

Octahedral Frames for Feature-Aligned Cross Fields

PAUL ZHANG, Massachusetts Institute of Technology
JOSH VEKHTER, University of Texas at Austin
EDWARD CHIEN, Massachusetts Institute of Technology
DAVID BOMMES, University of Bern
ETIENNE VOUGA, University of Texas at Austin
JUSTIN SOLOMON, Massachusetts Institute of Technology



Fig. 1. A variety of feature-aligned cross fields computed using our novel cross-field formulation.

We present a method for designing smooth cross fields on surfaces that automatically align to sharp features of an underlying geometry. Our approach introduces a novel class of energies based on a representation of cross fields in the spherical harmonic basis. We provide theoretical analysis of these energies in the smooth setting, showing that they penalize deviations from surface creases while otherwise promoting intrinsically smooth fields. We demonstrate the applicability of our method to quad meshing and include an extensive benchmark comparing our fields to other automatic approaches for generating feature-aligned cross fields on triangle meshes.

P. Zhang received generous support from the Department of Energy Computer Science Graduate Fellowship. E. Vouga received generous support from Adobe, SideFX, and NSF IIS-1910274. D. Bommes received generous support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (AlgoHex, grant agreement no. 853343). J. Solomon received generous support from Army Research Office grant W911NF-12-R-0011, National Science Foundation grant IIS-1838071, Air Force Office of Scientific Research award FA9550-19-1-0319, and a gift from Adobe Systems. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of these organizations.

Authors' addresses: P. Zhang, E. Chien, and J. Solomon, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139; emails: pzp1@mit.edu, eddchien@csail.mit.edu, jsolomon@mit.edu; J. Vekhter and E. Vouga, University of Texas at Austin, Austin, TX 78712; emails: {josh, evouga}@cs.utexas.edu; D. Bommes, University of Bern, Hochschulstrasse 6, Bern, 3012, Switzerland; email: david.bommes@inf.unibe.ch.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Link to use: <https://creativecommons.org/licenses/by/4.0/>

0730-0301/2020/04-ART25
<https://doi.org/10.1145/3374209>

CCS Concepts: • **Computing methodologies** → **Mesh models; Shape analysis;**

Additional Key Words and Phrases: Discrete differential geometry, geometry processing, total variation, singularities, feature alignment

ACM Reference format:

Paul Zhang, Josh Vekhter, Edward Chien, David Bommes, Etienne Vouga, and Justin Solomon. 2020. Octahedral Frames for Feature-Aligned Cross Fields. *ACM Trans. Graph.* 39, 3, Article 25 (April 2020), 13 pages.
<https://doi.org/10.1145/3374209>

1 INTRODUCTION

N -rotationally symmetric (RoSy) tangential vector fields over surfaces are ubiquitous in computer graphics. 2-RoSy fields can be used to generate stripe patterns due to their ambivalence to rotation by π about the normal. 4-RoSy fields (cross fields) are heavily used in both surface parameterization and quadrilateral (quad) meshing, thanks to their symmetry with respect to rotations by $\frac{\pi}{2}$ about the surface normal.

Depending on the application, n -RoSy field design algorithms must trade off between several desirable properties of the field. In almost all cases, n -RoSy fields are expected to be as smooth as possible. For surfaces with boundary, constraints on how the field aligns to the boundary are common, and for artistic applications, users may wish to prescribe a sparse set of streamlines that the field must follow. For meshing applications, alignment of n -RoSy fields to salient geometric features is also desirable as a means to identify or preserve mesh detail. Our focus will be on improving this latter aspect for the important case of 4-RoSy fields.




Fig. 2. Two surfaces (the three-cylinder-intersection and wavy-box) whose maximal curvature directions (blue lines) contradict its feature curves (red lines).

There are two broad strategies for achieving feature alignment. The first is to optimize only for smoothness, under the assumption that a well-chosen functional for measuring cross-field smoothness will automatically penalize fields that fail to align to geometric features. The most commonly used smoothness functionals (including the Dirichlet energy and its variants) are *intrinsic* and recover solutions that are unique only up to rotation [Knöppel et al. 2013]. These are ambivalent to isometric deformations of the surface and ignore *extrinsic* features such as creased folds.

An alternative strategy is to include energy terms that explicitly enforce alignment to an input guiding field of principal curvature directions during cross-field design [Brandt et al. 2018; Knöppel et al. 2013]. Drawbacks include the difficulty of robustly computing principal curvature directions on noisy meshes, the fact that forcing alignment to a guiding field based on local geometry may exclude cross-field designs that are globally more optimal, and more critically the fact that principal curvature directions are often different from features (e.g., Figure 2).

Our main observation is that neither of the preceding strategies adequately identify those features most important to generating high-quality quad meshes. Often the surface being modeled is constructed from smooth patches that are joined along *sharp extrinsic feature curves* where the normal direction is discontinuous or changes rapidly. On the one hand, such features are invisible to intrinsic smoothness functionals; on the other hand, the orientation of the feature curves often contradict that of nearby curvature lines.

Consider the surfaces shown in Figure 2: neither existing strategy will promote alignment to the features curves shown in red. Both of these shapes are developable away from a sparse set of cone singularities at the corners; the Gaussian curvature is nearly zero at creased edges and curved facets, and so purely intrinsic approaches have no hope of aligning to the creases. Augmenting with a guiding field based on extrinsic curvature is counter-productive, as the curvature lines (blue) are not compatible with the surface's more important crease features curves (red).

We approach feature alignment in a new way, which detects and aligns cross fields to sharp features in a stable fashion. Our method is based on an extrinsic representation of cross fields using spherical harmonic (SH) basis functions. SH functions have been used successfully in volumetric octahedral field problems for hexahedral meshing [Huang et al. 2011; Ray et al. 2016;

Solomon et al. 2017], and we argue that this representation is well suited not only for computing octahedral fields in volumes but also for field computation on surfaces. In particular, we apply an SH representation of octahedral frames, or frames of three orthogonal directions in \mathbb{R}^3 , proposed by Huang et al. [2011]; when one of its directions is constrained to the surface normal, it exhibits the same symmetry as a two-dimensional cross. We use this fact to devise a class of cross-field energies that promote intrinsic smoothness in smooth regions of the surface. Over sharp creases, however, our energy aligns the field to the crease direction, achieving automatic feature alignment without the need for explicit computation of extrinsic curvature directions or feature curves.

Contributions. In this work, we

- introduce SH functions for the computation of surface cross fields,
- propose a family of field smoothness energies whose optima are feature-aligned cross fields,
- provide a theoretical analysis of the behavior of a few important members of this family, and
- introduce cross fields with soft normal alignment for increased versatility/robustness.

Our approach is able to extract feature-aligned fields with comparable levels of efficiency to those of purely intrinsic algorithms. We tested our algorithm extensively on more than 200 different meshes, with results presented in both Section 6 and the supplementary material. We leverage our algorithm to produce feature-aligned cross fields and demonstrate their usefulness for quad meshing.

2 RELATED WORK

The generation of tangential n -RoSy fields over surfaces has many applications in computer graphics ranging from surface BRDF modification [Brandt et al. 2018] to meshing [Bommes et al. 2009; Jakob et al. 2015] to texture synthesis [Knöppel et al. 2015] and sketch-based modeling [Bessmeltsev and Solomon 2019; Iarussi et al. 2015]. Surveys of n -RoSy field design methods are provided in Vaxman et al. [2016] and de Goes et al. [2015].

2.1 Cross-Field Design

Cross fields ($n = 4$) have been especially well studied since their $\frac{\pi}{2}$ -symmetry allows them to behave like local coordinate systems, resulting in intuitive seamless surface parameterizations.

Methods to compute intrinsically smooth cross fields with alignment and singularity constraints were studied by Ray et al. [2008], Crane et al. [2010], and Knöppel et al. [2013]. More similar to our work, Jakob et al. [2015] instead formulated an *extrinsic* smoothness functional on cross fields in an attempt to automatically align to surface features. Their method penalizes an extrinsic distance between neighboring crosses that does not use a shared tangent space or connection. The resulting energy is non-convex but is minimized to local optimality, often resulting in more singularities than necessary. Huang and Ju [2016] analyzed this extrinsic energy, finding that it can be decomposed into an energy expressed in terms of intrinsic twisting and alignment to extrinsic curvature directions. We performed a similar analysis of the energy we introduce in the supplemental materials.

When using cross fields for quad mesh parameterization or processing, methods [Bommes et al. 2009; Brandt et al. 2018; Campen et al. 2016; Knöppel et al. 2013] often promote feature alignment by including a loss term penalizing disagreement with curvature directions. However, as we argue in Section 1 and illustrate in Figure 2, alignment to curvature directions is often less important than alignment to sharp creases. Other parameterization methods such as those of Bommes et al. [2013, 2009] and Campen et al. [2015] allow feature alignment but just assume that such feature curves are provided as input.

