

The Vector Heat Method

NICHOLAS SHARP, YOUSUF SOLIMAN, and KEENAN CRANE, Carnegie Mellon University

This paper describes a method for efficiently computing parallel transport of tangent vectors on curved surfaces, or more generally, any vector-valued data on a curved manifold. More precisely, it extends a vector field defined over any region to the rest of the domain via parallel transport along shortest geodesics. This basic operation enables fast, robust algorithms for extrapolating level set velocities, inverting the exponential map, computing geometric medians and Karcher/Fréchet means of arbitrary distributions, constructing centroidal Voronoi diagrams, and finding consistently ordered landmarks. Rather than evaluate parallel transport by explicitly tracing geodesics, we show that it can be computed via a short-time heat flow involving the *connection Laplacian*. As a result, transport can be achieved by solving three prefactored linear systems, each akin to a standard Poisson problem. To implement the method we need only a discrete connection Laplacian, which we describe for a variety of geometric data structures (point clouds, polygon meshes, *etc.*). We also study the numerical behavior of our method, showing empirically that it converges under refinement, and augment the construction of intrinsic Delaunay triangulations (iDT) so that they can be used in the context of tangent vector field processing.

CCS: • **Mathematics of computing** → **Discretization; Partial differential equations**; • **Computing methodologies** → **Shape analysis**;

Additional Key Words and Phrases: discrete differential geometry, parallel transport, velocity extrapolation, logarithmic map, exponential map, Karcher mean, geometric median

ACM Reference Format:

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. The Vector Heat Method. *ACM Trans. Graph.* 38, 3, Article 00 (June 2019), 19 pages. <https://doi.org/00.0000/0000000.0000000>

1 INTRODUCTION

Given a vector at a point of a curved domain, how do we find the most parallel vector at all other points (as shown in Fig. 1)? This “most parallel” vector field—not typically considered in numerical algorithms—provides a surprisingly valuable starting point for a wide variety of tasks across geometric and scientific computing, from extrapolating level set velocity to computing centers of distributions. To compute this field, one idea is to transport the vector along explicit paths from the source x to all other points y , but even just constructing these paths is already quite expensive (Sec. 2). We instead leverage a little-used relationship between parallel transport and the *vector heat equation*, which describes the diffusion of a given vector field over a time t . As t goes to zero, the diffused field is related to the original one via parallel transport along minimal geodesics, *i.e.*, shortest paths along the curved domain (Sec. 3.4).

Authors’ address: Nicholas Sharp; Yousuf Soliman; Keenan Crane, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

0730-0301/2019/6-ART00

<https://doi.org/00.0000/0000000.0000000>




Fig. 1. Given a vector at a point, the vector heat method computes the most parallel vector at every other point. The method easily generalizes to other data (such as a velocity field along a curve), providing a novel and efficient way to implement fundamental algorithms across geometry and simulation.

The same principle applies not only to point sources, but also to vector fields over curves or other subsets of the domain. Since diffusion equations are expressed in terms of standard Laplace-like operators, we effectively reduce parallel transport tasks to sparse linear systems that are extremely well-studied in scientific computing—and can hence immediately benefit from mature, high-performance solvers. Moreover, since discrete Laplacians are available for a wide variety of shape representations (polygon meshes, point clouds, *etc.*), and generalize to many kinds of vector data (symmetric direction fields, differential forms, *etc.*), we can apply this same strategy to numerous applications. In particular, this paper introduces

- a fast method for computing parallel transport from a given source set (Sec. 4)
- an augmented intrinsic Delaunay algorithm for vector field processing (Sec. 5.4)
- the first method for computing a logarithmic map over the entire surface, rather than in a local patch (Sec. 8.2), and
- the first method for computing true Karcher/Fréchet means and geometric medians on general surfaces (Sec. 8.3).

We also describe how to discretize the connection Laplacian on several different geometric data structures and types of vector data (Sec. 6), and consider a variety of other applications including distance-preserving velocity extrapolation for level set methods, computing geodesic centroidal Voronoi tessellations (GCVT), and finding consistently ordered intrinsic landmarks (Sec. 8).

2 RELATED WORK

Discrete Parallel Transport. Parallel transport has a long history in the discrete setting. One of the earliest ideas, perhaps, is *Schild’s ladder* which approximates parallel transport via short geodesic segments; this technique has proven useful for parallel transport in high-dimensional spaces representing image data [Lorenzi and Pennec 2014], but is not directly related to parallel transport of vectors on discrete surfaces. A more natural predecessor to the type of parallel transport encountered in geometry processing is the simplicial calculus of Regge [1961], largely used for problems in general relativity [Gentle 2002]. On surfaces, this approach essentially amounts to the notion of *discrete connections* studied by Crane et al. [2010]. To discretize our method on triangle meshes, we will instead consider vectors at vertices, building on ideas from Polthier and Schmies [1998] and Knöppel et al. [2013]. Finally, Azencot et al. [2015] explore a spectral approach to parallel transport, though here the goal is different from ours: transporting one vector field along another, rather than transporting vectors along shortest geodesics.

Discrete Geodesics. A seemingly natural solution to our problem is to explicitly transport vectors along geodesic paths—in the case of triangle meshes, one could unfold the triangles along the path and apply a simple translation in the plane (*à la* Polthier and Schmies [1998]). However, finding these paths is not straightforward: one can either compute exact polyhedral geodesics via expensive window-based methods [Mitchell et al. 1987; Chen and Han 1990] that demand sophisticated acceleration schemes [Surazhsky et al. 2005; Ying et al. 2013; Qin et al. 2016]; or trace integral curves of a piecewise linear geodesic distance function [Kimmel and Sethian 1998; Crane et al. 2013b], which may have very different behavior from true geodesics [Tricoche et al. 2000]. Our approach is far simpler: just build Laplace matrices and solve linear systems. It is also more efficient: even if there were no cost associated with computing geodesics, each shortest path on a discrete surface with $O(n)$ elements has length $O(\sqrt{n})$, yielding an overall cost in $O(n^{1.5})$. In our method, the cost is dominated by solving sparse diffusion equations, which has complexity approaching $O(n)$ for both iterative and direct methods [Spielman and Teng 2004; Gillman and Martinsson 2014]; prefactorization can be used to further reduce amortized cost across many different source points or sets (Sec. 7.2). In practice we observe that merely *extracting* paths from a given piecewise constant vector field is more than an order of magnitude slower than executing our entire algorithm (Fig. 2). Moreover, the diffusion-based approach also provides an accurate and reliable solution (Sec. 7.3).




Fig. 2. Even just *tracing* all the geodesics to a given source point (without accounting for the cost of distance computation) is already an order of magnitude more expensive than applying the vector heat method—here, 15x slower.

Relationship to Scalar Heat Method. The original, scalar heat method [Crane et al. 2013b] computes a related, but fundamentally different quantity from the vector heat method: the former computes geodesic distance; the latter computes parallel transport along shortest geodesics. Computationally, these methods share some basic features: rather than directly solve a difficult nonlinear hyperbolic problem (wavefront propagation from a source), they reformulate computation in terms of much easier linear elliptic PDEs (local averaging); all nonlinearity is captured by simple pointwise operations. However, the structure of the vector version is different: unlike the scalar heat method, there is no dependence among linear equations (STEP I–STEP III of Algorithm 1), making error behavior easier to analyze, and providing additional opportunities for acceleration. Moreover, the vector heat method does not require discrete divergence or gradient operators, making it easier to apply to data structures like point clouds, or even (in principle) data on general graphs [El Karoui and Wu 2015].

Connection Laplacians. On flat domains like the plane, vector diffusion amounts to diffusion of individual scalar components. On curved domains things are not so simple: there is typically no global coordinate system, and one must therefore apply a diffusion process that accounts for parallel transport, achieved via the *connection Laplacian* (Sec. 3). Singer and Wu [2012] use a similar process to obtain a *vector diffusion distance*, motivated by tasks in data analysis and machine learning. Lin et al. [2014] likewise consider vector diffusion in the learning context; we leverage a similar technique in Algorithm 2, STEP II, deriving initial conditions that substantially improve accuracy. On triangle meshes, Knöppel et al. [2013, 2015] consider two connection Laplacians: one based on finite elements, and another in the spirit of *discrete exterior calculus* [Desbrun et al. 2006]; we build primarily on the latter. Algorithmically, fast solvers for connection Laplacians are an active area of research [Kyng et al. 2016]; applications built on top of the vector heat method can immediately benefit from new developments in this area.

Applications. Though we postpone detailed background on applications to Sec. 8, it is worth noting that the heat flow approach is the first practical way to compute an accurate *logarithmic map* (sometimes referred to by its inverse, the *exponential map*) over the entire domain rather than just a local patch—Fig. 20 provides a comparison with previous methods, which exhibit significant error over longer distances. This global accuracy in turn yields the first efficient and reliable method for computing *Karcher means* and *geometric medians* on arbitrary surfaces (see especially Fig. 23). Computationally, previous methods are Dijkstra-like and necessitate dynamic branches and different memory access patterns for each source point. In contrast, heat methods execute a fixed and hence highly predictable traversal of a minimal data structure (a matrix factorization). As a result, the constants involved tend to be much smaller—for instance, our log map computation is faster than even the basic method of Schmidt et al. [2006]. More broadly, tasks that depend on global integration of information (such as computing means or landmarks) benefit from the robust global nature of our algorithm.

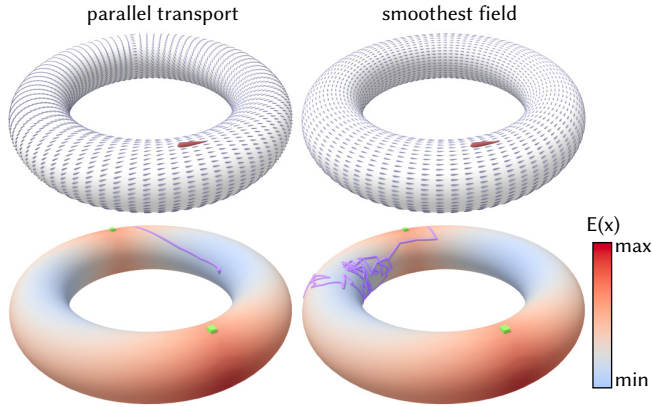


Fig. 3. The field obtained by parallel transport along shortest geodesics plays a special role in several basic geometric algorithms—replacing it with a field that is merely *smooth* can cause these algorithms to fail. Here for instance we use the field to locate the *Karcher mean* equidistant from the two green points (*left*). Attempting to instead use the smoothest possible field causes this search to fail, wandering randomly over the surface (*right*).

3 PRELIMINARIES

The basic idea of our method is to approximate parallel transport via short-time diffusion of vector-valued data. In the Euclidean setting, one can simply diffuse individual scalar components via the ordinary heat equation, but on curved domains this approach fails, since for vectors in different tangent spaces equality of coordinates has no geometric significance. We instead consider a particular vector heat equation which, for a short time t , keeps vectors parallel. We first provide some basic notation and definitions.

3.1 Notation

Throughout we consider a Riemannian manifold M with metric g . We use $d(x, y)$ to denote the corresponding geodesic distance, *i.e.*, the length of the shortest path between any two points $x, y \in M$. The *cut locus* of any subset $\Omega \subseteq M$ is the set of all points $p \in M$ for which there is not a unique closest point $q \in \Omega$ (Fig. 19, *bottom*). For a vector field X on M , we use $X|_p$ to denote the vector at a point p . We will use $i \in \mathbb{C}$ to denote the imaginary unit, *i.e.*, $i^2 = -1$. Finally, we use δ_x to denote the Dirac delta centered at $x \in M$.

3.2 Heat Diffusion

The most basic diffusion equation is the scalar *heat equation*, which describes how an initial heat distribution $\phi_0 : M \rightarrow \mathbb{R}$ looks after being diffused for a time $t > 0$:

$$\frac{d}{dt} \phi_t = \Delta \phi_t. \quad (1)$$

The operator Δ is the (negative semidefinite) Laplace-Beltrami operator on M ; in Euclidean \mathbb{R}^n , Δ is just the usual Laplacian.

