Applications of the Laplacian

Justin Solomon

6.8410: Shape Analysis

Spring 2023

Review:

Rough Intuition: Spectral Geometry

http://pngimg.com/upload/hammer_PNG3886.png

You can learn a lot about a shape by **hitting it** (lightly) **with a hammer**!

Review:

Rough Definition

What can you learn about its shape from **vibration frequencies and oscillation patterns?**

$$\Delta f = \lambda f$$

Review:

The Cotangent Laplacian

$$L_{ij} = \begin{cases} \frac{1}{2}\sum_{i\sim k}(\cot\alpha_{ik} + \cot\beta_{ik}) & \text{if } i = j \\ -\frac{1}{2}(\cot\alpha_{ij} + \cot\beta_{ij}) & \text{if } i \sim j \\ 0 & \text{otherwise} \end{cases}$$

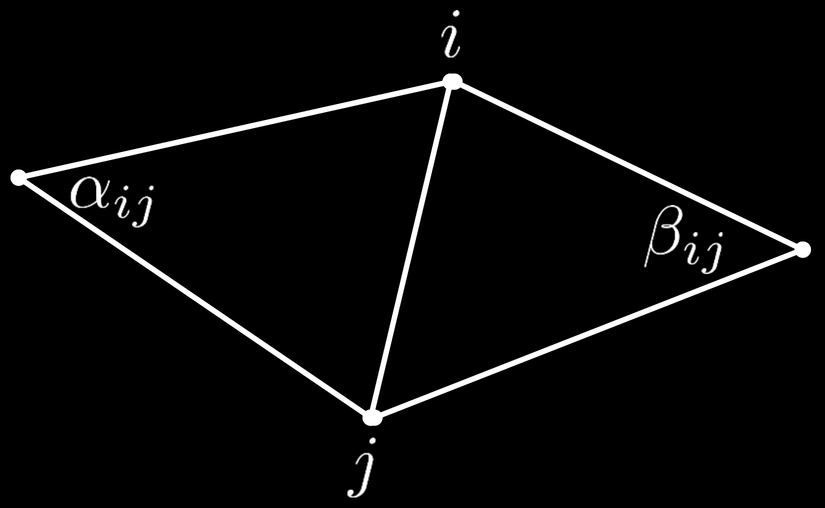

Key property:

Sparsity

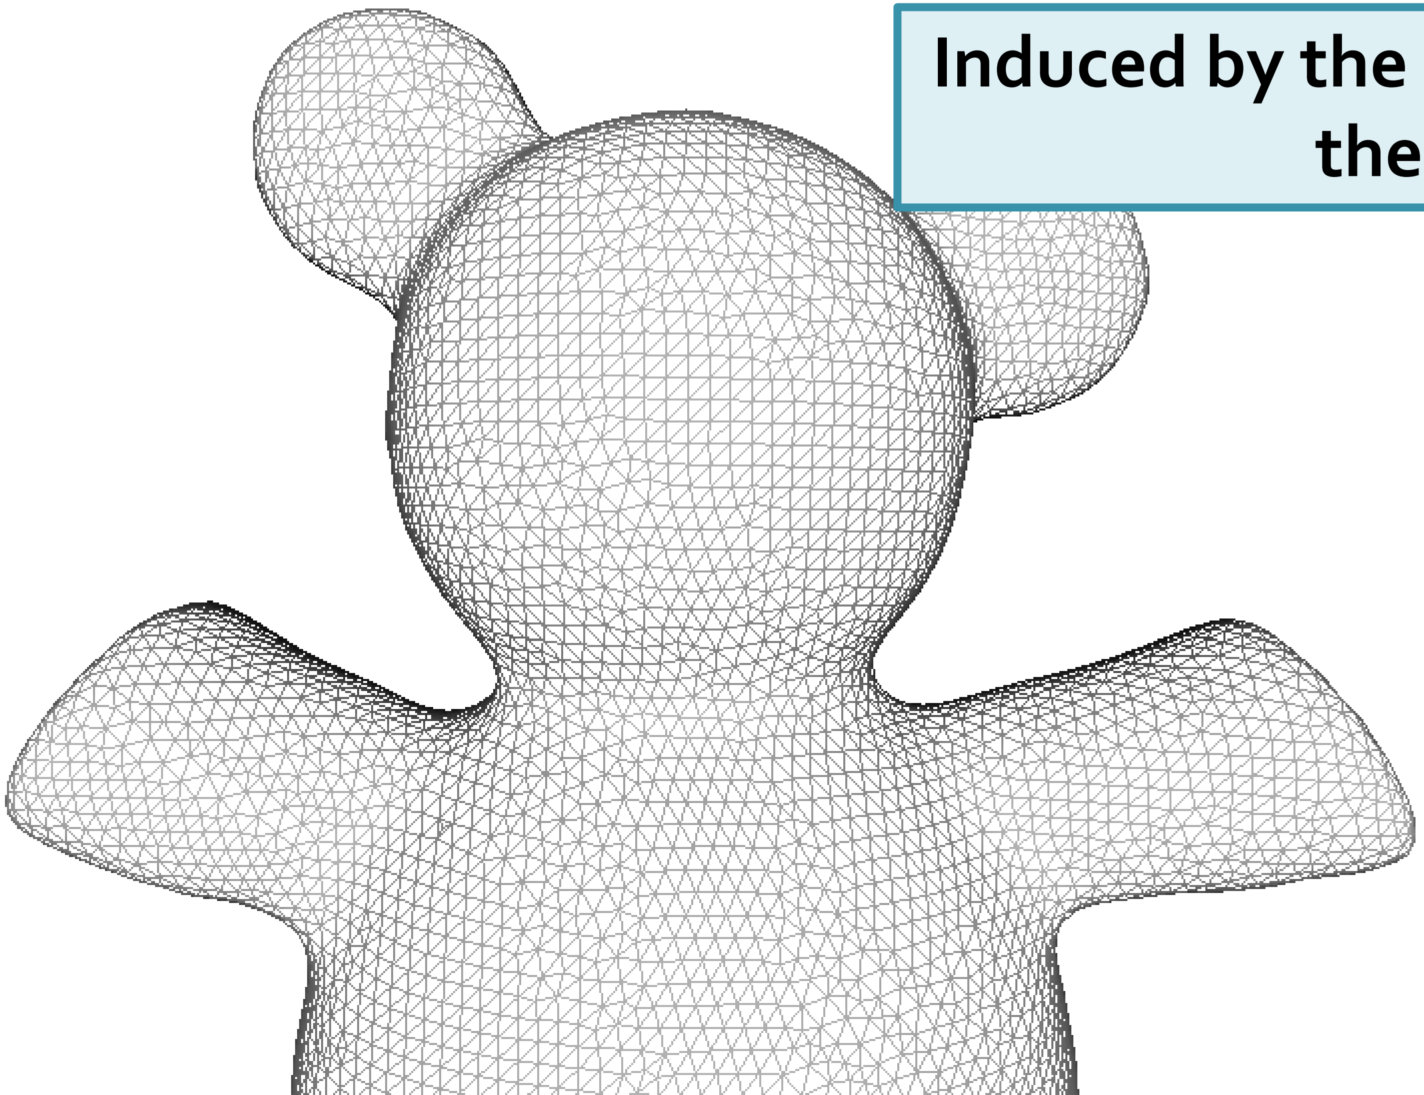

Induced by the connectivity of the triangle mesh.

Our Next Topic

Discrete Laplacian operators:

What are they good for?

- Useful properties of the Laplacian
- Applications in graphics/shape analysis
- Applications in machine learning

A quick survey:
A popular field!

Our Next Topic

Discrete Laplacian operators:

What are they good for?

- Useful properties of the Laplacian
- Applications in graphics/shape analysis
- Applications in machine learning

A quick survey:
A popular field!

One Object, Many Interpretations

$$L_{vw} = \overline{D} - A = \begin{cases} -1 & \text{if } v \sim w \\ \text{degree}(v) & \text{if } v = w \\ 0 & \text{otherwise} \end{cases}$$

<table>
<tr><th>Labeled graph</th><th>Degree matrix</th><th>Adjacency matrix</th><th>Laplacian matrix</th></tr>
<tr><td>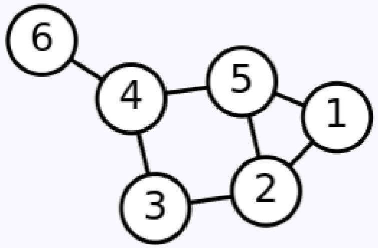
</td>
<td>$\begin{pmatrix} 2&0&0&0&0&0\\ 0&3&0&0&0&0\\ 0&0&2&0&0&0\\ 0&0&0&3&0&0\\ 0&0&0&0&3&0\\ 0&0&0&0&0&1 \end{pmatrix}$</td>
<td>$\begin{pmatrix} 0&1&0&0&1&0\\ 1&0&1&0&1&0\\ 0&1&0&1&0&0\\ 0&0&1&0&1&1\\ 1&1&0&1&0&0\\ 0&0&0&1&0&0 \end{pmatrix}$</td>
<td>$\begin{pmatrix} 2&-1&0&0&-1&0\\ -1&3&-1&0&-1&0\\ 0&-1&2&-1&0&0\\ 0&0&-1&3&-1&-1\\ -1&-1&0&-1&3&0\\ 0&0&0&-1&0&1 \end{pmatrix}$</td></tr>
</table>

https://en.wikipedia.org/wiki/Laplacian_matrix

Deviation from neighbors

One Object, Many Interpretations

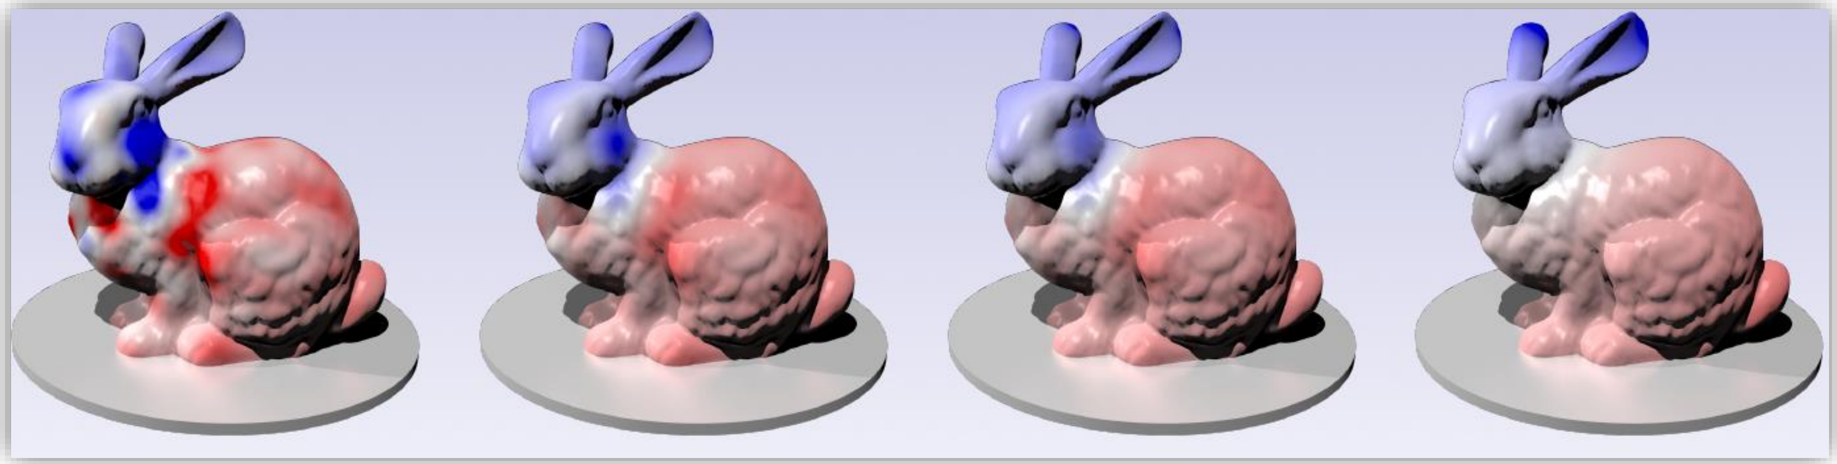

Decreasing *E*

$$E[f] := \frac{1}{2}\int_S \|\nabla f\|_2^2 \, dA = \frac{1}{2}\int_S f(x)\Delta f(x)\, dA(x)$$

Images made by E. Vouga

Dirichlet energy: Measures smoothness

One Object, Many Interpretations

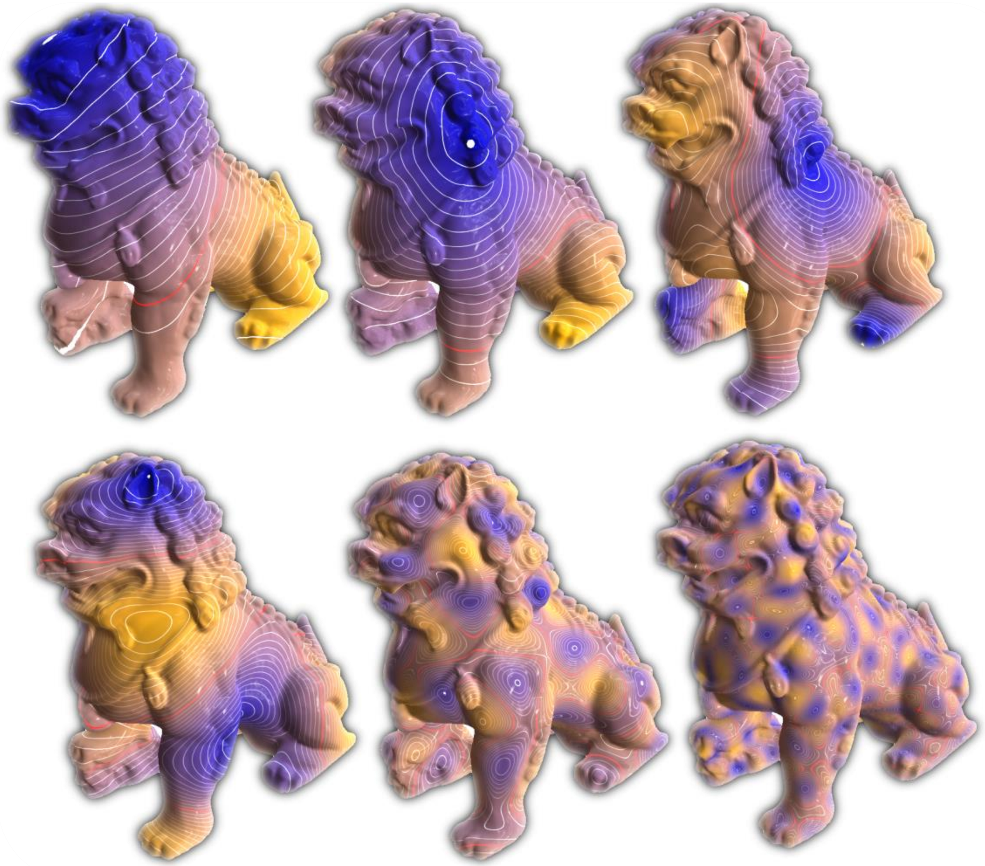

$$\Delta\psi_i = \lambda_i\psi_i$$

Vibration modes of surface (not volume!)

http://alice.loria.fr/publications/papers/2008/ManifoldHarmonics//photo/dragon_mhb.png

Vibration modes

Key Observation (in discrete case)

$$L_{ij} = \begin{cases} \frac{1}{2}\sum_{i\sim k}(\cot\alpha_{ik} + \cot\beta_{ik}) & \text{if } i = j \\ -\frac{1}{2}(\cot\alpha_{ij} + \cot\beta_{ij}) & \text{if } i \sim j \\ 0 & \text{otherwise} \end{cases}$$

$$M_{ij} = \begin{cases} \frac{\text{one-ring area}}{6} & \text{if } i = j \\ \frac{\text{adjacent area}}{12} & \text{if } i \neq j \end{cases}$$

Can be written in terms of angles and areas!