2.2 Octahedral Fields and Volumetric Representations

The three-dimensional generalization of a cross field is an *octahedral field*. Octahedral fields are often used in volumetric problems like hexahedral meshing [Nieser et al. 2011]. A single octahedral frame consists of three mutually orthogonal vectors and their negations. Huang et al. [2011] introduced a particularly convenient representation of an octahedral frame as a rotation of the spherical function $g(x, y, z) : (x, y, z) \in S^2 \mapsto x^4 + y^4 + z^4$, encoded by coefficients in the SH basis. Ray et al. [2016] used this representation to generate volumetric normal-aligned octahedral fields, and Solomon et al. [2017] combined the SH representation with the boundary element method (BEM) to remove the need for a volumetric mesh. Both of these methods use normal alignment constraints at the surface to enforce alignment of the octahedral frame with the volume's boundary.

Although there is no canonical three-dimensional generalization of arbitrary n -RoSy fields, the SH representation allows for frames that mimic the symmetries of all platonic solids [Shen et al. 2016], including octahedral fields [Corman and Crane 2019; Liu et al. 2018; Solomon et al. 2017]. Algebraic characterization of the orbit of $g(x, y, z)$ under the space of rotations as a subset of all possible SH coefficients was presented by Palmer et al. [2019] and Chemin et al. [2018].


3 PRELIMINARIES

Since our formulation relies heavily on both the SH representation of octahedral frames and vectorial total variation, we present a preliminary introduction to these topics.

3.1 SH Octahedral Frames

As introduced by Huang et al. [2011], the canonical axis-aligned octahedral frame can be represented by SHs as a function $g_0 : S^2 \rightarrow \mathbb{R}$ written as $g_0 = \sqrt{\frac{5}{12}} Y_{44} + \sqrt{\frac{7}{12}} Y_{40}$, where Y_{lm} denotes the basis for real SHs. The function g_0 can be understood as the scaled projection of $x^4 + y^4 + z^4$ onto the fourth band ($l = 4$) of SHs. Written differently, we can encode g_0 as a vector of coefficients in the full basis of fourth-band SHs $Y_{4(-4)}, \dots, Y_{44}$:

$$f_0 = \left[0, 0, 0, 0, \sqrt{\frac{7}{12}}, 0, 0, 0, \sqrt{\frac{5}{12}} \right]^T.$$




The space of octahedral frames can be described as all rotations of the canonical octahedral frame—that is, the orbit of f_0 under the group of 3D rotations $SO(3)$ (see Definition 3.5 in Palmer et al. [2019]). We write this via exponentiation of the Lie algebra elements: the set of octahedral frames is

$$\mathcal{V} = \left\{ f \mid \text{there exists } v \in \mathbb{R}^3 \text{ with } f = e^{v \cdot L} f_0 \right\},$$

where $v \cdot L = v_x L_x + v_y L_y + v_z L_z$ and L_x, L_y, L_z are the angular momentum operators expressed in the basis of band-four SHs. In this basis, L_x, L_y, L_z are each 9×9 matrices. The angular momentum operators are explicitly written in Section 4 of the supplementary material. In this description, v can be interpreted as an axis-angle representation of rotation, with corresponding rotation matrix $e^{[v]}$. $[v]$ denotes the skew-symmetric matrix that acts as $[v]u = v \times u$. Accordingly, $e^{v \cdot L} g_0$ encodes the octahedral frame whose directions are $\hat{x}, \hat{y}, \hat{z}$ rotated by $e^{[v]}$, where $^{\text{norm}}$ denotes normalization (see inset).

Using such SH rotations, we can present an alternative interpretation of the octahedral frame f_0 as the sum of three orthogonal SH lobe-shaped functions. The z -aligned lobe is $l = [0, 0, 0, 0, \sqrt{\frac{7}{12}}, 0, 0, 0, 0]$ and is depicted in the inset. Lobes can be rotated in the same way that frames can (i.e., by applying $e^{v \cdot L}$). The canonical octahedral frame can therefore be equivalently expressed as $f_0 = l + e^{\frac{\pi}{2} L_x} l + e^{\frac{\pi}{2} L_y} l$.



The space of octahedral frames that are aligned to a unit vector \hat{n} can be described by the set

$$\left\{ e^{v_n \cdot L} e^{\theta L_z} f_0 \mid \theta \in S^1 \right\},$$

where v_n is any axis-angle rotation taking \hat{z} to \hat{n} (e.g., the vector parallel to $\hat{z} \times \hat{n}$ and has magnitude equal to the angle from \hat{z} to \hat{n}), and θ encodes an additional twist of the frame about \hat{n} . The first rotation about \hat{z} can be written in explicit form [Huang et al. 2011] as

$$e^{\theta L_z} f_0 = \left[\sqrt{\frac{5}{12}} \cos 4\theta, 0, 0, 0, \sqrt{\frac{7}{12}}, 0, 0, 0, \sqrt{\frac{5}{12}} \sin 4\theta \right]^T.$$

The preceding allows us to formulate the set of all octahedral frames g aligned to a given direction \hat{n} in terms of two constraints:

$$\|f\|_2 = 1, \quad W_n f = u_0 = \left[0, 0, 0, \sqrt{\frac{7}{12}}, 0, 0, 0 \right]^T, \quad (1)$$

where W_n is the second through eighth rows of $e^{-v_n \cdot L}$. The linear constraint rotates the frame from normal alignment to \hat{z} alignment, and the norm constraint ensures that the first and last components are of the appropriate form.

Last, we will make use of the projection operator $\pi_{\mathcal{V}} : \mathbb{R}^9 \rightarrow \mathcal{V}$ onto the space of octahedral frames \mathcal{V} , as defined in Section 5.5 of Palmer et al. [2019].

3.2 Vectorial Total Variation

We will later make use of a total variation energy (among others) to analyze the behavior of our cross fields on creased surfaces. Here, we introduce total variation and vectorial total variation definitions in \mathbb{R}^n and provide intuition about their use. The extension to functions on a Riemannian manifold is straightforward, using the standard intrinsic gradient and divergence operators.

The total variation of a differentiable scalar function $h : \Omega \rightarrow \mathbb{R}$ is $TV[h] = \int_{\Omega} \|\nabla h\|_2 dA$ where $\Omega \subset \mathbb{R}^n$ [Ambrosio et al. 2000]. For non-differentiable h , the relevant definition is

$$TV[h] = \sup_{\phi \in C_c^1, \forall x \|\phi(x)\|_2 \leq 1} \left(\int_{\Omega} h \nabla \cdot \phi dA \right),$$

where C_c^1 denotes differentiable, compactly supported vector fields. For smooth h , equivalence to $\int_{\Omega} \|\nabla h\|$ follows from integration by parts and Stokes's theorem. In this case, the maximizing ϕ is $\frac{-\nabla h}{\|\nabla h\|}$. If h is the indicator function of a suitably regular (e.g., non-fractal) subset $A \subset \Omega$, then $TV[h]$ is the perimeter of A .

When $h : \Omega \rightarrow \mathbb{R}^m$ is vector valued rather than scalar valued, there are many different definitions for the vectorial total variation $VTV[h]$ [Sapiro 1996]. We use one proposed by Di Zenzo [1986], which for differentiable h is given by $VTV[h] = \int_{\Omega} \|\nabla h\|_F dA$, where $\|\cdot\|_F$ is the Frobenius norm. More generally, we can take

$$VTV[h] = \sup_{\phi \in C_c^1, \forall x \|\phi(x)\|_F \leq 1} \left(\sum_{i=1}^m \int_{\Omega} h_i \nabla \cdot \phi_i \right), \quad (2)$$

where $h = (h_1, h_2, \dots, h_m)$, and $\phi = (\phi_1, \phi_2, \dots, \phi_m)$ is a differentiable, compactly supported m -tuple of vector fields. This definition is *not* equivalent to a sum of m independent scalar total variations: the constraint on ϕ introduces non-trivial coupling between the dimensions. This definition of total vectorial variation is considered in the case where Ω is a surface in \mathbb{R}^3 by Bresson and Chan [2008] but without specific analysis for discontinuous h .

4 SH OCTAHEDRAL FRAMES AS CROSS FIELDS

We use normal-aligned octahedral fields to encode tangent cross fields on surfaces, with the goal of computing a smooth cross field on a surface aligned to sharp features. The SH representation will enable us to capture features even when they are purely extrinsic. To this end, our next task is to define a means of measuring smoothness by examining the gradient of a SH field along the surface.

