

Algebraic Representations for Volumetric Frame Fields

DAVID PALMER, Massachusetts Institute of Technology
DAVID BOMMES, University of Bern
JUSTIN SOLOMON, Massachusetts Institute of Technology




Fig. 1. Octahedral fields generated with our methods on various models.

Field-guided parameterization methods have proven effective for quad meshing of surfaces; these methods compute smooth *cross fields* to guide the meshing process and then integrate the fields to construct a discrete mesh. A key challenge in extending these methods to three dimensions, however, is *representation* of field values. Whereas cross fields can be represented by tangent vector fields that form a linear space, the 3D analog—an octahedral frame field—takes values in a nonlinear manifold. In this work, we describe the space of octahedral frames in the language of differential and algebraic geometry. With this understanding, we develop geometry-aware tools for optimization of octahedral fields, namely geodesic stepping and exact projection via semidefinite relaxation. Our algebraic approach not only provides an elegant and mathematically sound description of the space of octahedral frames but also suggests a generalization to frames whose three axes scale independently, better capturing the singular behavior we expect to see in volumetric frame fields. These new *odeco frames*, so called as they are represented by orthogonally decomposable tensors, also admit a semidefinite program-based projection operator. Our description of the spaces of octahedral and odeco frames suggests computing frame fields via manifold-based optimization algorithms; we show that these algorithms efficiently produce high-quality fields while maintaining stability and smoothness.

The work of D. Palmer was generously supported by the Fannie and John Hertz Foundation through the Hertz Graduate Fellowship. D. Bommes was generously supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (AlgoHex, grant agreement no. 853343). J. Solomon was generously supported by 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: D. Palmer and J. Solomon, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139; emails: {drp, jsolomon}@mit.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.

© 2020 Copyright held by the owner/author(s).
0730-0301/2020/03-ART16
<https://doi.org/10.1145/3366786>

CCS Concepts: • **Computing methodologies** → **Computer graphics**; **Mesh geometry models**; **Volumetric models**;

Additional Key Words and Phrases: Hexahedral meshing, octahedral frame fields, convex relaxations, convex algebraic geometry

ACM Reference format:

David Palmer, David Bommes, and Justin Solomon. 2020. Algebraic Representations for Volumetric Frame Fields. *ACM Trans. Graph.* 39, 2, Article 16 (March 2020), 17 pages.
<https://doi.org/10.1145/3366786>

1 INTRODUCTION

Inspired by the success of field-based approaches to quadrilateral meshing on surfaces (cf. Vaxman et al. [2016]), recent research in applied geometry has focused on developing an analogous approach to hexahedral meshing. Motivated by applications in finite-element modeling, hexahedral meshing is the problem of dividing a given volume into hexahedral elements (deformed cubes) with minimal distortion and such that mesh boundary faces are aligned to the boundary of the volume.

Hexahedral meshing couples a geometry problem—minimizing distortion of mesh elements—to a combinatorial problem, placing mesh elements to achieve a desired connectivity structure. As in 2D, field-based approaches first ignore combinatorial constraints and solve for a *frame field*, which represents the local alignment and singular structure of a mesh (see Figure 1). Then, they integrate that field to guide the placement of hex elements [Nieser et al. 2011]. So as not to impose unnatural constraints, the space of frame fields must be expressive enough to represent the range of possible singularities that may appear in hexahedral meshes. These singularities are described by gluing relations restricted to the symmetries of a cube (i.e., the octahedral group).

One might hope that 2D field-based methods would extend easily to 3D. There are at least two obstructions to transferring ideas from the 2D case. First, the singular structure of a 3D frame field can be much more complicated than that of a cross field,

comprising an embedded graph rather than a set of isolated points. Second, we must understand the geometry of the space of frames to measure and optimize smoothness of a frame field.

Complementing the recent results of Liu et al. [2018] on frame field singularities and necessary conditions for hex meshability, in the present work we study the second challenge, namely characterizing the geometry of the space of frames. A 3D frame field is, intuitively, an assignment of three mutually orthogonal directions to each point in a volume. An orthonormal basis of three vectors—comprising an orthogonal matrix—is sufficient to specify a frame, but thanks to octahedral symmetry, multiple orthogonal matrices can specify the same frame. Formally, the space of octahedral frames can be viewed as the quotient of the group of 3D rotations $SO(3)$ by the right action of the octahedral group O comprising the rotational symmetries of a cube (see Section 3). The noncommutativity of 3D rotations makes the geometry of this quotient more complicated than that of its 2D counterpart. Recent attempts to lift ideas and techniques from 2D have either ignored the geometry of the frame space entirely or treated it as a black box for nonlinear optimization.

Consider perhaps the simplest form of optimization on a space—projection—finding the closest point in the space to a given point in an ambient space. Previous work on frame fields has treated this projection problem as a nonlinear, nonconvex optimization problem over frames parameterized by Euler angles, with no guarantees on convergence or global optimality. Our description of the octahedral frame space as an algebraic variety suggests a different approach to projection based on semidefinite programming, which yields a certificate of global optimality in polynomial time. Our semidefinite relaxation of projection is exact in a neighborhood of the octahedral variety, and we conjecture—with strong empirical evidence—that it is so universally.

Even when conducting local optimization on the space of octahedral frames, parameterization by Euler angles may not be the best approach. We show that the map from $SO(3)$ to the octahedral variety is a local isometry, enabling us to compute geodesics on the octahedral variety in closed form. Manifold-based optimization that moves along geodesics can then be used to accelerate local optimization dramatically.

Beyond precisely characterizing the space of octahedral frames, our algebraic approach admits a generalization to frames whose axes scale independently. This larger space better captures frame field geometry—for example, allowing for a nonzero direction aligned to singular arcs even if the directions orthogonal to the arcs must vanish. We call these new objects *odeco frames*, thanks to their construction using orthogonally decomposable tensors, and we derive relevant projection operators.

Our experiments show how the theoretical objects we study enable volumetric frame field design in practice. In particular, we apply standard manifold-based optimization algorithms to field design, built on our differential and algebraic descriptions of octahedral and odeco frames (Figure 2). The end result is an efficient suite of techniques for producing smooth fields that obey typical constraints for our target application.

Outline. The rest of the article is organized as follows. In Section 2, we discuss related work. In Section 3.1, we reintroduce the spherical harmonic representation of octahedral frames and




Fig. 2. We compute frame fields as maps into the *octahedral* and *odeco* varieties, a projected slice of which is depicted here.

demonstrate how this amounts to an equivariant *isometric* embedding of the quotient $SO(3)/O$ into \mathbb{R}^9 . In Section 3.2, we introduce the *odeco* frames, whose axes can scale independently, and we exhibit the spaces of octahedral and odeco frames as nested varieties cut out by quadratic equations. In Section 4, we describe essential primitives for optimization over octahedral and odeco frames, namely geodesics and projection via semidefinite programming. In Section 5, we formulate the frame field optimization problem. In Section 6, we describe two algorithms for optimizing fields. Section 7 describes experiments using these algorithms. A discussion and conclusion follow in Section 8. Further experimental results are included in the supplemental material. In summary, our contributions are

- a proof of isometric embedding of $SO(3)/O$ in \mathbb{R}^9 ;
- descriptions of the spaces of octahedral and more general odeco frames as nested algebraic varieties; and
- new state-of-the-art optimization techniques for volumetric frame fields valued in both varieties, featuring geodesics and semidefinite program (SDP)-based projection as primitives.

2 RELATED WORK

2.1 2D Frame Fields and Quadrilateral Meshing

Cross fields, and their application to quad meshing, have been studied extensively in geometry processing (cf. Vaxman et al. [2016]). A useful insight from cross field research is that it is advantageous to replace field values defined up to some symmetry with a *representation vector*—that is, some function of the field value invariant under the relevant symmetry. In two dimensions, four vectors forming a right-angled cross can be represented in a unified manner by their common complex fourth power. This amounts to an embedding of a quotient manifold into a Euclidean space; we show that the octahedral variety generalizes this idea to 3D in a natural and isometric manner.

Recently, an effort has been made to unify and formalize the various algorithmic approaches to cross fields, borrowing from

the Ginzburg-Landau theory from physics [Beaufort et al. 2017; Osting and Wang 2017; Viertel and Osting 2019]. This amounts to replacing an ill-posed unit norm constraint with a penalty term and taking the limit as the penalty parameter goes to infinity. Viertel and Osting [2019] show that local optima of such a procedure have isolated singularities with indices $\pm 1/4$, as appropriate for quad meshing. They propose a diffusion-generated algorithm to compute such local optima. Inspired by the work of Merriman, Bence, and Osher (MBO) (1992) on mean curvature flow, this algorithm alternates between finite-time diffusion and pointwise normalization. Osting and Wang [2017] study an analogous method applied to orthogonal matrix-valued fields. Our projection operators enable us to develop similar MBO diffusion-generated methods for optimization of octahedral and odedo fields. In Section 7, we demonstrate that a modified strategy, where the diffusion parameter is adjusted on-the-fly, frequently helps to avoid local minima.

2.2 3D Frame Fields and Hexahedral Meshing

Huang et al. [2011] introduce a representation of frames as functions over the sphere that exhibit octahedral symmetry, parameterized by coefficients in the spherical harmonic basis. As an initialization step, they solve a Laplace equation, resulting in coefficients in the interior that do not correspond to octahedral frames. These must be projected via nonconvex optimization over frames parameterized by Euler angles. They further optimize the results by minimizing a discrete Dirichlet energy over Euler angles. Ray et al. [2016] refine the three stages of this approach, with improved boundary constraints in the initialization, an efficient projection technique, and an L-BFGS optimization algorithm.

Solomon et al. [2017] reformulate the initialization step of Ray et al. [2016] using the boundary element method (BEM). This provides a way to harmonically interpolate the boundary conditions exactly, ignoring the constraints in the interior. Finally, sampled interior values are approximately projected onto the constraints as in the previous methods. As the constraints are ignored at the Dirichlet energy-minimization stage, there is no sense in which the final frame fields are optimal.

A related but distinct problem is the computation of fields of symmetric second-order tensors (i.e., symmetric matrices) [Palacios et al. 2017]. Every symmetric matrix has an orthonormal basis of eigenvectors, which corresponds to an octahedral frame. One might thus think symmetric matrix fields can be used to parameterize octahedral frame fields. Although a symmetric matrix field corresponds to at least one frame field, the correspondence is not one-to-one—for example, a field of identity matrices corresponds to all frame fields. Moreover, symmetric matrix fields can only represent singularities whose indices are multiples of $\pm 1/2$, whereas indices $\pm 1/4$ are crucial for hex meshing [Liu et al. 2018]. Our fourth-order octahedral and odedo fields are rich enough to represent all indices.

The work of Fang et al. [2016] on generalized polycubes imposes an even-more-extreme restriction that singularities may only appear on the boundary. After cutting handles, the regular frame field in the interior of the volume may be represented by a field of rotation (orthogonal) matrices. This is further relaxed to a field of matrices, and orthogonality is approximately enforced via a penalty

term. The resulting field is used to construct a polycube map of the cut volume, through which a hex mesh is pulled back. By requiring regularity in the interior, this method may exclude fields that achieve lower distortion overall.