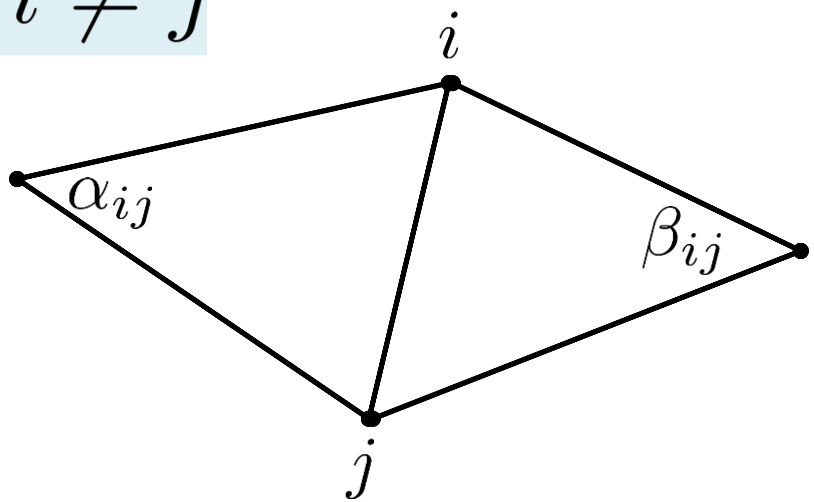

After (More) Trigonometry

$$L_{vw} = \frac{1}{8}\begin{cases} -\sum_{u\sim v} L_{uv} & \text{when } v = w \\ \mu(T)^{-1}(\ell_{vw}^2 - \ell_v^2 - \ell_w^2) & \\ \quad +\mu(T')^{-1}(\ell_{vw}^2 - \ell_v'^2 - \ell_w'^2) & \text{when } v \sim w \\ 0 & \text{otherwise} \end{cases}$$

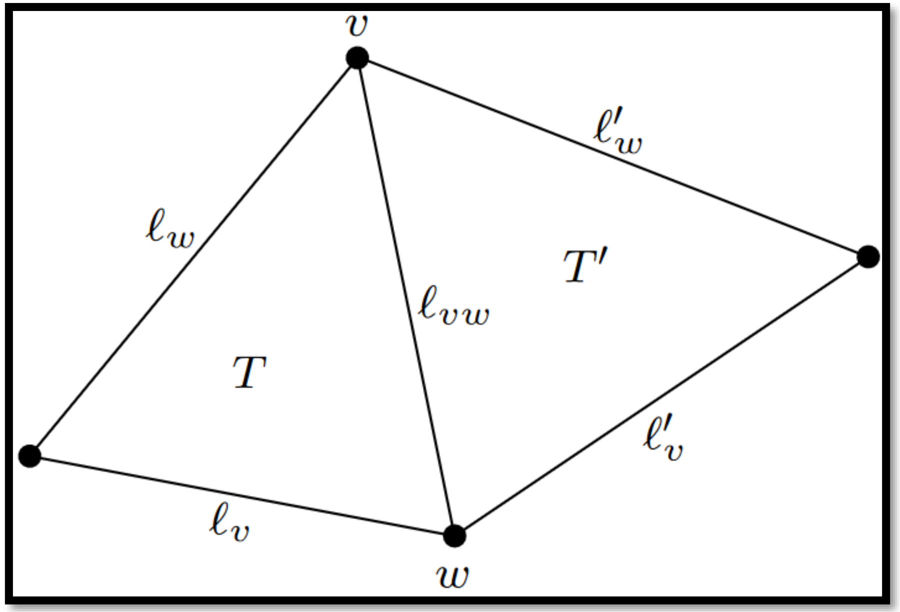

Image/formula in "Functional Characterization of Instrinsic and Extrinsic Geometry," TOG 2017 (Corman et al.)

Laplacian <u>only</u> depends on edge lengths

Isometry (for surfaces)

[ahy-**som**-i-tree]:

Bending without stretching.

Lots of Interpretations

Global isometry

$$d_1(x,y) = d_2(f(x), f(y))$$

Local isometry

$$\begin{aligned} g_1 &= f^* g_2 \\ g_1(v,w) &= g_2(f_* v, f_* w) \end{aligned}$$

Intrinsic Techniques

http://www.revedreams.com/crochet/yarncrochet/nonorientable-crochet/

Isometry invariant

Isometry Invariance: Hope

http://www.flickr.com/photos/melvinvoskuijl/galleries/72157624236168459

Isometry Invariance: Reality

“Rigidity”

http://www.4tnz.com/content/got-toilet-paper

Few shapes *can* deform isometrically

Isometry Invariance: Reality

Behavior for approximate isometry is important, too!

http://www.4tnz.com/content/got-toilet-paper

Few shapes *can* deform isometrically

Rigidity Properties

Contents lists available at SciVerse ScienceDirect

Graphical Models

journal homepage: www.elsevier.com/locate/gmod

ELSEVIER

Discrete heat kernel determines discrete Riemannian metric

Wei Zeng [a,*], Ren Guo [b], Feng Luo [c], Xianfeng Gu [a]

[a] Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA
[b] Department of Mathematics, Oregon State Univer
[c] Department of Mathematics, Rutgers University,

Theorem 3.5. *Suppose two Euclidean polyhedral surfaces* $(S,T,\mathbf{d}_1)$ *and* $(S,T,\mathbf{d}_2)$ *are given,*

$$L_1 = L_2,$$

if and only if $\mathbf{d}_1$ *and* $\mathbf{d}_2$ *differ by a scaling.*

ARTICLE INFO

Article history:
Received 5 March 2012
Accepted 28 March 2012
Available online 12 April 2012

Keywords:
Discrete heat kernel
Discrete Riemannian metric
Laplace–Beltrami operator
Legendre duality principle
Discrete curvature flow

crete heat kernel and the discrete Riemannian metric (unique up to a scaling) are mutually determined by each other. Given a Euclidean polyhedral surface, its Riemannian metric is represented as edge lengths, satisfying triangle inequalities on all faces. The Laplace–Beltrami operator is formulated using the cotangent formula, where the edge weight is defined as the sum of the cotangent of angles against the edge. We prove that the edge lengths can be determined by the edge weights unique up to a scaling using the variational approach.

The constructive proof leads to a computational algorithm that finds the unique metric on a triangle mesh from a discrete Laplace–Beltrami operator matrix.

Published by Elsevier Inc.

1. Introduction

Laplace–Beltrami operator plays a fundamental role in Riemannian geometry [26]. Discrete Laplace–Beltrami operators on triangulated surface meshes span the entire spectrum of geometry processing applications, including

1.1. Motivation

The Laplace–Beltrami operator on a Riemannian manifold plays an fundamental role in Riemannian geometry. The spectrum of its eigenvalues encodes the Riemannian metric information, the nodal lines of its eigenfunctions re-

Functional Characterization of Intrinsic and Extrinsic Geometry

ETIENNE CORMAN*
LIX, École Polytechnique
JUSTIN SOLOMON*
Massachusetts Institute of Technology
MIRELA BEN-CHEN
Technion — Israel Institute of Technology
LEONIDAS GUIBAS
Stanford University
MAKS OVSJANIKOV
LIX, École Polytechnique

We propose a novel way to capture and characterize distortion between pairs of shapes by extending the recently proposed framework of shape differences built on functional maps. We modify the original definition of shape differences slightly and prove that, after this change, the discrete metric is fully encoded in two shape difference operators and can be recovered by solving two *linear* systems of equations. Then, we introduce an extension of the shape difference operators using offset surfaces to capture extrinsic or embedding-dependent distortion, complementing the purely intrinsic nature of the original shape differences. Finally, we demonstrate that a set of four operators is *complete*, capturing intrinsic and extrinsic structure and fully encoding a shape up to rigid motion in both discrete and continuous settings. We highlight the usefulness of our constructions by showing the complementary nature of our extrinsic shape differences in capturing distortion ignored by previous approaches. We additionally provide examples where we recover local shape structure from the shape difference operators, suggesting shape editing and analysis tools based on manipulating shape differences.

along the surface, whereas extrinsic quantities are those that must be defined using surface normal vectors and/or an embedding into space. A crowning result of classical differential geometry describes local geometry in terms of two quantities: the first and second fundamental forms, which capture the intrinsic Gaussian and extrinsic mean curvatures, respectively [Bonnet 1867].

Considerable research in geometry processing has been dedicated to measuring intrinsic and extrinsic curvature in an attempt to replicate this attractive characterization of shape. From a practical standpoint, however, this task remains challenging for potentially noisy or irregular meshes considered in geometry processing. After all, surface curvature is a second-derivative quantity whose approximation on a piecewise-linear mesh requires discretization and mollification to deal with noise. Measurement of curvature aside, algorithms for recovering geometry from discrete curvatures remain difficult to formulate for many discretizations.

In this paper, we formulate an alternative characterization of ...ed for analysis, comparison, and synthesis ...tting. Several desiderata inform our design; a ...r representing shape should

...uish intrinsic and extrinsic geometry,
...erties in a multiscale fashion to distinguish ... detail from large-scale structure,
...h theory of shape,
...y to tessellation,
...sible on continuous surfaces and on triangle ..., and
—admit an inverse operator for reconstructing the embedded shape

PROPOSITION 1. *Suppose* M *has a boundary or at least one interior vertex with odd valence. Then,* $A(\mu)$ *uniquely determines* μ*, recoverable via a linear solve.*

PROPOSITION 2. *Assume that the mesh* M *is manifold without boundary. Then, for almost all choices of areas* μ*, the map* $C(\ell^2;\mu)$ *uniquely determines* ℓ*, which is recoverable via a linear solve.*

Beware

Figure 1: Deformations of a glove (left) and a solid hand (right) are an illustration of the difference between boundary and volume isometries.

Image from: Raviv et al. "Volumetric Heat Kernel Signatures." 3DOR 2010.

But calculations on a volume are expensive!

(changing!)

Not the same.

Why Study the Laplacian?

- **Encodes intrinsic geometry**
  Edge lengths on triangle mesh, Riemannian metric on manifold
- **Multi-scale**
  Filter based on frequency
- **Geometry through linear algebra**
  Linear/eigenvalue problems, sparse positive definite matrices
- **Connection to physics**
  Heat equation, wave equation, vibration, …

Our Next Topic

Discrete Laplacian operators:

What are they good for?

- Useful properties of the Laplacian
- Applications in graphics/shape analysis
- Applications in machine learning

A quick survey:
A popular field!

Example Task: Shape Descriptors (Features)

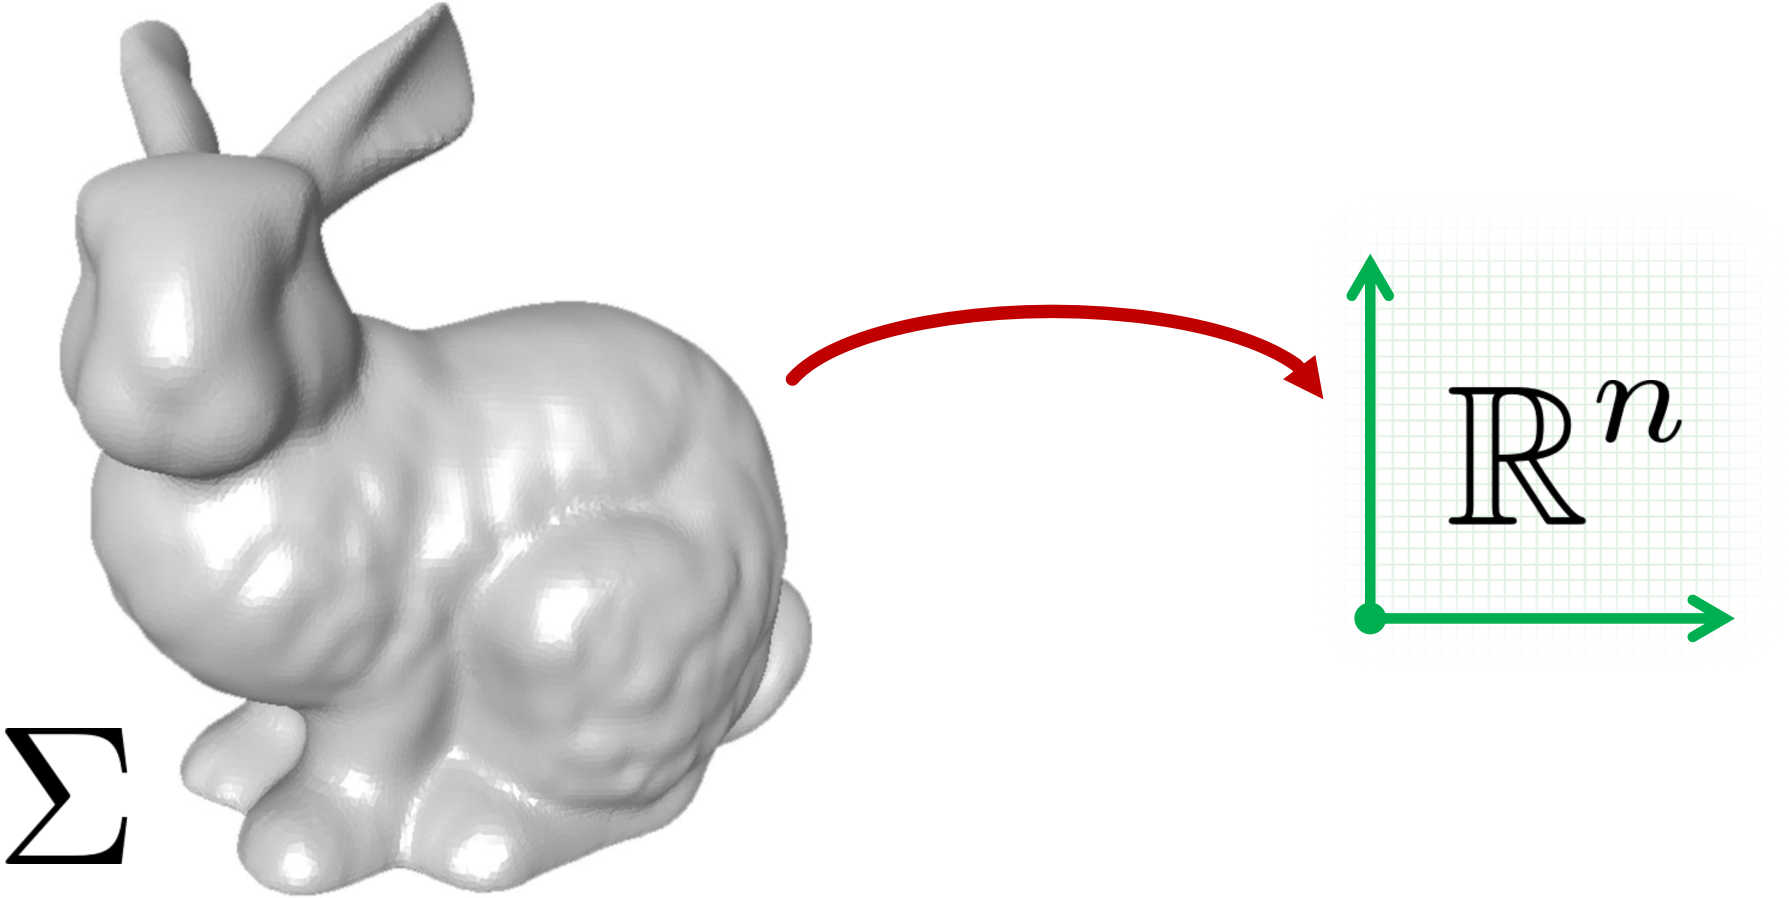

http://liris.cnrs.fr/meshbenchmark/images/fig_attacks.jpg

Pointwise quantity

Descriptor Tasks

- **Characterize local geometry**
  Feature/anomaly detection
- **Describe point's role on surface**
  Symmetry detection, correspondence