Heat Kernel. When the initial heat distribution ϕ_0 is just a spike δ_x at a single point x , the solution to Eqn. 1 is referred to as the *heat kernel* k_t . The heat kernel is the *fundamental solution*, in the sense

that convolution of the initial data ϕ_0 with k_t yields the solution to the heat equation at time t . When the domain is Euclidean (*i.e.*, $M = \mathbb{R}^n$), this fundamental solution is just a Gaussian of constant total mass, centered at a point $x \in \mathbb{R}^n$:

$$G_t(x, y) := \frac{1}{(4\pi t)^{n/2}} e^{-d(x, y)^2/4t}. \quad (2)$$


Here n denotes the dimension of the domain (*e.g.*, $n = 2$ for the plane). More generally, the heat kernel has the asymptotic expansion

$$k_t(x, y) \sim \frac{e^{-d(x, y)^2/4t}}{(4\pi t)^{n/2}} j(x, y)^{-1/2} \left(1 + \sum_{i=1}^{\infty} t^i \Phi_i(x, y) \right). \quad (3)$$

For our purposes the definition of the functions j and Φ_i will not be important, especially since we consider the limit as $t \rightarrow 0$ (see [Berline et al. 1992, Theorem 2.30] for further discussion). In practice, we obtain a numerical approximation of k_t by solving Eqn. 1 directly, *i.e.*, by placing a Dirac delta at a source point x and “smearing” it out via heat diffusion.

3.3 Parallel Transport and Connections

Given a tangent vector X at a point p of a curved surface M , which vector at another point q should be considered “parallel?” If we have a smooth curve $\gamma(s)$ going from p to q , one reasonable idea is that X should experience no “unnecessary turning,” *i.e.*, no change along the tangent direction $T := \frac{d}{ds} \gamma$; the vector we obtain at the end of the path is called the *parallel transport* of X along γ , which we will denote by $P_\gamma(X)$. An important fact about parallel transport is that it is *path dependent*, *i.e.*, for two different curves γ_1, γ_2 from p to q , it is not necessarily true that $P_{\gamma_1}(X) = P_{\gamma_2}(X)$. A good example is transporting a vector from the north to the south pole of the Earth along two different lines of longitude: at the south pole, the angle between the resulting vectors will be related to the difference in longitudes (see inset, *top*). However, as q gets closer and closer to p , only the outgoing direction of the path matters, since very short segments of paths with the same tangent become indistinguishable. We can therefore use parallel transport to define the derivative of one vector field Y along another vector field Z . In particular, at any point $p \in M$ the *covariant derivative* $\nabla_Z Y$, describes the change in Y as we travel an infinitesimally short distance along *any* curve γ with tangent Z at p . More formally, letting $p = \gamma(0)$ and $q = \gamma(s)$, $\nabla_Z Y|_p := \lim_{s \rightarrow 0} (P_{q \rightarrow p}(Y|_q) - Y|_p)/s$, where $P_{q \rightarrow p}$ denotes parallel transport from q back to p (see inset, *bottom*). The operator ∇ is referred to as the *Levi-Civita connection*.



3.4 Connection Laplacian

The *connection Laplacian* Δ^∇ is a second derivative on vector fields with many of the same basic properties as the ordinary Laplacian Δ : it is negative semidefinite, self-adjoint, and elliptic. Just as the ordinary negative semidefinite Laplacian can be expressed

as the trace of the Hessian, or as the divergence of the gradient ($\Delta f = \text{tr}(H(f)) = \text{div} \circ \text{grad} f$), the connection Laplacian associated with a connection ∇ is given by the trace of the second covariant derivative, or by the composition of the covariant derivative with its adjoint ($\Delta^\nabla X = \text{tr}(\nabla^2 X) = -\nabla^* \nabla X$). Some intuition can be obtained by relating the connection Laplacian to the *vector heat equation*

$$\frac{d}{dt} X_t = \Delta^\nabla X_t. \quad (4)$$

Intuitively, the evolution of the vector field X_t over time will look like a “smearing out” of an initial vector field X (Fig. 4). We can make this statement more precise by considering the associated heat kernel $k_t^\nabla(x, y)$, which describes how a vector at a single point x will diffuse to all other points y over time t . For points y that are not on the cut locus of x , this kernel has the asymptotic expansion

$$k_t^\nabla(x, y) \sim \frac{e^{-d(x,y)^2/4t}}{(4\pi t)^{n/2}} j(x, y)^{-1/2} \left(\sum_{i=0}^{\infty} t^i \Psi_i(x, y) \right), \quad (5)$$

where the functions Φ_i from the scalar heat kernel have been replaced by maps Ψ_i taking vectors at x to vectors at y . Most importantly, the first function in this series is given by

$$\Psi_0(x, y) = P_{\gamma_{x \rightarrow y}}, \quad (6)$$

where $\gamma_{x \rightarrow y}$ is the shortest curve from x to y , *i.e.*, the *shortest geodesic* [Berline et al. 1992, Theorem 2.30]. In other words, as $t \rightarrow 0$, the vector heat kernel behaves like parallel transport along shortest paths, along with a decay in magnitude that is identical to the decay of the scalar heat kernel. (As a side note: for $t \rightarrow \infty$, X_t approaches the smoothest possible vector field—independent of initial conditions—since the vector heat equation corresponds to gradient descent on the vector Dirichlet energy; normalizing this field would yield the optimal direction field in the sense of Knöppel et al. [2013].)

Note that not all vector diffusion equations yield the same behavior: for instance, a vector diffusion equation formulated in terms of the *Hodge Laplace operator* (discussed in Sec. 6.1.1) will exhibit different behavior with respect to parallel transport. The discrete picture also provides some useful intuition for the connection Laplacian—see Sec. 5.3.




Fig. 4. Similar to the scalar heat kernel k_t (top), the vector heat kernel k_t^∇ (bottom) “smears out” vectors over time. In the flat Euclidean case one can simply diffuse each scalar component independently, but on a curved domain the *connection Laplacian* is needed to diffuse vectors from one tangent space to another. (Figures not to scale.)




Fig. 5. Numerically, just taking the quotient of short-time solutions to the vector and scalar heat equations yields a poor approximation of magnitude (left), which in this case should be piecewise constant. By separately computing magnitudes (right), we obtain a more accurate solution.

4 SMOOTH FORMULATION

The relationship between the vector heat kernel and the parallel transport map (Eqn. 6) is critical to our method, since it allows us to compute parallel transport by solving diffusion problems—which in turn amount to easy linear systems. For instance, to transport a single unit vector to the rest of the surface, one could simply compute the vector heat kernel for small time t , then normalize the resulting vectors. In the general case, this strategy will not work: consider three vectors of different magnitudes, or a vector field of varying magnitude along a curve (Fig. 8). One way to account for this varying magnitude is to observe that the scalar heat kernel k_t (Eqn. 3) and the vector heat kernel k_t^∇ (Eqn. 5) have identical leading coefficients. Therefore, as $t \rightarrow 0$ the higher-order terms vanish and we can recover the parallel transport map as a simple quotient:

$$\lim_{t \rightarrow 0} \frac{k_t^\nabla(x, y)}{k_t(x, y)} = P_{\gamma_{x \rightarrow y}}. \quad (7)$$

More generally, suppose we diffuse a given vector field X supported (*i.e.*, nonzero) on a set Ω , and diffuse the corresponding scalar indicator function $\phi_0 = \mathbb{1}_\Omega$ (formally, a Hausdorff measure of appropriate dimension). Since diffusion is equivalent to convolution with the heat kernel, these quantities approach the same magnitude at each point, *i.e.*,

$$\lim_{t \rightarrow 0} |X_t| - \phi_t = 0.$$

Hence, the quotient X_t/ϕ_t should exactly factor out any decay in magnitude, leaving only the result of parallel transport along shortest geodesics. Numerically, however, the situation is not so simple: even for fairly small values of t , diffused vectors pointing in different directions will yield small cancellation errors, further reducing the magnitude of the numerator X_t (Fig. 5, left). To get reliable numerical results we will need to consider an alternative approach: use scalar diffusion to obtain the magnitude of the transported vectors (Sec. 4.1); use vector diffusion to obtain their direction (Sec. 4.2). Together these operations define our basic algorithm (Algorithm 1), though nothing restricts the method to two dimensional surfaces, nor to the tangent bundle: everything we state in the smooth setting immediately applies to any vector bundle over a Riemannian manifold of any dimension, as we will discuss in Sec. 6.




Fig. 6. A simple example of interpolation by diffusion: a Gaussian weighted combination of values at two points p_1, p_2 , yields a closest point interpolation \bar{u}_t as the Gaussian width t goes to zero. Here u_t is the linear combination prior to normalization.

4.1 Scalar Interpolation

Suppose we have a pair of source points $p_1, p_2 \in \mathbb{R}^2$, with associated values $u_1, u_2 \in \mathbb{R}$. How can we find a function \bar{u} over the rest of the plane whose value at each point is equal to the value at the closest source p_i ? For this particular example the answer is obvious (just find the line separating p_1 and p_2), but we can obtain it in an interesting way that will naturally generalize. Suppose we use the Gaussian kernel (Eqn. 2) to define a weighted average

$$\bar{u}_t = \frac{u_1 G_{t,p_1} + u_2 G_{t,p_2}}{G_{t,p_1} + G_{t,p_2}}.$$

As t goes to zero, this weighted average provides a closest point interpolation (Fig. 6, *bottom*), since for points closer to p_1 than p_2 , the numerator $u_t := u_1 G_{p_1,t} + u_2 G_{p_2,t}$ is dominated by the first term, and vice versa (Fig. 6, *top*).

This basic idea is easily generalized to curved domains: interpolation is again achieved by dividing a weighted sum by the sum of weights, except that we replace the Gaussian kernel G_t with the scalar heat kernel k_t (Eqn. 3). In particular, given a collection of sources $p_1, \dots, p_n \in M$ and associated values $u_1, \dots, u_n \in \mathbb{R}$, we solve two independent heat equations for functions u and ϕ , using initial conditions

$$\begin{aligned} u_0 &= \sum_{i=1}^n u_i \delta_{p_i}, \\ \phi_0 &= \sum_{i=1}^n \delta_{p_i}. \end{aligned}$$

The interpolant is then simply the limit as t goes to zero of the normalized function

$$\bar{u}(t) := \frac{u(t)}{\phi(t)}.$$

The intuition is the same as in the planar case: for points closest to p_i , the weighted sum will be dominated by the u_i term. Points exactly on the cut locus will approach an average of values; though a precise analysis of this behavior becomes more difficult [Grigor'yan 2009], in practice these values are well-behaved.

More generally, the source set can be any subset $\Omega \subset M$ —on a surface, for instance, Ω can be a collection of points, curves, and regions (see for instance Fig. 7). In this case, the initial conditions are essentially a Dirac-type measure concentrated on Ω (or more formally, a sum of Hausdorff measures of appropriate dimension); in the discrete setting we can integrate basis functions with respect to this measure to obtain initial conditions (as in App. A).

ALGORITHM 1: Vector Heat Method

Input: A vector field X supported on a subset $\Omega \subset M$ of the domain M .

Output: A vector field \bar{X} on all of M .

- I. Integrate the vector heat flow $\frac{d}{dt} Y_t = \Delta^\nabla Y_t$ for time t , with $Y_0 = X$.
 - II. Integrate the scalar heat flow $\frac{d}{dt} u_t = \Delta u_t$ for time t , with $u_0 = |X|$.
 - III. Integrate the scalar heat flow $\frac{d}{dt} \phi_t = \Delta \phi_t$ for time t , with $\phi_0 = \mathbb{1}_\Omega$.
 - IV. Evaluate the vector field $\bar{X}_t = u_t Y_t / \phi_t |Y_t|$.
-

4.2 Vector Heat Method

We now define our main algorithm, the **vector heat method**, which is summarized in Algorithm 1. The basic idea is to first diffuse a given vector field via the vector heat equation (Eqn. 4). For small time t , the resulting vectors will have essentially the right direction, but the wrong magnitude. To obtain the right magnitude, we interpolate the magnitudes of the source vectors (as in Sec. 4.1), and scale the normalized vectors by these magnitudes. The result is a field where the vector at each point q closely approximates the parallel vector at the closest point p . More precisely, for any given vector field X supported on a subset $\Omega \subset M$ of the domain M , we obtain a vector field \bar{X}_t such that at each point $q \in M$ not in the cut locus of Ω ,

$$\lim_{t \rightarrow 0} \bar{X}_t|_q = P_{Y_{p \rightarrow q}} X|_p,$$

where p is the point of Ω closest to q , and $\gamma_{p \rightarrow q}$ is the shortest geodesic from p to q . In practice we cannot evaluate the limit field directly; instead, we pick a small time step $t > 0$ (as detailed in Sec. 7.3) and solve the vector diffusion equation

$$\frac{d}{dt} Y_t = \Delta^\nabla Y_t \quad \text{with} \quad Y_0 = X,$$

to obtain a diffused vector field Y_t . For a small time $t > 0$, each vector $Y_t|_q$ will closely match the *direction* of the closest source vector $X|_p$, but will have the wrong *magnitude*, i.e., $|Y_t|_q \neq |X|_p$. To get the right magnitude, we solve the two scalar diffusion equations

$$\begin{aligned} \frac{d}{dt} u_t &= \Delta u_t & \text{with} & \quad u_0 = |X|, & \quad \text{and} \\ \frac{d}{dt} \phi_t &= \Delta \phi_t & \text{with} & \quad \phi_0 = \mathbb{1}_\Omega. \end{aligned}$$

The quotient $\bar{u}_t := u_t / \phi_t$ then gives the magnitude of the vector at the closest point, and the final vector field is hence just

$$\bar{X}_t = \bar{u}_t Y_t / |Y_t|.$$

This strategy resembles the scalar heat method, where one cannot simply apply Varadhan's formula, but must instead normalize the gradient to obtain the correct magnitude. Likewise, in the vector heat method we cannot simply divide the vector heat kernel by the scalar heat kernel (as in Eqn. 7), but must carefully interpolate magnitudes. This strategy provides numerical robustness: even if there are small errors in direction, the magnitudes are essentially perfect (Fig. 5, *right*).