The main driving force for research on 3D frame fields has been hexahedral mesh generation. For a broader overview, we refer the reader to the surveys of Armstrong et al. [2015] and Yu et al. [2015], and we limit the subsequent discussion to techniques involving frame fields. The idea of such methods is to construct a volumetric integer-grid map [Liu et al. 2018], through which portions of the Cartesian integer grid are pulled back to a structure-aligned hexahedral mesh in the input domain. Nieser et al. [2011] introduced a parameterization technique that turns a given frame field into an integer-grid map by solving a mixed-integer Poisson problem. Extraction of the hexahedral mesh from the map is hampered by degeneracies in the map, motivating the sanitization technique of Lyon et al. [2016] to improve robustness. Several refinements of the preceding ideas have been proposed, including different guidance for the frame field [Li et al. 2012] and heuristics to improve the singularity structure by decimation or splitting [Jiang et al. 2014; Li et al. 2012] to avoid degeneracies of the map.

Although robust hexahedral meshing based on frame fields remains an open problem, recently several hex-dominant meshing algorithms have been proposed [Gao et al. 2017; Sokolov et al. 2016] that also use frame fields but circumvent the problem of non-meshable singularities. Gao et al. [2017] propose a hierarchical optimizer for frame fields that is based on local relaxation.

However, many practical applications demand pure hexahedral meshes, such as those for the construction of volumetric spline spaces in isogeometric analysis, and consequently a full understanding of meshable field topology is required. To this end, Liu et al. [2018] enumerate the singular vertex types that may occur in a hex mesh with bounded edge valence; they also develop a topological index formula analogous to the Poincaré-Hopf formula for vector fields on surfaces. These local and global constraints being established, they propose an algorithm to generate a (meshable) frame field from a prescribed (meshable) singular structure. The theoretical portion of our work complements the topological work of Liu et al. [2018] with a closer study of the geometry of spaces of frames.

In concurrent work, Corman and Crane [2019] propose an alternative approach to computing a frame field with prescribed singular structure via a discretized connection on a frame bundle. This approach does not immediately extend to a method for *de novo* frame field design. However, the connection associated to a field is a natural object for the study of such properties as integrability. We leave to future work the study of connections associated to fields valued in the octahedral and odedo varieties and the extraction of such connections directly from the field coefficients.

2.3 Alternative Frame Representations

Complementing the spherical harmonic-based representation in the preceding works on octahedral frame fields, Chemin et al. [2018] propose an equivalent representation of octahedral frames as certain symmetric tensors of order four. They introduce algebraic equations characterizing the octahedral frames among




Fig. 3. An odeco field on a triangular prism. The norm of the second band coefficients indicates degree of anisotropy (left). Odeco frames close to the approximate position of the singular curve, indicated by a dot (right), scale toward zero in the directions normal to the curve but remain nondegenerate in the direction along the curve.

symmetric fourth-order tensors. These equations are equivalent under a change of basis to our defining equations for the octahedral variety. However, by using a basis suggested by the structure of $SO(3)$, we are able to not only present the defining equations of the octahedral variety but also illuminate how it is an isometric embedding of the quotient space $SO(3)/O$, enabling us to compute geodesics. Additionally, we use our algebraic description of the octahedral variety to introduce a semidefinite relaxation of projection and to place it in the context of the more general odeco variety.

In concurrent work, Golovaty et al. [2019] also represent frames by fourth-order symmetric tensors subject to some algebraic conditions, and they propose gradient flow on a Ginzburg-Landau-type energy to optimize for smooth fields.

An additional alternative representation is proposed by Gao et al. [2017], who use quaternions to encode octahedral frames. Their representation uses relatively few values, but a matching procedure has to be embedded in their optimization objective to account for the fact that 48 quaternions correspond to the same frame. The concurrent work [Beaufort et al. 2019] uses quaternions to derive another parameterization of octahedral frames by points of a variety in three complex dimensions.

All previous work on 3D frame fields has only considered octahedral frames, which do not capture the unidirectional behavior of frame fields near singular curves (cf. Figure 3). In the algebraic geometry community, Robeva [2016] and Boralevi et al. [2017] have completely characterized a family of algebraic varieties of *orthogonally decomposable* (odeco) tensors. We show how the octahedral variety is embedded in one of the odeco varieties, and we introduce a technique for optimization over the odeco frames. For our purposes, odeco frames generalize octahedral frames by allowing independent scaling of the “axes” of a frame, including degeneration to unidirectionality at singular curves and to zero at singular nodes.

Finally, Shen et al. [2016] consider fields having different local discrete symmetry groups, such as the tetrahedral and icosahedral groups. They extend the methods of Huang et al. [2011] and Ray et al. [2016] to compute such fields.

2.4 Semidefinite Relaxations

Relaxation of algebraic optimization problems to SDPs has been studied extensively in the field of real algebraic geometry and optimization. Blekherman et al. [2012] provide an introduction to this discipline. The efficacy of semidefinite relaxation in computer science was demonstrated dramatically in the seminal paper of Goemans and Williamson [1995] on the maximum cut problem. Since then, semidefinite relaxation has continued to be employed to solve both combinatorial and continuous optimization problems, such as angular synchronization [Singer 2011]. In geometry processing and graphics, this machinery has been applied to problems such as correspondence [Kezurer et al. 2015], consistent mapping [Huang and Guibas 2013], registration [Maron et al. 2016], camera motion estimation/calibration [Agrawal and Davis 2003; Ozyesil et al. 2015], and deformation [Kovalsky et al. 2014].

Most relevant to the present work are relaxations of the Euclidean projection problem onto a variety defined by quadratic equations, an example of a quadratically constrained quadratic program (QCQP). Cifuentes et al. [2017] have recently shown a stability result implying that the semidefinite relaxation of Euclidean projection onto a smooth, quadratically defined variety is exact in a neighborhood of the variety. Cifuentes et al. [2018] have additionally shown that the region in which the relaxation is exact is a *semialgebraic set*, and they have provided a formula for the degree of its algebraic boundary. Unfortunately, computing this boundary is generally intractable for interesting varieties. A deeper theoretical understanding of when semidefinite relaxations of Euclidean projection are globally exact is still lacking.

3 SPACES OF FRAMES

As studied previously in Chemin et al. [2018], Huang et al. [2011], Ray et al. [2016], and Solomon et al. [2017], the basic unknown in the volumetric frame field problem is a tuple of three mutually orthogonal directions at each point in a volume. These directions may be represented by an orthonormal basis of vectors, but the signs and order of the vectors are irrelevant thanks to octahedral symmetry. This redundancy makes detecting smoothness of a field of tuples difficult. Hence, it pays to use a unified representation invariant to the symmetries of the frame.

In the following, we apply machinery from differential and algebraic geometry to derive a succinct description of this basic octahedral frame and show how it is related to rotations of the function $x_1^4 + x_2^4 + x_3^4$ on the unit sphere expressed in the spherical harmonic basis (as used to represent frames in Huang et al. [2011], Ray et al. [2016], and Solomon et al. [2017]) and to tensorial representations (as used in Chemin et al. [2018]). Algebraic language not only provides a succinct description of previous representations but also suggests a means of generalizing to frames whose three axes scale independently (e.g., rotations of the function $\sum_i \lambda_i x_i^4$ for varying $\lambda \in \mathbb{R}^3$), whereas previous work requires them to have the same length ($\lambda_1 = \lambda_2 = \lambda_3$). This broader set better aligns with the realities of the frame field problem, since singular edges have nontrivial directionality that cannot be captured by existing representations. Relevant background material may be found in Appendix A.

3.1 Octahedral Variety

Intuitively, an octahedral frame field is a smooth assignment of (unoriented) orthonormal coordinate axes to each point x in a region $\Omega \subset \mathbb{R}^3$. Such frames are called *octahedral* because they exhibit symmetries described by the octahedral group O .¹ The space of such octahedral frames can be identified with the quotient space $\mathcal{F} = \text{SO}(3)/O$ —that is, the quotient of the group of oriented rotations by the right action of the octahedral group. Since O is not a normal subgroup of $\text{SO}(3)$ (indeed $\text{SO}(3)$ is simple), \mathcal{F} is not a group. However, as O is a finite group acting freely on $\text{SO}(3)$, \mathcal{F} is a manifold, and the surjective map $\text{SO}(3) \rightarrow \mathcal{F}$ is a covering map. The universal cover of \mathcal{F} is that of $\text{SO}(3)$, namely $\text{SU}(2)$, and so its fundamental group $\pi_1(\mathcal{F})$ is the lift of O to $\text{SU}(2)$, the binary octahedral group BO . In particular, BO classifies singular curves of octahedral fields [Mermin 1979].

$\text{SO}(3)$ admits a bi-invariant Riemannian metric (23), which means that the action of O by right multiplication is isometric. Thus, the Riemannian metric on $\text{SO}(3)$ descends to \mathcal{F} , making it a Riemannian manifold. In particular, \mathcal{F} has geodesics that lift to geodesics of $\text{SO}(3)$. We will employ these geodesics to compute optimization substeps on \mathcal{F} (cf. Section 4.1).

We have introduced \mathcal{F} as an abstract smooth manifold, but an embedding of \mathcal{F} in a Euclidean space is necessary to make it amenable to computation. Previous works have introduced this equivariant embedding via the action of $\text{SO}(3)$ on polynomials and spherical harmonic coefficients. For completeness, we reintroduce it here, although in a slightly different way that highlights the role of representation theory and the orbit-stabilizer theorem. Later, we will show that the embedding of \mathcal{F} in \mathbb{R}^9 is an *equivariant isometric embedding* (cf. Moore [1976]).

Consider the irreducible representation $\rho : \text{SO}(3) \rightarrow \text{SO}(9)$ corresponding to the fourth band of spherical harmonics. The 9×9 orthogonal matrices $\rho(r)$ for $r \in \text{SO}(3)$ are sometimes referred to as Wigner D -matrices. Form a linear operator $H \in \mathbb{R}^{9 \times 9}$ as

$$H = \frac{1}{|O|} \sum_{o \in O} \rho(o).$$

This H is a projection operator onto the subspace of \mathbb{R}^9 invariant under all octahedral rotations.

LEMMA 3.1. For any $o' \in O$, $\rho(o')H = H$.

PROOF.

$$\rho(o')H = \frac{1}{|O|} \sum_{o \in O} \rho(o')\rho(o) = \frac{1}{|O|} \sum_{o \in O} \rho(o'o) = \frac{1}{|O|} \sum_{o \in O} \rho(o) = H. \quad \square$$

COROLLARY 3.2. For $q \in \mathbb{R}^9$, $\rho(O) \cdot q = q$ if and only if $q \in \text{Im}(H)$.

