creating Up vectors (blue arrows) from a camera path and LookAt vectors (red arrows) along the path (far-left, green dot is the focal point): projection of a global Z vector (middle-left), using cross-product of successive LookAt vectors (middle), parallel-transporting an initial Up vector (middle-right), and an extrinsically smooth Up vector field (far-right). Note that the last field is smooth and minimizes both rolling and pitching motion of the camera.

curve, but the resulting camera may have large “pitching” motion (Figure 9 middle-right). In comparison, an extrinsically smooth Up vector field minimizes both rolling and pitching of the camera (Figure 9 far-right).

Next, consider a 3-dimensional volume M with a smooth direction field f within the volume. A potential application for an extrinsically smooth direction field g orthogonal to f is enhancing the visualization of f . While f is commonly visualized by its samples or streamlines [32], these 1-dimensional geometric primitives do not easily convey the 3-dimensional shape of f (Figure 8 top). We can turn these lines or curves into 2-dimensional tiles or ribbons using the g direction. In the simplest case, we can visualize each sample vector of f as a rectangular tile whose long and short sides are aligned with f and g , respectively (Figure 8 bottom). The resulting tiles not only enhance the 3-dimensional feel of the visualization but also reveals the variation of f itself, since g strives to stay orthogonal to the direction of greatest variation of f .

Performance. Our method was implemented in C++ on a 2.5 GHz Intel Core i7. Initializing the field (Section 4.3) takes between 0.047 and 4.406 seconds for M ranging from 1200 to 61844 vertices. Iterative optimization (Section 4.2) takes between 0.192 and 10.276 seconds.

6. Conclusion

In this paper, we advocate *extrinsic smoothness* in designing direction (i.e., unit vector) fields for a variety of problem settings. We give a simple geometric explanation of the relation of extrinsic variation and the shape of the underlying domain in two motivating problems, finding normal fields over a curve and finding tangent fields over a surface. We enrich, justify, and improve an existing optimization framework to compute extrinsically smooth fields over curves, surfaces, and volumes.

In the future, we would like to explore the extension of our analysis in several directions, including general n -vector fields and direction fields in domains of higher dimensions (e.g., for use in visualizing time-varying flow

fields). It would also be interesting to investigate how vector field topology can be incorporated as an additional constraint to produce fields that are not only smooth and conforming to shape, but also having a prescribed topological structure (e.g., location and index of singularities, a given network of separatrices, etc.).

References

- [1] Siltanen P, Woodward C. Normal orientation methods for 3d offset curves, sweep surfaces and skinning. *Comput Graph Forum* 1992;11(3):449–57.
- [2] Wang W, Joe B. Robust computation of the rotation minimizing frame for sweep surface modeling. *Computer-Aided Design* 1997;29(5):379–91.
- [3] Hanson AJ, Ma H. Quaternion frame approach to streamline visualization. *IEEE Trans Vis Comput Graph* 1995;1(2):164–74.
- [4] Turk G. Texture synthesis on surfaces. In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. ISBN 1-58113-374-X; 2001, p. 347–54.
- [5] Wei LY, Levoy M. Texture synthesis over arbitrary manifold surfaces. In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. ISBN 1-58113-374-X; 2001, p. 355–60.
- [6] Praun E, Hoppe H, Webb M, Finkelstein A. Real-time hatching. In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. ISBN 1-58113-374-X; 2001, p. 581.
- [7] Hertzmann A, Zorin D. Illustrating smooth surfaces. In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '00. ISBN 1-58113-208-5; 2000, p. 517–26.
- [8] Xu K, Cohn-Or D, Ju T, Liu L, Zhang H, Zhou S, et al. Feature-aligned shape texturing. *ACM Transactions on Graphics (Proceedings SIGGRAPH Asia 2009)* 2009;28(5):1–7.
- [9] Jakob W, Tarini M, Panozzo D, Sorkine-Hornung O. Instant field-aligned meshes. *ACM Trans Graph* 2015;34(6):189:1–189:15.
- [10] Bishop RL. There is more than one way to frame a curve. *Amer Math Monthly* 1975;82(3):246–51.
- [11] Wang W, Jüttler B, Zheng D, Liu Y. Computation of rotation minimizing frames. *ACM Trans Graph* 2008;27(1).
- [12] Hanson AJ. Constrained optimal framings of curves and surfaces using quaternion gauss maps. In: *IEEE Visualization*. 1998, p. 375–82.
- [13] Lazarus SC, Jancene P. Axial deformations: an intuitive deformation technique. *Computer-Aided Design* 1994;26(8):607–13.