Note that for points on the cut locus, \bar{X}_t essentially approaches an average of all closest vectors (see for instance Ludewig [2018]). Importantly, we have no particular interest in approximating the cut locus itself; the presence of a cut locus is merely a natural feature of any globally accurate approximation of the true (smooth) solution to our problem. Global accuracy turns out to be essential for applications—see for instance Fig. 3 and discussion in Sec. 8.




Fig. 7. Interpolation of scalar values at points (left) and curves (right).

5 DISCRETE FORMULATION

Fundamentally, the vector heat method is an algorithm formulated in the smooth setting—so far we have not assumed that we work with any particular discretization (such as point clouds or polygon meshes). In this section we discretize the method on triangle meshes; other possibilities are explored in Sec. 6.2.

5.1 Discrete Surface

Topology. Throughout we consider a manifold triangle mesh $K = (V, E, F)$, with or without boundary. In principle our method applies to nonorientable domains, but to simplify exposition (and implementation) it will be easier to assume that K is oriented. We use tuples of vertex indices to specify simplices—for instance, ijk is a triangle with vertices $i, j, k \in V$. Indices appearing on both sides of an equation are held fixed in sums, for instance, $a_{ij} = \sum_{ijk \in F} b_{ijk}$ denotes a sum over only those triangles $ijk \in F$ containing edge ij .

Geometry. The only geometric information we need to formulate our algorithm is positive edge lengths $\ell : E \rightarrow \mathbb{R}_{>0}$ satisfying the triangle inequality in each face; from this data one can easily determine the area A_{ijk} of each triangle, and the interior angle θ_i^{jk} at each corner i of each triangle ijk (via Heron’s formula and the law of cosines, *resp.*). For problems involving tangent vector fields, this purely *intrinsic* point of view has some attractive consequences—in particular, it enables us to talk about tangent vector fields on an *intrinsic Delaunay triangulation* (Sec. 5.4), which in practice can significantly improve accuracy and reliability (Fig. 10).

5.2 Intrinsic Tangent Spaces

At each vertex $i \in V$, we encode tangent vectors X_i in local polar coordinates (r_i, φ_i) , à la Knöppel et al. [2013]. Conceptually, one can imagine isometrically mapping a small neighborhood of the vertex onto a circular cone whose base has a radius $r = 1$ (see inset); the direction of any tangent vector can then be expressed as an angle $\varphi \in [0, 2\pi)$, equal to the arc length along the cone boundary. Concretely, we pick a canonical reference edge ij_0 to represent the direction $\varphi = 0$; all other directions are expressed as a counter-clockwise rotation relative to this edge. In particular, letting

$$\Theta_i := \sum_{ijk \in F} \theta_i^{jk}$$





Fig. 8. Parallel transport of vectors from points (left) and curves (right).

be the total interior angle at vertex i , we define normalized angles

$$\tilde{\theta}_i^{jk} := 2\pi\theta_i^{jk}/\Theta_i, \quad (8)$$

which sum to 2π . The direction of the outgoing edges ij_0, ij_1, \dots (in counter-clockwise order) are then given by the cumulative sums

$$\varphi_{ija} := \sum_{p=0}^{a-1} \tilde{\theta}_i^{j_p, j_{p+1}}. \quad (9)$$

A tangent vector X_i at any vertex i is specified by an angle φ_i and magnitude r_i in this coordinate system. In practice, we will encode this data as a complex number $r_i e^{i\varphi_i} \in \mathbb{C}$.


Discrete Parallel Transport. For any edge $ij \in E$, the angles φ_{ij} and $\pi + \varphi_{ji}$ encode the edge direction relative to the coordinate systems at vertices i and j , *resp.* To keep a vector parallel as we go from i to j , we must therefore rotate by the angle

$$\rho_{ij} := (\varphi_{ji} + \pi) - \varphi_{ij}.$$

We encode the corresponding rotations as unit complex numbers

$$r_{ij} := e^{i\rho_{ij}}, \quad (10)$$

which will help to construct our discrete connection Laplacian.



5.3 Discrete Laplace Operators


For triangle meshes, the Laplace-Beltrami operator Δ can be discretized as a weighted graph Laplacian $L \in \mathbb{R}^{|V| \times |V|}$, given by

$$(L\phi)_i = \frac{1}{2} \sum_{ij \in E} \underbrace{(\cot \theta_k^{ij} + \cot \theta_l^{ji})}_{=: w_{ij}} (\phi_j - \phi_i)$$

at each vertex $i \in V$, where k, l denote the vertices opposite edge ij ; the *cotan weights* w_{ij} simply account for the shape of the triangles (see Crane et al. [2013a, Chapter 6]). To be concrete, let $a := \cot \theta_i^{jk}$, $b := \cot \theta_j^{ki}$ and $c := \cot \theta_k^{ij}$ be the cotangents of the angles at the corners of a triangle $ijk \in F$. One way to build L is to accumulate, for each triangle, the local 3×3 matrix

$$-\frac{1}{2} \begin{bmatrix} b+c & -c & -b \\ -c & c+a & -a \\ -b & -a & a+b \end{bmatrix}$$

into the corresponding entries of L .



The connection Laplacian Δ^∇ is given by a nearly identical *complex* matrix $L^\nabla \in \mathbb{C}^{|V| \times |V|}$; the only change is that the off-diagonal entries are multiplied by the edge rotations r_{ij} :

$$-\frac{1}{2} \begin{bmatrix} b+c & -cr_{ij} & -br_{ik} \\ -cr_{ji} & c+a & -ar_{jk} \\ -br_{ki} & -ar_{kj} & a+b \end{bmatrix}.$$

This matrix is Hermitian since r_{ij} and r_{ji} are unit complex numbers encoding equal and opposite rotations; hence, $r_{ji} = r_{ij}^{-1} = \bar{r}_{ij}$. This matrix naturally arises as the Hessian of the vector Dirichlet energy $\sum_{ij \in E} w_{ij} |X_j - r_{ij} X_i|^2$, which quantifies the “straightness” of a given vector field X (see Knöppel et al. [2015, Section 3.2]). Both L and L^∇ effectively encode zero Neumann boundary conditions; zero Dirichlet conditions will yield similar results (see discussion in Crane et al. [2013b, Section 3.4]). We also have a diagonal $|V| \times |V|$ *lumped mass matrix* with entries

$$M_{ii} = \frac{1}{3} \sum_{ijk} A_{ijk}.$$

This matrix is either real or complex depending on whether we are building the scalar or vector heat equation (*resp.*). In particular, we apply a one-step backward Euler approximation to our short-time heat equations (STEPS I–III of Algorithm 1) to obtain

$$\begin{aligned} (M - tL^\nabla)Y &= Y_0, \\ (M - tL)u &= u_0, \\ (M - tL)\phi &= \phi_0. \end{aligned}$$

Here the vectors $Y_0 \in \mathbb{C}^{|V|}$, $u_0, \phi_0 \in \mathbb{R}^{|V|}$ describe the source data. For instance, if the source is a collection of points $i_1, \dots, i_k \in V$ with associated vectors X_{i_1}, \dots, X_{i_k} , then

$$Y_0 = \sum_{i=1}^k X_k \delta_k, \quad u_0 = \sum_{i=1}^k |X_k| \delta_k, \quad \phi_0 = \sum_{i=1}^k \delta_k,$$

where δ_k is the Kronecker delta at vertex k . The final result is obtained by evaluating $u_i Y_i / \phi_i |Y_i|$ at each vertex $i \in V$ (STEP IV).





Fig. 9. *Left*: the intrinsic Delaunay algorithm performs edge flips to obtain a mesh where edge angle sums are no greater than π . *Right*: we augment this algorithm to update the local tangent spaces on each flip.

5.4 Tangent Intrinsic Delaunay

Although this discretization already works quite well, we can improve robustness and accuracy by building the matrices L , L^∇ , and M with respect to the *intrinsic Delaunay triangulation* of the given mesh [Bobenko and Springborn 2007]. The basic idea is to construct a different triangulation of the same piecewise Euclidean surface (*i.e.*, without changing the geometry) which leads to better numerical behavior for operators like the Laplacian. Conceptually, edges may cross multiple faces of the original mesh (see inset), but in practice we need only keep track of the usual connectivity information: which edges are connected to which vertices? Since the vertex set is preserved, any solution computed on the Delaunay mesh can be directly copied back to the vertices of the original mesh. However, we will need to augment this construction to keep track of tangent-valued data.



Let k, l denote the vertices opposite a given edge ij . The basic intrinsic Delaunay algorithm iteratively flips any edge ij where the angle sum $\theta_k^{ij} + \theta_l^{ji}$ is greater than π , or equivalently, where the cotan weight w_{ij} is negative. After a flip, edge kl is assigned a new length ℓ_{kl} , equal to the distance between k and l along the previous triangles ijk and jil . This length can be computed via the law of cosines—see Fisher et al. [2006] for further details. We make a small but important modification to this algorithm: after each flip, we also compute the angles φ encoding the outgoing direction of the flipped edge kl relative to its endpoints. In particular, we set

$$\begin{aligned} \varphi_{lk} &\leftarrow \varphi_{lj} + 2\pi\theta_l^{jk} / \Theta_l, \\ \varphi_{kl} &\leftarrow \varphi_{ki} + 2\pi\theta_k^{il} / \Theta_k. \end{aligned}$$

In other words we add the (normalized) angle between the preceding edge and the new edge to the angle of the preceding edge; here the angles $\theta_l^{jk}, \theta_k^{il}$ can be computed from the updated lengths. This way, we preserve the local polar coordinate systems throughout the flipping process, and therefore know how to map tangent data computed on the intrinsic Delaunay triangulation back to the original mesh: simply copy the polar coordinates (r_i, φ_i) .

This procedure is useful not only for our algorithm, but any algorithm that seeks to improve the quality and reliability of tangent vector field processing without altering the geometry of the input mesh. For instance, the discrete connection Laplacian L^∇ may be indefinite since (unlike the cotan Laplacian L) it cannot simply be interpreted as the restriction of the smooth connection Laplacian Δ^∇ to the subspace of piecewise linear functions. The intrinsic Delaunay condition ensures that L^∇ is semidefinite, since the intrinsic cotan weights w_{ij} are guaranteed to be nonnegative. It also ensures that no vector X_i can be “flipped,” in the sense that it will always be a positive linear combination of the neighbors $r_{ji} X_j$ (and hence in their convex cone). In practice the intrinsic Delaunay condition is not strictly necessary to obtain high-quality results, but helps to guarantee that problems will not occur, even on pathological inputs (Fig. 10). Formally understanding further properties of the intrinsic connection Laplacian and associated objects (*e.g.*, special intrinsic vector fields) is an enticing question for future work.