PROOF. The forward implication follows by writing $q = Hq$. The reverse implication follows directly from Lemma 3.1. \square

Because O is a maximal subgroup of $\text{SO}(3)$, we have the following corollary. Here, $\text{stab } q$ denotes the stabilizer of q —that is, the subgroup of $\text{SO}(3)$ leaving q invariant (see Appendix A).

COROLLARY 3.3. Let $q \in \text{Im}(H)$ and nonzero. Then $\text{stab}(q) = O$.

¹Technically, their stabilizers are conjugate to O .

It happens that $H \in \mathbb{R}^{9 \times 9}$ has rank one (i.e., its image is 1D), motivating the following definitions.

Definition 3.4. The *canonical octahedral frame* is the normalized vector

$$q_0 = \left(0, 0, 0, 0, \sqrt{\frac{7}{12}}, 0, 0, 0, \sqrt{\frac{5}{12}} \right)^\top \in \mathbb{R}^9$$

such that $H = q_0 q_0^\top$.

Our canonical frame is the same as the normalized projection of the polynomial $\sum_i x_i^4$ into the fourth band of spherical harmonics, such as in Solomon et al. [2017].

Definition 3.5. The *octahedral variety* is the orbit of q_0 under the action of $\text{SO}(3)$,

$$F = \rho(\text{SO}(3))q_0 = \{\rho(r)q_0 \mid r \in \text{SO}(3)\}.$$

To summarize, F is the orbit of q_0 , whose stabilizer is O . A smooth version of the orbit-stabilizer theorem (Theorem 21.18 of Lee [2012]) shows that there is an equivariant diffeomorphism $\phi : \mathcal{F} \rightarrow F$.

$$\begin{array}{ccc} \text{SO}(3) & \xrightarrow{\rho} & \text{SO}(9) \\ \downarrow & & \downarrow \cdot q_0 \\ \mathcal{F} & \xrightarrow{\phi} & F \end{array}$$

Next, we will characterize the Riemannian geometry of F by showing that ϕ is an isometry up to a uniform scale factor. To our knowledge, this observation has not appeared in previous work.

PROPOSITION 3.6. Let $\alpha = \sqrt{3/20}$ and $F_\alpha := \alpha F = \{\alpha q : q \in F\}$. Let $\pi_\alpha : \mathbb{R}^{9 \times 9} \rightarrow \mathbb{R}^9$ denote matrix multiplication by the scaled canonical octahedral frame $q_\alpha := \alpha q_0$. In other words, $\pi_\alpha(A) = Aq_\alpha$. Then the map

$$\pi_\alpha \circ \rho : \text{SO}(3) \rightarrow F_\alpha$$

is a local isometry, making the induced diffeomorphism $\phi_\alpha : \mathcal{F} \rightarrow F_\alpha$ an isometry.

PROOF. Taking the differential of ρ at the identity yields the associated Lie algebra representation

$$(D\rho)_I : \mathfrak{so}(3) = T_I \text{SO}(3) \rightarrow \mathfrak{so}(9) \subset \mathbb{R}^{9 \times 9}.$$

$(D\rho)_I$ is characterized by $L_i := (D\rho)_I(l_i)$, the images of the Lie algebra generators l_i (see Appendix A for definitions and supplemental material for explicit expressions). π_α is a linear map, so its differential is also multiplication by q_α .

To see that $\pi_\alpha \circ \rho$ is a local isometry, first recall that the metric on $\text{SO}(3)$ is bi-invariant. In particular, for each $g \in \text{SO}(3)$, an orthonormal basis for $T_g \text{SO}(3)$ is given by the right-translated Lie algebra generators

$$\{(DR_g)_I l_i\}_{i=1}^3.$$

Thus, it suffices to prove that their images under $\pi_\alpha \circ \rho$ form an orthonormal basis in $T_{\rho(g)q_\alpha} F_\alpha$. But

$$\begin{aligned} D(\pi_\alpha \circ \rho)_g (DR_g)_I l_i &= (D\pi_\alpha)(D\rho)_g (DR_g)_I l_i \\ &= (D\pi_\alpha)(DR_{\rho(g)})_I (D\rho)_I l_i \\ &= (D\pi_\alpha) L_i \rho(g) \\ &= L_i \rho(g) q_\alpha, \end{aligned} \quad (1)$$

where the second equality follows by differentiating the representation property $\rho \circ R_g = R_{\rho(g)} \circ \rho$ at the identity. Therefore, the equation we need to prove is

$$\langle L_i q, L_j q \rangle = \delta_{ij} \quad \forall q \in F_\alpha. \quad (2)$$

This isometry condition can be checked explicitly at q_α :

$$\langle L_i q_\alpha, L_j q_\alpha \rangle = \delta_{ij} = \langle l_i, l_j \rangle. \quad (3)$$

For a general $q = \rho(g)q_\alpha \in F_\alpha$, we compute

$$\begin{aligned} \langle L_i \rho(g)q_\alpha, L_j \rho(g)q_\alpha \rangle &= \langle \rho(g)^\top L_i \rho(g)q_\alpha, \rho(g)^\top L_j \rho(g)q_\alpha \rangle \\ &= \langle D\rho_I(g^\top l_i g)q_\alpha, D\rho_I(g^\top l_j g)q_\alpha \rangle. \end{aligned} \quad (4)$$

The first equality follows because $\rho(g) \in \text{SO}(9)$, and the second follows by Lemma A.1. Let $a_{ik} = a_{ik}(g)$ be the coefficients of the adjoint representation of $\text{SO}(3)$ (see Appendix A)—that is, $g^\top l_i g =: \sum_k a_{ik} l_k$; these coefficients form an orthogonal matrix. Now using linearity of $D\rho$ along with (3) and (4), we obtain

$$\begin{aligned} \langle L_i \rho(g)q_\alpha, L_j \rho(g)q_\alpha \rangle &= \left\langle \sum_r a_{ir} L_r q_\alpha, \sum_s a_{js} L_s q_\alpha \right\rangle \\ &= \sum_{r,s} a_{ir} a_{js} \langle L_r q_\alpha, L_s q_\alpha \rangle \\ &= \sum_k a_{ik} a_{jk} = \delta_{ij} \end{aligned} \quad (5)$$

as required. \square

In summary, we have described the octahedral variety as an embedded submanifold in \mathbb{R}^9 isometric to $\mathcal{F} = \text{SO}(3)/O$. This isometry means that we can do manifold optimization over frames by working in the embedding (cf. Section 6.1). To show that F is really an algebraic variety, we will need to exhibit equations that cut it out. We will delay this until Section 3.2.1, when we can give a unified description of the octahedral and odeco varieties.

3.2 Odeco Variety

The previous section provides more insight into the octahedral frames used in all previous work. Octahedral frames are suitable for representing singularity-free frame fields. However, frame fields commonly encountered in applications have singularities comprising an embedded graph. Consider the triangular prism shown in Figure 3. Near the singular curve, a smooth octahedral field would rotate infinitely quickly, and it would not have a well-defined value along the curve. This issue is analogous to the case of *unit* cross fields—note that for cross fields, the hairy ball theorem requires singularities on simply connected surfaces.

One solution to this is to replace the hard constraint that the field values lie in the octahedral variety with a penalty term, as motivates the MBO method for octahedral fields detailed in the following (Section 6.2). Another solution is homogenization—that is, allowing field values to scale, replacing singularities by zeros, as is considered for cross fields in Knöppel et al. [2013]. But consider the triangular prism again: a scaled octahedral field would vanish completely at the singular curve since all three orthogonal axes must scale uniformly. This makes the octahedral representation unable to capture the alignment of the field to the singular curve. To cope with this problem and to show the value of our algebraic

approach, we now describe a superset of the octahedral frames. This set allows the axes to scale independently. For instance, as shown in Figure 3, the frame axes orthogonal to the singular curve scale toward zero while the axis tangent to the singular curve remains nonzero.

The symmetric orthogonally decomposable (*odeco*) tensor varieties, introduced in [Robeva 2016], parameterize symmetric tensors $T \in \text{Sym}^d \mathbb{R}^n$ that can be written as

$$T = \sum_{i=1}^n \lambda_i (v_i)^{\otimes d}$$

for some set of n orthonormal vectors $v_i \in \mathbb{R}^n$, where $v^{\otimes d}$ denotes the d -wise tensor power of the vector v . In other words, an odeco tensor encodes a set of orthogonal vectors up to permutation. If d is even, T is also invariant under sign changes to the v_i . Moreover, an odeco tensor assigns *weights* λ_i to the vectors v_i . This property allows us to encode frames whose axes scale independently.

There is a one-to-one correspondence between symmetric tensors of order d over \mathbb{R}^n and homogeneous polynomials of degree d in n variables (cf. Section 1.2 of Robeva [2016]). This correspondence is given in one direction by taking a polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ to its tensor of d th derivatives and in the other direction by symbolic evaluation on the vector of formal variables $x = (x_1, \dots, x_n)$:

$$T \in \text{Sym}^d \mathbb{R}^n \mapsto p_T(x) = \frac{1}{d!} T(x, \dots, x) \in \mathbb{R}[x].$$

This is a generalization of the correspondence between symmetric bilinear forms (i.e., symmetric 2-tensors) and quadratic forms (i.e., homogeneous quadratic polynomials). Note that T is odeco if and only if p_T can be written as a sum of d th powers of linear forms

$$p_T(x) = \sum_i \lambda_i (v_i^\top x)^d,$$

where the v_i are orthonormal as earlier. In this case, we also refer to p_T as an odeco polynomial.

The defining equations of the odeco varieties are quadrics—homogeneous quadratic equations—in the tensor coefficients (Theorem 4 of Boralevi et al. [2017]) or equivalently in the coefficients of the associated polynomial p_T . In other words, a homogeneous polynomial

$$p(x) = \sum_{d_1 + \dots + d_n = d} u_{d_1, \dots, d_n} x_1^{d_1} \dots x_n^{d_n}$$

is odeco if and only if the coefficients u satisfy

$$u^\top A_i u = 0 \quad (6)$$

for a finite set of symmetric matrices A_i . In the case relevant to us, where $n = 3$ and $d = 4$, there are exactly 27 such defining equations. Although dimension counting might suggest otherwise, these equations are not redundant—as can be seen by computing a Gröbner basis of the ideal they generate. The matrices A_i are listed explicitly in the supplementary material. We will henceforth refer to this particular odeco variety simply as the *odeco variety* \tilde{F} . Figure 4 plots several fourth-order odeco polynomials over the unit sphere.




Fig. 4. Examples of z -aligned odeco polynomials, plotted over the sphere.

3.2.1 Octahedral Variety as a Subvariety of Odeco Variety. Recall that our octahedral frames were represented by coefficients in an irreducible representation of $SO(3)$, whereas the odeco variety was defined using monomial or tensor coefficients. To see the relationship between the two varieties, it is beneficial to recast the odeco variety in the irreducible representation basis. This corresponds to looking at the coefficients of odeco polynomials in the basis of spherical harmonics.