Descriptors We've Seen Before

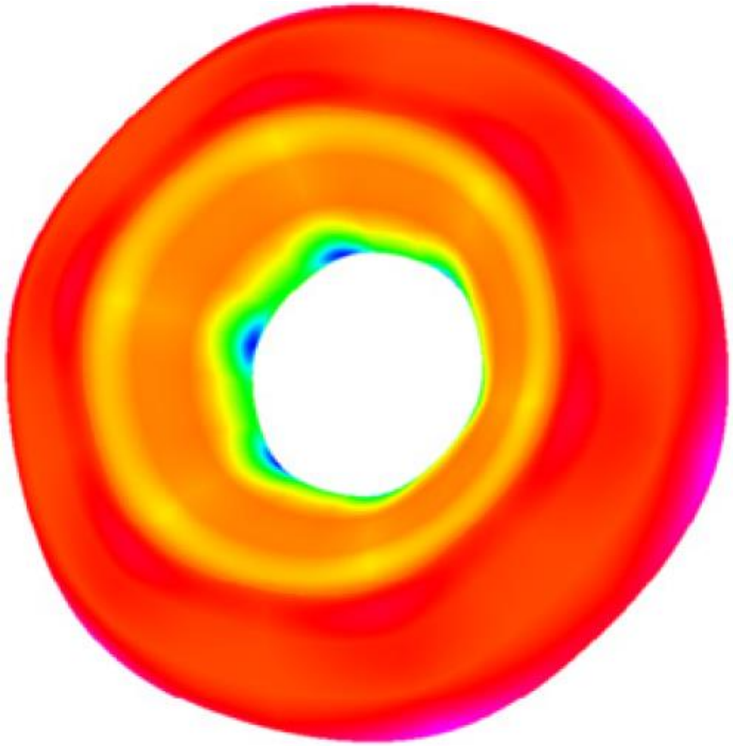

$$K := \kappa_1 \kappa_2 = \det \mathrm{I\!I}$$

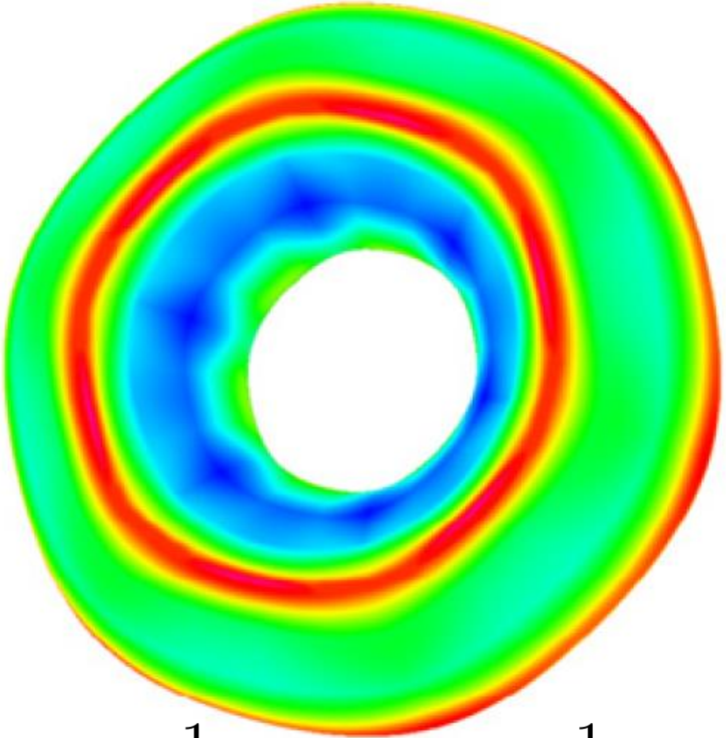

$$H := \frac{1}{2}(\kappa_1 + \kappa_2) = \frac{1}{2}\mathrm{tr}\,\mathrm{I\!I}$$

http://www.sciencedirect.com/science/article/pii/S0010448510001983

Gaussian and mean curvature

Desirable Properties

- **Distinguishing**

  Provides useful information about a point

- **Stable**

  Numerically and geometrically

- **Intrinsic**

  No dependence on embedding

  Sometimes undesirable!

Intrinsic Descriptors

Invariant under

- Rigid motion
- Bending without stretching

Intrinsic Descriptor

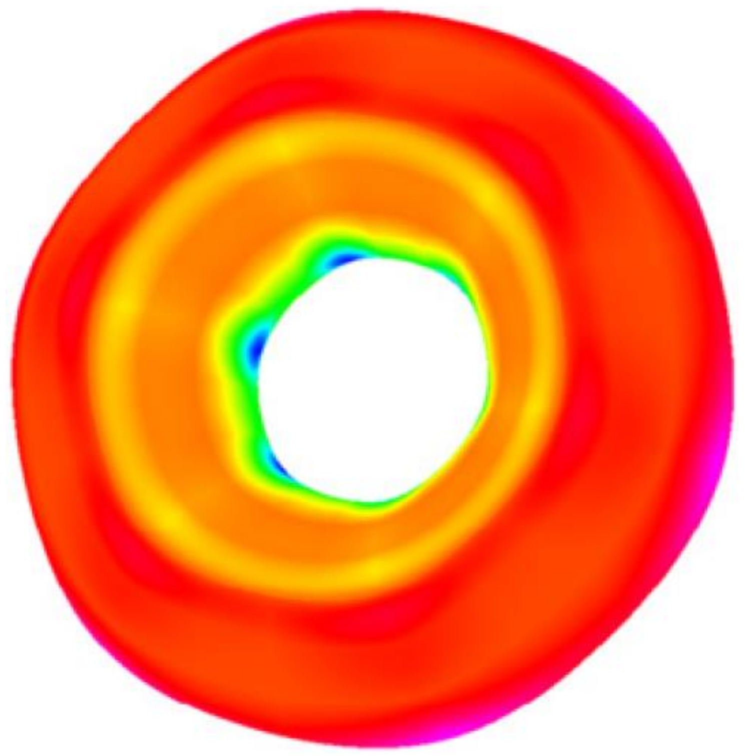

$$K := \kappa_1 \kappa_2 = \det \mathrm{I\!I}$$

Theorema Egregium

("Totally Awesome Theorem"):

Gaussian curvature is intrinsic.

http://www.sciencedirect.com/science/article/pii/S0010448510001983

Gaussian curvature

End of the Story?

$$K = \kappa_1 \kappa_2$$

Second derivative quantity

End of the Story?

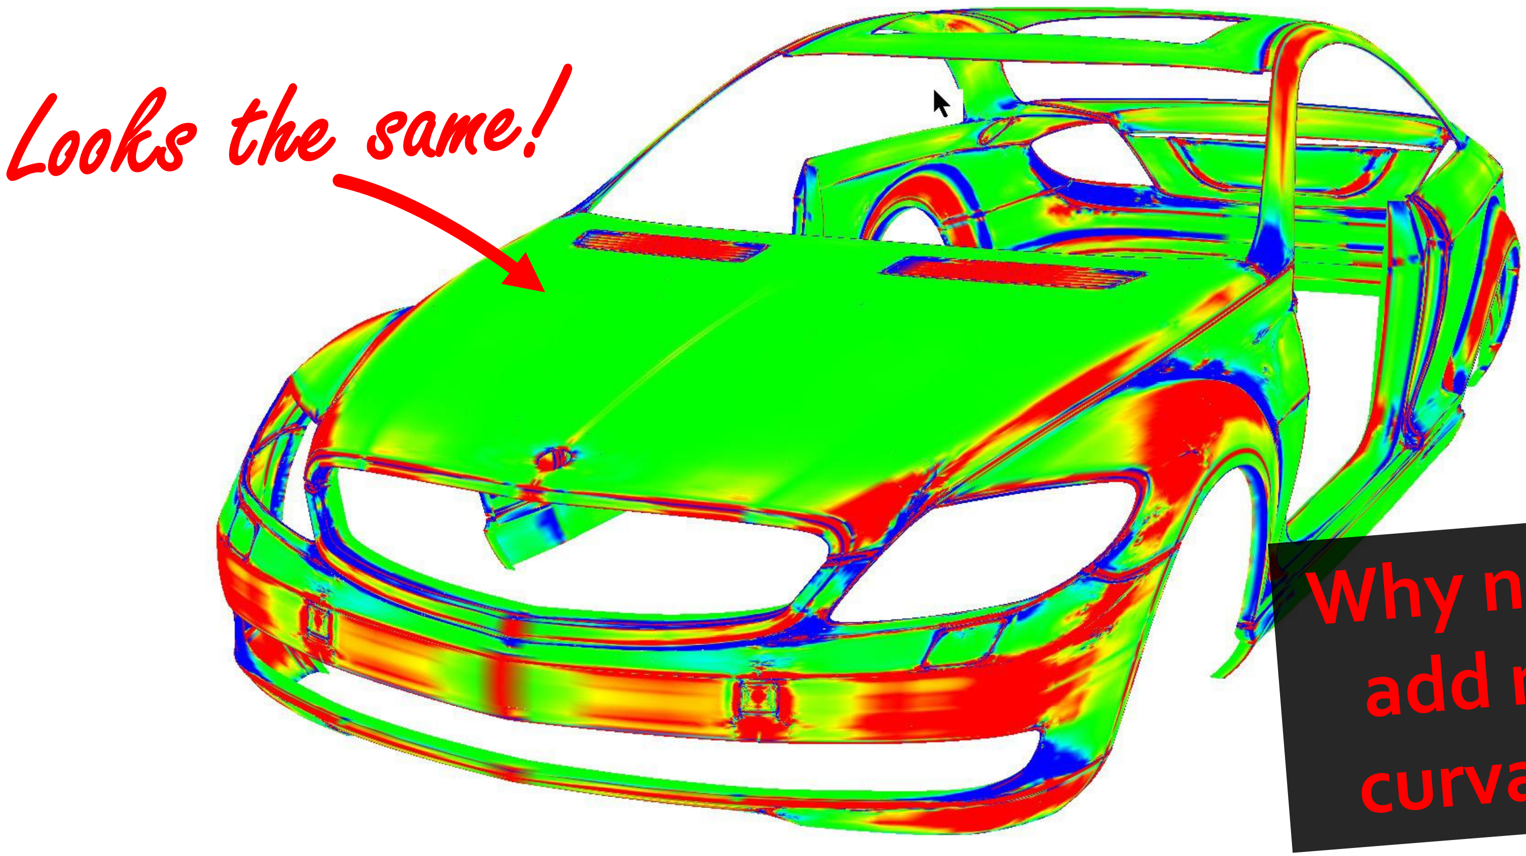

Why not just add mean curvature?

http://www.integrityware.com/images/MerceedesGaussianCurvature.jpg

Non-unique

Desirable Properties

Incorporates neighborhood information in an intrinsic fashion

Stable under small deformation

Global Point Signature

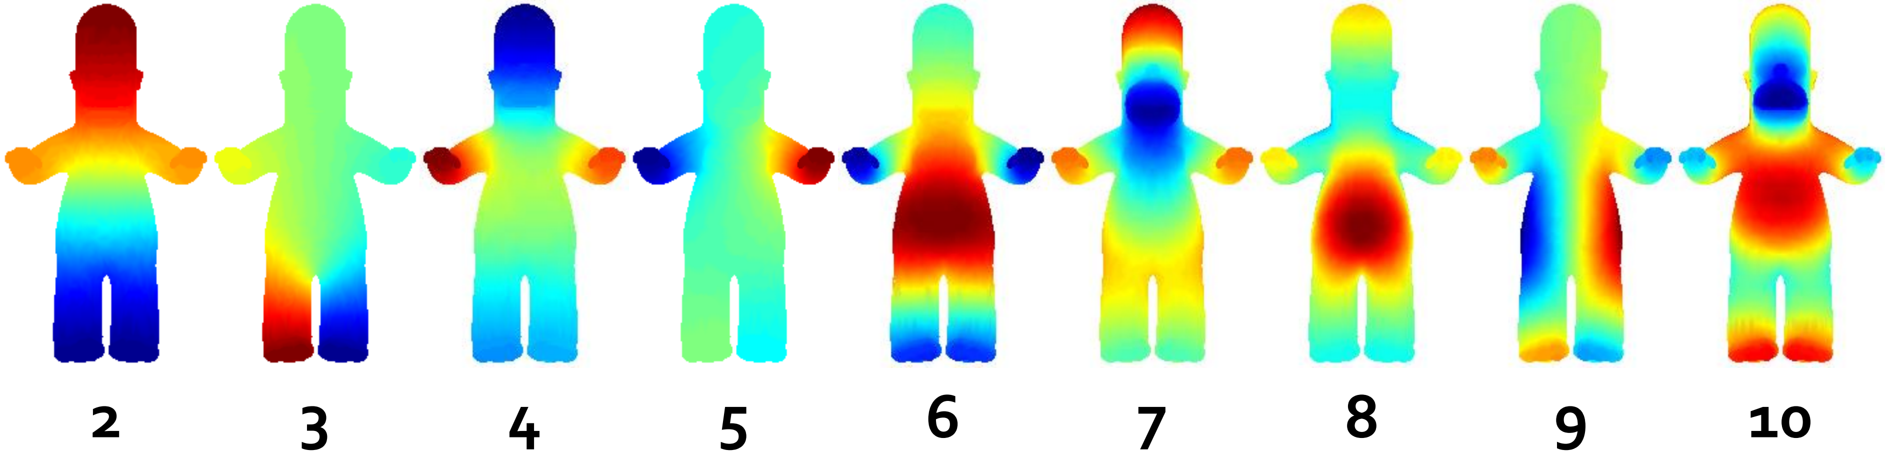

$$\mathrm{GPS}(p) := \left(-\frac{1}{\sqrt{\lambda_1}}\phi_1(p), -\frac{1}{\sqrt{\lambda_2}}\phi_2(p), -\frac{1}{\sqrt{\lambda_3}}\phi_3(p), \cdots \right)$$

"Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation"
Rustamov, SGP 2007

Global Point Signature

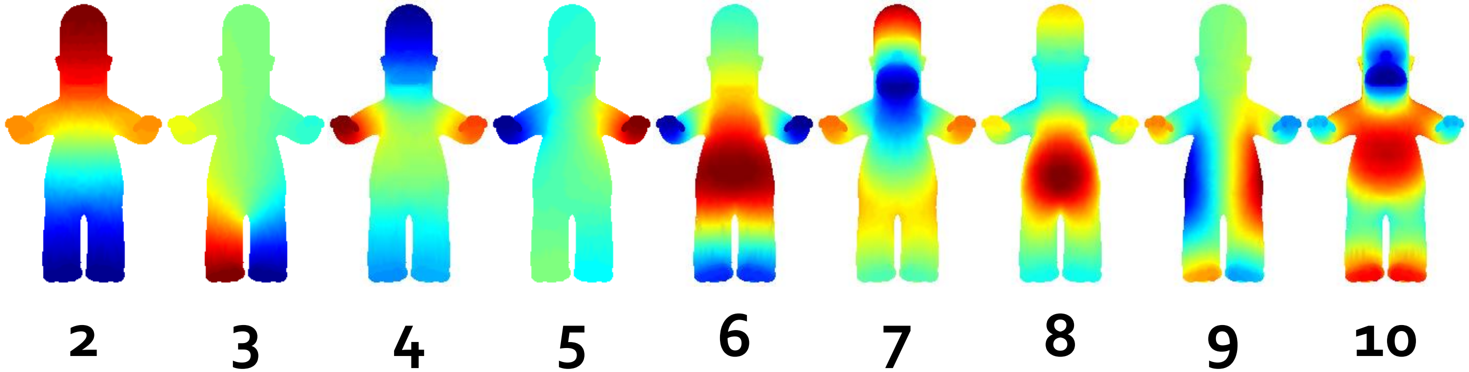

$$\mathrm{GPS}(p) := \left(-\frac{1}{\sqrt{\lambda_1}}\phi_1(p), -\frac{1}{\sqrt{\lambda_2}}\phi_2(p), -\frac{1}{\sqrt{\lambda_3}}\phi_3(p), \cdots\right)$$

If surface does not self-intersect, neither does the GPS embedding.

Proof: Laplacian eigenfunctions span $L^2(\Sigma)$; if GPS(p)=GPS(q), then all functions on Σ would be equal at p and q.

Global Point Signature

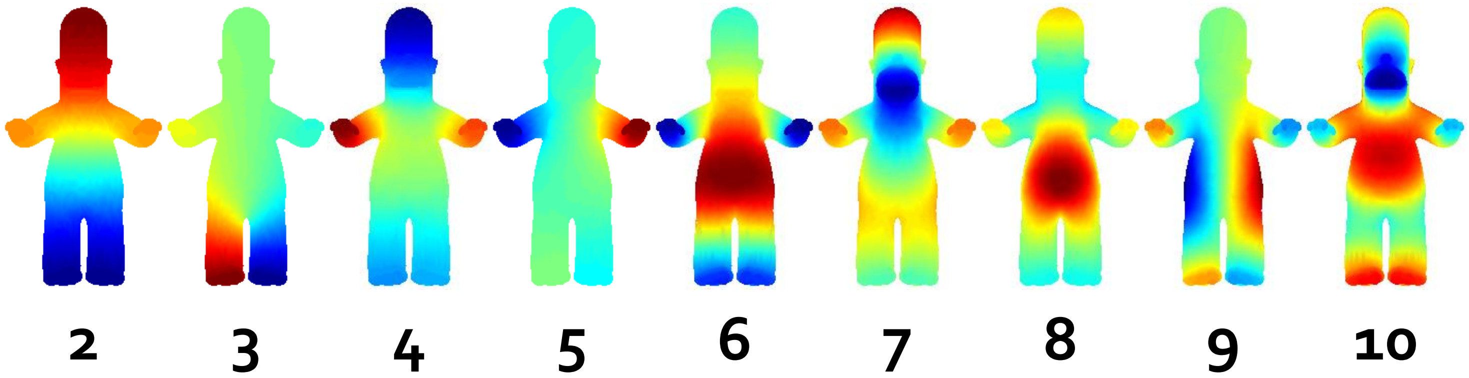

$$\mathrm{GPS}(p) := \left(-\frac{1}{\sqrt{\lambda_1}}\phi_1(p), -\frac{1}{\sqrt{\lambda_2}}\phi_2(p), -\frac{1}{\sqrt{\lambda_3}}\phi_3(p), \cdots \right)$$

GPS is isometry-invariant.

Proof: Comes from the Laplacian.

New inspiration:

Physics Applications of the Laplacian

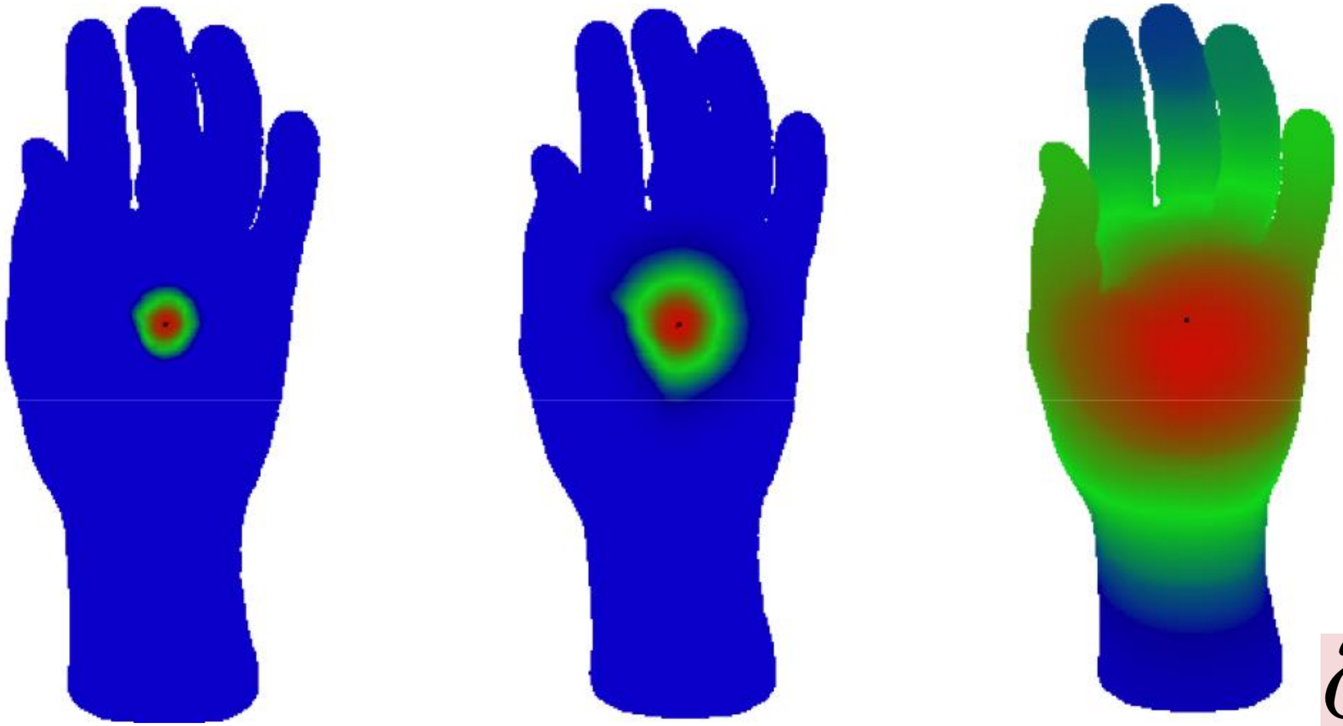

$$\frac{\partial u}{\partial t} = -\Delta u$$

http://graphics.stanford.edu/courses/cs468-10-fall/LectureSlides/11_shape_matching.pdf

Heat equation

Physics Applications of the Laplacian

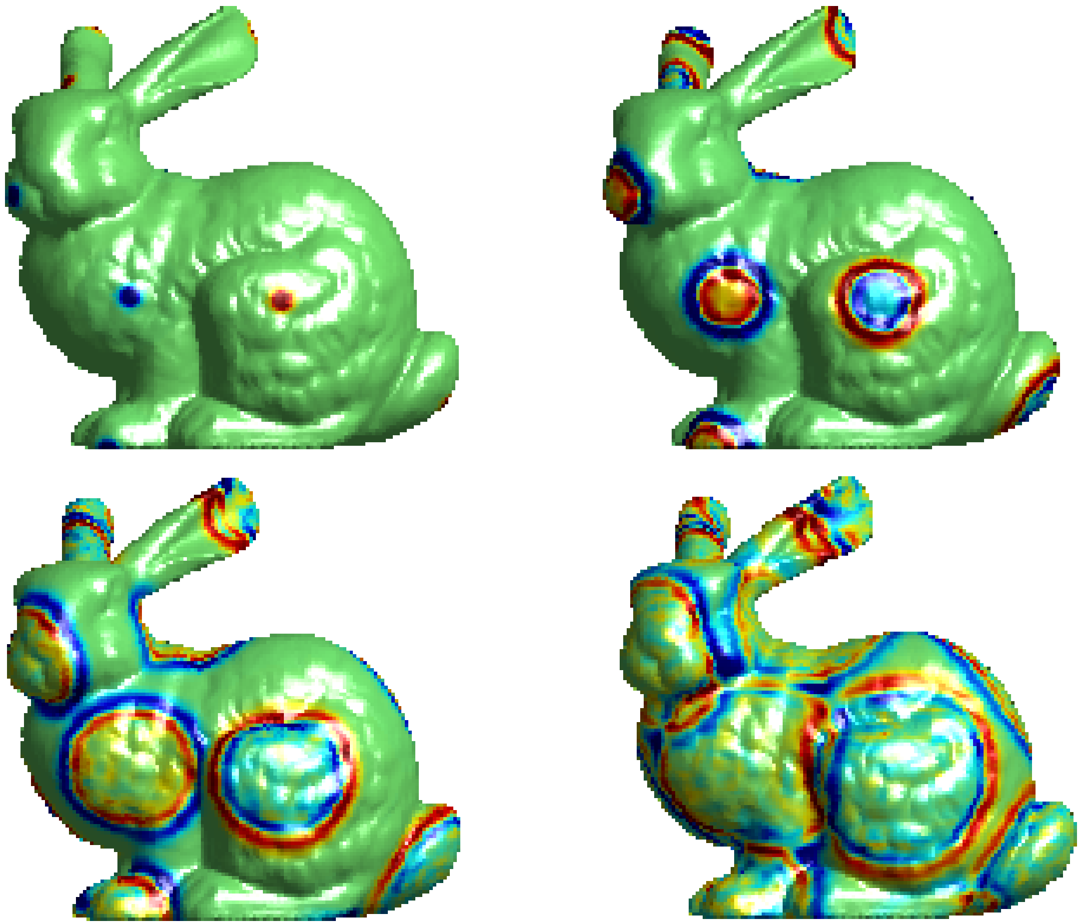

$$\frac{\partial^2 u}{\partial t^2} = -\Delta u$$

Image courtesy G. Peyré

Wave equation

Solutions in the LB Basis

$$\frac{\partial u}{\partial t} = -\Delta u$$

Heat equation

$$u = \sum_{n=0}^{\infty} a^n e^{-\lambda_n t} \phi_n(x)$$

$$a^n = \int_{\Sigma} u_0(x) \cdot \phi_n(x)\, dA = \langle u_0, \phi_n \rangle$$

Heat Kernel Signature (HKS)

$$k_t(x, x) = \sum_{n=0}^{\infty} e^{-\lambda_i t} \phi_n(x)^2$$

Continuous function of $t \in [0, \infty)$

How much heat diffuses from *x* to itself in time *t*?

Heat Kernel Signature (HKS)

$$k_t(x,x) = \sum_{n=0}^{\infty} e^{-\lambda_i t} \phi_n(x)^2$$

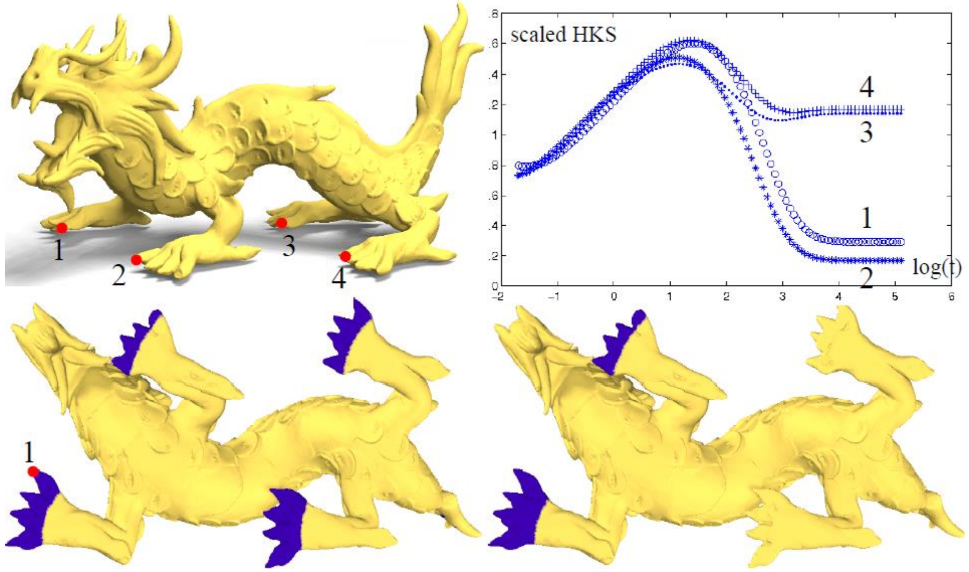

"A concise and provably informative multi-scale signature based on heat diffusion"
Sun, Ovsjanikov, and Guibas; SGP 2009

Heat Kernel Signature (HKS)

$$k_t(x,x) = \sum_{n=0}^{\infty} e^{-\lambda_i t} \phi_n(x)^2$$

Good properties:

- Isometry-invariant
- Multiscale
- Not subject to switching
- Easy to compute
- Related to curvature at small scales

Heat Kernel Signature (HKS)

$$k_t(x,x) = \sum_{n=0}^{\infty} e^{-\lambda_i t} \phi_n(x)^2$$

Bad properties:

- Issues remain with repeated eigenvalues
- Theoretical guarantees require (near-)isometry

Wave Kernel Signature (WKS)

$$\mathrm{WKS}(E, x) = \lim_{T \to \infty} \frac{1}{T} \int_0^T |\psi_E(x,t)|^2 \, dt = \sum_{n=0}^{\infty} \phi_n(x)^2 f_E(\lambda_n)^2$$

Initial energy distribution

Average probability over time that particle is at *x*.