4.1 Derivatives of SH Octahedral Frames

To calculate $\|\nabla f\|^2$, we first express it in an appropriate local coordinate system that simplifies the formulas in coordinates and better reveals the structure. Following the notation in Section 3.1, an octahedral field $f(r) : \Omega \rightarrow \mathcal{V} \subset \mathbb{R}^9$ can be parameterized relative to a point r^* by $v(r) : \Omega \rightarrow \mathbb{R}^3$, where $v(r)$ is the axis-angle rotation from $f(r^*)$ to $f(r)$. This implies that $v(r^*) = [0, 0, 0]$. Without loss of generality, we rotate the surface so that the normal of Ω at r^* is \hat{z} . We can then compute the gradient ∇f at the point r^* from the formula $f(r) = e^{v(r) \cdot L} f(r^*)$:

$$\nabla f(r)|_{r^*} = \begin{bmatrix} | & | & | \\ L_x f(r^*) & L_y f(r^*) & L_z f(r^*) \\ | & | & | \end{bmatrix} \begin{bmatrix} | & | & | \\ \nabla_x v & \nabla_y v & \nabla_z v \\ | & | & | \end{bmatrix}_{r^*}. \quad (3)$$

As the field $f(r)$ encodes an extrinsically embedded frame at each point, we take the gradient ∇ to be the component-wise derivative of the field's nine scalar functions rather than a covariant or Lie derivative along the surface to capture the extrinsic geometry of the surface. We use Section 1.2.5 in Rossmann [2002] and the fact that $v(r^*) = [0, 0, 0]$ to derive Equation (3).

By combining facts about the SH representation and standard results in differential geometry, we show that the squared norm $\|\nabla f(r)\|^2$ at r^* can then be expressed in the following more intuitive way.

PROPOSITION 4.1. *Let $f(r) : \Omega \rightarrow \mathcal{V} \subset \mathbb{R}^9$ be a normal-aligned octahedral field over a smooth surface Ω . Then at every point $r^* \in \Omega$, $\|\nabla f\|_F^2 = k_1^2 + k_2^2 + w$, where k_1 and k_2 are the principal curvatures and w measures the intrinsic tangential twist of the octahedral field. Using mean and Gauss curvatures H and K , we can write $\|\nabla f\|_F^2 = 2H^2 - K + w$.*

We leave the full proof of this formula to the supplementary material. Proposition 4.1 gives a more intuitive form for Equation (3) and relates the SH representation of an octahedral frame to properties of the frame it represents. Most notably, the Dirichlet energy of the SH representation can be effectively decoupled into extrinsic dependence of $\|\nabla f\|_F^2$ on the surface Ω and the intrinsic tangential twisting of the normal-aligned octahedral field $f(r)$. The values of H and K simply contribute a fixed quantity depending on Ω rather than the field. Therefore, the influence of f on $\|\nabla f\|_F^2$ is just in w , the intrinsic twist of the cross field it represents. We stress that this behavior is quite different from the behavior of the component-wise derivative evaluated on vectors, as studied in Huang and Ju [2016], where their smoothness energy promotes alignment to extrinsic curvature directions.

4.2 L^p Smoothness Energy of SH Cross Fields

Suppose we wish to measure smoothness of a normal-aligned octahedral field in the SH representation. We define the following class of convex smoothness energies using the L^p -norm of $\|\nabla f\|_F$ over the surface Ω for $p \geq 1$:

$$E_p(\Omega, f) = \left(\int_{\Omega} \|\nabla f\|_F^p dA \right)^{\frac{1}{p}}. \quad (4)$$

We now analyze the behavior of the E_p energy for cross fields in several select cases.

4.2.1 Case $p = 2$: Dirichlet Energy. We begin with a common choice in geometry processing when smoothness is desirable: the Dirichlet energy E_2 . Given Proposition 4.1, we can write the Dirichlet energy as $\int_{\Omega} 2H^2 - K + w$. Since H and K are independent of the octahedral field f , they have no influence over the f that minimizes E_2 . Therefore, on smooth Ω , we recover intrinsically smooth cross fields.

Since the Dirichlet energy may diverge at singularities [Knöppel et al. 2013], this choice of energy has the theoretical drawback of diverging for all f in the neighborhood of creases that break octahedral symmetry. In the discretized setting, however, the behavior of E_2 is dependent on mesh resolution and empirically leads to strong feature alignment as demonstrated in Section 6. It also leads to an easily-solved optimization problem described in Section 5.2.

4.2.2 Case $p = 1$: Vectorial Total Variation. As noted in the previous section, the conventional means of measuring field smoothness fails to be well defined for our field representation on creased surfaces. We show here that the E_1 energy is not only finite across sharp edges and around singular points but also provides an intuitive measure of field quality that captures both smoothness and feature alignment. It is also known as the *vectorial total variation*.

Consider a function $f : \Omega \rightarrow \mathbb{R}^9$ that is piecewise smooth on n closed patches Ω_j intersecting in a finite-length curve network $\Gamma = \bigcup_{k=1}^s \gamma_k$, where each γ_k is a C^1 curve. Equivalently, $\Gamma = \bigcup_{j=1}^n \partial\Omega_j$, and the vectorial total variation can be decomposed into integrals over each patch and Γ .

PROPOSITION 4.2. *For compact Ω and f as earlier, $VTV[f]$ is finite and given by the following equation:*

$$VTV[f] = \sum_{j=1}^n \int_{\Omega_j} \|\nabla f\|_F dA + \sum_{k=1}^s \int_{\gamma_k} \|f^+ - f^-\|_2 dL, \quad (5)$$

where f^+ and f^- refer to the limiting values of f on either side of γ_k , and Ω_j denotes the interior of Ω_j .

The basic argument starts from Equation (2), splits it into integrals over the patches, applies integration by parts, and utilizes partitions of unity to construct a maximizing sequence of ϕ 's. The full argument is contained in supplementary materials Section 2. An analogous result, which applies to arbitrary functions on \mathbb{R}^n with bounded variation, is contained in Ambrosio et al. [2000], with the addition of a third term representing the Cantor part of f . Since our f is piecewise smooth, however, we can safely ignore the Cantor part. The second term is often referred to as the *jump* part in the total variation literature.

The formula (5) provides an intuitive description of the total variation of an octahedral field in the SH basis as a measure of intrinsic smoothness with extra jump terms. Letting f represent a normal-aligned octahedral field, we obtain

$$VTV[f] = \sum_{j=1}^n \int_{\Omega_j} \sqrt{2H^2 - K + w} dA + \sum_{k=1}^s \int_{\gamma_k} \|f^+ - f^-\|_2 dL. \quad (6)$$

Generalizing to creased surfaces. Although the preceding result is derived for smooth surfaces Ω and discontinuous f , we can further generalize the result to a surface Ω constructed from smooth open patches Ω_j joined along a network of sharp creases $\Gamma = \bigcup_{k=1}^s \gamma_k$. As there is neither a consistent metric nor a consistent tangent space on Ω across Γ , there is no well-defined choice of gradient. We therefore use Equation (6) as the *definition* for E_1 on such a creased surface. Since f is a normal-aligned octahedral field, it is necessarily discontinuous across creases, resulting in contributions to the jump term.


The jump $\|f^+ - f^-\|_2$, where f^+ and f^- represent octahedral frames aligned to different normal directions, is minimized if f^+ and f^- are both aligned to the axis of rotation from one normal to the other. We formalize this property by Proposition 4.3.

PROPOSITION 4.3. *Let Ω^+ and Ω^- be smooth patches of a surface with normal directions \hat{n}^+ and \hat{n}^- that meet at a crease. Let \hat{d} denote the intersection of their tangent spaces at the crease. Let f_θ^+ and f_ϕ^- be the octahedral frames on either side of the crease aligned to \hat{n}^+*

and \hat{n}^- , respectively. θ and ϕ denote their deviation from alignment to \hat{d} . The cost $\|f_\theta^+ - f_\phi^-\|_2$ is minimized by $\theta = \phi = 0$.

Proof of this proposition is left to Section 3 of the supplementary material.

The setup is depicted on the right, showing discontinuous normal directions \hat{n}^+ and \hat{n}^- as the left and right red arrows, respectively. The crease direction \hat{d} is shown by the middle red arrow.



Octahedral frames near crease

We emphasize that this proposition implies (locally) crease alignment always minimizes the VTV. We extensively test and show in supplemental material that this crease alignment tends to globally hold on surfaces with complicated geometry and topology as well.