The quartic polynomials comprising the odeco variety decompose as linear combinations of the spherical harmonics of bands 0, 2, and 4. Consider an odeco polynomial $\sum_{i=1}^3 \lambda_i (v_i^\top x_i)^4$, and let

$$q = (q_0, q_2, q_4) \in V_0 \times V_2 \times V_4$$

be its coefficients in this basis of even spherical harmonics. These 15 coefficients give us a different representation of odeco frames in \mathbb{R}^{15} , where each band has a clear meaning. In particular, $q_0 = C_0 \sum_i \lambda_i$, where C_0 is a constant, and similarly

$$\|q_2\|^2 = C_2 \left(\sum_i \lambda_i^2 - \sum_{i < j} \lambda_i \lambda_j \right).$$

The parenthesized expression is the squared distance between $(\lambda_1, \lambda_2, \lambda_3)$ and the line $\lambda_1 = \lambda_2 = \lambda_3$. Thus, $q_2 = 0$ if and only if q_4 is a scalar multiple of an octahedral frame. The set

$$\{q \in \tilde{F} \mid q_2 = 0\}$$

consists of scalar multiples of the octahedral variety indexed by q_0 . Fix $q_2 = 0$ and let q_0 be a constant such that $\|q_4\| = 1$. The octahedral variety is the intersection of this affine subspace with the odeco variety. Reducing the Equations (6) of the odeco variety with respect to this subspace yields 15 inhomogeneous quadratic equations

$$\begin{pmatrix} 1 \\ q_4 \end{pmatrix}^\top P_i \begin{pmatrix} 1 \\ q_4 \end{pmatrix} = 0, \quad i = 1, \dots, 15 \quad (7)$$

cutting out the octahedral variety F . As was the case for the odeco variety, these equations are not redundant. The symmetric matrices P_i are listed explicitly in the supplementary material.

4 GEODESICS AND PROJECTION

We have introduced two spaces, the octahedral and odeco varieties, that can serve as target sets for frame fields. To compute such fields, we will have to solve optimization problems over products of many copies of these varieties. Naively, one might plug the quadratic constraints in (6) and (7) directly into a generic quadratic optimization solver. However, the A_i or P_i are not positive semidefinite, nor can their span be rewritten as the span of positive semidefinite matrices. Thus, the constraints are nonconvex and challenging to enforce.

As an alternative, we use optimization algorithms that are tailored to the manifold-valued variables in our problem. These algorithms employ substeps such as geodesic traversal and projection. We derive these operations for a single frame in the following.

4.1 Octahedral Geodesics

By Proposition 3.6, the scaled octahedral variety F_α is locally isometric to $SO(3)$ via the map $r \mapsto \rho(r) \cdot q_\alpha$. It follows that geodesics of $SO(3)$ push forward to geodesics of F_α . The relation (17) in Appendix A then allows us to compute geodesics on F_α in closed form without evaluating the representation map ρ explicitly. Let $v \in T_q F_\alpha$. v can be written in a basis induced by the $SO(3)$ action, $v = \sum_{i=1}^3 v_i L_i q$. Here, the coefficient vector $v = (v_1, v_2, v_3) \in \mathbb{R}^3$ is the “axis-angle” representation of a rotation, and the $SO(3)$ exponential maps it to the corresponding rotation matrix. This can be computed by conjugation with a rotation r taking the unit vector $v/\|v\|$ to $(0, 0, 1)$. Composing with the representation map, we have

$$\exp_q(v) = \rho(r^\top \exp(\|v\|L_3)r)q = \rho(r)^\top \exp(\|v\|L_3)\rho(r)q, \quad (8)$$

where (unsubscripted) \exp denotes the ordinary matrix exponential.

To compute $\rho(r)$, define spherical coordinates θ, φ such that

$$v = \|v\|(\cos \theta \sin \varphi, \sin \theta \sin \varphi, \cos \varphi).$$

Then one possible choice for r is

$$r = \exp(-\phi L_2) \exp(-\theta L_3) = r_{23}^\top \exp(-\phi L_3) r_{23} \exp(-\theta L_3),$$

where $r_{23} = \exp((\pi/2)L_1)$. Thus,

$$\rho(r) = R_{23}^\top \exp(-\phi L_3) R_{23} \exp(-\theta L_3), \quad (9)$$

where $R_{23} = \exp((\pi/2)L_1)$. Combining (9) with (8), we can compute geodesics by products of two simple ingredients: R_{23} and $\exp(tL_3)$ for $t \in \mathbb{R}$. Closed-form expressions for both appear in Section 2 of the supplementary material. Note that a formula similar to (9) is common in the graphics literature, such as for rotating BRDFs expressed in spherical harmonic coefficients (cf. the appendix in Kautz et al. [2002]).

4.2 Projection via Semidefinite Relaxation

F and \tilde{F} are both varieties defined by quadratic equations. $F \subset \mathbb{R}^9$ is cut out by 15 inhomogeneous quadratic equations (7), whereas $\tilde{F} \subset \mathbb{R}^{15}$ is cut out by 27 homogeneous quadratic equations (6). Consider the problem of finding the closest point in F to a given point $y \in \mathbb{R}^9$:

$$\Pi_F(y) = \arg \min_{q \in F} \|q - y\|^2. \quad (P_F)$$

This problem is referred to as Euclidean projection onto a quadratic variety. It is an example of a QCQP, and the general recipe for semidefinite relaxation detailed in Section A.2 automatically applies. The SDP will have the form

$$\begin{aligned} & \arg \min_{Q \in \mathbb{R}^{10 \times 10}} \langle Y, Q \rangle \\ & \text{subject to} \quad Q_{11} = 1 \\ & \quad \langle P_i, Q \rangle = 0, \quad i = 1, \dots, 15 \\ & \quad Q \geq 0, \end{aligned} \quad (\text{SDP}_F)$$

where P_i are the symmetric matrices from (7), and

$$Y = \begin{pmatrix} \|y\|^2 & -y^\top \\ -y & I_{9 \times 9} \end{pmatrix}.$$

Let Q^* be an optimal solution to (SDP_F) . Since (SDP_F) is a relaxation of (P_F) , $\langle Y, Q^* \rangle$ is a lower bound on the objective value of (P_F) . If it happens that $\text{rank}(Q^*) = 1$, then Q^* can be written in the form

$$Q^* = \begin{pmatrix} 1 \\ q^* \end{pmatrix} \begin{pmatrix} 1 \\ q^* \end{pmatrix}^\top,$$

from which it follows that

$$\begin{pmatrix} 1 \\ q^* \end{pmatrix}^\top P_i \begin{pmatrix} 1 \\ q^* \end{pmatrix} = 0, \quad i = 1, \dots, 15$$

(i.e., $q^* \in F$). In this case, q^* is the globally optimal solution to (P_F) . This situation is called *exact recovery*. The foregoing discussion also applies, *mutatis mutandis*, to \tilde{F} .

In our MBO algorithm presented in the following (cf. Section 6.2), we alternate projection with stepping off the variety. We refer to the following theorem, which suggests that when taking small enough steps from smooth points of a variety, projection will be generically exact.

THEOREM 4.1 (CIFUENTES ET AL. [2017], THEOREM 1.2). *Consider the Euclidean projection problem*

$$\arg \min_{q \in \mathcal{V}} \|q - y\|, \quad y \in \mathbb{R}^n,$$

where $\mathcal{V} \subset \mathbb{R}^n$ is a real variety defined by quadratic equations $f_1(q) = \dots = f_m(q) = 0$. Let $\tilde{y} \in \mathcal{V}$ be a point at which the rank of the Jacobian $\nabla f(\tilde{y})$ is equal to the codimension of \mathcal{V} . Then the semidefinite relaxation of projection is exact for $y \in \mathbb{R}^n$ sufficiently close to \tilde{y} .

In addition to this theoretical motivation, we have ample empirical evidence that our relaxations are exact under much more general conditions. The blue histogram in Figure 5 shows the results of attempting to project 10^6 random points onto F via semidefinite relaxation. We use the ratio of the second largest eigenvalue $\lambda_2(Q^*)$ to the largest eigenvalue $\lambda_1(Q^*)$ as a proxy for exactness, as it indicates whether Q^* was rank-one up to machine precision. For all octahedral projections, $\lambda_2(Q^*)/\lambda_1(Q^*)$ was $\approx 10^{-8}$ or less. Based on these results, we make the following conjecture.

CONJECTURE 4.2. *For a generic point $q_0 \in \mathbb{R}^9$, the solution to (SDP_F) has rank one and therefore yields the exact projection $\Pi_F(q_0)$.*

We also tried projecting 10^6 random quartic polynomials onto the odeco variety, as shown in the red histogram in Figure 5. Most of the odeco projections were also exact up to numerical precision: out of 10^6 solutions, only 60 had an eigenvalue ratio greater than 10^{-8} . Theorem 4.1 gives us some intuition as to why this might happen. The stability result only holds near smooth points of the variety, but whereas the octahedral variety is a smooth manifold, the odeco variety has a singularity at the origin, separating polynomials of different signs.

CONJECTURE 4.3. *For a generic point $q_0 \in \mathbb{R}^{15}$ representing a positive polynomial, the SDP relaxation yields the exact projection $\Pi_{\tilde{F}}(q_0)$.*




Fig. 5. Histogram of eigenvalue ratio $\lambda_2(Q^*)/\lambda_1(Q^*)$ for solutions to the SDP relaxations of Euclidean projection onto F and \tilde{F} . Projections of 10^6 random points were tested for each. The maximum ratio for octahedral projection was 2.41×10^{-8} . The maximum ratio for odeco projection of positive quartic polynomials was 1.54×10^{-9} . See Section 4.2.




Fig. 6. For some query polynomials, our globally optimal SDP-based projection onto the octahedral variety (blue) yields dramatically different results from the approximate projection of Ray et al. [2016] (red). Our projections are closer to the query points y (distances shown in blue) compared to theirs (distances in red). The query polynomials are plotted on the sphere, with color and distance from the center proportional to magnitude.

Indeed, the green histogram in Figure 5 shows the results of odeco projection on 10^6 random positive quartic polynomials, generated as sums of squares of random quadratic polynomials. For all such positive initial points, the SDP solution had $\lambda_2(Q^*)/\lambda_1(Q^*) < 10^{-8}$, supporting Conjecture 4.3.

For octahedral frames, we also compare our SDP-based projection to the previous state-of-the-art method proposed by Ray et al. [2016]. Because their method is based on nonconvex optimization, we would expect it to get stuck in local minima. Indeed, out of 100,000 trials, we found 600 cases for which the result of their method was at least 10^{-3} further from the initial point than our projection—a nearly 1% error rate. Moreover, the difference between the computed projections can be substantial, as illustrated in Figure 6.

5 FROM FRAMES TO FRAME FIELDS

The overall aspiration of this work is to compute smooth frame fields in a region $\Omega \subseteq \mathbb{R}^3$ aligned to the normal \mathbf{n} on its boundary $\partial\Omega$. We assume Ω to be compact with $\partial\Omega$ a union of smooth,

embedded surfaces. We think of a frame field as a map $\phi : \Omega \rightarrow \mathcal{V}$ into some space of frames \mathcal{V} satisfying alignment boundary conditions.