"The Wave Kernel Signature: A Quantum Mechanical Approach to Shape Analysis"
Aubry, Schlickewei, and Cremers; ICCV Workshops 2012

Wave Kernel Signature (WKS)

$$\mathrm{WKS}(E,x) = \lim_{T\to\infty} \frac{1}{T}\int_0^T |\psi_E(x,t)|^2\,dt = \sum_{n=0}^{\infty} \phi_n(x)^2 f_E(\lambda_n)^2$$

HKS

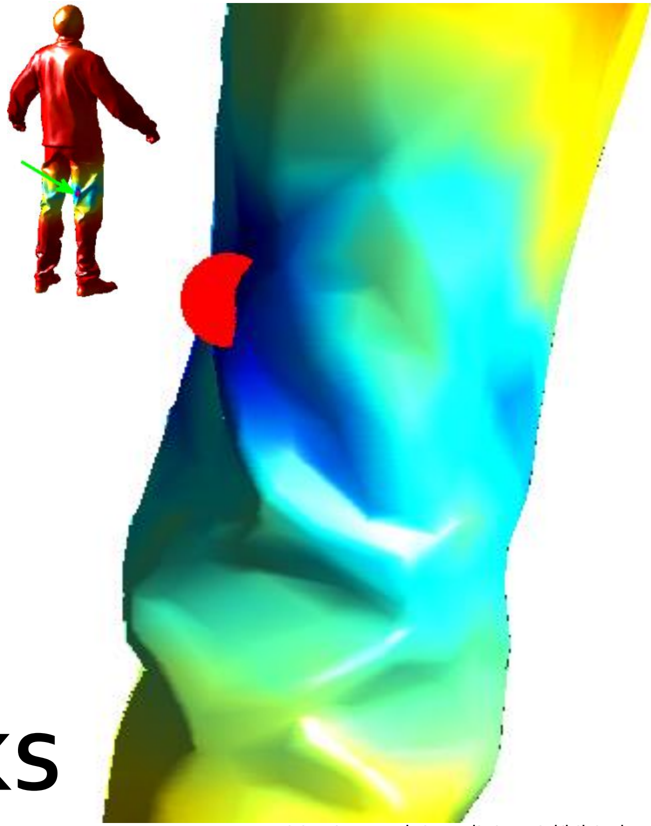

WKS

vision.in.tum.de/_media/spezial/bib/aubry-et-al-4dmod11.pdf

Wave Kernel Signature (WKS)

$$\mathrm{WKS}(E,x) = \lim_{T\to\infty} \frac{1}{T}\int_0^T |\psi_E(x,t)|^2\, dt = \sum_{n=0}^{\infty} \phi_n(x)^2 f_E(\lambda_n)^2$$

Good properties:

- [Similar to HKS]
- Localized in frequency
- Stable under some non-isometric deformation
- Some multi-scale properties

Wave Kernel Signature (WKS)

$$\mathrm{WKS}(E,x) = \lim_{T\to\infty} \frac{1}{T}\int_0^T |\psi_E(x,t)|^2\, dt = \sum_{n=0}^{\infty} \phi_n(x)^2 f_E(\lambda_n)^2$$

Bad properties:

- [Similar to HKS]
- Can filter out *large*-scale features

Many Others

Lots of spectral descriptors in terms of Laplacian eigenstructure.

Combination with Machine Learning

$$p(x) = \sum_k f(\lambda_k)\phi_k^2(x)$$

Learn *f* rather than defining it

Fig. 3. Correspondences computed on TOSCA shapes using the spectral matching algorithm [30]. Shown are the matches with geodesic distance distortion below 10 percent of the shape diameter, from left to right: HKS (34 matches), WKS (30 matches), and trained descriptor (54 matches).

Learning Spectral Descriptors for Deformable Shape Correspondence

Litman and Bronstein; PAMI 2014

Even More Generic Version

HodgeNet: Learning Spectral Geometry on Triangle Meshes

DMITRIY SMIRNOV and JUSTIN SOLOMON, Massachusetts Institute of Technology, USA

Fig. 1. Mesh segmentation results on the full-resolution MIT animation dataset. Each mesh in the dataset contains 20,000 faces (10,000 vertices). We show an example ground truth segmentation in the bottom-left. In contrast to previous works, which downsample each mesh by more than 10×, we efficiently process dense meshes both at train and test time.

Constrained by the limitations of learning toolkits engineered for other applications, such as those in image processing, many mesh-based learning algorithms employ data flows that would be atypical from the perspective of conventional geometry processing. As an alternative, we present a technique for learning from meshes built from standard geometry processing modules and operations. We show that low-order eigenvalue/eigenvector computation from operators parameterized using discrete exterior calculus is amenable to efficient approximate backpropagation, yielding spectral per-element or per-mesh features with similar formulas to classical descriptors like the heat/wave kernel signatures. Our model uses few parameters, generalizes to high-resolution meshes, and exhibits performance and time complexity on par with past work.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; **Mesh geometry models**; **Neural networks**.

Additional Key Words and Phrases: Machine learning, meshes, operators

ACM Reference Format:
Dmitriy Smirnov and Justin Solomon. 2021. HodgeNet: Learning Spectral Geometry on Triangle Meshes. *ACM Trans. Graph.* 40, 4, Article 166 (August 2021), 11 pages. https://doi.org/10.1145/3450626.3459797

Numerous technical challenges preclude modern learning methods from being adopted for meshes. Deep learning—arguably the most popular recent learning methodology—relies on *regularity* of the data and *differentiability* of the objective function for efficiency. For example, convolutional neural network (CNN) training is built on high-throughput processing of images through convolution and per-pixel computations to obtain gradients with respect to network weights, required for stochastic gradient descent.

Meshes, a primary means of representing geometry in graphics, defy the considerations above. They come as sparse, irregular networks of vertices varying in number; the same piece of geometry easily can be represented by multiple meshes and at multiple resolutions/densities. Advances in graph neural networks (GNNs) have as a byproduct helped advance mesh processing, but typical graphs in geometry processing are fundamentally different from those in network science—vertices have low valence, are related through long chains of edges, can be connected in many roughly-equivalent ways, and can be deformed through rigid motions and isometries.

The end result is that mesh-based learning architectures often

Application: Feature Extraction

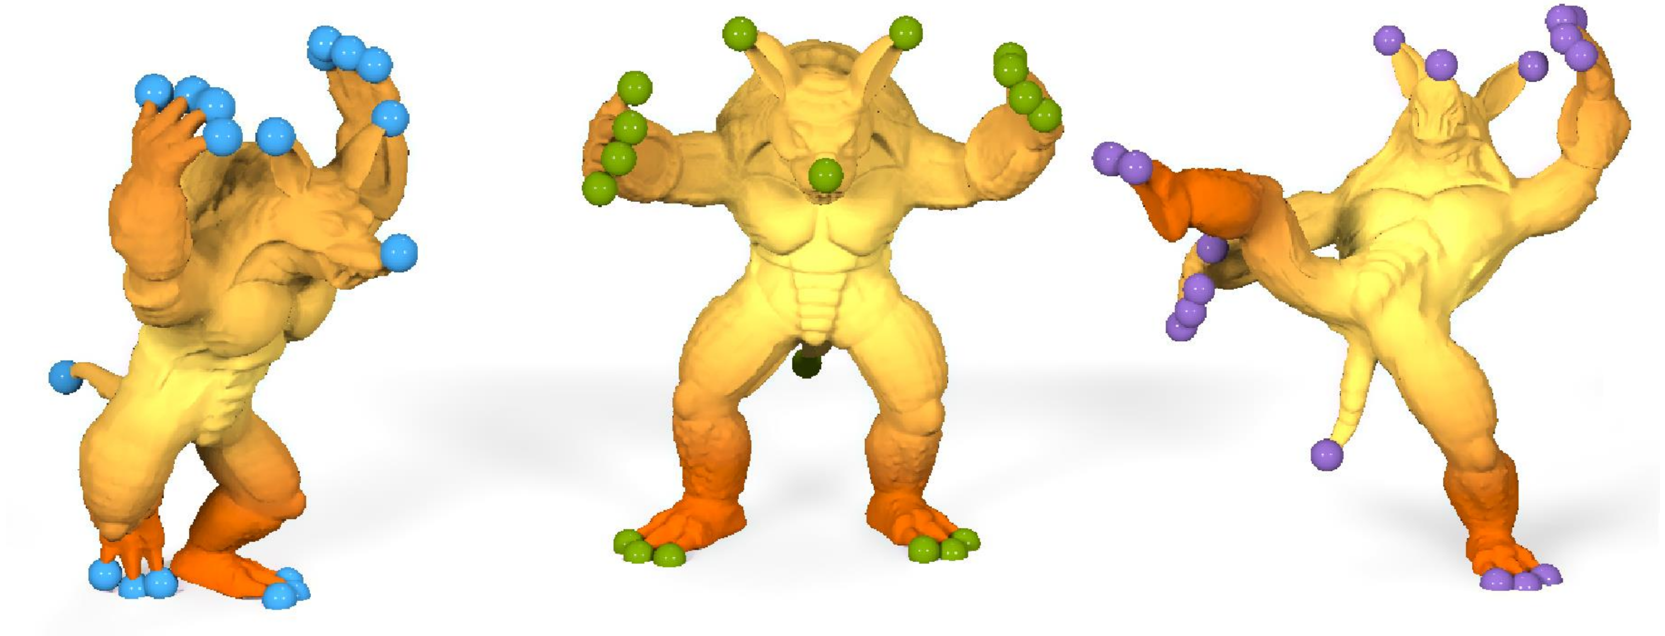

Maxima of $k_t(x,x)$ over x for large t.

A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion
Sun, Ovsjanikov, and Guibas; SGP 2009

Feature points

Preview: Correspondence

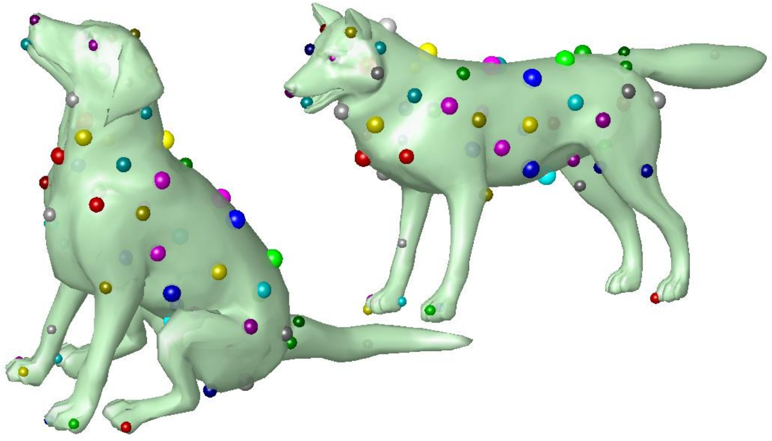

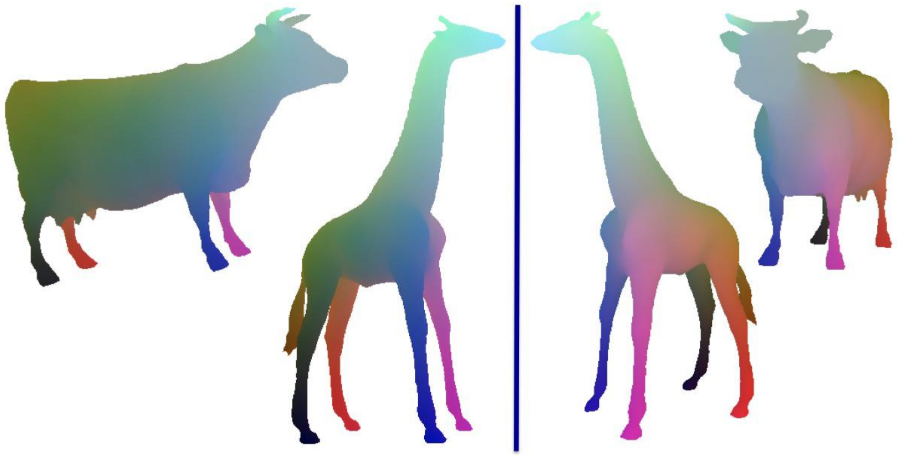

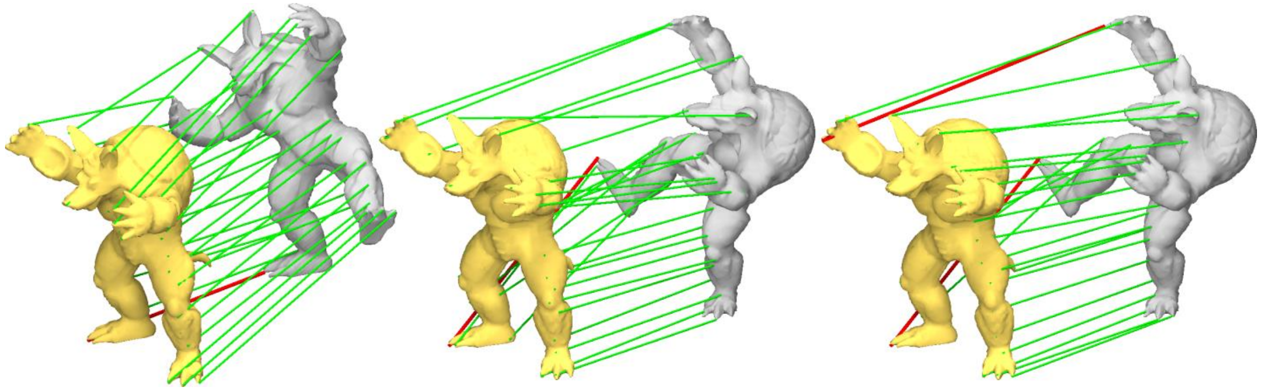

http://graphics.stanford.edu/projects/lgl/papers/ommg-opimhk-10/ommg-opimhk-10.pdf
http://www.cs.princeton.edu/~funk/sig11.pdf
http://gfx.cs.princeton.edu/pubs/Lipman_2009_MVF/mobius.pdf

Descriptor Matching

Simply match **closest points** in descriptor space.