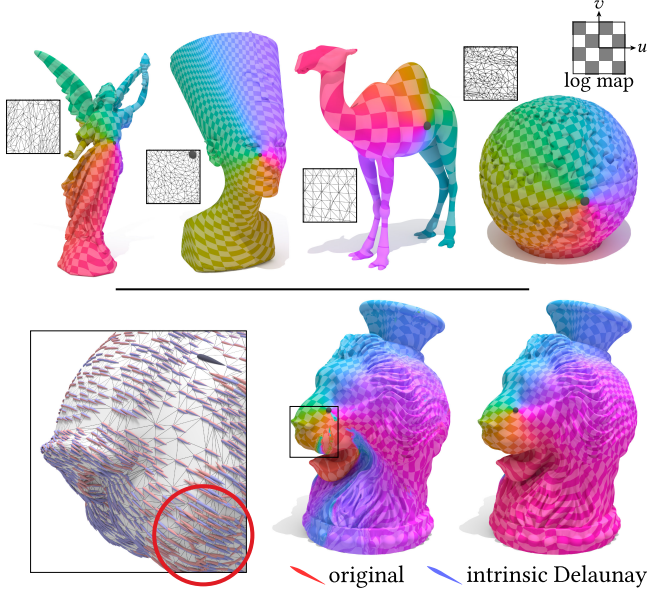


Fig. 10. We often obtain high-quality results even on non-Delaunay meshes (top). Occasionally, however, transported fields can have improperly oriented vectors (inset), here causing errors in the log map (bottom center). By keeping track of tangent spaces during intrinsic Delaunay flips, we obtain a high-quality solution (bottom right) without having to change the input geometry.

6 GENERALIZATIONS

The vector heat method easily generalizes to other domains and other kinds of vector-valued data (Sec. 6.1), and can easily be implemented on many data structures beyond triangle meshes (Sec. 6.2).

6.1 Other Vector Bundles

As stated, the algorithm described in Sec. 4 already applies to any *vector bundle*. Loosely speaking, a vector bundle is a manifold M with a copy of the same vector space V at each point. A choice of vector at each point of M is called a *section* of the bundle—for instance, a tangent vector field X is a section of the tangent bundle TM . As hinted at in Sec. 3.3, the connection ∇ defines what it means for nearby vectors to be parallel. In general we may want to change the domain (i.e., the choice of manifold M), the type of vector data (i.e., the choice of vector space V), or the notion of what it means for vectors to be parallel (i.e., the choice of connection ∇). These choices ultimately determine the operators Δ and Δ^∇ , which are all we need to formulate Algorithm 1. An elementary example is the *trivial real line bundle*, where the vector space is just $V = \mathbb{R}$, i.e., sections are just real-valued functions, and parallel transport simply copies values from one point to another. In this case the vector heat method reduces to the scalar interpolation scheme described in Sec. 4.1—some more interesting examples are given below.

6.1.1 Differential 1-Forms. One vector bundle common in applications is the *cotangent bundle* T^*M , whose sections are called *differential 1-forms*. In this context, it is often easiest to discretize the Hodge Laplacian $\Delta_1 := d\delta + \delta d$. To obtain a corresponding

connection Laplacian Δ^∇_1 , we will employ the *Weitzenböck identity*, which for a smooth surface M with Gaussian curvature K says that

$$\Delta^\nabla_1 \alpha = \Delta_1 \alpha + \frac{1}{2} K \alpha$$

for any 1-form α . We can therefore obtain a discrete connection Laplacian by adding a matrix representing the curvature term to an existing discretization of the Hodge Laplacian. In particular, let

$$\Omega_{ijk} := \pi - (\tilde{\theta}_i^{jk} + \tilde{\theta}_j^{ki} + \tilde{\theta}_k^{ij})$$

be the total Gaussian curvature of triangle ijk (see Sharp and Crane [2018, Section 5.2]). Similarly, let

$$\Omega_{ij} := \frac{1}{3}(\Omega_{ijk} + \Omega_{jil})$$


be the total curvature in the barycentric diamond around an edge ij with opposite vertices $k, l \in V$, which has area

$$A_{ij} := \frac{1}{3}(A_{ijk} + A_{jil}).$$

One discrete Hodge Laplacian

is provided by *discrete exterior calculus* [Desbrun et al. 2006], defined in terms of simplicial boundary/coboundary operators, in conjunction with diagonal mass matrices \star_k for differential k -forms. One can likewise add a diagonal matrix with entries Ω_{ij}/A_{ij} to approximate the connection Laplacian—unfortunately, this approach does not yield the correct result (inset, top); seemingly, the highly local region of support is not sufficient to capture transport. We therefore replace \star_1 with an approximate Galerkin mass matrix \star_1^B obtained via one point barycentric quadrature [Mohamed et al. 2016], which has larger support and appears to capture sufficient directional information. For any two edges ij and jk contained in a common face $ijk \in F$, this matrix has a nonzero entry

$$(\star_1^B)_{ij,jk} := \frac{|e_{ij}^*| \cos(\xi)}{|e_{jk}^*| \cos(\zeta)},$$




where ξ is the angle between the barycentric dual edges e_{ij}^* and e_{jk}^* (i.e., the vectors from the edge midpoints to the face barycenter), and ζ is the angle between the normal n of edge jk and the dual edge e_{jk}^* (see inset figure). To discretize the term $\frac{1}{2}K$ we then build a matrix $K \in \mathbb{R}^{|E| \times |E|}$ with diagonal entries

$$K_{ij,ij} = \frac{\Omega_{ij}}{2A_{ij}} (\star_1^B)_{ij,ij}$$

for each edge $ij \in E$, and off-diagonal entries

$$K_{ij,jk} = \frac{\Omega_{ijk}}{2A_{ijk}} (\star_1^B)_{ij,jk}$$

for all pairs of edges $ij, jk \in E$ that share a triangle. When used in Algorithm 1, this discretization appears to yield the correct solution—for instance, the solution on the sphere above closely matches the analytical solution (compare with Fig. 16, bottom).



6.1.2 Symmetric Direction Fields.

Another important example of vector-valued data in computer graphics and geometry processing are *symmetric direction fields* such as line fields, cross fields, *etc.* [Vaxman et al. 2016], which play a key role in applications like surface shading [Knöppel et al. 2015] and

quadrilateral remeshing [Bommes et al. 2013]. Formally, such fields are sections of the k th tensor power $TM^{\otimes k}$ of the tangent bundle, where k determines the degree of symmetry. In this setting, one can build the connection Laplacian exactly as described in Knöppel et al. [2013]—in particular, all one has to do is raise the coefficients r_{ij} from Eqn. 10 to the k th power, and apply Algorithm 1 as usual. The final vectors are given by the k th complex roots at each vertex, as described in Knöppel et al. [2013, Sec. 2]. An example is shown in the inset figure.


6.1.3 Different Connection. Another possibility is to change the connection ∇ itself. In this case, Algorithm 1 will compute parallel transport along ordinary geodesics, but the notion of what it means to parallel will change. For instance, setting $r_{ij} = 1 + 0i$ for all edges ij simply yields closest point interpolation of complex functions. A more interesting connection is considered by Knöppel et al. [2015], who compute a global parameterization aligned to a given vector field Z . Here the Levi-Civita connection ∇ is replaced by the connection $\tilde{\nabla} = d - iZ^b$; in practice this just means that the rotations r_{ij} are larger for edges that align with Z (see Knöppel et al. [2015, Section 3.2]). Using the corresponding connection Laplacian $\tilde{\Delta}^{\nabla}$ in Algorithm 1 will yield a *local* field aligned parameterization centered around a given point x , since for any shortest geodesic γ starting at x the augmented parallel transport map is given by

$$\tilde{P}_{\gamma} = e^{i \int_{\gamma} \langle Z, \gamma' \rangle ds},$$

i.e., a rotation determined by how much the tangent of γ lines up with Z . Hence, the angle of the transported field $\eta(y) := \arg(\tilde{P}_{\gamma_{x \rightarrow y}} 1)$ will be a scalar function increasing along Z (see inset). The gradient of η will therefore be closely aligned with Z near the source point x ; any non-integrability is effectively dealt with by pushing it out toward the cut locus (dashed line), rather than inserting new singularities (as in Knöppel et al. [2015]) or globally projecting onto a more integrable field (as in Ray et al. [2006]).

6.2 Other Discretizations

There is no fundamental reason why we must use triangle meshes to discretize the vector heat method: any geometric representation



that admits a discretization of the scalar Laplace-Beltrami operator Δ and the connection Laplacian Δ^{∇} will suffice. Discrete Laplacians have been developed for a wide variety of domains, including point clouds [Liu et al. 2012], polygon meshes [Alexa and Wardetzky 2011], subdivision surfaces [de Goes et al. 2016], tetrahedral meshes [Belyaev and Fayolle 2015], spline surfaces [Nguyen et al. 2016], and *digital surfaces*, *i.e.*, voxelizations [Caissard et al. 2017], all of which have been used to implement the scalar heat method (see either the preceding references, or Crane et al. [2013b]).

Given a scalar Laplacian, a connection Laplacian can be obtained by following the same strategy used for triangle meshes (Sec. 5.3): for any pair of nearby nodes i and j (representing vertices, points, *etc.*), determine the transformation between tangent spaces. For a surface embedded in \mathbb{R}^n , this transformation is just the smallest rotation between tangent planes, and can be encoded by a unit complex number r_{ij} . Then simply multiply the off-diagonal entries L_{ij} by the values r_{ij} to obtain the connection Laplacian L^{∇} . If L was symmetric, L_{ij} will be Hermitian (assuming $r_{ji} = r_{ij}^{-1}$), and will hence exhibit the properties discussed in Sec. 7.1). We consider three specific cases in detail.

6.2.1 Polygon Meshes. For surface meshes comprised of general, possibly non-planar polygons, we augment the discrete Laplacian of Alexa and Wardetzky [2011] (and use the same mass matrix M). In this setting we need a transport coefficient for all pairs of vertices i, j contained in each polygon—not just those connected by an edge. We therefore define extrinsic tangent planes, by picking any reasonable normal direction N_i at each vertex $i \in V$, and any direction E_i orthogonal to N_i that serves as the zero direction. Letting $R_{ij} \in \text{SO}(3)$ denote the smallest rotation taking plane i to plane j , the rotations r_{ij} are then determined by the angle from $R_{ij}E_i$ to E_j in plane j . These values are used to modify the off-diagonal entries as described above. Fig. 11, *center right* shows an example on a quad-dominant mesh containing non-planar quads and pentagons, of the type commonly used in numerical simulation.

6.2.2 Point Clouds. As in Crane et al. [2013b, Section 3.2.3], we use the positive semidefinite point cloud Laplacian of Liu et al. [2012] to implement our algorithm on unstructured point cloud data. In this setting we must already estimate tangent planes at each point in order to build the scalar Laplacian; the transport coefficients can therefore be computed exactly as in the polygonal case: pick a direction E_i at each tangent plane and compute the rotations r_{ij} from $R_{ij}E_i$ to E_j . As in the scalar case, the mass matrix M is given by the Voronoi areas associated with points. An example is shown in Fig. 11, *center left*.

6.2.3 Voxelizations. Finally, for a voxelized or *digital surface* a discrete Laplacian was recently developed by Caissard et al. [2017]. Here, values are associated with faces (*i.e.*, quads) on the voxelization boundary; the basic principle is the same as in the point cloud case, but normals and areas are carefully estimated based on the voxelization geometry. (If the voxelization arises from an Eulerian signed distance function, these normals might also be used.) An example is shown in Fig. 11, *far right*.




Fig. 11. To implement our method on a given geometric data structure, we essentially just need a scalar Laplacian and a notion of tangent spaces at each point or vertex. Here we show the solution for three sources of different magnitudes on a triangle mesh, point cloud, polygon mesh, and voxelization.

MODEL	TRIANGLES	PRECOMPUTE	SOLVE
Chair	11k	0.06 s	0.002 s
Bust	100k	0.73 s	0.031 s
Children	100k	0.74 s	0.027 s
Seahorse	145k	1.25 s	0.047 s
Snake	293k	2.74 s	0.101 s
Rhino	310k	4.02 s	0.105 s

Table 1. Timings for Figure 12, including intrinsic Delaunay preprocessing.

7 DISCUSSION AND VALIDATION

Here we study numerical properties of our method, and compare with other candidate approaches. Importantly, the main task we consider (parallel transport along shortest geodesics to a given set) is not one previously considered in geometry processing, and hence there are no standard algorithms. The closest analogy, perhaps, is a recent numerical integrator for closed-form Riemannian metrics rather than discrete meshes [Louis et al. 2018]. For surfaces of revolution we use the exact solution (computed via Clairaut’s relation) as a basis for comparison; for more complicated models we compute exact polyhedral geodesics (*à la* Surazhsky et al. [2005]), and apply parallel transport via unfolding, as described by Polthier and Schmieß [1998]. Note that even the polyhedral approach does not yield the true (smooth) solution on coarse meshes or near the cut locus; on fine meshes and away from the cut locus it nonetheless provides a useful benchmark for comparison. Overall we find that the vector heat method provides an excellent performance/quality trade off, making it well-suited for practical geometry processing tasks; application-specific comparisons are explored in Sec. 8.