We have examined the geometry of two candidates for the fiber \mathcal{V} —the octahedral variety F and the odeco variety \tilde{F} .² It remains to describe the boundary conditions and to define an energy we want ϕ to minimize.

5.1 Boundary Conditions

In either F or \tilde{F} , the frames aligned to a given direction $\mathbf{n} \in \mathbb{R}^3$ form the intersection of an affine subspace with the respective variety. In each case, this intersection is a lower-dimensional variety defined by a different set of quadratic equations. We can impose alignment boundary conditions by working over this lower-dimensional variety. For boundaries with sharp creases, we can optionally exclude creased vertices, where the normal direction is ill defined, from the alignment constraint. Other constraints can also be imposed at creases—for example, alignment to the crease tangent.

5.1.1 Octahedral Boundary Conditions. Let $\rho : \text{SO}(3) \rightarrow \text{SO}(9)$ be the irreducible representation as in Section 3. First consider the octahedral frames aligned to the vertical $\mathbf{z} = (0, 0, 1)^\top$. The space of vertical-aligned frames $F_z := \rho(A_z)q_0$, where A_z is the coaxial subgroup consisting of all rotations about \mathbf{z} —for instance,

$$A_z = \{\exp(tL_3) \mid t \in \mathbb{R}\},$$

and q_0 is the canonical octahedral frame. As derived in Huang et al. [2011] and Ray et al. [2016], such frames have the form

$$\begin{aligned} \exp(tL_3)q_0 &= \left(\sqrt{\frac{5}{12}} \sin(4t), 0, 0, 0, \sqrt{\frac{7}{12}}, 0, 0, 0, \sqrt{\frac{5}{12}} \cos(4t) \right)^\top \\ &= q_z + B_z s_z, \end{aligned} \quad (10)$$

where $q_z = (0, 0, 0, 0, \sqrt{7/12}, 0, 0, 0, 0)^\top$, $s_z = (\cos(4t), \sin(4t))^\top$, and B_z is the obvious 9×2 matrix. Given $q \in \mathbb{R}^9$, the closest vertical-aligned frame is

$$\Pi_{F_z}(q) = q_z + B_z \frac{B_z^\top q}{\|B_z^\top q\|}.$$

For a general unit normal \mathbf{n} , the aligned frames can be written as

$$\rho(r_n)(q_z + B_z s_z), \quad (11)$$

where $r_n \in \text{SO}(3)$ is any rotation taking \mathbf{z} to \mathbf{n} . The projection of $q \in \mathbb{R}^9$ onto the \mathbf{n} -aligned frames is then

$$\Pi_{F_n} = \rho(r_n) \Pi_{F_z} (\rho(r_n)^\top q). \quad (12)$$

During optimization, this projection is used for boundary-aligned frames.

²The term *fiber* is intended to be suggestive. It may be fruitful to consider a bundle $\pi : P \rightarrow \Omega$ with fiber \mathcal{V} , whose restriction $\partial P \rightarrow \partial \Omega$ is nontrivial and encodes the boundary alignment condition. The map ϕ would be replaced with a section of P .

5.1.2 Odeco Boundary Conditions. Consider the standard odeco frame $\sum_i \lambda_i x_i^4$ rotated by some $r \in A_z$, and fix $\lambda_3 = 1$. As described in Section 3.2.1, it has coefficients in the even-numbered irreducible representations (bands of spherical harmonics)

$$q = (q_0, q_2, q_4) \in V_0 \times V_2 \times V_4 = \mathbb{R} \times \mathbb{R}^5 \times \mathbb{R}^9.$$

Just as in the octahedral case, q can be decomposed into a fixed part and a rotating part parameterized by a lower-dimensional vector

$$q = q_z + B_z s_z, \quad (13)$$

where now $s_z \in \mathbb{R}^5$ and $B_z \in \mathbb{R}^{15 \times 5}$. The Equations (6) reduce to three quadratic equations in the coefficients of s_z , which can be used to construct an SDP to project onto the vertical-aligned odeco frames, as in Section 4.2. The details are given in the supplementary material. For frames aligned to a direction \mathbf{n} ,

$$q = \rho(r_n)(q_z + B_z s_z),$$

where ρ is now the direct sum representation on $V_0 \times V_2 \times V_4$. Projection onto the \mathbf{n} -aligned odeco frames can be constructed from projection onto the \mathbf{z} -aligned odeco frames as in (12).

5.2 Objective Function

As we are searching for smoothest fields, a natural choice for the energy is Dirichlet energy $\frac{1}{2} \int_\Omega \|\nabla \phi\|^2 dx$, where the norm is taken with respect to a metric on \mathcal{V} . There are multiple problems that this formulation will need to confront. As we have seen, F is a smooth manifold. But being a quotient of the 3-sphere, it has positive curvature, making mere existence of harmonic maps $\Omega \rightarrow F$ a hard problem. Even more fundamentally, it is not clear how to represent singularities of ϕ —the map cannot be consistently defined along singular curves. \tilde{F} attempts to solve this problem by representing singular frames with only one direction while allowing the other two to vanish. Along a singular curve, we would expect an odeco field to take rank-one values of the form $\lambda v^{\otimes 4}$, where v is tangent to the singular curve.

5.3 Discretization

Let $\mathcal{T} = (V, T)$ be a tetrahedral mesh of Ω with vertices V and tetrahedra T . The set of boundary vertices will be denoted by ∂V . A discrete octahedral frame field on \mathcal{T} is specified by a map $q : V \rightarrow \mathcal{V} = F$ or \tilde{F} , giving a frame q_i for each vertex i . As the octahedral variety $F \subset \mathbb{R}^9$ and the odeco variety $\tilde{F} \subset \mathbb{R}^{15}$, we can think of q as a $d \times n$ matrix, where $n = |V|$ and $d = 9$ or 15 , respectively.

We then define a discretized Dirichlet energy using the Euclidean metric in the spherical harmonic basis to compare elements of \mathcal{V} . This is equivalent to measuring L^2 distance between the corresponding polynomials over the sphere, as employed in Huang et al. [2011], Ray et al. [2016], and Solomon et al. [2017]. Note that this would not be the case if we compared odeco frames in the monomial basis. The discrete energy is

$$E(q) = \frac{1}{2} \text{tr}(qLq^\top),$$

where L is the linear finite element Laplacian on \mathcal{T} , with the usual cotangent weights.

6 ALGORITHMS

We now have two *variety-constrained optimization* problems of the form

$$\begin{aligned} \min & \frac{1}{2} \operatorname{tr}(qLq^\top) \\ \text{s.t. } & q_i \in \mathcal{V}, \quad \forall i \in V \\ & q_i \in \mathcal{V}_{n_i}, \quad \forall i \in \partial V, \end{aligned} \quad (14)$$

where the variety \mathcal{V} is either F or \tilde{F} , q_i are the columns of q , ∂V denotes the set of boundary vertices, and \mathcal{V}_{n_i} is the variety of frames aligned to the normal at boundary vertex i .

6.1 Manifold Optimization Methods

In the case $\mathcal{V} = F$, (14) becomes a manifold-constrained optimization problem. The octahedral variety is a Riemannian manifold, and we can compute geodesics on it as described in Section 4.1. Thus, we can apply the standard Riemannian trust-region (RTR) algorithm [Absil et al. 2007] as implemented in the Manopt toolbox for MATLAB [Boumal et al. 2014]. We consider the field to be a point in the product manifold

$$q \in \prod_{i \in V \setminus \partial V} F \times \prod_{i \in \partial V} F_{n_i}.$$

In addition to the geodesics computed in Section 4.1, we also need a way to project vectors (e.g., gradients) in the ambient space \mathbb{R}^9 into the tangent space at a point q , $\operatorname{proj}_q : \mathbb{R}^9 \rightarrow T_q F$. The vectors $\{L_i q\}_{i=1}^3$ form an orthogonal basis for this tangent space (cf. Proposition 3.6), so projection is just given by taking the inner product with each of these vectors.

The odeco variety \tilde{F} is not a manifold, but it is smooth almost everywhere. We know of no way to compute exact geodesics on \tilde{F} , but we can implement a version of RTR by replacing the exact exponential map with a *retraction*—that is, a first-order approximation (cf. [Absil et al. 2007, Definition 2.1]). In doing so, we give up the superlinear local convergence guarantees of RTR (cf. [Absil et al. 2007, Theorem 4.12]) but retain a global convergence guarantee. In practice, we find that odeco RTR converges at a reasonable rate.

At a smooth point of \tilde{F} , it is easy to project vectors onto the tangent space using the fact that \tilde{F} is defined by *quadratic* equations. Let $q \in \tilde{F}$, and assume the coefficients of q are expressed in spherical harmonic coefficients. Then differentiating the equations $q^\top \tilde{A}_i q = 0$ shows that the normal space at q is

$$N_q \tilde{F} = \operatorname{span}\{\tilde{A}_i q\},$$

where \tilde{A}_i are expressed in the spherical harmonic basis. Then $T_q \tilde{F} = (N_q \tilde{F})^\perp$, where the orthogonal complement is taken with respect to the standard metric under which the spherical harmonic functions are orthonormal.

We compute retractions as follows. The tangent space to the odeco variety decomposes into a *rotational part* and a *scaling part*:

$$T_q^r \tilde{F} = \operatorname{span}\{\tilde{L}_i q\}_{i=1}^3, \quad T_q^s \tilde{F} = (T_q^r \tilde{F})^\perp,$$

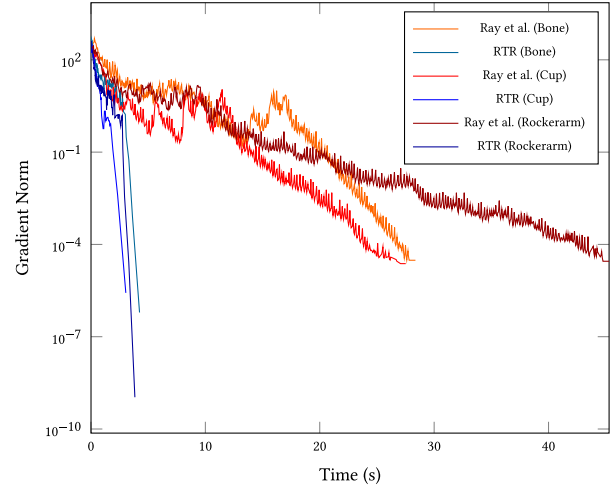


Fig. 7. Convergence behavior of octahedral RTR and the method of Ray et al. [2016] on various models, starting from their initialization and using the combinatorial Laplacian.

where the orthogonal complement is taken within $T_q \tilde{F}$. We then set

$$\operatorname{retr}_q(v) = e^{v_r \cdot \tilde{L}}(q + v_s),$$

where v_s is the component of v in $T_q^s \tilde{F}$, v_r is the component of v in $T_q^r \tilde{F}$, and

$$v_r \cdot \tilde{L} := \sum_j (v_r)_j \tilde{L}_j,$$

where $v_r = \sum_j (v_r)_j \tilde{L}_j q$. It is simple to verify that $\operatorname{retr}_q(v) \in \tilde{F}$ and that


$$\left. \frac{\partial}{\partial t} \operatorname{retr}_q(tv) \right|_{t=0} = v.$$

We compare the convergence behavior of octahedral RTR to that of the method of Ray et al. [2016] under two sets of conditions—their initialization and combinatorial Laplacian (Figure 7), and our random initialization and linear finite element Laplacian (Figure 8). The quadratic local convergence of RTR stands in stark contrast to the slower linear convergence behavior of their method. RTR converges faster but finds solutions of similar Dirichlet energy to previous work (see Table 1 in the supplementary material).