4.2.3 General $p \geq 2$. By Equation (4) and Proposition 4.1, E_p incentivizes intrinsic smoothness for all p on smooth domains. On creased domains, we have demonstrated (local) crease alignment for the $p = 1$ case. For $p \geq 2$, the value of E_p diverges for a creased surface. However, we find empirically that minimizing E_p (by recovering solutions to Equation (7)) on a discretized surface leads to stronger feature alignment as p increases. This behavior may be explained by Proposition 4.3, which affects all edges regardless of p . The p simply exponentiates the energy across each edge before accumulating it into the total E_p . Local to a single edge, the energy-minimizing configuration is unaffected by p . Based on our experiments, we further conjecture that the sequence of fields obtained by minimizing E_2 on an increasingly dense discrete approximation of Ω converges to a feature-aligned cross field. This intrinsically smooth feature alignment is empirically shown in Figure 6. We leave proof of this conjecture to future work.

4.2.4 Relation to Polycube Surfaces. We achieve an additional property for all values of p through our use of SH octahedral frames. Consider the case of Ω being a cube: minimizers of E_p will have zero energy, despite the cube's sharp corners, since the field's octahedral symmetry allows it to simultaneously align to all three creases at each corner. Effectively, a surface with many angle- $\frac{\pi}{2}$ turns and cube corners can have just as low of an energy as one with no creases at all. More generally, if Ω is a polycube surface, $E_p(\Omega, f) = 0$ by choosing f to be a facet-aligned uniform frame field.

5 OPTIMIZING FOR AN OCTAHEDRAL FRAME FIELD

Our preceding discussion provides a new class of energies based on the SH representation of cross fields, which naturally promote both intrinsic smoothness and extrinsic crease alignment without the need for feature curve detection or reliance on potentially noisy local curvature estimates. For this reason, we propose solving the following variational problem to find a cross-field f^* over a surface:

$$\begin{aligned} f^* &= \arg \min_f E_p(\Omega, f) \\ \text{subject to } W_{n(x)} f(x) &= u_0. \end{aligned} \quad (7)$$

Recall that the constraint encodes normal alignment of the frame field (see Equation (1)). Some past algorithms have an extra

$\|f(x)\|_2 = 1$ constraint that results in uniform-scale isotropic octahedral fields over Ω . This constraint makes the problem non-convex and causes the functional to diverge in the neighborhood of field singularities, which are unavoidable on generic surfaces by the Poincaré–Hopf theorem. Accordingly, a relaxation is naturally required; we drop the constraint $\|f(x)\|_2 = 1$, yielding a convex problem with globally optimal solution. Dropping the $\|f(x)\|_2 = 1$ constraint allows the frame’s two tangential components to scale independently from its normal component, resulting in anisotropic octahedral fields. We obtain octahedral fields with uniform-magnitude normal lobes and varying scale in the magnitude of the tangential cross field. This relaxation is similar in spirit to the one that appears in Knöppel et al. [2013] and has similar benefits, including automatic placement of singularities, and bounded-energy minimizers in the smooth limit (which is necessary for the field to be insensitive to the underlying mesh).

5.1 Soft Normal Alignment

It is sometimes beneficial to relax the normal alignment constraint, such as in cases where the mesh contains sliver triangles with unstable normal directions. In these cases, a smoother cross field can be obtained by deviating slightly from exact normal alignment. This relaxation changes the optimization problem from Equation (7) into the following:


$$\begin{aligned} f^* = \operatorname{argmin}_f \quad & E_p(\Omega, f) \\ \text{subject to} \quad & \|W_{n(x)}f(x) - u_0\|_2 \leq \epsilon. \end{aligned} \quad (8)$$

This problem imposes a pointwise normal alignment constraint with tolerance ϵ . When $\epsilon = 0$, we recover the hard normal alignment formulation (7). On the opposite side of the spectrum, as $\epsilon \rightarrow \|u_0\|_2 = \sqrt{\frac{7}{12}} \approx 0.76$, the solution to (8) approaches a constant octahedral field. This is the case where normal alignment has relaxed so far that the octahedral frames are effectively unconstrained.

For values in between, we perform the following experiment to obtain a rough correspondence between soft normal alignment parameter ϵ and maximum angle deviation from normal alignment: for each value of ϵ between 0 and .7 (at intervals of .05), we sample 100,000 ϵ -perturbations of a \hat{z} -aligned frame, extract the frame they represent, and compute its maximum angle deviation from the \hat{z} -axis. Results are shown in the inset.

We highlight that this parameter encodes a pointwise constraint uniformly applied over the mesh. As such, its interpretation does not change with different meshes. Please see the supplemental material for results on more than 200 different meshes using a variety of values of ϵ .

The benefit of soft normal alignment is demonstrated in Figure 5. Due to the influence of a sliver triangle in the buste mesh



with unstable normal direction, the hard-normal-aligned cross field is forced to create a localized artifact. By using soft normal alignment, the sliver triangle’s unstable normal direction has less influence over the resulting cross field, therefore increasing the quality of the result. A similar benefit is demonstrated on the duck and armchair meshes shown in the supplementary material.

Additionally, we test soft normal alignment on a cube mesh with artificial noise added in Figure 5. With hard normal alignment the cross fields exhibit undesirable alignment to noise that increases with p . With soft normal alignment, the cross fields show significantly decreased sensitivity to noise.

5.2 Discretization

Now we describe how to construct smooth cross fields by numerically optimizing a discretization of E_p . We assume the surface Ω has been triangulated into a manifold mesh $\mathcal{M} = (V, E, F)$. Let n_t be the normal direction of triangle $t \in F$. We represent a cross on M as a normally aligned octahedral frame $f_t \in \mathcal{V} \subset \mathbb{R}^9$ per triangle. We use the shorthand f to denote the concatenation of all f_t into a single $9|F| \times 1$ vector. V is a $n_v \times 3$ matrix of vertex positions, where n_v is the number of vertices. E denotes the $n_e \times 2$ matrix of edges, where n_e is the number of edges. The energy E_p can be discretized as

$$E_p = \left(\sum_{e \in E} w_e \|f_{t_1} - f_{t_2}\|_2^p \right)^{\frac{1}{p}}, \quad (9)$$

where t_1 and t_2 are triangles adjacent to edge e , and w_e are weights corresponding to the dual Laplacian. We use $w_e = \frac{\|e\|}{\|e^*\|}$, where $\|e\|$ is the length of edge e and $\|e^*\|$ is the distance between barycenters of t_1 and t_2 .

For $\epsilon = 0$, the normal alignment constraint is discretized by the linear constraint $Wf = u$, where W is a sparse block-diagonal matrix with a block W_{n_t} for each triangle. It has dimensions $7|F| \times 9|F|$. The vector u is a repetition of u_0 for each triangle, resulting in a $7|F| \times 1$ vector. For $\epsilon > 0$, the normal alignment constraint is discretized by a second-order cone constraint: $\|W_{n_t}f_t - u_0\|_2 \leq \epsilon$ per triangle.

For the case $p = 1$ and a completely flat surface Ω , our discretization agrees with the standard discretization of total variation in image processing [Chambolle et al. 2010; Rudin et al. 1992].

5.3 Meshes with Boundary

When the mesh contains boundaries, we do not enforce any boundary conditions; in PDE parlance, this choice corresponds to “natural boundary conditions.” Although it is tempting to expect the natural boundary conditions for $p = 2$ to imply zero Neumann boundary conditions [Stein et al. 2018], the SH representation vector is complicated by being constrained to a spatially varying linear subspace. We simply allow the cross on the boundary to be that which minimizes total energy. If desired, one can enforce a constraint that the cross field on the boundary be aligned to the boundary through the method described in Section 5.4.