7.1 Basic Properties

Which properties of smooth parallel transport are preserved by our discrete algorithm? For a single point source, one can easily argue that we exactly preserve elementary properties such as linearity ($P_Y(aX + Y) = aP_Y X + P_Y Y$), conservation of magnitude ($|P_Y X| = |X|$), and covariance with respect to rotation, *i.e.*, rotating the initial

vector is equivalent to rotating the final solution by the same angle. A more interesting property is symmetry: in the smooth setting, $P_{Y \rightarrow X} \circ P_{X \rightarrow Y} = \text{id}$, *i.e.*, transporting from x to y and back again should yield the original vector. This property turns out to be true in the discrete case as well: since the matrix $L^V \in \mathbb{C}^{|V| \times |V|}$ encoding the connection Laplacian is Hermitian, the solution operator

$$A := (M + tL^V)^{-1}$$

is also Hermitian. Letting δ_i denote a Kronecker delta at the source vertex $i \in V$, we can write the parallel vector field corresponding to the vector $b := z\delta_i$ (for $z \in \mathbb{C}$) as $X = Ab$. The transported vector at any vertex $j \in V$ can be written as $X_j = (\delta_j^T X)\delta_j$; when we transport this vector back to i , we therefore get

$$\delta_i^T (AX_j) = \delta_i^T A \delta_j (\delta_j^T X) = z (\delta_i^T A \delta_j) (\delta_j^T A \delta_i) = z A_{ij} A_{ji}.$$

Since A is Hermitian, $A_{ij} = \bar{A}_{ji}$, hence $A_{ij} A_{ji}$ is real, *i.e.*, the initial vector z experiences a scaling and no rotation. But since the overall process preserves scale, the final vector is the same as the initial one. Further properties of parallel transport (such as equivalence between curvature and monodromy around closed loops) may not hold exactly; likewise, these properties may not hold exactly in the case of multiple sources or curves, since for $t > 0$ vectors pointing in different directions may result in cancellation of magnitude. In general, we expect that many properties will at least be preserved under refinement—see Sec. 7.3 for further discussion.

7.2 Implementation and Performance

We implemented our method in C++ using double precision; all timings were taken on a single thread of an Intel Core i7 3.5GHz CPU. To solve linear systems we prefactor matrices via CHOLMOD [Chen et al. 2008] or UMFPACK [Davis 2004] and applied backsubstitution for each subsequent problem. For the basic algorithm (Algorithm 1) we need to pre-factor two $|V| \times |V|$ Laplace matrices (one real, one complex); for each new source set we need only three backsolves, and trivial per-vertex multiplication/division operations. The (optional) intrinsic Delaunay mesh can be constructed as a preprocess;



Fig. 12. Results of the vector heat method on a variety of models; source is marked by a large arrow. For visualization, vectors are sampled *à la* Bridson [2007].

in practice we find the cost is about the same as one matrix factorization. On a mesh of 100k triangles, preprocessing takes about 1 second overall; computing parallel transport from any subsequent source to all points on the surface then takes about 30ms. We did not carefully optimize our code, though further accelerations are relatively straightforward: for instance, since both matrices have the same sparsity pattern one could re-use the symbolic factorization; moreover, since there is no dependence among the linear systems, backsubstitution could be applied in parallel. (See also discussion of fast preconditioners in Sec. 2.) In contrast, computing exact polyhedral geodesics and applying parallel transport via unfolding takes about 40x longer on the same mesh (using Kirsanov’s implementation of [Surazhsky et al. 2005]). One could significantly improve performance of the polyhedral strategy via any number of recent acceleration schemes (such as [Ying et al. 2013]), or, at the cost of accuracy, by extracting geodesics from a cheaper piecewise linear distance function. However, even just *tracing* the geodesics from the source point to each vertex (whether using the polyhedral scheme or fast marching) already takes about 10–20x longer than executing our entire method; in general, it would seem quite difficult to develop a polyhedral strategy that is competitive with the diffusion-based approach.

7.3 Convergence and Accuracy

The accuracy of any method for computing parallel transport will depend on the resolution and quality of the surface tessellation. For the vector heat method, we find that using an intrinsic Delaunay triangulation improves quality, and hence apply this technique throughout our examples. Exact polyhedral schemes also re-tessellate the input by slicing it up into polygonal “windows” relative to the given source. This point of view helps to explain the relative trade offs of the two approaches: the vector heat method re-tessellates the domain at most once (as an optional pre-process), whereas polyhedral schemes must re-tessellate for each new source point. The vector heat method hence has far better amortized performance, whereas window-based schemes can provide greater accuracy since the domain is effectively meshed along characteristics of the equation being solved (*i.e.*, along geodesics). Empirically, we observe a convergence rate of roughly $O(h)$ and $O(h^2)$ for the two methods, *resp.*, relative to the mean edge length h (Fig. 13). Fields computed via these two methods mainly differ near the cut locus, where *neither* approach can guarantee accurate results—in fact, it is well-known that the exact polyhedral cut locus is a poor approximation of the smooth one [Itoh and Sinclair 2004]. On the whole, accuracy and rates of convergence are in line with the scalar heat method—for a more in-depth analysis, see Crane et al. [2013b, Section 4.2].

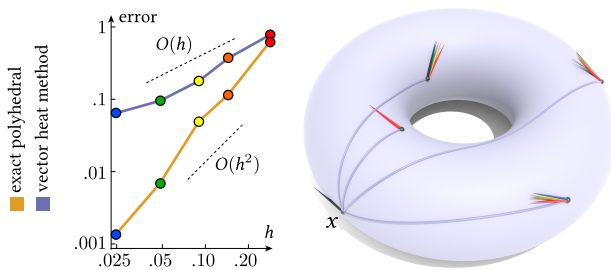


Fig. 13. The vector heat method appears to converge linearly. *Left*: On progressively finer meshes (ranging from 1250–11250 triangles), we transport a vector from a source x to several other points and measure the error relative to the true solution on the smooth surface. *Right*: visualization of vectors transported via the vector heat method; black is reference solution.




Fig. 14. Our algorithm yields very similar results to the brute-force approach of explicitly unfolding triangles along exact polyhedral geodesics. Away from the cut locus, the difference is typically just a few degrees (*left, right*).




Fig. 15. Our method enables fast velocity extension for level set methods. Here we show a sequence of frames from a simple curve-shortening flow, plus a constant tangential term. In each frame, our scalar interpolation scheme provides a closest-point interpolation of normal velocity (*top left*), resulting in excellent preservation of the signed distance function over long integration times (*bottom left*); note that we never apply explicit redistancing. Since Algorithm 1 provides accurate extrapolation of vector data over the entire domain (*top right*), we can track particles even very far from the interface (*bottom right*).

Choice of t . The accuracy of the vector heat method will be affected by the choice of the parameter t . Here we observe the exact same behavior as for the scalar heat method: if h is the average spacing between nodes (e.g., the mean edge length in a triangle mesh) then setting

$$t = h^2$$

for both scalar and vector heat flow tends to yield the best accuracy (see Fig. 26 and [Crane et al. 2013b, Section 3.2.4]). We use this value throughout all of our examples.

8 APPLICATIONS

Fast parallel transport along shortest paths provides a basic foundation on top of which many algorithms can easily be built—here we consider several important examples from geometry processing and simulation. Implementation of these algorithms via the vector heat method is often much simpler than existing alternatives: mainly just setting up and solving linear systems. In each case one enjoys a common set of benefits, such as low amortized cost (due to prefactorization) and the ability to generalize to many different geometric data structures (point clouds, polygon meshes, *etc.*). To keep discussion concrete, we will describe algorithms primarily in terms of triangulated surfaces.

8.1 Velocity Extrapolation

Perhaps the most straightforward application of our method is extrapolation of scalar or vector velocity in the context of free boundary problems; beyond physical simulation, such methods are increasingly used for tasks ranging from shape optimization to semantic shape analysis. If a signed distance function needs to be updated, one can simply extrapolate the scalar velocity in the normal direction, using the approach described in Sec. 4.1. To advect auxiliary quantities (color, temperature, particles, *etc.*), one also needs to extrapolate the tangential velocity, which can be done using Algorithm 1. Fig. 15 illustrates several features of our extrapolation strategy: for instance, since we solve a global problem we get a well-behaved velocity field far from the interface; in fact, the closest point property ensures that a signed distance function will be nearly preserved even over long integration times. Note that due to the use of a finite value t , data directly on the interface may not be exactly preserved, but will generally be very close. In the scalar

case, the performance comparison with fast marching is identical to the comparison found in [Crane et al. 2013b, Section 4.1]: the cost of our method is dominated by two backsolves, whereas fast marching executes a Dijkstra-like traversal. Note that one does *not* need to refactor matrices for a changing boundary. In the vector case, there is not a clear comparison: vector extrapolation is well-studied for Euclidean domains (e.g., [Adalsteinsson and Sethian 1999]), but these methods do not immediately generalize to curved surfaces due to the added complexity of parallel transport (especially on data structures like point clouds).

8.2 Logarithmic Map

At any point x of a closed surface M , the *exponential map* $\exp_x(rv)$ yields the point $y \in M$ obtained by walking in the unit tangent direction v and continuing along a geodesic for a distance r . The *logarithmic map* \log_x is the inverse operation: given any point y (away from the cut locus), it finds the smallest distance r and corresponding unit vector v such that $\exp_x(rv) = y$, in analogy with the ordinary logarithm and exponential. If we encode v as an angle φ , then we essentially have polar coordinates (r, φ) relative to an origin x . As observed by Schmidt et al. [2006], the logarithmic map (referred to there as the exponential map) is useful for a large number of tasks in geometry processing, such as interactive shape editing [Schmidt and Singh 2010] and texture decaling (Fig. 17). A




Fig. 16. *Left*: the *logarithmic map* provides a coordinate system on the surface, relative to a chosen origin x . *Right*: on the sphere we easily compute the correct log map; note that far from the source even the analytical log map can be highly skewed.

global logarithmic map also helps translate algorithms from Euclidean space to curved domains—for instance, in Sec. 8.3 the log map enables us to easily compute generalized centers of mass, by identifying points on the surface with vectors in the tangent plane.

How can we compute the log map on a surface? For polar coordinates (r, φ) in the Euclidean plane, φ can be expressed as the angle between a horizontal direction H and a radial vector field R emanating from the origin. Likewise, on a curved surface the radial vector field R is given by the gradient of the geodesic distance to a source point x , and the “horizontal” vector field H is obtained by transporting any unit vector at x to every other point. The angular coordinate of the log map is then the angle from H to R ; the radial component is just the geodesic distance. We compute the horizontal field by applying Algorithm 1 as usual, where the choice of initial vector determines the zero direction ($\varphi = 0$). Obtaining the radial field R is more challenging: if we simply take derivatives of a per-vertex distance function, we get numerical noise (Fig. 18, *left*). Explicitly smoothing this field is not an attractive option, since it will distort features like the cut locus, and generally degrade the accuracy of subsequent computations (*e.g.*, when computing Karcher means).

Instead, we can apply our parallel transport algorithm (Algorithm 1) to a small circle of outward pointing normals around the source vertex i . Here, care must be taken in formulating the initial conditions R^0 for the vector diffusion step: simply setting R_j^0 to the outward pointing direction $-e_{ji}/|e_{ji}|$ at each neighbor j can lead to anisotropy in the resulting map (Fig. 18, *middle*). We instead derive initial conditions by carefully projecting unit normals on a circle of radius ε around the source vertex onto piecewise linear basis functions (as discussed in App. A). This approach yields a log map which is both accurate and smooth (Fig. 18, *right*). Note that since both H and R are unit vector fields we do not need to interpolate magnitudes. The radial coordinate r (corresponding to the geodesic distance) can be computed using any method; we simply integrate R by solving the Poisson equation $\Delta r = \nabla \cdot R$, requiring only one additional prefactorization and backsubstitution. In particular, to evaluate the right hand side of this equation, we first map vectors R_i at vertices to integrated values $R_{ij} \in \mathbb{R}$ per edge, by averaging the inner product with the edge vector:

$$R_{ij} := \frac{1}{2} (\langle e_{ij}, X_i \rangle + \langle -e_{ji}, X_j \rangle)$$

(see Knöppel et al. [2015, Section 3.2] for further discussion). Keeping in mind that $R_{ij} = -R_{ji}$, the total divergence $\nabla \cdot R$ for vertex i can then be expressed as

$$(\nabla \cdot R)_i = \sum_{ij \in E} w_{ij} R_{ij},$$

where w_{ij} are the cotan weights from Sec. 5.3. The final log map is encoded via Cartesian coordinates $(u, v) := r_i(\cos \varphi_i, \sin \varphi_i)$ at each vertex $i \in V$.