6.2 Generalized MBO Methods

As we have observed, it is possible to do optimization on the octahedral and odeco varieties by moving along curves that stay on the varieties exactly. However, this hard constraint sometimes causes pure manifold optimization to get stuck in local minima. To avoid such local minima, an approach that allows “tunneling” is required.

The MBO method (1992) is a diffusion-generated method for computing (possibly singular) maps into a target space embedded in a Euclidean space. Following Osting and Wang [2017] and Viertel and Osting [2019], we first replace the hard constraint that the field values lie in the variety with a penalty






Fig. 8. Convergence behavior of octahedral RTR and the method of Ray et al. [2016] on various models, starting from random initialization and using the geometric Laplacian.

term, yielding an energy of Ginzburg-Landau type,

$$E_\epsilon(q) = \frac{1}{2} \int_\Omega \|\nabla q\|^2 dx + \frac{1}{2\epsilon^2} \int_\Omega \text{dist}(q, \mathcal{V})^2 dx, \quad (15)$$

where \mathcal{V} is our variety. Consider this energy in the limit $\epsilon \rightarrow 0$, as in Section 4 of Viertel and Osting [2019]. The MBO method consists of alternating descent on the two terms of E_ϵ . Gradient flow on the first term—Dirichlet energy—yields componentwise heat diffusion,

$$\begin{aligned} \partial_t q(x, t) &= \Delta q(x, t) \\ q(x, 0) &= q_{k-1}, \end{aligned} \quad (16)$$

with the boundary constraints given in Section 5.1. In practice, we do one step of implicit Euler integration per overall algorithm step. Gradient flow on the second term of (15) in the limit $\epsilon \rightarrow 0$ amounts to projection onto the variety. The overall algorithm is shown in Algorithm 1.

We follow the suggestion of van Gennip et al. [2014] and Viertel and Osting [2019] to set τ_0 relative to the inverse of the smallest nonzero eigenvalue of the Laplacian. In ordinary MBO, $\beta(k) = 1$ —that is, τ does not change over the course of the algorithm. In Figure 9, we test the effects of different values of τ . As we decrease τ , the field is able to relax more and reduce the Dirichlet energy. However, larger values of τ allow more tunneling so that the field can escape local minima. This suggests a modified MBO scheme (mMBO) in which we start with a large τ and progressively reduce it over the course of optimization—analogously to what happens to the temperature in simulated annealing. Our mMBO uses an optimization schedule $\beta(k) = 50k^{-3}$. This optimization schedule starts with a large diffusion time for robustness to random initialization, then sweeps τ through various orders of magnitude quickly. We have found that this heuristic produces a good balance between quick convergence and field quality.

Figure 10 depicts the convergence behavior of RTR, MBO with various τ values, and mMBO starting from a projection of an approximately harmonic field. The energy value achieved by ordinary MBO is limited for each τ ; however, mMBO achieves a lower




Fig. 9. Results of octahedral MBO on a torus. With a relatively large diffusion time τ , MBO produces a field with tightly packed singular regions (left). At a smaller value of τ , singular curves relax toward the boundary, reducing the Dirichlet energy (center). RTR reduces the energy even further (right), pushing the singular curves to the boundary.

ALGORITHM 1: MBO over a variety $\mathcal{V} \subset \mathbb{R}^d$

input: initial $d \times n$ field q_0 , diffusion parameter τ_0 , optimization schedule $\beta(k)$

result: q_k

$k \leftarrow 1$

repeat

$\tau_k \leftarrow \beta(k)\tau_0$

Diffusion step: Solve the linear system $(M - \tau_k L)\bar{q}_k^\top = Mq_{k-1}^\top$ with columns $(\bar{q}_k)_i$ constrained to be in the affine span (11) or (13) for each $i \in \partial V$.

Projection step: Project \bar{q}_k into the variety:

$$(q_k)_i \leftarrow \begin{cases} \Pi_{\mathcal{V}}((\bar{q}_k)_i) & i \notin \partial V \\ \Pi_{\mathcal{V}_{n_i}}((\bar{q}_k)_i) & i \in \partial V. \end{cases}$$

$$\Delta_k \leftarrow \frac{\text{tr}((q_k - q_{k-1})M(q_k - q_{k-1})^\top)}{\text{tr}(q_k M q_k^\top)}$$

$$E_k \leftarrow \text{tr}(q_k L q_k^\top)$$

$$\Delta E_k \leftarrow E_k - E_{k-1}$$

$k \leftarrow k + 1$

until $\Delta E_k / E_k < \delta$ or $\Delta_k < \delta$

value by sweeping through various values of τ . Note that the iteration counts shown do not reflect wall-clock time; in particular, RTR runs at least an order of magnitude faster than the other algorithms.

7 EXPERIMENTS

We initialize our algorithms with vertexwise random octahedral fields, generated by—separately for each vertex—starting with the canonical frame, rotating it by a random angle between 0 and 2π about the positive z -axis, and then rotating the positive z -axis to a random direction. We do this even for optimization of odedo fields to avoid encountering odedo frames that have negative weights λ_i (cf. Conjecture 4.3). When starting from octahedral initialization, the odedo frames we compute always have nonnegative weights.

In Figure 11, we demonstrate the robustness of MBO to random initialization. On the sphere, global rotations of a field do not change the objective value. Initializing odedo MBO with




Fig. 10. Energy convergence on the rockerarm_91k model. MBO’s convergence to lower energies is limited by the diffusion time τ . Decreasing τ according to the optimization schedule used in mMBO achieves the best results overall.




Fig. 11. Robustness to random initialization. Odeco MBO transforms vertexwise random initial fields into qualitatively identical results up to the symmetry of the sphere. All results were computed on a sphere with 48950 vertices. Integral curves are depicted in regions of rapid rotation (i.e., in proximity to singularities).

different fields of vertexwise random frames, we find that it converges to randomly rotated copies of a qualitatively identical, symmetric field.

Figure 12 illustrates the behavior of octahedral and odeco fields as the density of the underlying tetrahedral mesh increases. Due to the unit-norm constraint, the gradient of an octahedral field becomes unbounded close to its singularities. This leads to logarithmic divergence of the total energy as the tet mesh becomes finer.

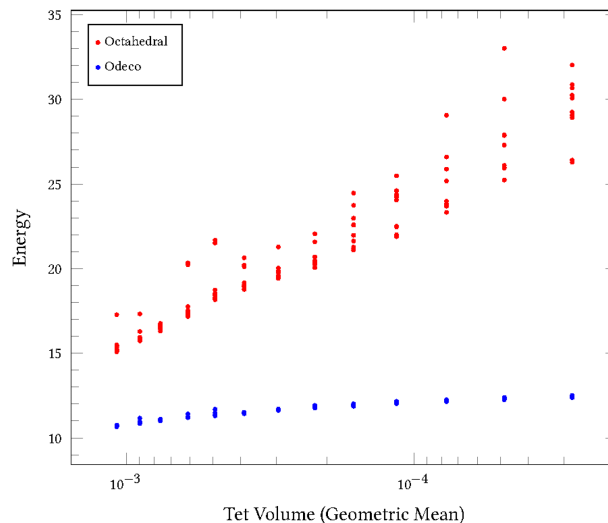


Fig. 12. Energy density diverges for octahedral fields, but plateaus for odeco fields, as mesh density increases. Also notice the higher variance for octahedral fields. Ten fields of each type were computed by RTR on each of 13 tetrahedral meshes of the unit sphere of various densities.




Fig. 13. The energy density at singularities of an octahedral field dominates the total energy (left). In contrast, an odeco field’s energy is distributed more uniformly (right). Results computed using MBO + RTR.

In contrast, the additional scaling degrees of freedom available to odeco fields allow renormalization of singularities, as shown in Figure 13. Thus, odeco field energy plateaus as mesh density increases. Note also the much smaller variance in energy between runs for odeco fields, quantitatively illustrating robustness to random initialization.

Table 1 (in the supplementary material) compares timings and energy values for our methods and the method of Ray et al. [2016], comprising 18 types of experiments on 15 different models, for 270 total runs. Direct comparison of energy values with Ray et al. [2016] is not possible, as their method uses the graph Laplacian and initializes by solving a linear system, whereas our method uses the geometric (finite element) Laplacian and random initialization. To attempt a fair comparison, we report results of their method alone, their method followed by RTR with the geometric Laplacian, and their method substituted with the geometric Laplacian and random initialization. All experiments in the table were conducted on an Ubuntu workstation with a four-core, 3.6-GHz Intel Core i7-7700




Fig. 14. Twisted singular curves on the handle of the cup in the result of Ray et al. [2016] lead to twisting in the resulting mesh (top). Our field yields a mesh with a more regular structure throughout (bottom).




Fig. 15. On the bone model, twisted singular curves in a field computed by the method of Ray et al. [2016] lead to a degenerate integer grid map, producing highly distorted hexahedra (left). Our field yields a mesh without this degeneracy (right).

and 32 GB of RAM. A MATLAB implementation of our algorithms accompanies this article and also includes our implementation of Ray et al. [2016].

The results in Table 1 show that RTR converges very quickly but, like previous work, easily gets stuck in local minima. In contrast, mMBO takes longer to converge but produces higher-quality fields that reflect the symmetries of the volume. Running RTR to polish the results of mMBO produces the best energies.

Despite sometimes getting stuck in local minima (see Figure 6), we observe that the approximation of projection by Ray et al. [2016] can be useful in practice. Table 1 (in the supplementary material) includes experiments with MBO and mMBO (see Section 6.2) substituting the projection of Ray et al. [2016] for the true projection. In most cases, the resulting energy is very similar, suggesting that the diffusion iterations in MBO smooth out any errors resulting from incorrect projection. This hybrid MBO can be a useful tradeoff between correctness and speed.

In Figures 14 through 17, we compare fields computed by octahedral mMBO + RTR to those computed by the method of Ray et al. [2016]. Our fields not only have lower Dirichlet energy but also better singular structures. To visualize singular structures, we use the visualization technique of Liu et al. [2018]. To illustrate the importance of singular structure, we have generated hexahe-




Fig. 16. Due to the twisting of singular curves in the field generated by the method of Ray et al. [2016] (top), the mesh is highly distorted on the right arm. In contrast, our field (bottom) has a regular cubic singular structure, leading to fewer mesh degeneracies.