5.4 Manual Guidance

To support manual guidance of the octahedral frame field, we can prescribe alignment of the frame field to streamlines. Streamline

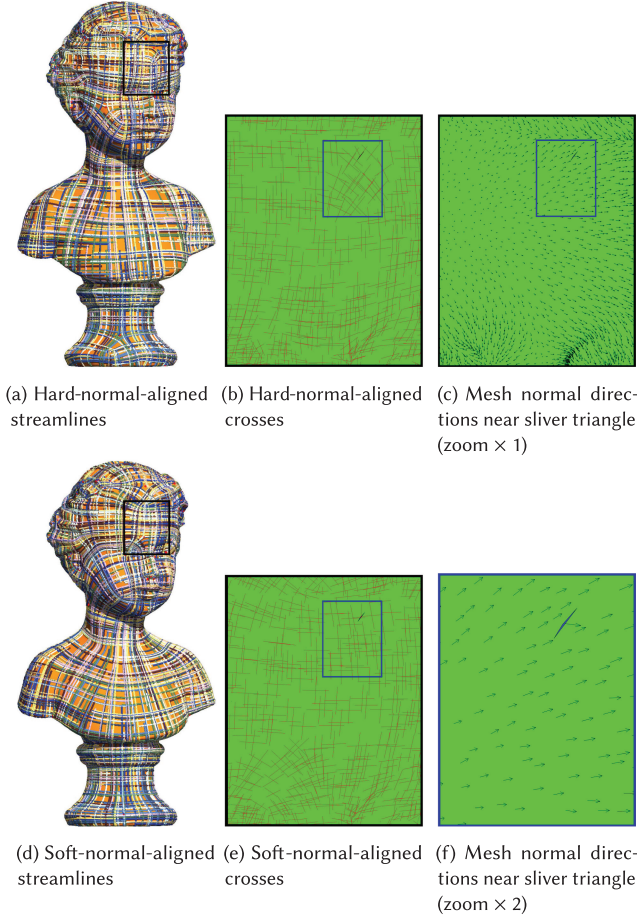


Fig. 3. Soft normal alignment increases quality of the cross field and decreases the influence of mesh artifacts. The buste mesh is shown with $p = \infty$ and varying normal alignment: $\epsilon = 0$ for the top figure and $\epsilon = 0.5$ for the bottom. (a) Hard-normal-aligned streamlines. (b) Magnified crosses show a small patch of diagonal crosses in an otherwise regular region. (c) Magnified triangle normals visualized with sliver triangle 4611 shaded in blue. (d) Soft normal-aligned streamlines. (e) Magnified crosses no longer show diagonal artifacts. (f) Extra-magnified triangle normals visualized with sliver triangle 4611 shaded in blue. While the normal direction of the region points diagonally up and right, the sliver triangle's normal direction points almost completely to the right.

constraints combined with normal alignment result in a fully determined frame. Therefore, prescribing streamlines is equivalent to prescribing the value of f_t on a subset of triangles T_p . Denote the prescribed octahedral frame on triangle t as F_t . We then add a new linear constraint that

$$\forall t \in T_p, f_t = F_t. \quad (10)$$

This technique is demonstrated in Figure 4.

5.5 Non-Triviality Constraint

As a result of dropping the unit-norm constraint from Equation (7), we have no explicit guarantee that the tangential components of octahedral frames do not degenerate to zero. On a surface with a crease, however, the normal alignment constraint on one side of




Fig. 4. Octahedral fields obtained by minimizing E_∞ on the hand mesh before and after adding manual direction. The manually added streamline is shown by the inset black arrow. This constraint removes a singularity from the original octahedral field.

the crease imposes that the magnitude of the tangential component on the other side of the crease is close to one. As a result, we observe empirically that most of our octahedral frames do not degenerate.

In the case that octahedral frames do degenerate significantly, their norms can be too small to project robustly. We locate these by using the octahedral projection from Palmer et al. [2019] to measure the distance from f_t to the octahedral variety $\mathcal{V} : d(f_t) = \|\pi_{\mathcal{V}}(f_t) - f_t\|_2$, and thresholding by $d(f_t) > .665$. If such frames are found, we run the optimization again while holding non-degenerate frames to their projected values. In our experiments, just one round of re-solving results in 99.8% non-degenerate frames.

5.6 Solving for an Octahedral Field

In its most general form, our problem formulation consists of minimizing a mixed-norm objective, with both linear and second-order cone constraints. This results in a convex problem that we solve with Mosek 9 [MOSEK ApS 2017]. The normal alignment constraint becomes $[\epsilon, (W_{n_t} f_t - u_0)^T] \in \mathcal{L}^8$, where \mathcal{L}^8 is the eight-dimensional Lorentz cone. Likewise, the energy is formulated using a single p -norm cone. Our code is written in Matlab with a mex interface to Mosek; it builds cross platform. Since our problem is convex, any dependence on initialization would entirely be due to non-unique solutions, which we do not observe in practice. Furthermore, we use the interior point method, which does not accept manual initialization. In the specific case of $\epsilon = 0, p = 2$, solving this optimization is equivalent to solving a linear system.

6 RESULTS

We begin with a comparison of the behavior of our energy for different values of p . This experiment is depicted in Figure 7. We observe that our cross fields naturally align to features with increasing strength for higher p . In the case $p = 1$, our cross field is discontinuous over all creases, but although it is provably incentivized to align, it sometimes deviates due to the influence of neighboring creases (e.g., on the top surface of the fandisk). For $p = 2$, our cross fields achieve close alignment to the upper half of the shallow crease, as well as alignment on the top face where




Fig. 5. As ϵ increases or as p decreases, the cross fields become less sensitive to noise added to the cube mesh.

the $p = 1$ case failed. Finally, for $p = \infty$, our fields align down the entirety of the shallow crease. Although in theory E_p for $p \geq 2$ diverges on creases, we observe that its discretization yields empirically strong crease alignment. This may be due to application of Proposition 4.3 over all edges of the mesh. We show our fields for different discretizations of the same geometry in Figure 8 and observe that in all cases we achieve crease alignment.

Supplementary material. In our supplementary document, we perform an empirical study to evaluate the performance of our method. We evaluate our method on several models drawn from the Thingi10k [Zhou and Jacobson 2016] dataset, as well as several other commonly used benchmark models to demonstrate effective crease alignment on real-world models. We also compare our approach to several baseline methods [Brandt et al. 2018; Jakob et al. 2015; Knöppel et al. 2013] by generating fields on the models in the “Robust Field-Aligned Global Parametrization” dataset [Myles et al. 2014], taking care to sample the relevant parameter space for each formulation. Although it is difficult to precisely quantify the quality of a vector field, we highlight several cases where our method recovers fields that more faithfully conform to mesh features than baseline methods on real-world models.

Our runtimes are shown for a set of meshes with 240 to 76K vertices and 480 to 152K faces in Figure 11. Runtimes naturally increase with mesh size and appear to grow linearly with number of

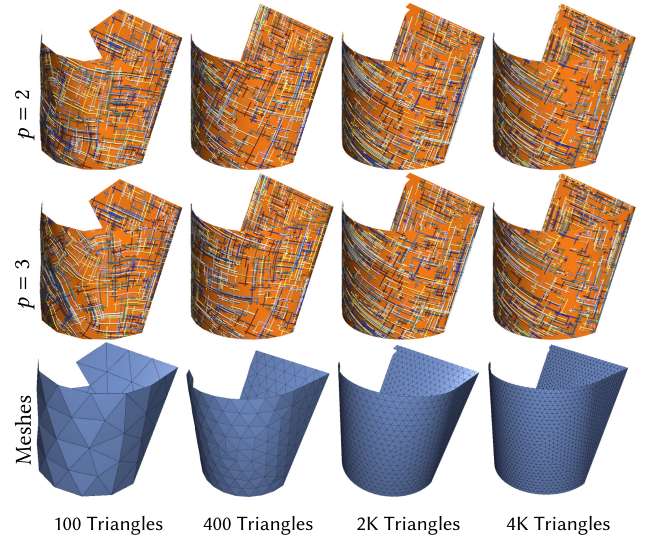


Fig. 6. On this developable surface, our cross fields are intrinsically smooth in the limit of refinement but exhibit some mesh sensitivity on coarse meshes, particularly for higher p values. They are crease aligned for all resolutions. Note that the extrinsic curvature of the cylindrical bend has no effect on the cross fields at higher resolutions.

triangles in our mesh test set. Memory costs are incurred to store a W_{n_t} per triangle, a single u_0 , w_e per edge, and f_t per triangle. Hence, storage is linear in size of the mesh. More detailed information regarding parameter choices and runtimes is provided in the supplementary material. Table 1 shows a summary of our runtimes in comparison to that of other methods. Our runtimes are on the same scale as Knöppel et al. [2013] and to the bases setup step in Brandt et al. [2018].

Comparison to explicit feature curves. Next, we compare our feature-aligned cross fields to those produced with the help of explicitly computed feature curves. We obtain feature curves on the 1904-triangle Moai mesh from Figure 6 in Gehre et al. [2016]. We compute a cross field with additional hard constraints as described in Section 5.4 to enforce alignment to the pre-computed feature curves. We compare the resulting field with and without explicit feature curve alignment in Figure 4. Although the feature curves help guide the cross field, just a few artifacts in the computed features drastically influence the resulting cross field to have more singularities and be less smooth without clear benefit. The Moai is shown from an angle where these differences are most pronounced.

Effect of mesh resolution on crease alignment vs extrinsic curvature. In this experiment, we test on a geometry where a sharp crease is mis-aligned to extrinsic curvature directions. We generate meshes of this geometry at varying resolution to see how crease alignment interacts with extrinsic curvature. Results of this experiment are depicted in Figure 6. As mesh resolution increases, our cross fields become crease aligned and intrinsically smooth, agreeing with the theory. For very low mesh resolution, the cross fields are more sensitive to the underlying meshing pattern.