Heat Kernel Map

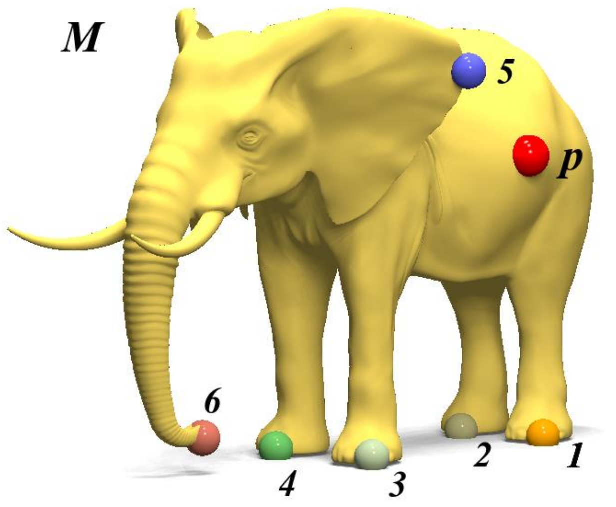

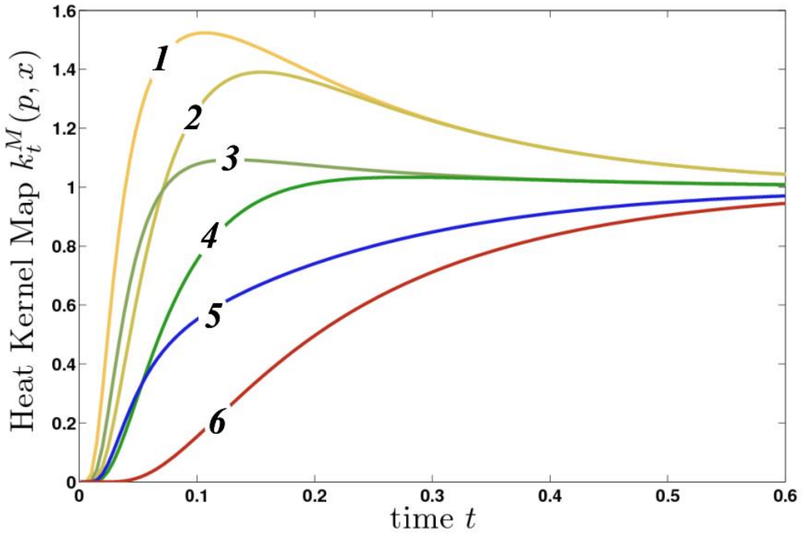

$$\mathrm{HKM}_p(x,t) := k_t(p,x)$$

How much heat diffuses from *p* to *x* in time *t*?

One Point Isometric Matching with the Heat Kernel
Ovsjanikov et al. 2010

Heat Kernel Map

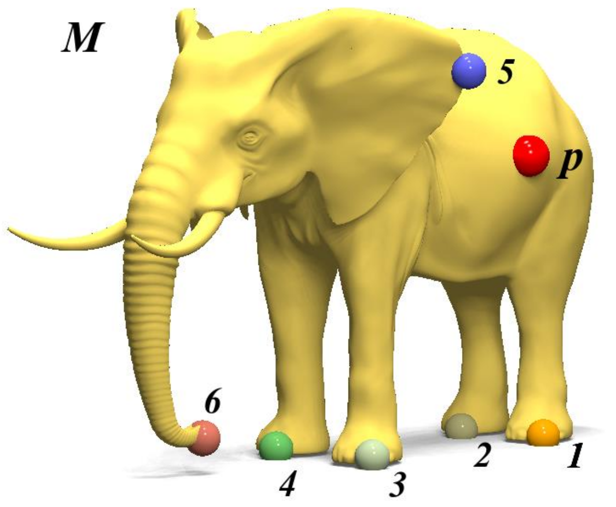

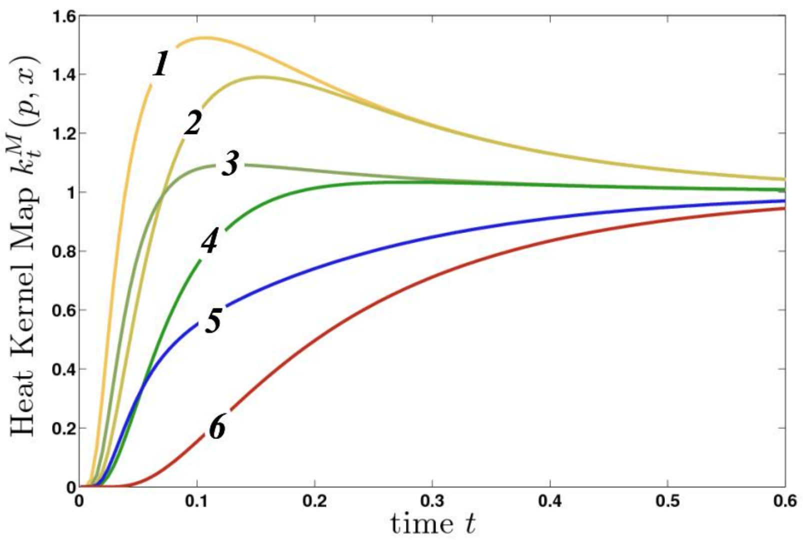

$$\mathrm{HKM}_p(x,t) := k_t(p,x)$$

Theorem: Only have to match one point!

KNN

One Point Isometric Matching with the Heat Kernel

Ovsjanikov et al. 2010

Descriptor Matching Problem

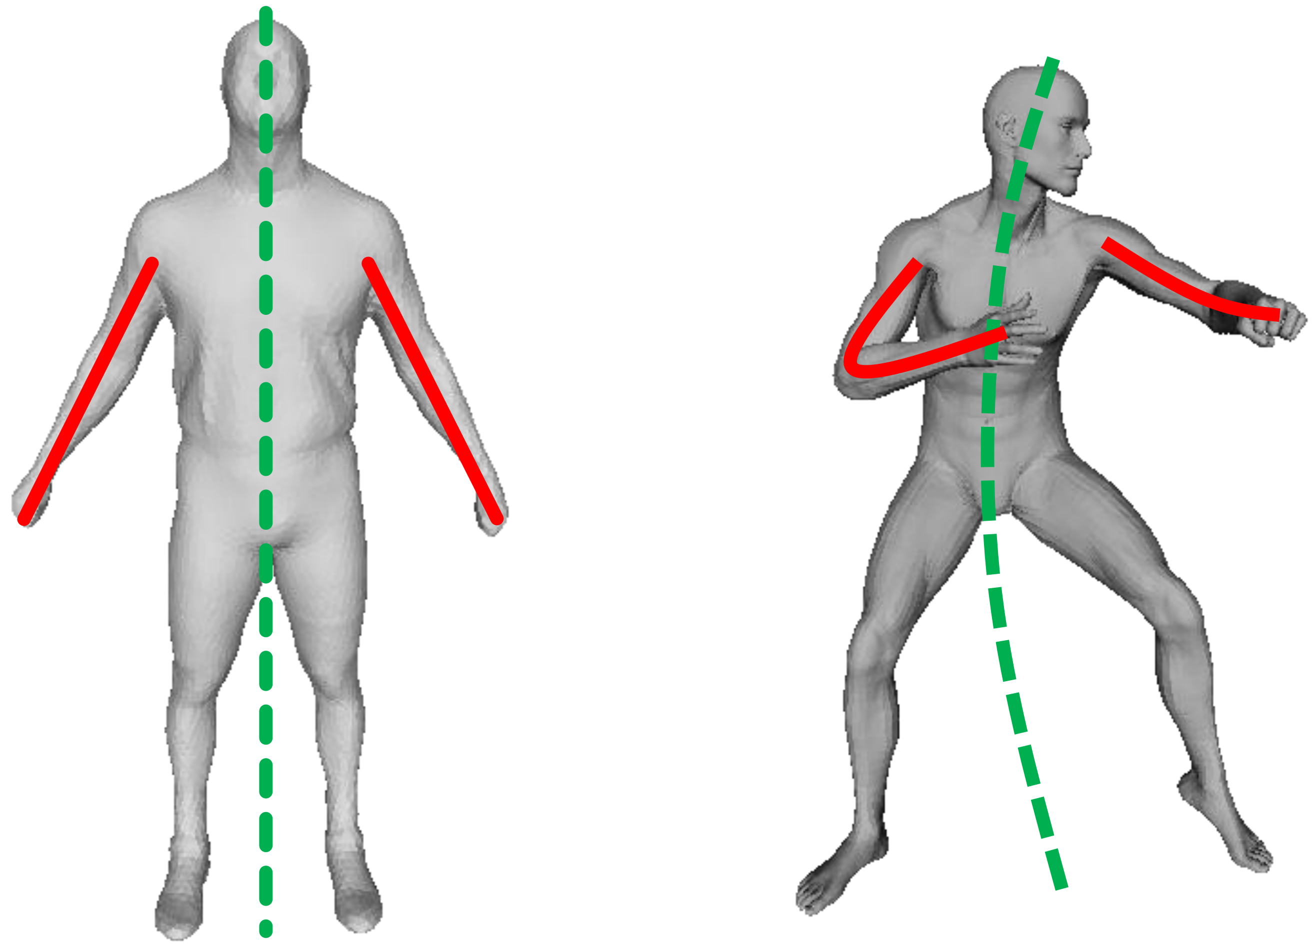

Symmetry

Self-Map: Symmetry

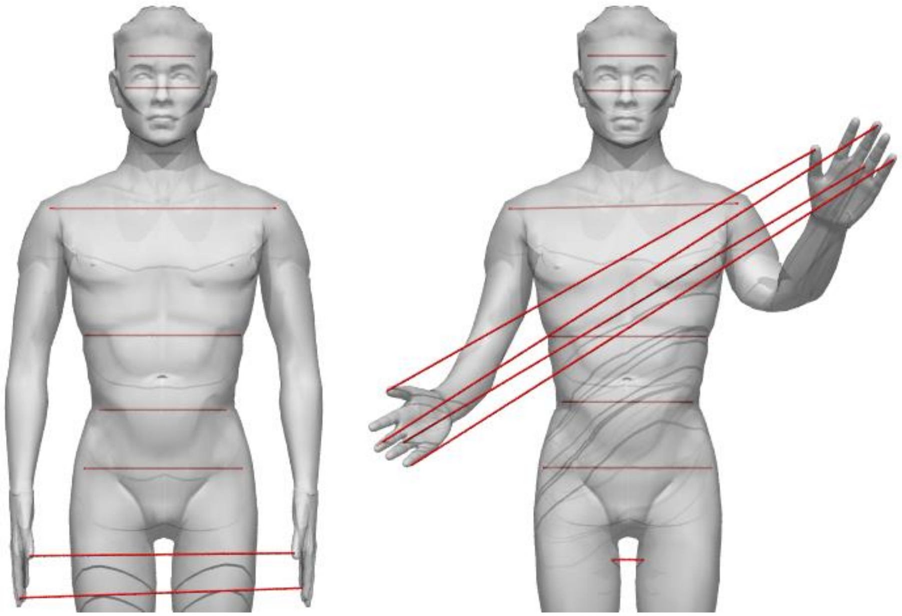

Intrinsic symmetries become extrinsic in GPS space!

Global Intrinsic Symmetries of Shapes
Ovsjanikov, Sun, and Guibas 2008

“Discrete intrinsic” symmetries

All Over the Place

Laplacians appear everywhere in shape analysis and geometry processing.

Biharmonic Distances

$$d_b(p,q) := \|g_p - g_q\|_2, \text{ where } \Delta g_p = \delta_p$$

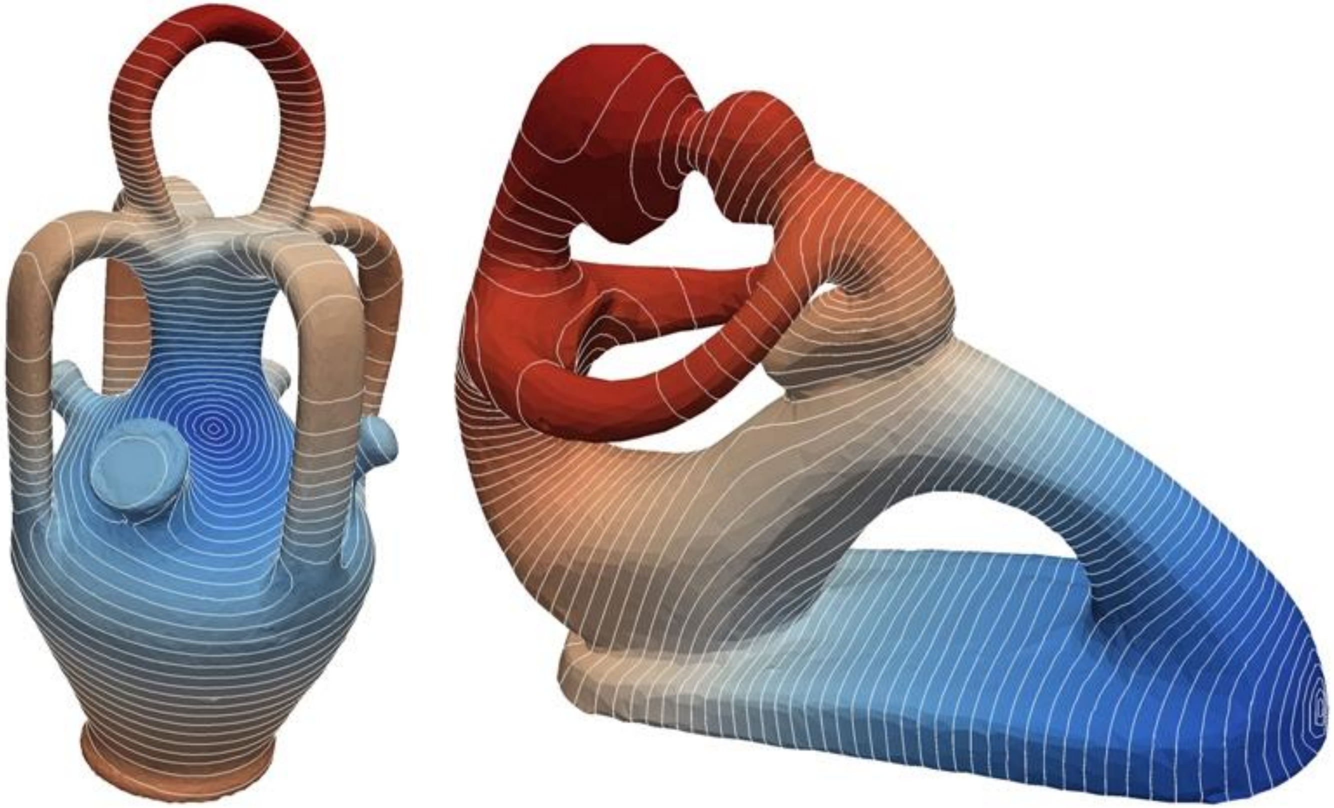

"Biharmonic distance"
Lipman, Rustamov & Funkhouser, 2010

Geodesic Distances

$$d_g(p,q) = \lim_{t\to 0} \sqrt{-4t \log k_{t,p}(q)}$$

"Varadhan's Theorem"