Fig. 17. Note the simpler, more regular singular curves in our result as compared to the result of Ray et al. [2016]—especially on the bunny’s nose. The degeneracy in their result leads to a collapse of the integer grid map on the head.

dral meshes from both sets of fields, following the methods laid out in Nieser et al. [2011] and Lyon et al. [2016]. The meshes are computed from the raw field data, with no preprocessing other than tetrahedral mesh refinement to resolve and localize singular curves. In particular, we do not “correct” singularities—thus, both sets of meshes show some degeneracies resulting from collapsed or flipped tetrahedra during the parameterization step. However, our fields yield meshes with fewer and smaller degeneracies, less distortion, and more regular structure.

Figures 18 and 19 compare fields produced by mMBO + RTR to fields produced by the method of Gao et al. [2017]. Our fields




Fig. 18. Comparison of octahedral field algorithms on the gear model. A field produced with mMBO + RTR (right) exhibits lower energy and greater symmetry than one produced by the method of Gao et al. [2017] (left).




Fig. 19. Comparison of octahedral field algorithms on the space-filling torus model. A field produced with mMBO + RTR (right) exhibits lower energy and greater symmetry than one produced by the method of Gao et al. [2017] (left).

better respect the symmetries of the underlying models. We note that mMBO + RTR yields a higher-quality result without requiring a multiscale method like the one Gao et al. [2017] employ.

8 DISCUSSION AND FUTURE WORK

A stronger understanding of the unknowns in the volumetric frame field problem enables both theoretical and practical developments. On the theoretical side, our reexamination of the typical representation of octahedral frames not only yields useful geodesic formulas and projection operators but also suggests generalization to odeco frames. For both frame representations, optimization algorithms designed explicitly to traverse the manifold of unknowns yield gains in efficiency and in the quality of computed results.

Our work is intended not only to improve on existing frame field computation algorithms but also to inspire additional research into the structure of spaces of frames and to highlight the significance of this structure to practical aspects of field computation and meshing. Most immediately, a few open questions are left from our discussion. On the theoretical side, conjectures 4.2 and 4.3 remain to be proven. We anticipate that both can be embedded in a more general framework explaining when relaxations of projection problems are exact. On the algorithmic side, RTR seems to get stuck in local minima much more often than MBO, especially on denser




Fig. 20. A field of lower Dirichlet energy (bottom) may still result in more mesh degeneracies than a field of higher Dirichlet energy (top) due to topological impediments to meshability, namely the presence of an additional valence 3–5 junction.

meshes. We hypothesize that this is because RTR strictly moves along the manifold, whereas MBO is free to tunnel through the ambient space. However, RTR converges much faster and yields high-quality fields on coarse meshes. Given these observations, we anticipate that it may be possible to incorporate RTR into a multiscale method that leverages its efficiency while avoiding local minima that appear at the finest scales.

As with most existing frame field computation algorithms, even when mMBO+RTR converges to a smooth field, the topology of the field is not always hex meshable. Figure 20 shows a case where the method of Ray et al. [2016] yields a field of higher Dirichlet energy but leads to fewer degeneracies than our field. In particular, the presence of a singular curve whose type changes from valence 3 to 5 leads to a degeneracy in the integer-grid map. Although our method succeeds at finding fields of lower Dirichlet energy—the objective function of our method and that of Ray et al. [2016]—minimizing this energy is an incomplete proxy for our ultimate goal, namely to obtain smooth, meshable fields. We would like to investigate additional metrics, such as integrability of frame fields, that might make it possible to express meshability constraints rigorously.

To define such additional metrics, it might be fruitful to consider further alternative frame field representations. For example, given our use of Lie algebra representations, a logical next step might be to introduce an $SO(3)$ -principal bundle and work with connections on that bundle as variables. In this theory, quantities such as curvature, torsion, or the Chern-Simons functional might encode important features of frame fields. The discretization of connections in Corman and Crane [2019] represents a promising first step.

Finally, although our algorithms do not explicitly take account of symmetry, we find that fields computed by MBO consistently reproduce the symmetries of the volume. It would be interesting to develop a better theoretical understanding of this behavior and to develop machinery for explicitly promoting conformation to near-symmetry in deformed domains.

These and other challenges appear when extending well-known machinery from geometry processing on surfaces to volumetric domains, as demanded by applications in engineering, simulation,

and other disciplines. Over the longer term, careful consideration of problems like the ones we consider in this article will lead to a versatile and generalizable approach to geometry processing.

APPENDIX

A MATHEMATICAL BACKGROUND

Representation theory and algebraic geometry provide a concise language for the discussion of the preceding octahedral and odedo varieties. We begin with definitions of a few terms used throughout the article. A detailed introduction to either of these subjects is outside the scope of our discussion; however, we refer the reader to Stillwell [2008] for an introduction to matrix Lie groups and to Blekherman et al. [2012] for a primer on real algebraic geometry and optimization.

A.1 Lie Groups and Representations

A *Lie group* is a group G that is also a smooth manifold, and such that the multiplication $\cdot : G \times G \rightarrow G$ and inversion $^{-1} : G \rightarrow G$ are smooth maps. A *matrix Lie group* is a subgroup of the invertible matrices $GL(n) \subset \mathbb{R}^{n \times n}$ that is simultaneously a submanifold. Examples include the orthogonal groups $O(n)$ and the special orthogonal groups $SO(n)$. The primary Lie group of interest to us in this work will be $SO(3)$.

For a matrix Lie group $G \subset \mathbb{R}^{n \times n}$, the *Lie algebra* is a linear subspace of $\mathbb{R}^{n \times n}$ identified with the tangent space to G at the identity I , denoted $\mathfrak{g} = T_I G$. The *inner automorphism* $Ad_g : G \rightarrow G$ taking $h \mapsto ghg^{-1}$ preserves the identity, and its derivative at the identity induces the *Lie bracket* $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$. In the case of a matrix Lie group, this is just the matrix commutator.

As left translation $L_g : h \mapsto gh$ is a diffeomorphism, every element of \mathfrak{g} is also associated to a left-invariant vector field on G . The exponential map $\exp : \mathfrak{g} \rightarrow G$ is given by starting at I and integrating this vector field for one unit of time. For a matrix Lie group, \exp is the ordinary matrix exponential.

A *representation* of a matrix Lie group G is a smooth group homomorphism $\rho : G \rightarrow GL(n)$ for some n . A representation of G induces a representation of each subgroup $H \subset G$. A Lie group representation comes with a *Lie algebra representation* $(D\rho)_I : \mathfrak{g} \rightarrow \mathbb{R}^{n \times n}$, which preserves the Lie bracket. Moreover, representations commute with exponentials:

$$\rho \circ \exp = \exp \circ (D\rho)_I. \quad (17)$$

An *irreducible* representation is one that cannot be decomposed as a direct sum of subrepresentations. The irreducible representations of *compact* Lie groups such as $SO(3)$ are finite dimensional. The *adjoint representation* is the irreducible representation of G on its own Lie algebra given by conjugation:

$$\begin{aligned} Ad : G &\rightarrow \text{Aut}(\mathfrak{g}) \\ Ad(g)a &:= g^{-1}ag. \end{aligned} \quad (18)$$

LEMMA A.1. *Ad commutes with representations of G :*

$$(D\rho)_I(g^{-1}ag) = \rho(g)^{-1}((D\rho)_I a)\rho(g). \quad (19)$$

PROOF. Using (17) and conjugating by g , we obtain

$$\rho(g^{-1} \exp(ta)g) = \rho(g)^{-1} \exp(t(D\rho)_I a)\rho(g). \quad (20)$$

Differentiating in t at $t = 0$ yields the result. \square

The *stabilizer* $\text{stab}(v)$ of a vector v is the subgroup consisting of all elements $g \in G$ that preserve v —that is, such that $\rho(g)v = v$.

The Lie algebra $\mathfrak{so}(3)$ associated to $SO(3)$ consists of the skew-symmetric 3×3 matrices. $\mathfrak{so}(3)$ has a basis consisting of infinitesimal rotations about the three coordinate axes:

$$l_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix} \quad l_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix} \quad l_3 = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Their Lie brackets are

$$[l_i, l_i] = 0 \quad [l_1, l_2] = l_3 \quad [l_2, l_3] = l_1 \quad [l_3, l_1] = l_2. \quad (21)$$

These generators might be familiar as the angular momentum operators from quantum mechanics. Indeed, $SO(3)$ acts on functions on the sphere by rotating them, and its Lie algebra elements act as rotational derivatives. This action decomposes into irreducible representations. A basis for each irreducible representation of $SO(3)$ is provided by *spherical harmonics*; each irreducible representation is called a *band* of spherical harmonics. The real representations are odd-dimensional vector spaces, and the representation matrices $\rho(\cdot)$ are also referred to as Wigner D -Matrices.

$SO(3)$ admits a *bi-invariant Riemannian metric* whose value at the identity is given by

$$\langle u, v \rangle = \frac{1}{2} \sum_{i,j} u_{ij}v_{ij}. \quad (22)$$

Note that the generators l_i form an orthonormal basis under this metric. Bi-invariance means that

$$\langle (DL_g)_I v, (DL_g)_I w \rangle = \langle v, w \rangle = \langle (DR_g)_I v, (DR_g)_I w \rangle, \quad (23)$$

where $v, w \in \mathfrak{so}(3)$ and R_g is right translation by $g \in SO(3)$, defined analogously to left translation. In addition, (23) completely describes the Riemannian metric on $SO(3)$. Under this metric, the matrix exponential is also the Riemannian exponential map at the identity, giving rise to geodesics.

The *octahedral group* O is a discrete subgroup of $SO(3)$ generated by right-angle rotations about the three axes,

$$\left\{ r_i = \exp\left(\frac{\pi}{2} l_i\right) \right\}_{i=1}^3.$$

O has 24 elements comprising all symmetries of the cube or, equivalently, of its dual, the octahedron.

A.2 Semidefinite Relaxations

A *(real) algebraic variety* is the set of solutions of a system of polynomial equations over n real variables:

$$V(p_1, \dots, p_k) = \{x \in \mathbb{R}^n \mid p_i(x) = 0, i = 1, \dots, k\},$$

where p_1, \dots, p_k are polynomials.