Fig. 7. Cross fields generated by minimizing E_p for $p = 1, 2, \infty$ on the fandisk mesh. The shallow crease of the fandisk mesh is marked in red. Our cross fields naturally align to the shallow crease with increasing strength for higher p .

Table 1. Runtimes in Seconds for Computing Cross-fields Using Different Methods on Meshes with a Varying Number of Triangles

Number of triangles	Bases setup	Biharmonic solve	Instant meshes	Globally optimal	Ours
3K	5	.005	.026	.85	2.8
12K	24	.005	.053	20.46	15.058
20K	44	.005	.080	21.913	25.895
69K	170	.006	.141	62.733	135.09
80K	181	.006	.222	71.15	112.3

Methods listed are those of Brandt et al. [2018], Jakob et al. [2015], Knöppel et al. [2013], and our own. Runtimes for fields from [Brandt et al. 2018] are split into time needed for the setup of 500 bases eigenfields and the field computation separately because of drastically differing timescales.

Comparison to three-dimensional octahedral fields. Due to similarity of frame representation, we compare our method to surface cross fields obtained by optimizing a volumetric octahedral field. Algorithms like those of Huang et al. [2011] and Ray et al. [2016] can generate surface cross fields by approximating the surface with the limiting behavior of a thin layer of tetrahedra or prism elements. However, prism elements are non-standard and both element types will be poorly conditioned without introducing further restrictions such as zero normal gradient to mimic a triangle mesh. We instead opt to compare with the BEM [Solomon et al. 2017], which acts directly on surface triangle meshes. We use the 2,500-triangle fandisk mesh for this comparison. As observed earlier, our method has increasing feature alignment with increased values of p . In comparison, Figure 10 shows that the BEM field fully ignores the shallow crease of the fandisk, running through it at a 45° offset. Moreover, despite the fact that the BEM only needs boundary data as input, its runtime is close to 50 times slower than ours.

Challenging test cases. We compare feature alignment of our cross fields with that of existing methods on several meshes illustrative examples in Figure 12. As pointed out in Section 1, a key advantage of our technique is that it recovers crease-aligned fields on models whose maximal curvature directions *disagree* with their creases. This occurs naturally when models are specified by the intersections of developable patches—a very common primitive in CAD tools. We introduce two benchmark models for testing crease




Fig. 8. Cross fields generated by minimizing E_2 on different meshings of the three-cylinder-intersection with number of faces indicated. Cross field (d) is computed on the multi-resolution mesh (e). Notice that we obtain the same feature-aligned cross field each time.

alignment when creases disagree with intrinsic notions of curvature. The three-cylinder-intersection mesh is composed of 12 quadrilateral patches, where each patch is a subset of a cylinder and has maximal curvature directions making $\frac{\pi}{4}$ angles with its boundary creases. The wavey-box example has the same creases as a standard cube, with the modification that each of its faces has a sine wave ripple running diagonally through it. These two cases are shown in Figure 2. The fandisk mesh is another example of a challenging case for feature alignment due to its shallow crease with strong non-aligning neighboring creases, which is representative of one way that such features arise in real-world models.

Our cross fields on these test cases are shown in Figure 12. We observe proper feature alignment in our fields, and although other methods can sometimes be tuned per model to achieve the same feature alignment, there is no choice of parameters that worked on all test cases. In particular:

- Fields from Jakob et al. [2015] are distracted by extrinsic curvature on the three-cylinder-intersection and entirely pave over the shallow crease of the fandisk. Their results on wavey-box and wedge are successfully aligned to the creases.
- Fields from Brandt et al. [2018] are challenging to tune with λ representing alignment to a guiding extrinsic curvature field. We show their method for the biharmonic energy ($m = 2$) as a point of contrast to Dirichlet energy. We choose two values of λ , $\lambda = -.0001$ for slight extrinsic curvature alignment and $\lambda = -.1$ for stronger extrinsic curvature alignment. Their fields are unable to align to features of the three-cylinder-intersection in both cases, and specifically for




Fig. 9. Comparison of our feature-aligned cross fields to those generated when adding additional explicit feature curve alignment constraints. Explicit feature curves were obtained from Figure 7 in Gehre et al. [2016]. Despite the extra cost of pre-computing explicit feature curves, slight artifacts in the feature curves (most pronounced on the side) force the explicitly guided cross field to have lower quality.




Fig. 10. Cross field and runtime comparison of our method to a method optimizing volumetric octahedral frames [Solomon et al. 2017]. The fandisk used contains 2.5K triangles.

$\lambda = -0.1$ the field strongly aligns to noise on the flat upper face of the fandisk mesh. Their fields are successfully crease aligned for the wedge mesh.

- We compare against both the anti-holomorphic and Dirichlet energies of Knöppel et al. [2013] with the curvature alignment parameter λ set to -0.1 . This results in good alignment on the three-cylinder-intersection, but noisy or unaligned fields for the remaining test cases.

In contrast, our method for $p = 2$ achieves feature alignment on all test cases without unnecessary discontinuities in the field over flat faces. We show additional results for more than 200 different meshes with both smooth and creased geometries with varying values of p and ϵ in the supplementary material. The fields are crease aligned for all creased meshes and are otherwise intrinsically smooth. For comparison, we include fields from Brandt et al. [2018] and Knöppel et al. [2013] on a larger range of λ . We also include fields from Brandt et al. [2018] for $m = 1$ and fields from Jakob et al. [2015] in the supplementary material.

Quad meshing. Feature alignment is especially important when using cross fields to guide high-fidelity quad meshing. We generate quad meshes using Campen et al. [2015] to parameterize our cross fields. We compare against a standard quad meshing pipeline using cross fields from Bommers et al. [2009] and Campen et al. [2015] for parameterization. We also test against parameterization by Campen et al. [2016], which introduces extra guidance to encourage extrinsic curvature alignment.

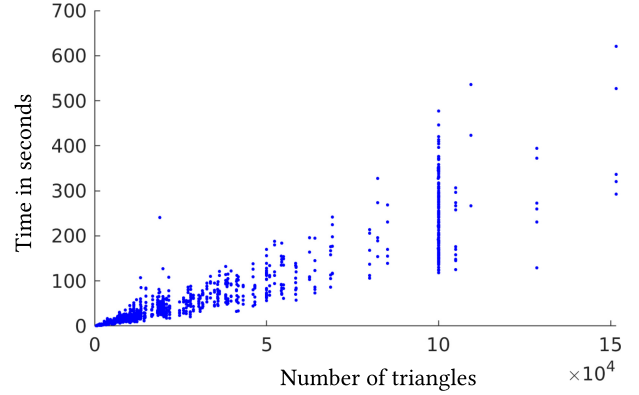


Fig. 11. Runtimes to compute cross fields over various mesh sizes.

For the fandisk mesh, prior methods generate quad meshes that are influenced by the shallow crease but do not manage to capture it sharply (see Figure 13). We observe that by placing singularities near the shallow crease of the fandisk, our quad meshes manage to align much more sharply. The quad mesh generated by minimizing E_∞ aligns even better than for E_2 .

We also compare quad meshes generated from our cross fields against the prior art on the anchor, spot, moomoo, and three-cylinder-intersection meshes. These results are shown in Figure 14. We observe generally better alignment in the quad meshes generated from our method. By placing singularities on the cylindrical region of the anchor, our quad meshing manages to align better to its creases. On the spot mesh, we see a straighter connection between the ear and the head. For the three-cylinder-intersection, the quad mesh generated from our fields clearly aligns better. Since the moomoo is a relatively smooth mesh, we do not see particularly defining differences in quality.

7 DISCUSSION AND CONCLUSION

Feature alignment is a desirable property in many geometry processing applications. In the context of cross fields and remeshing, we consider features to be creases where the surface changes non-smoothly. Quality of feature detection and alignment can significantly impact quality of the remeshing and the usefulness of the resulting cross fields. Although significant effort has been put into extrinsic alignment of cross fields to curvature directions, they are not always appropriate substitutes for crease alignment. By specifically targeting discontinuities of the surface, we have created a new class of octahedral frame field energies parameterized by $p \geq 1$ for computing crease-aligned cross fields on surfaces. The resulting fields are intrinsically smooth over smooth surfaces and can be used for crease-aligned quad meshing. Moreover, alignment is fully automatic and does not rely on explicit extraction of feature curves, itself an open problem and active area of research.