Fig. 17. A noise-free log map allows us to place delicious decals (*left*) or tantalizing tattoos (*right*) on a surface without having to worry about difficult issues like where to place cuts.

Previous approaches aim to compute a log map only in a neighborhood around the source using *extrinsic* approximations [Schmidt et al. 2006; Brun 2007; Melvær and Reimers 2012], making them inaccurate over longer distances, and precluding isometry invariance. For instance, Zhang et al. [2006] simply project the closest 25% of vertices onto the local tangent plane which can fail badly for highly curved surfaces; Dijkstra-like algorithms can also deviate wildly on highly curved domains (see Fig. 20 for comparisons). As a result, such approximations cannot reliably be used for algorithms that require global information, such as computation of Karcher means (Sec. 8.3) or intrinsic landmarks (Sec. 8.5). Our method nicely resolves the map over the whole surface, all the way up to the cut locus (Fig. 19). It is also quite competitive in terms of performance since it simply needs to execute highly optimized backsubstitution operations, rather than a Dijkstra-like traversal (see discussion in Sec. 2). When a parameterization is desired only in a small region (*e.g.*, when computing centroidal Voronoi diagrams), further speedups might be achieved by applying the localized Cholesky strategy of Herholz et al. [2017a], which fits perfectly into our framework. Fig. 17 shows some simple examples where our log map is used to add texture decals to a surface. Note that in the smooth setting the log map is only well-defined for closed surfaces—nonetheless, our method still works nicely on surfaces with boundary, especially within the image of the exponential map.




Fig. 18. To get a high-quality log map (here, from a source point x) one must carefully discretize the distance gradient. *Left*: simply taking the gradient of a given distance function yields numerical noise. *Middle*: naive parallel transport of vectors emanating from the source yields global anisotropy. *Right*: proper discretization of initial conditions yields a smooth and accurate map, where the only remaining noise is near the cut locus (dashed line).

ALGORITHM 2: Logarithmic Map**Input:** A point $x \in M$ and zero direction $H_0 \in T_x M$.**Output:** A function $(u, v) : M \rightarrow \mathbb{R}^2$.

- I. Transport H_0 to all other points to get H , via Algorithm 1.
- II. Compute R using Algorithm 1, with initial conditions from App. A.
- III. Compute the angle φ from H to R at each point.
- IV. Solve the Poisson equation $\Delta r = \nabla \cdot R$.
- V. Evaluate $(u, v) := r(\cos \varphi, \sin \varphi)$ at each point.




Fig. 19. Our log map is globally accurate—even on long skinny models where most points are reached by traveling in nearly identical directions (far right). On a high-resolution mesh we nicely resolve the cut locus (dashed line).





Fig. 20. Since the vector heat method is based on a straightforward discretization of a smooth formulation, it leads to a close approximation of the true log map. Previous approximations exhibit not only small local errors, but large global inaccuracies—especially in the angular coordinate φ . (Reference solution is computed via the exact polyhedral scheme.)

Point cloud log map. As an example of how the diffusion-based approach easily extends to other geometric data structures, we implemented the log map directly on unstructured point clouds. Here, computation of the horizontal vector field H is straightforward: just apply Algorithm 1 to a unit vector at the source point x ; the radial field R can be computed using the same initial conditions described by Lin et al. [2014]: transport the vector from x to the tangent space at each neighbor y (via an extrinsic rotation) and project onto the tangent space. The angular component is then given by the angle from H to R ; the radial component is obtained by solving the Poisson equation $\Delta r = \nabla \cdot R$ exactly as described in Crane et al. [2013b, Section 3.2.3]. Two examples on scanned data are shown in the inset.




8.3 Karcher Means and Geometric Medians

Given a set of points $y_1, \dots, y_n \in M$, any minimizer of the energy

$$E(x) := \frac{1}{2n} \sum_{i=1}^n d(x, y_i)^p \quad (11)$$

provides a notion of center. For $p = 2$, such minimizers are called *Karcher means*; in Euclidean space, just the centroid or arithmetic mean. (If the Karcher mean is unique, it is known as the *Fréchet mean*.) For $p = 1$, minimizers are known as *geometric medians*, and tend to be more robust to outliers.



Though algorithms have been developed for finding Karcher means on special geometries [Buss and Fillmore 2001] or other notions of weighted averages on surfaces [Panozzo et al. 2013], to date there has been no practical algorithm for accurately computing Karcher means on general surfaces. Likewise, the geometric median has been considered in the space of images [Fletcher et al. 2009], but no efficient algorithms are known on discrete geometric

Fig. 22. At any point $x \in M$, the gradient v of the Karcher mean energy is just the sum of the logarithms of all the points y_i .

ALGORITHM 3: Karcher Mean**Input:** A collection of points $y_1, \dots, y_n \in M$.**Output:** A point $m \in M$.

- I. Pick a random initial guess $m^0 \in M$.
- II. Until the vector v has sufficiently small norm:
 - (a) Compute the log map at m^k via Algorithm 2.
 - (b) Evaluate the update vector $v = \frac{1}{n} \sum_i \log_{m^k}(y_i)$.
 - (c) Compute $m^{k+1} = \exp_{m^k}(\tau v)$, i.e., walk forward along v for time τ .




Fig. 21. Our globally accurate logarithmic map can be used to compute centers of data on a curved surface. *Left*: given a collection of points, we iteratively compute the Karcher mean; on simple geometries we reproduce the expected solution in just 2-3 iterations, and the algorithm generalizes to more complex geometries while still needing < 20 iterations. *Right*: We can also efficiently compute the centers of distributions; here we show both the Karcher mean (purple) and the geometric median (green). The Karcher mean is significantly influenced by outliers, while the geometric median is not.

domains (meshes, point clouds, *etc.*). Our algorithm for computing the log map (Sec. 8.2) enables a straightforward and efficient strategy for minimizing the energy E . In the case of the Karcher mean ($p = 2$), we just iteratively evaluate the update vector

$$v \leftarrow \frac{1}{n} \sum_{i=1}^n \log_{m^k}(y_i), \quad m^{k+1} \leftarrow \exp_{m^k}(\tau v), \quad (12)$$

where $\tau > 0$ is a step size. (For instance, if the domain is \mathbb{R}^N , this algorithm immediately converges to the centroid for $\tau = 1$.) For the geometric median ($p = 1$) we simply need to replace the expression for v with a convex combination $\sum_i \omega_i \log_{m^k}(y_i) / \sum_i \omega_i$, where the coefficients $\omega_i := 1/d(m^k, y_i)$ can also be computed from the log map—this strategy is known as the *Weiszfeld algorithm* [Weiszfeld 1937]. The log map is computed once per iteration, via the algorithm described in Sec. 8.2. To evaluate the exponential map, we simply trace a geodesic at m^k along the surface in the direction v for time τ (*à la* Polthier and Schmies [1998], in the case of triangle meshes). In practice this scheme tends to converge in no more than 20 iterations (and far fewer on simple models); in all our examples, the initial guess $m^0 \in M$ is chosen completely at random. The cost of each iteration is dominated by the backsolves to compute the log map. Line search can slightly reduce the number of steps, at the cost of additional solves; for most examples we use $\tau = 1$ and no line search.

In fact, since we know the log map over the entire surface, the number of points p_i has a negligible effect on the cost of computation (just taking a weighted sum). We can therefore apply the same method directly to arbitrary distributions $\rho : M \rightarrow \mathbb{R}_{>0}$ (as depicted in Fig. 21, *right*); in this case, we can define a center

$$m(\rho) = \arg \min_x \int_M \rho(y) d(x, y)^p dy. \quad (13)$$

The only change to the algorithm is that we now take a weighted average over *all* vertices, using weights $M_{ii} \rho_i$, where $\rho_i \in \mathbb{R}_{>0}$ is the density at each vertex, and M is the mass matrix.

The globally accurate approximation of the log map provided by the vector heat method allows us to reliably obtain the correct result (Fig. 23, *far right*). Although previous methods for *locally* approximating the log map can also be used to implement this algorithm [Zhang et al. 2006; Schmidt et al. 2006; Schmidt 2013], they are not accurate enough *globally* to produce the desired behavior:

iterates either wander around randomly and fail to converge, or converge only because line search eventually steers an inaccurate guess toward a local minimum (Fig. 23). The method of Panozzo et al. [2013] takes a different, non-iterative approach to computing averages on surfaces, but does not produce the true Karcher mean, and may not even respect basic symmetries (Fig. 23, *center right*). Likewise, simply using a smooth, somewhat parallel vector field to construct the log map also does not work, as shown in Fig. 3, *right*. To date, the vector heat method appears to be the only way to compute Karcher means on surfaces.

8.4 Geodesic Centroidal Voronoi Diagrams

A Voronoi diagram partitions a domain M into regions $U_i \subset M$ comprised of those points closest to a given collection of *sites* $s_i \in M$. In a centroidal Voronoi diagram, each site is located at the centroid of its associated region. Our fast Karcher mean algorithm (Sec. 8.3) provides an effective way to compute *geodesic centroidal Voronoi tessellations* (GCVT), where the centroid is defined as the cell's Karcher mean. Our method provides, to our knowledge, the first efficient approach for computing the true GCVT: using the Euclidean centroid (*à la* CCVT [Du et al. 2003]) does not work well for large, curved cells; likewise, a diffusion-based centroid (*à la* Herholz et al. [2017b]) can yield a diagram very different from the GCVT, especially when the Karcher mean is outside the cell.

To compute a GCVT, we simply apply Lloyd's algorithm, updating cell centers via the Karcher mean. More specifically, we consider a distribution associated with each cell, defined via the scalar heat kernel k_t as

$$\rho_i := \frac{k_t(s_i, \cdot)}{\sum_j k_t(s_j, \cdot)} \quad (14)$$




Fig. 24. Fast Karcher means allow us to compute centroidal geodesic Voronoi tessellations with large, possibly multiply-connected cells.

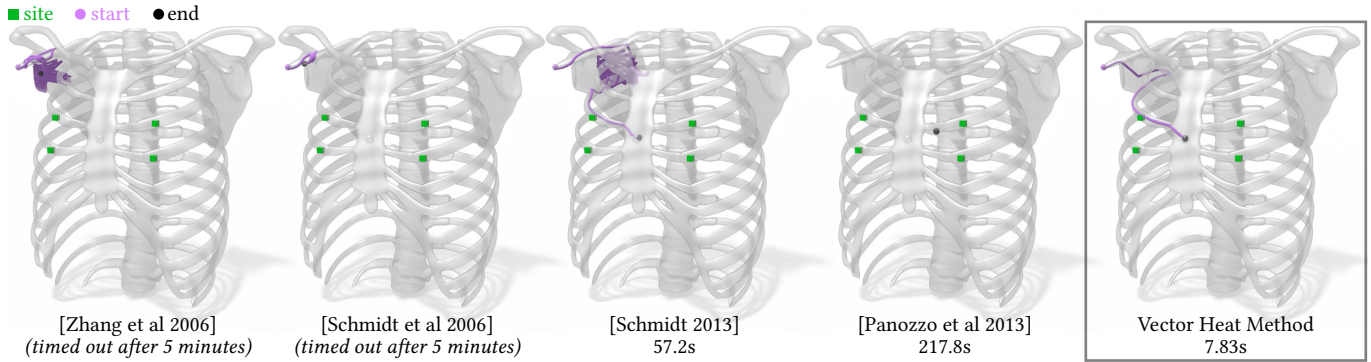


Fig. 23. Here we compute the Karcher mean of the four green points, which should appear at the center of the sternum. Previous local approximations of the log map yield poor behavior, with iterates failing to converge (*far left, center left*) or wandering around randomly until line search pushes the solution toward a local minimum (*center*). Other algorithms compute only approximate averages, which may not respect symmetries of the problem (*center right*). The global accuracy of the log map provided by the heat method guides the algorithm to the correct solution in just a few iterations (*far right*).

for a small time t . This distribution is an indicator function for each cell U_i ; k_i is computed as usual (by solving a discrete heat equation). We then apply Algorithm 3 to move each site s_i to the center of the distribution ρ_i , and repeat until convergence. In practice, we find it is more efficient overall to take just a single step of the Karcher algorithm, even though it results in more Lloyd iterations. We do not have to worry about topological issues like features separated by small extrinsic distances, and can handle multiply connected cells since we need only integrate the log map over each region U_i . Faster convergence might be achieved by replacing Lloyd’s algorithm with a more sophisticated optimization strategy—Liu et al. [2016, Section 2] provides a nice discussion in the context of GCVT.