If p_1, \dots, p_k are quadratic in x , they may be written as

$$p_i(x) = (1 \quad x^\top) A_i \begin{pmatrix} 1 \\ x \end{pmatrix} = \langle A_i, X \rangle,$$

where A_i are symmetric matrices, $\langle \cdot, \cdot \rangle$ denotes the usual inner product on matrices inducing the Frobenius norm, and

$$X = \begin{pmatrix} 1 \\ x \end{pmatrix} (1 \quad x^\top) = \begin{pmatrix} 1 & x^\top \\ x & xx^\top \end{pmatrix}.$$

A quadratically constrained quadratic program (QCQP) is of the form

$$\arg \min_x \{f(x) \mid x \in V(p_1, \dots, p_k)\},$$

where f, p_1, \dots, p_k are all quadratic. A QCQP may be rewritten in the more illuminating form

$$\begin{aligned} & \arg \min_X \quad \langle C, X \rangle \\ & \text{s.t.} \quad X_{11} = 1 \\ & \quad (A_i, X) = 0, \quad i = 1, \dots, k \\ & \quad X \geq 0 \\ & \quad \text{rank}(X) = 1. \end{aligned} \quad (24)$$

Absent the rank constraint, (24) would be a semidefinite program (SDP). SDPs may be solved in polynomial time by interior point methods, implemented in common optimization software packages like MOSEK (2019). Thus, it is natural to ignore the rank constraint and solve the associated SDP; this technique is called *semidefinite relaxation*. The optimal objective value of the relaxed problem is a lower bound for the globally optimal value of the QCQP (24), and if the recovered solution to the SDP has rank one, the solution is said to be *exact*, as it is the globally optimal solution of (24).

ACKNOWLEDGMENTS

The authors would like to thank Elina Robeva for helping us understand her work on the odeco varieties, Pablo Parrilo and Diego Cifuentes for sharing their insights into semidefinite relaxation, and Nicolas Boumal for helping us get started with manifold optimization. We thank Heng Liu for providing tetrahedral refinement, parameterization, and hexahedral mesh extraction code. We are grateful to Xifeng Gao and Daniele Panozzo for guiding us in using their code from Gao et al. [2017]. Finally, we thank Mirela Ben-Chen, Amit Bermano, Edward Chien, Etienne Corman, Keenan Crane, Fernando de Goes, Steven Gortler, Leif Kobbelt, Braxton Osting, Nir Sharon, Nir Sochen, Amir Vaxman, Josh Vekhter, and Paul Zhang for valuable discussions.

REFERENCES

- P.-A. Absil, C. G. Baker, and K. A. Gallivan. 2007. Trust-region methods on Riemannian manifolds. *Foundations of Computational Mathematics* 7, 3 (July 2007), 303–330. DOI: <https://doi.org/10.1007/s10208-005-0179-9>
- Motilal Agrawal and Larry S. Davis. 2003. Camera calibration using spheres: A semidefinite programming approach. In *Proceedings of the 9th IEEE International Conference on Computer Vision (ICCV'03)*. IEEE, Los Alamitos, CA, 782.
- Cecil G. Armstrong, Harold J. Fogg, Christopher M. Tierney, and Trevor T. Robinson. 2015. Common themes in multi-block structured quad/hex mesh generation. *Procedia Engineering* 124, suppl. C (2015), 70–82.
- Pierre-Alexandre Beaufort, Jonathan Lambrechts, Christophe Geuzaine, and Jean-Francois Remacle. 2019. Quaternionic octahedral fields: SU (2) parameterization of 3D frames. arXiv:1910.06240.
- Pierre-Alexandre Beaufort, Jonathan Lambrechts, François Henrotte, Christophe Geuzaine, and Jean-Francois Remacle. 2017. Computing cross fields: A PDE approach based on the Ginzburg–Landau theory. *Procedia Engineering* 203 (2017), 219–231.
- Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas. 2012. *Semidefinite Optimization and Convex Algebraic Geometry*. SIAM.
- Ada Boralevi, Jan Draisma, Emil Horobeț, and Elina Robeva. 2017. Orthogonal and unitary tensor decomposition from an algebraic perspective. *Israel Journal of Mathematics* 222, 1 (Oct. 2017), 223–260. DOI: <https://doi.org/10.1007/s11856-017-1588-6>
- N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. 2014. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research* 15 (2014), 1455–1459. <http://www.manopt.org>.
- A. Chemin, F. Henrotte, J.-F. Remacle, and J. Van Schaftingen. 2018. Representing three-dimensional cross fields using 4th order tensors. In *International Meshing Roundtable*. Lecture Notes in Computational Science and Engineering, Vol. 127. Springer, 89–108.
- Diego Cifuentes, Sameer Agarwal, Pablo A. Parrilo, and Rekha R. Thomas. 2017. On the local stability of semidefinite relaxations. arXiv:1710.04287.
- Diego Cifuentes, Corey Harris, and Bernd Sturmfels. 2018. The geometry of SDP-exactness in quadratic optimization. arXiv:1804.01796.
- Etienne Corman and Keenan Crane. 2019. Symmetric moving frames. *ACM Transactions on Graphics* 38, 4 (2019), Article 87.
- Xianzhong Fang, Weiwei Xu, Hujun Bao, and Jin Huang. 2016. All-hex meshing using closed-form induced polycube. *ACM Transactions on Graphics* 35, 4 (July 2016), Article 124, 9 pages. DOI: <https://doi.org/10.1145/2897824.2925957>
- Xifeng Gao, Wenzel Jakob, Marco Tarini, and Daniele Panozzo. 2017. Robust hex-dominant mesh generation using field-guided polyhedral agglomeration. *ACM Transactions on Graphics* 36, 4 (July 2017), Article 114, 13 pages.
- Michel X. Goemans and David P. Williamson. 1995. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM* 42, 6 (1995), 1115–1145.
- Dmitry Golovaty, Jose Alberto Montero, and Daniel Spirn. 2019. A variational method for generating n -cross fields using higher-order Q -tensors. arXiv:1909.00922.
- Jin Huang, Yiyong Tong, Hongyu Wei, and Hujun Bao. 2011. Boundary aligned smooth 3D cross-frame field. *ACM Transactions on Graphics* 30, 6 (Dec. 2011), Article 143, 8 pages. DOI: <https://doi.org/10.1145/2070781.2024177>
- Qi-Xing Huang and Leonidas Guibas. 2013. Consistent shape maps via semidefinite programming. In *Proceedings of the Symposium on Geometry Processing*. 177–186.
- Tengfei Jiang, Jin Huang, Yuanzhen Wang, Yiyong Tong, and Hujun Bao. 2014. Frame field singularity correction for automatic hexahedralization. *IEEE Transactions on Visualization & Computer Graphics* 20, 8 (Aug. 2014), 1189–1199.
- Jan Kautz, John Snyder, and Peter-Pike J. Sloan. 2002. Fast arbitrary BRDF shading for low-frequency lighting using spherical harmonics. *Rendering Techniques* 2, 1 (2002), 291–296.
- Itay Kezurer, Shahar Z. Kovalsky, Ronen Basri, and Yaron Lipman. 2015. Tight relaxation of quadratic matching. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 115–128.
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally optimal direction fields. *ACM Transactions on Graphics* 32, 4 (2013), 59.
- Shahar Z. Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. 2014. Controlling singular values with semidefinite programming. *ACM Transactions on Graphics* 33, 4 (2014), 68–71.
- J. Lee. 2012. *Introduction to Smooth Manifolds (2nd ed.)*. Springer, New York, NY. 2012945172
- Yufei Li, Yang Liu, Weiwei Xu, Wenping Wang, and Baining Guo. 2012. All-hex meshing using singularity-restricted field. *ACM Transactions on Graphics* 31, 6 (2012), 177.
- 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>
- Max Lyon, David Bommes, and Leif Kobbelt. 2016. HexEx: Robust hexahedral mesh extraction. *ACM Transactions on Graphics* 35, 4 (July 2016), Article 123, 11 pages.
- Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. 2016. Point registration via efficient convex relaxation. *ACM Transactions on Graphics* 35, 4 (2016), 73.
- N. D. Mermin. 1979. The topological theory of defects in ordered media. *Reviews of Modern Physics* 51, 3 (July 1979), 591–648. DOI: <https://doi.org/10.1103/RevModPhys.51.591>
- Barry Merriman, James Kenyard Bence, and Stanley Osher. 1992. *Diffusion Generated Motion by Mean Curvature*. Department of Mathematics, University of California, Los Angeles.
- John Douglas Moore. 1976. Equivariant embeddings of Riemannian homogeneous spaces. *Indiana University Mathematics Journal* 25, 3 (1976), 271–279. <http://www.jstor.org/stable/24891007>
- MOSEK ApS. 2019. *The MOSEK Fusion API for C++ Manual. Version 9.1*. Retrieved February 21, 2020 from <https://docs.mosek.com/9.1/cxxfusion/index.html>.
- Matthias Nieser, Ulrich Reitebuch, and Konrad Polthier. 2011. CUBECOVER—parameterization of 3d volumes. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1397–1406.
- Braxton Osting and Dong Wang. 2017. A generalized MBO diffusion generated motion for orthogonal matrix-valued fields. arXiv:1711.01365.
- Onur Ozyesil, Amit Singer, and Ronen Basri. 2015. Stable camera motion estimation using convex programming. *SIAM Journal on Imaging Sciences* 8, 2 (2015), 1220–1262.
- Jonathan Palacios, Lawrence Roy, Prashant Kumar, Chen-Yuan Hsu, Weikai Chen, Chongyang Ma, Li-Yi Wei, and Eugene Zhang. 2017. Tensor field design in volumes. *ACM Transactions on Graphics* 36, 6 (Nov. 2016), Article 188, 15 pages. DOI: <https://doi.org/10.1145/3130800.3130844>

- 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>
- E. Robeva. 2016. Orthogonal decomposition of symmetric tensors. *SIAM Journal on Matrix Analysis and Applications* 37, 1 (2016), 86–102. DOI: <https://doi.org/10.1137/140989340>
- Zhongwei Shen, Xianzhong Fang, Xinguo Liu, Hujun Bao, and Jin Huang. 2016. Harmonic functions for rotational symmetry vector fields. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 507–516.
- Amit Singer. 2011. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis* 30, 1 (2011), 20.
- Dmitry Sokolov, Nicolas Ray, Lionel Untereiner, and Bruno Lévy. 2016. Hexahedral-dominant meshing. *ACM Transactions on Graphics* 35, 5 (June 2016), Article 157, 23 pages.
- Justin Solomon, Amir Vaxman, and David Bommes. 2017. Boundary element octahedral fields in volumes. *ACM Transactions on Graphics* 36, 3 (May 2017), Article 114b. DOI: <https://doi.org/10.1145/3065254>
- John Stillwell. 2008. *Naive Lie Theory*. Springer Science & Business Media.
- Yves van Gennip, Nestor Guillen, Braxton Osting, and Andrea L. Bertozzi. 2014. Mean curvature, threshold dynamics, and phase field theory on finite graphs. *Milan Journal of Mathematics* 82, 1 (June 2014), 3–65. DOI: <https://doi.org/10.1007/s00032-014-0216-8>
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional field synthesis, design, and processing. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 545–572.
- R. Viertel and B. Osting. 2019. An approach to quad meshing based on harmonic cross-valued maps and the Ginzburg–Landau theory. *SIAM Journal on Scientific Computing* 41, 1 (2019), A452–A479. DOI: <https://doi.org/10.1137/17M1142703>
- W. Yu, K. Zhang, and X. Li. 2015. Recent algorithms on automatic hexahedral mesh generation. In *Proceedings of the 2015 10th International Conference on Computer Science Education (ICCSE'15)*. 697–702.

Received August 2019; revised December 2019; accepted December 2019