We find the behavior of E_p for $p \geq 2$ over creases of a discretization to be an interesting point for further exploration since all practical computations on surfaces are necessarily discrete and we observe strong feature alignment, despite the problem being ill posed in the smooth setting. Theoretical analysis of anisotropic normally aligned octahedral frame fields combined with




Fig. 12. Various cross field methods compared on several meshes with complex features and geometry. We test on the three-cylinder-intersection, wavey-box, wedge, and fandisk meshes and compare against the following works with various parameters: Brandt et al. [2018], Jakob et al. [2015], and Knöppel et al. [2013]. We use normal-aligned octahedral fields generated by minimizing E_2 . We achieve crease alignment on all test cases where other methods succeed sporadically.

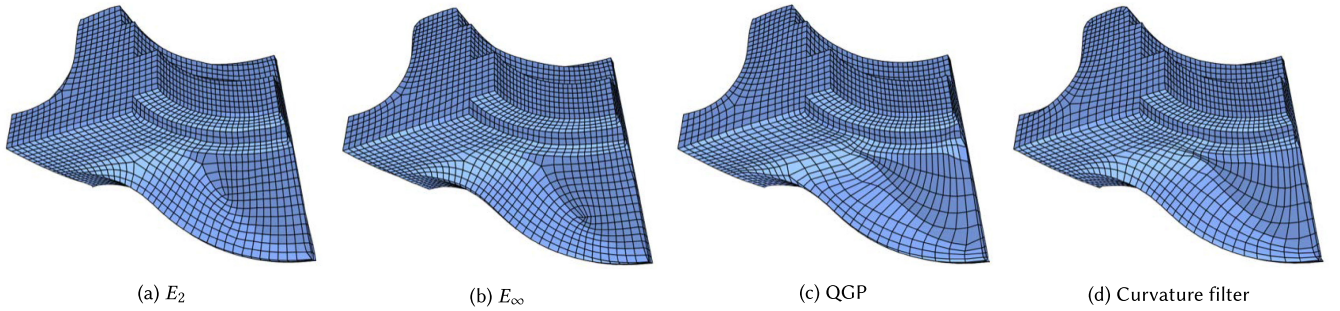


Fig. 13. Quad meshes of the fandisk mesh generated using cross fields from E_2 , E_∞ , MIQ + QGP [Campen et al. 2015], and MIQ + Curvature filter [Campen et al. 2016], respectively. Our methods achieve sharp alignment to the shallow crease with increased depth for higher p . Alternative methods are influenced by the crease only to a shallower extent.

Proposition 4.3 may be able to explain this behavior. Since all edges of a mesh are creases of a piecewise linear domain, the behavior of geometry processing algorithms on creased domains merits further study.

There are also further applications of soft-normal-aligned octahedral frame fields. Although in this article we fix ϵ as a single

parameter per mesh, it could also be defined as a scalar field representing “trust” in the quality of a mesh. It would be interesting to explore a spatially varying ϵ dependent on triangle quality or other metrics in the future. If we treat the mesh itself as variables, soft normal alignment enables a surface flow toward meshes with lower cross-field energy. Our analysis can be further extended to

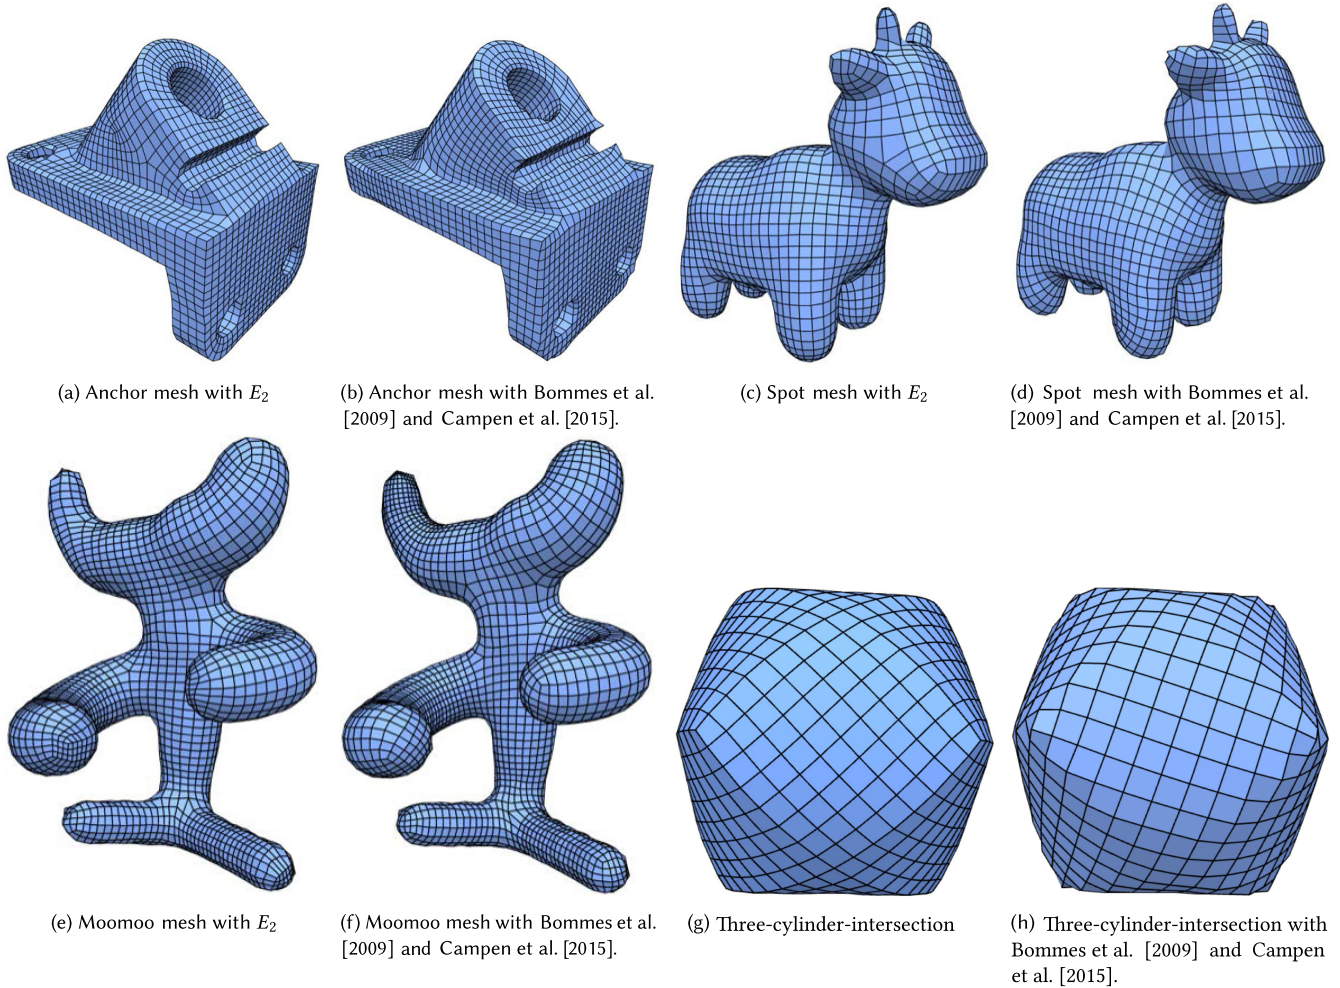


Fig. 14. Quad meshes of the anchor, spot, moomoo, and three-cylinder-intersection meshes. We compare quad meshes generated using cross fields from our E_2 energy with quad meshes generated through Campen et al. [2015] and Bommes et al. [2009]. Our methods achieve sharper feature alignment on the anchor, spot (on the ear), and three-cylinder-intersection meshes.

SH representations of n -RoSy fields or even platonic solid symmetries [Shen et al. 2016]. We also conjecture that with mild assumptions, the solution to our problem is unique, but a proof is outside the scope of this work; we leave exploration of these ideas to future work.

Even without these extensions, our method provides a practical solution to a challenging problem. By using a new representation of cross fields, we achieve crease-aligned cross fields on surfaces.

ACKNOWLEDGMENTS

The authors would like to thank Christopher Brandt for help obtaining comparison fields, Amir Vaxman for support with Directional and rendering, Michal Adamaszek for help debugging Mosek, and David Palmer for discussions and help with his code base. We thank Wenzel Jacob, Keenan Crane, and Qingnan Zhou for their open source code implementations. Finally, we thank Ryan Viertel for many valuable discussions and Panini Pals for the panini press.