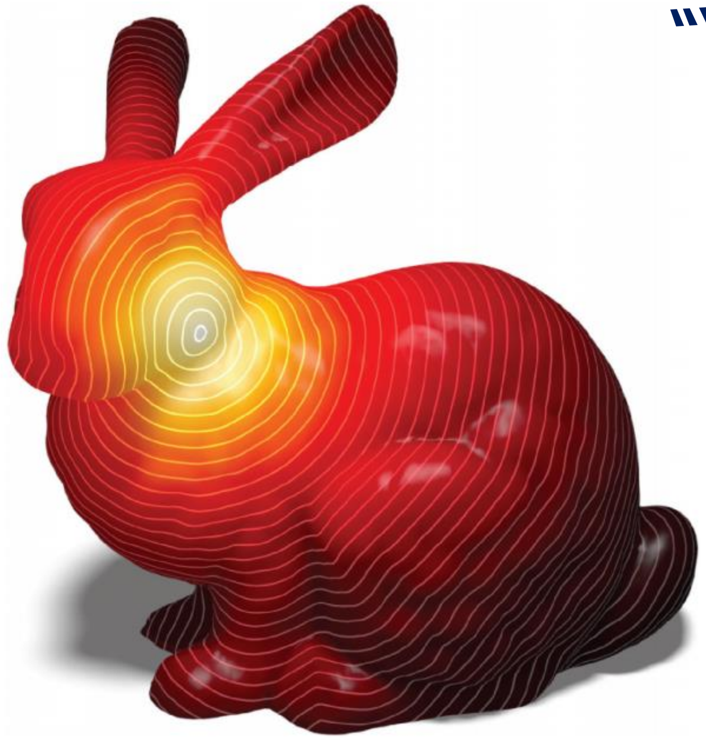

"Geodesics in heat"
Crane, Weischedel, and Wardetzky; TOG 2013

Alternative to Eikonal Equation

Algorithm 1 The Heat Method

I. Integrate the heat flow $\dot{u} = \Delta u$ for time t.
II. Evaluate the vector field $X = -\nabla u / |\nabla u|$.
III. Solve the Poisson equation $\Delta \phi = \nabla \cdot X$.

Crane, Weischedel, and Wardetzky. "Geodesics in Heat." TOG, 2013.

Implicit Fairing: Mean Curvature Flow

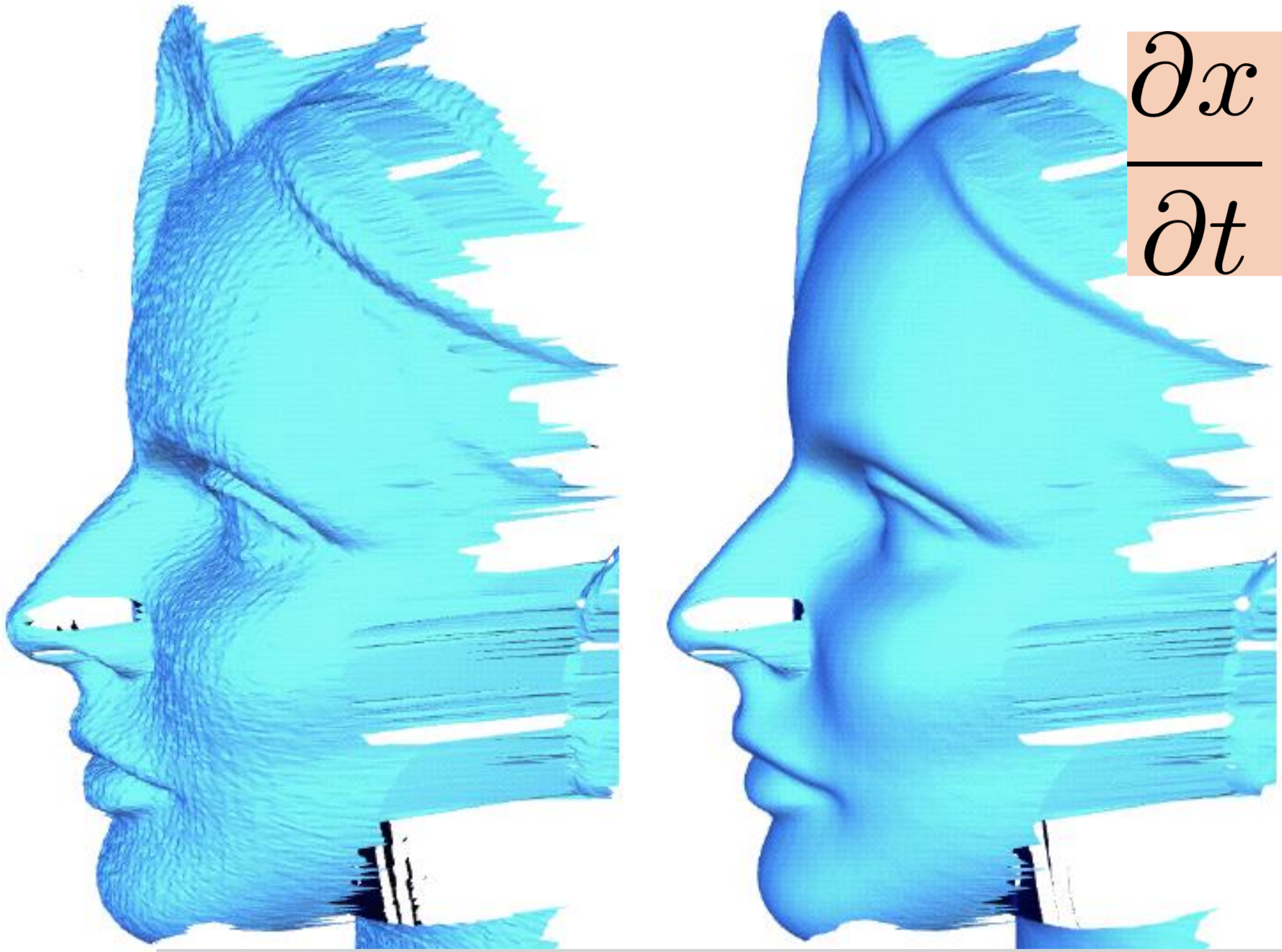

$$\frac{\partial x}{\partial t} = \Delta(x) \cdot x$$

"Implicit fairing of irregular meshes using diffusion and curvature flow"
Desbrun et al., 1999

Useful Technique

$$\frac{\partial f}{\partial t} = -\Delta f \text{ (heat equation)}$$

$$\to M\frac{\partial f}{\partial t} = Lf \text{ after discretization in space}$$

$$\to M\frac{f_T - f_0}{T} = Lf_T \text{ after time discretization}$$

Choice: Evaluate at time T

Unconditionally stable, but not necessarily accurate for large T!

(Semi-)Implicit time stepping

Parameterization: Harmonic Map

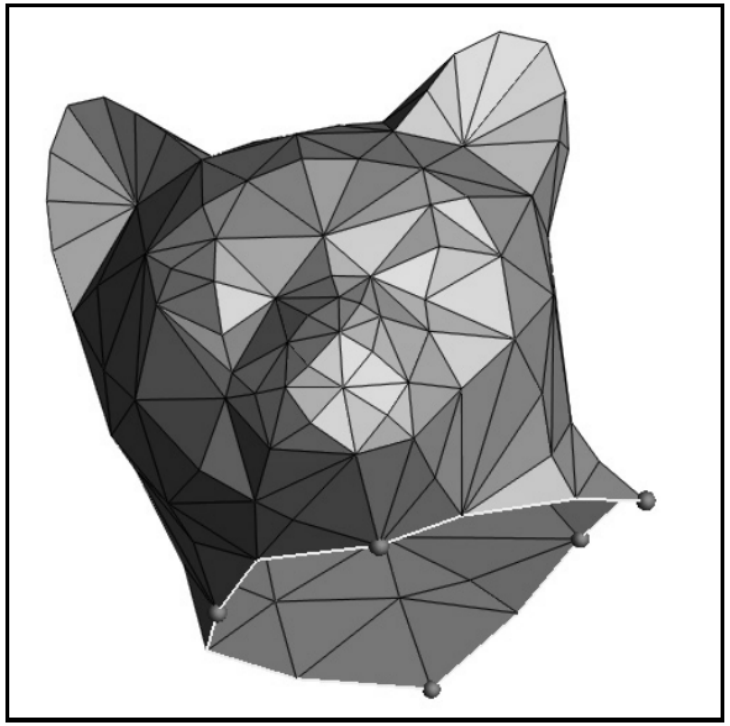

(a) Original mesh tile

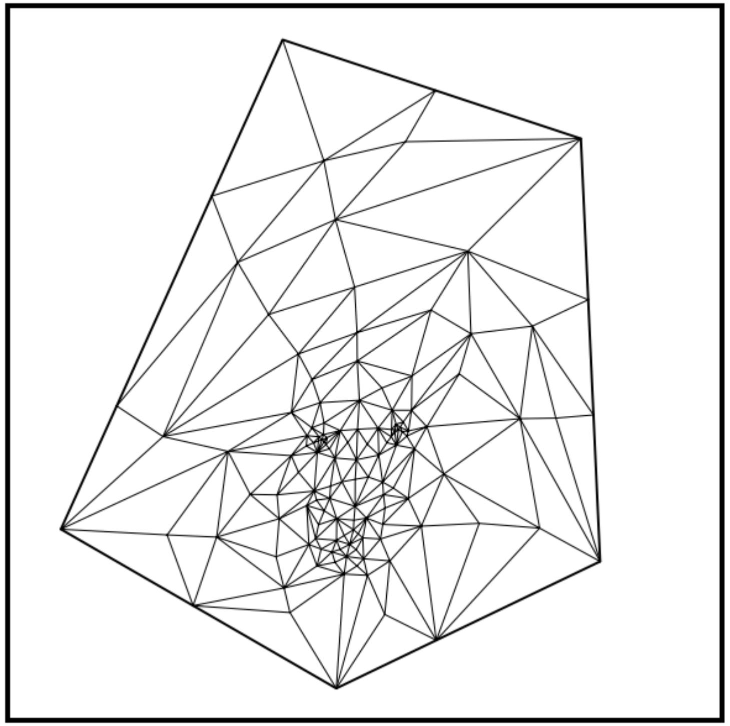

(b) Harmonic embedding

Recall:
Mean value principle

"Multiresolution analysis of arbitrary meshes"
Eck et al., 1995 (and many others!)

Others

- **Shape retrieval from Laplacian eigenvalues**
  "Shape DNA" [Reuter et al., 2006]

- **Quadrangulation**
  Nodal domains [Dong et al., 2006]

- **Surface deformation**
  "As-rigid-as-possible" [Sorkine & Alexa, 2007]

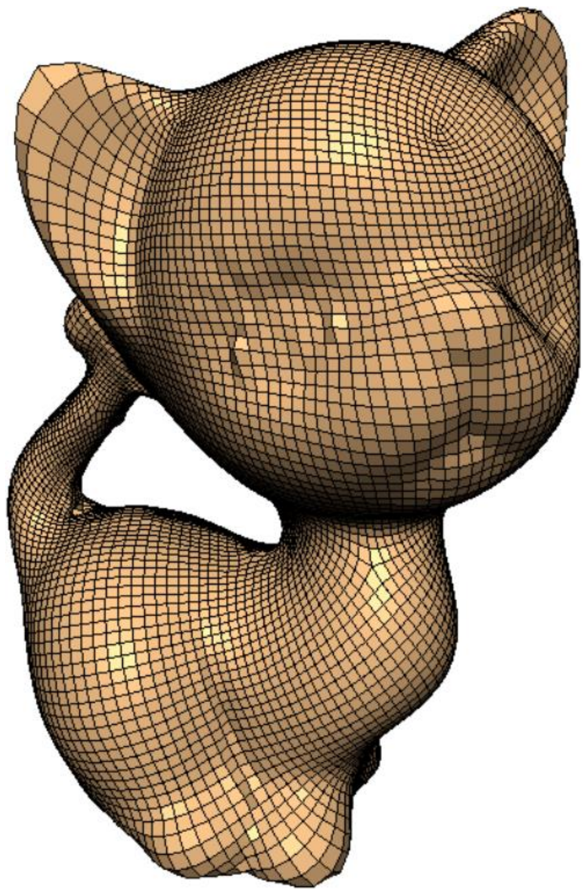

Our Next Topic

Discrete Laplacian operators:

What are they good for?

- Useful properties of the Laplacian
- Applications in graphics/shape analysis
- Applications in machine learning

A quick survey:
A popular field!