- [14] Llamas I, Powell A, Rossignac J, Shaw CD. Bender: A virtual ribbon for deforming 3d shapes in biomedical and styling applications. In: Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling. SPM '05. ISBN 1-59593-015-9; 2005, p. 89–99.
- [15] Bergou M, Wardetzky M, Robinson S, Audoly B, Grinspun E. Discrete elastic rods. *ACM Trans Graph* 2008;27(3):63:1–63:12.
- [16] Praun E, Finkelstein A, Hoppe H. Lapped textures. In: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '00. ISBN 1-58113-208-5; 2000, p. 465–70.
- [17] Fisher M, Schröder P, Desbrun M, Hoppe H. Design of tangent vector fields. *ACM Trans Graph* 2007;26(3).
- [18] Ray N, Vallet B, Li WC, Lévy B. N-symmetry direction field design. *ACM Trans Graph* 2008;27(2):10:1–10:13.
- [19] Ray N, Vallet B, Alonso L, Levy B. Geometry-aware direction field processing. *ACM Trans Graph* 2009;29(1):1:1–1:11.
- [20] Bommes D, Zimmer H, Kobbelt L. Mixed-integer quadrangulation. *ACM Trans Graph* 2009;28(3):77:1–77:10.
- [21] Crane K, Desbrun M, Schröder P. Trivial connections on discrete surfaces. *Comput Graph Forum* 2010;29(5):1525–33.
- [22] Lai Y, Jin M, Xie X, He Y, Palacios J, Zhang E, et al. Metric-driven rosy field design and remeshing. *IEEE Trans Vis Comput Graph* 2010;16(1):95–108.
- [23] Li G, Guo L, Nie J, Li K, Liu T. Direction field diffusion on cortical surface via graph cuts. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2010, San Francisco, CA, USA, 13-18 June, 2010. 2010, p. 95–102.
- [24] Knöppel F, Crane K, Pinkall U, Schröder P. Globally optimal direction fields. *ACM Trans Graph* 2013;32(4):59:1–59:10.
- [25] Iarussi E, Bommes D, Bousseau A. Bendfields: Regularized curvature fields from rough concept sketches. *ACM Trans Graph* 2015;34(3):24:1–24:16.
- [26] Gonzalez-Davila J, Vanhecke L. Energy and volume of unit vector fields on three-dimensional riemannian manifolds. *Differential Geometry and its Applications* 2002;16(3):225–44.
- [27] Azencot O, Ovsjanikov M, Chazal F, Ben-Chen M. Discrete derivatives of vector fields on surfaces – an operator approach. *ACM Trans Graph* 2015;34(3):29:1–29:13.
- [28] Liu B, Tong Y, Goes FD, Desbrun M. Discrete connection and covariant derivative for vector field analysis and design. *ACM Trans Graph* 2016;35(3):23:1–23:17.
- [29] Pinkall U, Polthier K. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 1993;2(1):15–36.
- [30] Ju T, Liepa P, Warren JD. A general geometric construction of coordinates in a convex simplicial polytope. *Computer Aided Geometric Design* 2007;24(3):161–78.
- [31] Knöppel F, Crane K, Pinkall U, Schröder P. Stripe patterns on surfaces. *ACM Trans Graph* 2015;34(4):39:1–39:11.
- [32] McLoughlin T, Laramee RS, Peikert R, Post FH, Chen M. Over two decades of integration-based, geometric flow visualization. In: Eurographics 2009 - State of the Art Reports, Munich, Germany, March 30 - April 3, 2009. 2009, p. 73–92.