REFERENCES

- Luigi Ambrosio, Nicola Fusco, and Diego Pallara. 2000. *Functions of Bounded Variation and Free Discontinuity Problems*. Clarendon Press, Oxford, UK.
- Mikhail Bessmeltsev and Justin Solomon. 2019. Vectorization of line drawings via polyvector fields. *ACM Transactions on Graphics* 38, 1 (2019), Article 9.
- David Bommes, Marcel Campen, Hans-Christian Ebke, Pierre Alliez, and Leif Kobbelt. 2013. Integer-grid maps for reliable quad meshing. *ACM Transactions on Graphics* 32, 4 (July 2013), Article 98. <https://hal.inria.fr/hal-00862648>.
- David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-integer quadrangulation. *ACM Transactions on Graphics* 28, 3 (July 2009), Article 77. DOI:<https://doi.org/10.1145/1531326.1531383>
- Christopher Brandt, Leonardo Scandolo, Elmar Eisemann, and Klaus Hildebrandt. 2018. Modeling n -symmetry vector fields using higher-order energies. *ACM Transactions on Graphics* 37, 2 (March 2018), Article 18, 18 pages. DOI:<https://doi.org/10.1145/3177750>
- Xavier Bresson and Tony Chan. 2008. Fast dual minimization of the vectorial total variation norm and applications to color image processing. *Inverse Problems and Imaging* 2, 4 (Nov. 2008), 455–484. DOI:<https://doi.org/10.3934/ipi.2008.2.455>
- Marcel Campen, David Bommes, and Leif Kobbelt. 2015. Quantized global parametrization. *ACM Transactions on Graphics* 34, 6 (Oct. 2015), Article 192, 12 pages. DOI:<https://doi.org/10.1145/2816795.2818140>
- Marcel Campen, Moritz Ibing, Hans-Christian Ebke, Denis Zorin, and Leif Kobbelt. 2016. Scale-invariant directional alignment of surface parametrizations. *Computer Graphics Forum* 35, 5 (Aug. 2016), 1–10. DOI:<https://doi.org/10.1111/cgf.12958>

- Antonin Chambolle, Vicent Caselles, Matteo Novaga, Daniel Cremers, and Thomas Pock. 2010. An introduction to total variation for image analysis. *Radon Series on Computational and Applied Mathematics* 9 (2010), 1–76. DOI: <https://doi.org/10.1515/9783110226157.263>
- Alexandre Chemin, François Henrotte, Jean-François Remacle, and Jean Van Schaftingen. 2018. Representing three-dimensional cross fields using 4th order tensors. In *Proceedings of the 2018 International Meshing Roundtables*. <http://arxiv.org/abs/1808.03999>.
- Etienne Corman and Keenan Crane. 2019. Symmetric moving frames. *ACM Transactions on Graphics* 38, 4 (July 2019), Article 87, 16 pages. DOI: <https://doi.org/10.1145/3306346.3323029>
- Keenan Crane, Mathieu Desbrun, and Peter Schröder. 2010. Trivial connections on discrete surfaces. *Computer Graphics Forum* 29, 5 (Sept. 2010), 1525–1533. DOI: <https://doi.org/10.1111/j.1467-8659.2010.01761.x>
- Fernando de Goes, Mathieu Desbrun, and Yiyi Tong. 2015. Vector field processing on triangle meshes. In *Proceedings of SIGGRAPH Asia 2015 Courses (SA '15)*. ACM, New York, NY, Article 17, 48 pages. DOI: <https://doi.org/10.1145/2818143.2818167>
- Silvano Di Zenzo. 1986. A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing* 33, 1 (Jan. 1986), 116–125. DOI: [https://doi.org/10.1016/0734-189X\(86\)90223-9](https://doi.org/10.1016/0734-189X(86)90223-9)
- Anne Gehre, Isaak Lim, and Leif Kobbelt. 2016. Adapting feature curve networks to a prescribed scale. *Computer Graphics Forum* 35, 2 (May 2016), 319–330. DOI: <https://doi.org/10.1111/cgf.12834>
- Jin Huang, Yiyi Tong, Hongyu Wei, and Hujun Bao. 2011. Boundary aligned smooth 3D cross-frame field. In *Proceedings of the 2011 SIGGRAPH Asia Conference (SA '11)*. ACM, New York, NY, Article 143, 8 pages. DOI: <https://doi.org/10.1145/2024156.2024177>
- Zhiyang Huang and Tao Ju. 2016. Extrinsically smooth direction fields. *Computers & Graphics* 58 (Aug. 2016), 109–117. DOI: <https://doi.org/10.1016/j.cag.2016.05.015>
- Emmanuel Iarussi, David Bommes, and Adrien Bousseau. 2015. BendFields: Regularized curvature fields from rough concept sketches. *ACM Transactions on Graphics* 34, 3 (May 2015), Article 24, 16 pages. DOI: <https://doi.org/10.1145/2710026>
- Wenzel Jakob, Marco Tarini, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Instant field-aligned meshes. *ACM Transactions on Graphics* 34, 6 (Oct. 2015), Article 189, 15 pages. DOI: <https://doi.org/10.1145/2816795.2818078>
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally optimal direction fields. *ACM Transactions on Graphics* 32, 4 (July 2013), Article 59. DOI: <https://doi.org/10.1145/2461912.2462005>
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe patterns on surfaces. *ACM Transactions on Graphics* 34, 4 (July 2015), Article 39, 11 pages. DOI: <https://doi.org/10.1145/2767000>
- Heng Liu, Paul Zhang, Edward Chien, Justin Solomon, and David Bommes. 2018. Singularity-constrained octahedral fields for hexahedral meshing. *ACM Transactions on Graphics* 37, 4 (July 2018), Article 93, 17 pages. DOI: <https://doi.org/10.1145/3197517.3201344>
- MOSEK ApS. 2017. *MOSEK Optimization Toolbox for MATLAB 8.1.0.82*. Retrieved March 14, 2020 from <http://docs.mosek.com/8.1/toolbox/index.html>.
- Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust field-aligned global parametrization. *ACM Transactions on Graphics* 33, 4 (July 2014), Article 135, 14 pages. DOI: <https://doi.org/10.1145/2601097.2601154>
- M. Nieser, U. Reitebuch, and K. Polthier. 2011. CubeCover—Parameterization of 3D volumes. *Computer Graphics Forum* 30, 5 (Aug. 2011), 1397–1406. DOI: <https://doi.org/10.1111/j.1467-8659.2011.02014.x>
- David Palmer, David Bommes, and Justin Solomon. 2019. Algebraic representations for volumetric frame fields. *arXiv:1908.05411*.
- Nicolas Ray, Dmitry Sokolov, and Bruno Lévy. 2016. Practical 3D frame field generation. *ACM Transactions on Graphics* 35, 6 (Nov. 2016), Article 233, 9 pages. DOI: <https://doi.org/10.1145/2980179.2982408>
- Nicolas Ray, Bruno Vallet, Wan Chiu Li, and Bruno Lévy. 2008. N-symmetry direction field design. *ACM Transactions on Graphics* 27, 2 (May 2008), Article 10, 13 pages. DOI: <https://doi.org/10.1145/1356682.1356683>
- Wulf Rossmann. 2002. *Lie Groups: An Introduction Through Linear Groups*. Oxford University Press.
- Leonid I. Rudin, Stanley Osher, and Emad Fatemi. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60, 1–4 (Nov. 1992), 259–268. DOI: [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F)
- G. Sapiro. 1996. Vector-valued active contours. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '96)*. IEEE, Los Alamitos, CA, 680–685. DOI: <https://doi.org/10.1109/CVPR.1996.517146>
- Zhongwei Shen, Xianzhong Fang, Xinguo Liu, Hujun Bao, and Jin Huang. 2016. Harmonic functions for rotational symmetry vector fields. *Computer Graphics Forum* 35, 7 (Oct. 2016), 507–516. DOI: <https://doi.org/10.1111/cgf.13047>
- Justin Solomon, Amir Vaxman, and David Bommes. 2017. Boundary element octahedral fields in volumes. *ACM Transactions on Graphics* 36, 3 (May 2017), Article 28, 16 pages. DOI: <https://doi.org/10.1145/3065254>
- Oded Stein, Eitan Grinspun, Max Wardetzky, and Alec Jacobson. 2018. Natural boundary conditions for smoothing in geometry processing. *ACM Transactions on Graphics* 37, 2 (May 2018), Article 23, 13 pages. DOI: <https://doi.org/10.1145/3186564>
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional field synthesis, design, and processing. In *Proceedings of the SIGGRAPH ASIA 2016 Courses (SA '16)*. Article 15, 30 pages.
- Qingnan Zhou and Alec Jacobson. 2016. Thingi10K: A dataset of 10,000 3D-printing models. *arXiv:1605.04797*.

Received September 2019; revised January 2020; accepted February 2020