8.5 Ordered Intrinsic Landmarks

An ongoing challenge in geometry processing is finding landmark points that provide correspondence between (near-)isometric surfaces. A significant difficulty is not only finding geometrically salient points, but also finding a consistent *ordering* for those points, *i.e.*, which landmark on the first surface corresponds to which landmark on the second surface? Trying to determine this ordering *a posteriori* leads to hard combinatorial *matching problems*; algorithms for efficiently computing such matchings are only just starting to be understood [Kezurer et al. 2015]. The ability to reliably compute geometric centers provides new opportunities for generating landmarks that are consistently ordered *a priori*. We explore a simple strategy




Fig. 25. Our robust geometric median facilitates computation of consistently ordered landmarks, helping to avoid difficult matching problems. (Models reconstructed from scans in Bogo et al. [2014].)

where we first compute the geometric medians *of the surface itself*; in other words, we apply the algorithm described in Sec. 8.3, setting $p = 1$ and using a uniform density $\rho = 1$ over the whole surface. We then progressively add points via furthest point sampling, *i.e.*, picking the point with greatest geodesic distance from our current set. Since in general there may be more than one geometric median (*e.g.*, on surfaces with symmetries) we sample a large set of initial guesses, and compute the furthest points relative to this *set* (discarding the medians themselves, since we cannot easily distinguish their order). Using the geometric median rather than the Karcher mean significantly reduces variability caused by values near the cut locus, which are effectively treated as outliers. As is common in shape correspondence problems, we can add a weak extrinsic factor (*e.g.*, the x -coordinate) to the density ρ in order to break symmetry. Though we have so far considered only the most basic implementation, results are already promising (Fig. 25); further refinement of such techniques would make for interesting future work.

9 LIMITATIONS

Potential challenges in applying our method stem mainly from two sources, namely (i) scalability of linear solvers, and (ii) numerical accuracy. Since we solve standard Poisson-like systems, the scalability issue is no different than for many other problems in scientific computing—for instance, if meshes become too big to factor, one can switch to more scalable solvers (of the kind discussed in Sec. 2). In practice, however, we find that modern direct solvers [Chen et al. 2008] provide an excellent solution up to millions of elements. In terms of accuracy, poor mesh quality can lead to problems such as indefinite matrices and spurious flipped vectors (Fig. 10). Using an intrinsic Delaunay triangulation helps significantly, though (as with any finite element method) meshes with too few elements or poorly distributed vertices can of course yield low-quality results. The use of low-order elements limits us to linear convergence; as with the scalar heat method, spline and subdivision bases might yield higher order accuracy [de Goes et al. 2016; Nguyen et al. 2016]. As discussed in Sec. 7.1 some basic properties of smooth parallel transport are not exactly preserved. Finally, the applications explored in Sec. 8

leave many open questions, such as preserving data along level sets, generalizing the definition of the log map for domains with boundary, and improving the robustness of landmark identification.

10 CONCLUSION

Vector fields arising from parallel transport along shortest geodesics enable the implementation of many fundamental algorithms in geometry processing, yet surprisingly have received little prior attention. We have presented a first method for efficiently and reliably computing such fields, though many questions remain. For instance: how to properly formulate boundary conditions, how to improve accuracy (e.g., using more sophisticated finite element discretizations [Arnold and Li 2017]), how to improve practical efficiency (e.g., via parallelism and local computation), and what additional properties might be guaranteed (expanding on Sec. 7.1). Another interesting question is whether the method can be augmented to compute transport along non-geodesic curves, or along a given vector field (*à la* [Azencot et al. 2015]). Overall, given that the starting point (vector diffusion via the connection Laplacian) is quite unlike many traditional methods from computational geometry, we expect our method will inspire new ways of looking at old problems and lead to very different computational trade offs. We are hopeful that the ease of implementation (just building and solving Laplace-like systems) will facilitate rapid adoption in real applications.

ACKNOWLEDGEMENTS

Thanks to Jooyoung Hahn for early discussions about signed distance computation, and Jim McCann for project feedback. This work was supported by a Packard Fellowship, NSF Award 1717320, an NSF Graduate Research Fellowship, and gifts from Autodesk, Adobe, and Facebook.

REFERENCES

- D. Adalsteinsson and J. Sethian. 1999. The Fast Construction of Extension Velocities in Level Set Methods. *J. Comput. Phys.* 148, 1 (1999).
- M. Alexa and M. Wardetzky. 2011. Discrete Laplacians on General Polygonal Meshes. (2011).
- D. Arnold and L. Li. 2017. Finite Element Exterior Calculus with Lower-Order Terms. *Math. Comput.* 86, 307 (2017), 2193–2212.
- O. Azencot, M. Ovsjanikov, F. Chazal, and M. Ben-Chen. 2015. Discrete Derivatives of Vector Fields on Surfaces—An Operator Approach. *ACM Trans. Graph.* 34, 3 (2015).
- A. Belyaev and P.A. Fayolle. 2015. On Variational and PDE-Based Distance Function Approximations. *Comp. Graph. Forum* 34, 8 (2015).
- N. Berline, E. Getzler, and M. Vergne. 1992. *Heat Kernels and Dirac Operators*.
- A. Bobenko and B. Springborn. 2007. A Discrete Laplace-Beltrami Operator for Simplicial Surfaces. *Disc. Comp. Geom.* 38, 4 (2007).
- F. Bogo, J. Romero, M. Loper, and M. Black. 2014. FAUST: Dataset and Evaluation for 3D Mesh Registration. *Comp. Vis. Patt. Recog.* (2014).
- D. Bommers, B. Lévy, N. Pietroni, E. Puppo, C. Silva, M. Tarini, and D. Zorin. 2013. Quad-Mesh Generation and Processing: A Survey. *Comp. Graph. Forum* 32, 6 (2013).
- R. Bridson. 2007. Fast Poisson Disk Sampling in Arbitrary Dimensions. In *ACM SIGGRAPH 2007 Sketches*.
- A. Brun. 2007. *Manifolds in Image Science and Visualization*. Ph.D. Dissertation. Institutionen för Medicinsk Teknik.
- S. Buss and J. Fillmore. 2001. Spherical Averages and Applications to Spherical Splines and Interpolation. *ACM Trans. Graph.* 20, 2 (2001).
- T. Caissard, D. Coeurjolly, J.O. Lachaud, and T. Roussillon. 2017. Heat Kernel Laplace-Beltrami Operator on Digital Surfaces. In *Disc. Geom. for Comp. Imag. (Lecture Notes in Computer Science)*.
- J. Chen and Y. Han. 1990. Shortest Paths on a Polyhedron. In *Symp. Comp. Geom.*
- Y. Chen, T. Davis, W. Hager, and S. Rajamanickam. 2008. Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate. (2008).
- K. Crane, F. de Goes, M. Desbrun, and P. Schröder. 2013a. Digital Geometry Processing with Discrete Exterior Calculus. In *ACM SIGGRAPH 2013 courses (SIGGRAPH '13)*.
- K. Crane, M. Desbrun, and P. Schröder. 2010. Trivial Connections on Discrete Surfaces. *Symp. Geom. Proc.* 29, 5 (2010).
- K. Crane, C. Weischedel, and M. Wardetzky. 2013b. Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow. *ACM Trans. Graph.* 32, 5 (2013).
- T. Davis. 2004. Algorithm 832: UMFPACK V4.3—an Unsymmetric-Pattern Multifrontal Method. *ACM Trans. Math. Softw.* 30, 2 (2004).
- F. de Goes, M. Desbrun, M. Meyer, and T. DeRose. 2016. Subdivision Exterior Calculus for Geometry Processing. *ACM Trans. Graph.* 35, 4 (2016).
- M. Desbrun, E. Kanso, and Y. Tong. 2006. Discrete Differential Forms for Computational Modeling. In *ACM SIGGRAPH 2006 Courses (SIGGRAPH '06)*.
- Q. Du, M. Gunzburger, and L. Ju. 2003. Constrained Centroidal Voronoi Tessellations for Surfaces. *SIAM J. Sci. Comp.* 24, 5 (2003).
- N. El Karoui and H. Wu. 2015. Graph Connection Laplacian and Random Matrices with Random Blocks. *Information and Inference* 4, 1 (03 2015), 1–44.
- M. Fisher, B. Springborn, A. Bobenko, and P. Schroder. 2006. An Algorithm for the Construction of Intrinsic Delaunay Triangulations with Applications to Digital Geometry Processing. In *ACM SIGGRAPH 2006 Courses*.
- P. Fletcher, S. Venkatasubramanian, and S. Joshi. 2009. The Geometric Median on Riemannian Manifolds with Application to Robust Atlas Estimation. *NeuroImage* 45, 1 (2009).
- A. Gentle. 2002. Regge Calculus: A Unique Tool for Numerical Relativity. *Gen. Rel. and Grav.* 34, 10 (2002).
- A. Gillman and P. G. Martinsson. 2014. A Direct Solver with $O(N)$ Complexity for Variable Coefficient Elliptic PDEs Discretized via a High-Order Composite Spectral Collocation Method. *SIAM J. Sci. Comp.* 36, 4 (2014).
- A. Grigor'yan. 2009. *Heat Kernel and Analysis on Manifolds*. Amer. Math. Soc.
- P. Herholz, T. Davis, and M. Alexa. 2017a. Localized Solutions of Sparse Linear Systems for Geometry Processing. *ACM Trans. Graph.* 36, 6 (2017).
- P. Herholz, F. Haase, and M. Alexa. 2017b. Diffusion diagrams: Voronoi cells and centroids from diffusion. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 163–175.
- J. Itoh and R. Sinclair. 2004. Thaw: A Tool for Approximating Cut Loci on a Triangulation of a Surface. *Experiment. Math.* 13, 3 (2004).
- I. Kezurer, S. Kovalsky, R. Basri, and Y. Lipman. 2015. Tight Relaxation of Quadratic Matching. *Symp. Geom. Proc.* (2015).
- R. Kimmel and J. Sethian. 1998. Fast Marching Methods on Triangulated Domains. *Proc. Nat. Acad. Sci.* 95 (1998).
- F. Knöppel, K. Crane, U. Pinkall, and P. Schröder. 2013. Globally Optimal Direction Fields. *ACM Trans. Graph.* 32, 4 (2013).
- F. Knöppel, K. Crane, U. Pinkall, and P. Schröder. 2015. Stripe Patterns on Surfaces. *ACM Trans. Graph.* 34, 4 (2015).
- R. Kyng, Y.T. Lee, R. Peng, S. Sachdeva, and D. Spielman. 2016. Sparsified Cholesky and Multigrid Solvers for Connection Laplacians. In *Symp. Theor. Comp.*
- B. Lin, J. Yang, X. He, and J. Ye. 2014. Geodesic Distance Function Learning via Heat Flow on Vector Fields. In *Int. Conf. Mach. Learn.*
- Y. Liu, B. Prabhakaran, and X. Guo. 2012. Point-Based Manifold Harmonics. *IEEE Trans. Vis. Comp. Graph.* 18, 10 (2012).
- Y.J. Liu, C.X. Xu, R. Yi, D. Fan, and Y. He. 2016. Manifold Differential Evolution (MDE): A Global Optimization Method for Geodesic Centroidal Voronoi Tessellations on Meshes. *ACM Trans. Graph.* 35, 6 (2016).
- M. Lorenzi and X. Pennec. 2014. Efficient Parallel Transport of Deformations in Time Series of Images: From Schild's to Pole Ladder. *Journal of Mathematical Imaging and Vision* 50, 1 (2014).
- M. Louis, B. Charlier, P. Jusselin, S. Pal, and S. Durrleman. 2018. A Fanning Scheme for the Parallel Transport Along Geodesics on Riemannian Manifolds. *SIAM J. Numer. Anal.* 56, 4 (2018), 2563–2584.
- M. Ludewig. 2018. Heat Kernel Asymptotics, Path Integrals and Infinite-Dimensional Determinants. *J. Geom. Phys.* 131 (2018), 66–88.
- E. Melvør and M. Reimers. 2012. Geodesic Polar Coordinates on Polygonal Meshes. In *Comp. Graph. Forum*, Vol. 31.
- J. Mitchell, D. Mount, and C. Papadimitriou. 1987. The Discrete Geodesic Problem. *SIAM J. Comput.* 16, 4 (1987).
- M. Mohamed, A. Hirani, and R. Samtaney. 2016. Comparison of Discrete Hodge Star Operators for Surfaces. *Comput. Aided Des.* 78, C (2016).
- A. Myles, N. Pietroni, and D. Zorin. 2014. Robust Field-Aligned Global Parametrization. *ACM Trans. Graph.* 33, 4 (2014).
- T. Nguyen, K. Karčiauskas, and J. Peters. 2016. C1 Finite Elements on Non-Tensor-Product 2d and 3d Manifolds. *Appl. Math. Comput.* 272 (2016).
- D. Panozzo, I. Baran, O. Diamanti, and O. Sorkine-Hornung. 2013. Weighted Averages on Surfaces. *ACM Trans. Graph.* 32, 4 (2013).
- K. Polthier and M. Schmies. 1998. *Straightest Geodesics on Polyhedral Surfaces*. Springer Verlag.
- Y. Qin, X. Han, H. Yu, Y. Yu, and J. Zhang. 2016. Fast and Exact Discrete Geodesic Computation Based on Triangle-oriented Wavefront Propagation. *ACM Trans. Graph.* 35, 4 (2016).