Semi-Supervised Learning

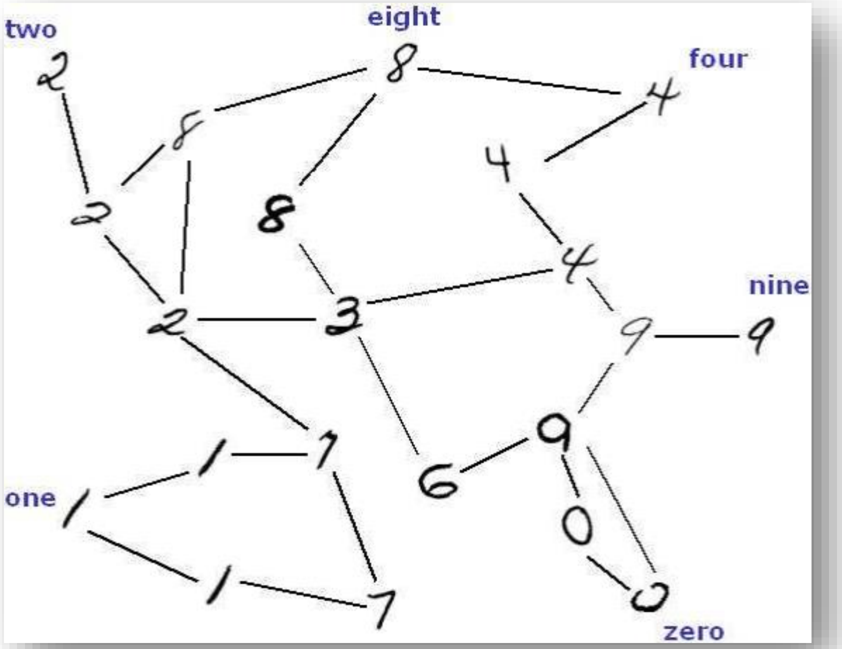

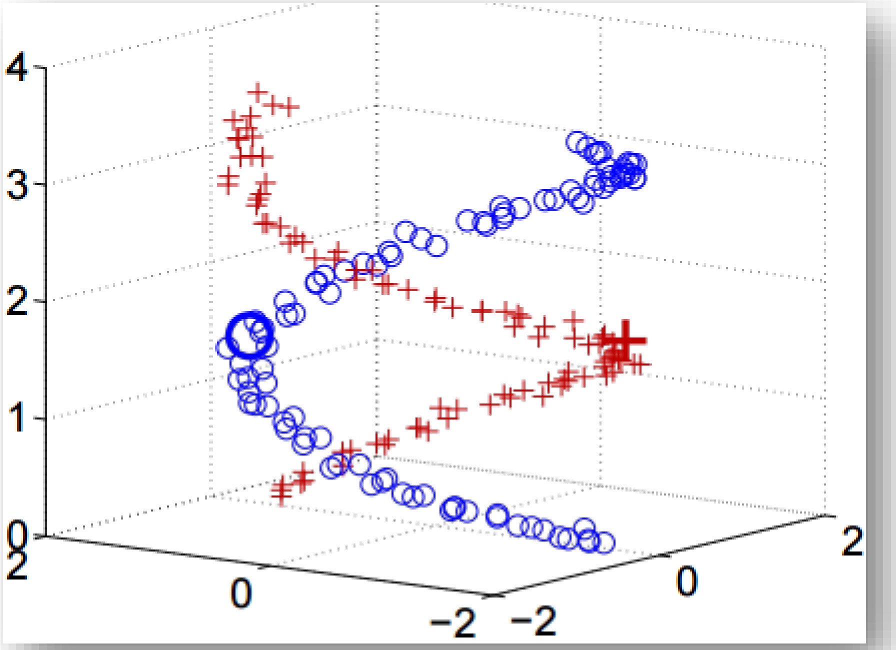

"Semi-supervised learning using Gaussian fields and harmonic functions"
Zhu, Ghahramani, & Lafferty 2003

Semi-Supervised Technique

$$\text{Given: } \ell \text{ labeled points } (x_1, y_1), \ldots, (x_\ell, y_\ell); y_i \in \{0, 1\}$$
$$u \text{ unlabeled points } x_{\ell+1}, \ldots, x_{\ell+u}; \ell \ll u$$

$$\min \frac{1}{2} \sum_{ij} w_{ij} (f(i) - f(j))^2$$
$$\text{s.t. } f(k) \text{ fixed } \forall k \leq \ell$$

Dirichlet energy → Linear system of equations (Poisson)

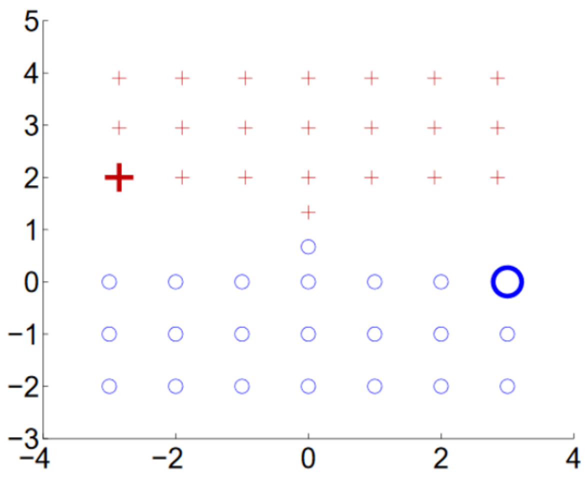

Related Method

- **Step 1:**
  Build *k*-NN graph
- **Step 2:**
  Compute *p* smallest Laplacian eigenvectors
- **Step 3:**
  Solve semi-supervised problem in subspace

"Using Manifold Structure for Partially Labelled Classification"
Belkin and Niyogi; NIPS 2002

Recent Related Method

Post-processing for Individual Fairness

Felix Petersen*
University of Konstanz
felix.petersen@uni.kn

Debarghya Mukherjee*
University of Michigan
mdeb@umich.edu

Yuekai Sun
University of Michigan
yuekai@umich.edu

Mikhail Yurochkin
IBM Research, MIT-IBM Watson AI Lab
mikhail.yurochkin@ibm.com

Abstract

Post-processing in algorithmic fairness is a versatile approach for correcting bias in ML systems that are already used in production. The main appeal of post-processing is that it avoids expensive retraining. In this work, we propose general post-processing algorithms for individual fairness (IF). We consider a setting where the learner only has access to the predictions of the original model and a similarity graph between individuals guiding the desired fairness constraints. We cast the IF post-processing problem as a graph smoothing problem corresponding to graph Laplacian regularization that preserves the desired "treat similar individuals similarly" interpretation. Our theoretical results demonstrate the connection of the new objective function to a local relaxation of the original individual fairness. Empirically, our post-processing algorithms correct individual biases in large-scale NLP models such as BERT, while preserving accuracy.

1 Introduction

There are many instances of algorithmic bias in machine learning (ML) models [1]–[4], which has led to the development of methods for *quantifying* and *correcting* algorithmic bias. To quantify algorithmic bias, researchers have proposed numerous mathematical definitions of algorithmic fairness. Broadly speaking, these definitions fall into two categories: *group fairness* [5] and *individual fairness* [6]. The former formalizes the idea that ML system should treat certain *groups* of individuals similarly, e.g., requiring the average loan approval rate for applicants of different ethnicities be similar [7]. The latter asks for similar treatment of similar *individuals*, e.g., same outcome for applicants with resumes that differ only in names [8]. Researchers have also developed many ways of correcting algorithmic bias. These fairness interventions broadly fall into three categories: pre-processing the

Buyer Beware: Ill-Posed in Limit?

Semi-Supervised Learning with the Graph Laplacian: The Limit of Infinite Unlabelled Data

Boaz Nadler
Dept. of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot, Israel 76100
boaz.nadler@weizmann.ac.il

Nathan Srebro
Toyota Technological Institute
Chicago, IL 60637
nati@uchicago.edu

Xueyuan Zhou
Dept. of Computer Science
University of Chicago
Chicago, IL 60637
zhouxy@cs.uchicago.edu

Abstract

We study the behavior of the popular Laplacian Regularization method for Semi-Supervised Learning at the regime of a fixed number of labeled points but a large

Potential fix:

Higher-order operators

Aside:

Common Misconception

$$\min_f \; E[f] \text{ s.t. } f(p) = \text{const.}$$

Point constraints are ill-advised

Manifold Regularization

$$\text{Regularized learning:} \quad \arg\min_{f \in \mathcal{H}} \frac{1}{\ell} \sum_{i=1}^{\ell} \underbrace{V(f(x_i), y_i)}_{\text{Loss function}} + \gamma \underbrace{\|f\|^2}_{\text{Regularizer}}$$

$$\|f\|_I^2 := \int \overbrace{\|\nabla f(x)\|^2}^{\text{Dirichlet energy}} \, dx \approx f^\top L f$$

"Manifold Regularization:
A Geometric Framework for Learning from Labeled and Unlabeled Examples"
Belkin, Niyogi, and Sindhwani; JMLR 2006

Examples of Manifold Regularization

- **Laplacian-regularized least squares (LapRLS)**

$$\arg\min_{f\in\mathcal{H}} \frac{1}{\ell}\sum_{i=1}^{\ell}(f(x_i)-y_i)^2 + \gamma\|f\|_I^2 + \mathrm{Other}[f]$$

- **Laplacian support vector machine (LapSVM)**

$$\arg\min_{f\in\mathcal{H}} \frac{1}{\ell}\sum_{i=1}^{\ell}\max(0, 1-y_i f(x_i)) + \gamma\|f\|_I^2 + \mathrm{Other}[f]$$

"On Manifold Regularization"
Belkin, Niyogi, Sindhwani; AISTATS 2005

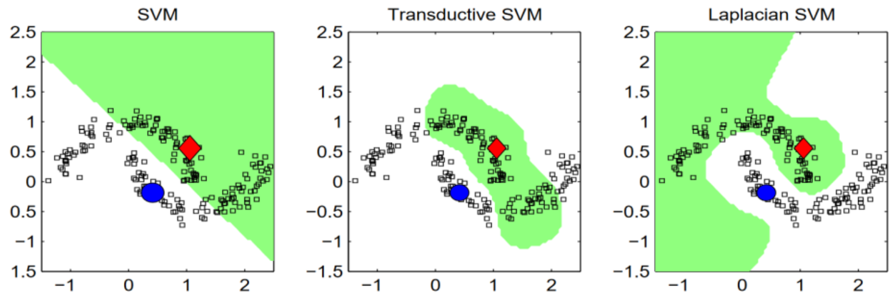

Diffusion Maps

Embedding from first *k* eigenvalues/vectors:

$$\Psi_t(x) := \left(\lambda_1^t \psi_1(x), \lambda_2^t \psi_2(x), \ldots, \lambda_k^t \psi_k(x)\right)$$

Roughly:

$|\boldsymbol{\Psi}_{\mathbf{t}}(\mathbf{x}) - \boldsymbol{\Psi}_t(\boldsymbol{y})|$ **is probability that *x, y* diffuse to the same point in time *t*.**

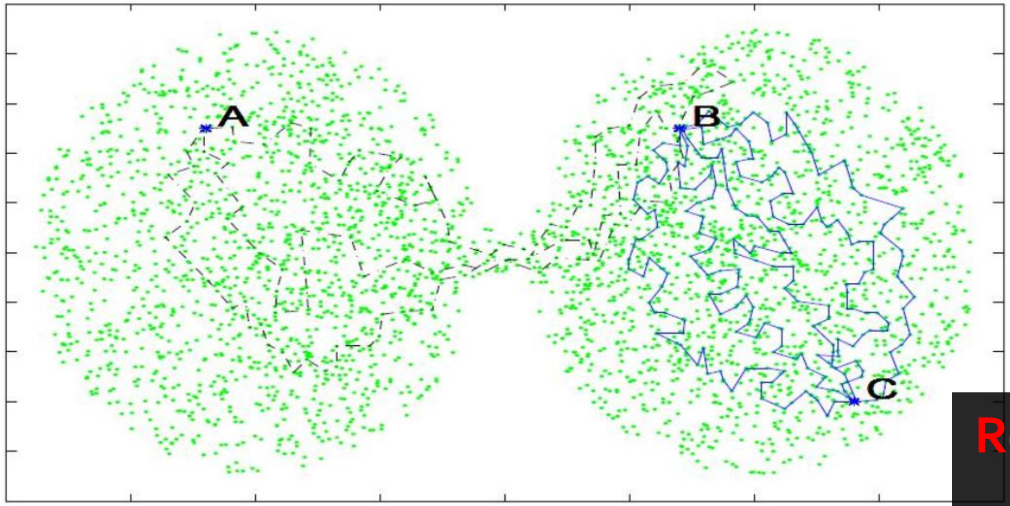

Robust to sampling and noise

"Diffusion Maps"
Coifman and Lafon; Applied and Computational Harmonic Analysis, 2006

Image from http://cpsc445.guywolf.org/Slides/CPSC445%20-%20Topic%2010%20-%20Diffusion%20Maps.pdf (nice slides!)

Graph Convolutional Networks

Spectral Networks and Deep Locally Connected Networks on Graphs

Joan Bruna
New York University
bruna@cims.nyu.edu

Wojciech Zaremba
New York University
woj.zaremba@gmail.com

Arthur Szlam
The City College of New York
aszlam@ccny.cuny.edu

Yann LeCun
New York University
yann@cs.nyu.edu

Abstract

Convolutional Neural Networks are extremely efficient architectures in image and audio recognition tasks, thanks to their ability to exploit the local translational invariance of signal classes over their domain. In this paper we consider possible generalizations of CNNs to signals defined on more general domains without the action of a translation group. In particular, we propose two constructions, one based upon a hierarchical clustering of the domain, and another based on the spectrum of the graph Laplacian. We show through experiments that for low-dimensional graphs it is possible to learn convolutional layers with a number of parameters independent of the input size, resulting in efficient deep architectures.

1 Introduction

Convolutional Neural Networks (CNNs) have been extremely succesful in machine learning problems where the coordinates of the underlying data representation have a grid structure (in 1, 2 and 3 dimensions), and the data to be studied in those coordinates has translational equivariance/invariance with respect to this grid. Speech [11], images [14, 20, 22] or video [23, 18] are prominent examples that fall into this category.

Convolution theorem for functions on $\mathbb{R}^n$:

$$\boldsymbol{f} * \boldsymbol{g} = \boldsymbol{\mathcal{F}}^{-1}[\boldsymbol{F} \cdot \boldsymbol{G}]$$

$$x_{k+1,j} = h\left(V \sum_{i=1}^{f_{k-1}} F_{kij} V^\top x_{ki}\right)$$

V contains eigenvectors of graph Laplacian

Useful Survey

Geometric deep learning: going beyond Euclidean data

Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, Pierre Vandergheynst

Many scientific fields study data with an underlying structure that is a non-Euclidean space. Some examples include social networks in computational social sciences, sensor networks in communications, functional networks in brain imaging, regulatory networks in genetics, and meshed surfaces in computer graphics. In many applications, such geometric data are large and complex (in the case of social networks, on the scale of billions), and are natural targets for machine learning techniques. In particular, we would like to use deep neural networks, which have recently proven to be powerful tools for a broad range of problems from computer vision, natural language processing, and audio analysis. However, these tools have been most successful on data with an underlying Euclidean or grid-like structure, and in cases where the invariances of these structures are built into networks used to model them.

Geometric deep learning is an umbrella term for emerging techniques attempting to generalize (structured) deep neural models to non-Euclidean domains such as graphs and manifolds. The purpose of this paper is to overview different

the data such as stationarity and compositionality through local statistics, which are present in natural images, video, and speech [14], [15]. These statistical properties have been related to physics [16] and formalized in specific classes of convolutional neural networks (CNNs) [17], [18], [19]. In image analysis applications, one can consider images as functions on the Euclidean space (plane), sampled on a grid. In this setting, stationarity is owed to shift-invariance, locality is due to the local connectivity, and compositionality stems from the multi-resolution structure of the grid. These properties are exploited by convolutional architectures [20], which are built of alternating convolutional and downsampling (pooling) layers. The use of convolutions has a two-fold effect. First, it allows extracting local features that are shared across the image domain and greatly reduces the number of parameters in the network with respect to generic deep architectures (and thus also the risk of overfitting), without sacrificing the expressive capacity of the network. Second, the convolutional architecture itself imposes some priors about the data, which appear very suitable especially for natural images [21], [18],

Applications of the Laplacian

Justin Solomon

6.8410: Shape Analysis

Spring 2023

MIT EECS