- N. Ray, W. C. Li, B. Lévy, A. Sheffer, and P. Alliez. 2006. Periodic Global Parameterization. *ACM Trans. Graph.* 25, 4 (2006).
- T. Regge. 1961. General relativity without coordinates. *Il Nuovo Cimento* 19, 3 (1961).
- R. Schmidt. 2013. Stroke Parameterization. *Comp. Graph. Forum* 32, 2 (2013).
- R. Schmidt, C. Grimm, and B. Wyvill. 2006. Interactive Decal Compositing with Discrete Exponential Maps. In *ACM Trans. Graph.*, Vol. 25.
- R. Schmidt and K. Singh. 2010. Meshmixer: An Interface for Rapid Mesh Composition. In *SIGGRAPH 2010 Talks*.
- N. Sharp and K. Crane. 2018. Variational Surface Cutting. *ACM Trans. Graph.* 37, 4 (2018).
- A. Singer and H.-T. Wu. 2012. Vector Diffusion Maps and the Connection Laplacian. *Comm. on Pure and Appl. Math.* 65, 8 (2012).
- D. Spielman and S.-H. Teng. 2004. Nearly-Linear Time Algorithms for Graph Partitioning, Graph Sparsification, and Solving Linear Systems. In *Symp. Theor. Comp.*
- V. Surazhsky, T. Surazhsky, D. Kirsanov, S. Gortler, and H. Hoppe. 2005. Fast Exact and Approximate Geodesics on Meshes. *ACM Trans. Graph.* 24, 3 (2005).
- X. Tricoche, G. Scheuermann, and H. Hagen. 2000. Higher Order Singularities in Piecewise Linear Vector Fields. *Math. of Surf. IX* (2000).
- A. Vaxman, M. Campen, O. Diamanti, D. Panozzo, D. Bommes, K. Hildebrandt, and M. Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing. *Computer Graphics Forum* (2016).
- E. Weiszfeld. 1937. Sur le point pour lequel la somme des distances de n points donnés est minimum. *Tohoku Mathematical Journal* 43 (1937).
- X. Ying, X. Wang, and Y. He. 2013. Saddle Vertex Graph (SVG): A Novel Solution to the Discrete Geodesic Problem. *ACM Trans. Graph.* 32, 6 (2013).
- E. Zhang, K. Mischaikow, and G. Turk. 2006. Vector Field Design on Surfaces. *ACM Trans. Graph.* 25, 4 (2006).

A DISTANCE GRADIENT DISCRETIZATION

For the log map (Sec. 8.2), we need a discretization of the radial vector field R , which in the smooth setting is the gradient of the distance function at x , *i.e.*, $R|_y := \nabla d(x, y)$. Away from the cut locus, R is a unit vector field tangent to the shortest geodesics emanating from x . To get an accurate discretization, we can therefore transport unit vectors in a small neighborhood around x to every other point—but must be careful about initial conditions. Simply *sampling* initial conditions onto vertices will not yield well-behaved solutions (Fig. 18, *center*); the dangers of intermingling sampling and finite-elements are well-known. We instead take a finite element approach, leading to reliable and accurate initial conditions (Fig. 18, *right*). We work in a flat Euclidean domain where we can use a single coordinate system for all tangent spaces; the resulting expressions also yield accurate results on curved domains due to the normalization by angle sums Θ_i (Sec. 5.2), which effectively “flattens out” each vertex tangent space.

Consider a small circle C_ε of radius $\varepsilon > 0$ centered around a source vertex $i \in V$, and let n denote its outward unit normal field. For a point in the plane—or at the tip of a cone, as depicted in Sec. 5.1—the normals n are exactly the gradient of geodesic distance along C_ε . Just as one might treat a point source as a measure of unit mass supported on a point (*i.e.*, a Dirac delta), we consider a measure of unit mass supported on C_ε , namely

$$\mu_\varepsilon := \frac{1}{\varepsilon^2} \mathcal{H}_{C_\varepsilon}^1,$$

where $\mathcal{H}_{C_\varepsilon}^1$ is the Hausdorff measure on the circle. Our initial conditions are then given by the vector-valued measure $X_\varepsilon := n\mu_\varepsilon$.

Now let $F := \text{span}\{\psi_v : v \in V\}$ denote the finite element space of piecewise linear hat functions ψ_v at vertices v . To discretize a solution to the PDE $(\text{id} + \Delta^\nabla)Y = X_\varepsilon$ in the space F , we solve

$$(M + A)y = x,$$

where M_{ij} is the mass matrix and A is the stiffness matrix discretizing the connection Laplacian—in our case we use the matrices defined in Sec. 5.3. We therefore just need to discretize the right-hand side, which is achieved by integrating each of the basis functions ψ_i with respect to the measure X_ε , *i.e.*, by evaluating the integrals

$$x_v := \int \psi_v dX_\varepsilon$$

over the whole domain. Since X_ε is supported only on triangles containing the source vertex i , and since each basis function ψ_v is supported only on the triangles containing v , we need only evaluate this expression for immediate neighbors of i .

For each neighbor j , we work in a polar coordinate system (r, θ) with i at the origin and edge ij along the horizontal axis (see inset). We treat points and vectors as complex numbers, using e to denote Euler’s number and i to denote the imaginary unit (which applies a 90-degree counter-clockwise rotation); we use $\langle \cdot, \cdot \rangle$ to denote the usual (real) inner product of vectors. For brevity we will also let $\alpha := \theta_i^k$ and $\beta := \theta_i^l$. In this coordinate system, we can express the unit normal as $n(\theta) := (\cos \theta, \sin \theta) = e^{i\theta}$. At any point x in triangle ijk , the piecewise linear hat function ψ_j can be expressed as

$$\psi_j(x) = \frac{1}{h_j^{ki}} \langle x, -ie_{ik}/\ell_{ik} \rangle = \frac{1}{2A_{ijk}} \langle x, -ie_{ik} \rangle$$

where e_{ik} is the edge vector from i to k , $\ell_{ik} := |e_{ik}|$ is its length, h_j^{ki} is the height of the triangle with apex j and base ki , and A_{ijk} is its area. In other words, to get a linear function with value 1 at j and value 0 at i and k , we take the dot product with the unit vector orthogonal to e_{ik} , and divide by height. Integrating this function over the triangle ijk with respect to the measure X_ε then yields

$$\begin{aligned} \int_{ijk} \psi_j dX_\varepsilon &= \frac{1}{\varepsilon^2} \int_0^\alpha \psi_j(\varepsilon e^{i\theta}) n(\theta) \varepsilon d\theta \\ &= \frac{\ell_{ik}}{2A_{ijk}} \int_0^\alpha \langle e^{i\theta}, e^{i(\alpha-\pi/2)} \rangle e^{i\theta} d\theta, \end{aligned}$$


where we have used the relationship $e_{ik} = \ell_{ik} e^{i\alpha}$. Integrating this expression and repeating a nearly identical calculation for triangle jil we find that the overall integral of ψ_j is given by

$$\tilde{x}_j = \frac{\ell_{ik} (\alpha \sin \alpha, \sin \alpha - \alpha \cos \alpha) / (4A_{ijk}) + \ell_{il} (\beta \sin \beta, \beta \cos \beta - \sin \beta) / (4A_{jil})}{4A_{ijk} + 4A_{jil}}. \quad (15)$$

To get the final entry for the right-hand side at vertex j , we rotate this value back into the tangent coordinate system of j :

$$x_j = -e^{i\varphi_{ji}} \tilde{x}_j. \quad (16)$$

Since Equations 15 and 16 involve only lengths, angles, and areas, they can easily be evaluated on any triangle mesh. Note that the angles φ_{ji} are exactly the same as those given in Eqn. 9.



Importantly, the initial value x_i at the source vertex i should *not* necessarily be zero: this value does not represent the pointwise value of the initial data, but is rather just one of several nodes that determines the best piecewise linear approximation. The calculation of this value is similar to the neighboring values except that we now take a sum over all triangles ijk containing vertex i . For each such triangle we will again construct a coordinate system with i at the origin and ij along the horizontal axis. In this coordinate frame, the outward unit vector orthogonal to edge kj can be expressed as

$$v := (\ell_{ij}e^{i0} - \ell_{ik}e^{i\alpha})/\ell_{jk},$$

where $\alpha := \theta_i^{jk}$. Within triangle ijk , the basis function ψ_i is then

$$\psi_i(x) = 1 - \langle x, v \rangle / h_i^{jk},$$

and the integral over the triangle is

$$\begin{aligned} \int_{ijk} \psi_i dX_\epsilon &= \frac{1}{\epsilon^2} \int_0^\alpha \psi_i(\epsilon e^{i\theta}) n(\theta) \epsilon d\theta \\ &= \int_0^\alpha e^{i\theta} d\theta - \frac{1}{2A_{ijk}} \int_0^\alpha \langle e^{i\theta}, (\ell_{ij} - \ell_{ik}e^{i\alpha}) \rangle e^{i\theta} d\theta. \end{aligned}$$

We can ignore the first term since over the whole circle it integrates to zero; the remaining term integrates to

$$\tilde{x}_{i,j} := \frac{1}{4A_{ijk}} \begin{bmatrix} -\sin \alpha (\ell_{ik}\alpha + \ell_{ij} \sin \alpha) \\ \ell_{ij}(\cos \alpha \sin \alpha - \alpha) + \ell_{ik}(\alpha \cos \alpha - \sin \alpha) \end{bmatrix}. \quad (17)$$

The initial value at the source is then given by the sum of these values, rotated back to the original coordinate system:

$$x_i := \sum_{ijk \in F} e^{i\theta_{ij}} \tilde{x}_{i,j}.$$

Again, these initial conditions can be easily computed on any triangle mesh. Notice that the initial conditions do not depend on the choice of radius ϵ . To obtain the radial field R we now run Algorithm 1 with initial conditions x , but can simply normalize the resulting vectors rather than solving for magnitudes. At the source vertex i we set R_i to zero.

Received May 2018; final version June 2018; accepted July 2018




Fig. 26. To determine a good choice for the parameter t , we measured the error of the vector heat method for varying values of $t = mh^2$, where h is the mean edge length. Each curve represents a mesh from the data set of Myles et al. [2014]. For each value of m , the error is the mean error relative to the exact polyhedral solution over the whole surface. Each curve is independently normalized by the error at its optimal m value, indicated by a black dot. Overall, we find that $t = h^2$ (i.e., $m = 1$) is a reasonable choice across a wide variety of examples.

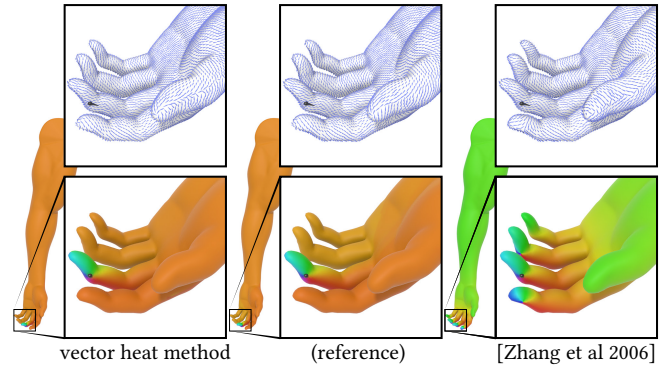


Fig. 27. Other local approximations of parallel transport, such as the method of Zhang et al [2006], do not provide a good global approximation over the larger domain. Here the field produced by the vector heat method closely matches the reference (exact polyhedral surface) allowing us to compute a globally accurate log map; the local approximation of Zhang et al deviates significantly away from the source.