MVA

Geometry Processing and Geometric Deep Learning

Today

- Surface and Shape Analysis
  - Surface features
  - Discrete representations
  - Discrete Laplace-Beltrami operator
  - Applications in shape comparison and shape analysis

Parametrized Surfaces

A parametrized surface is a map from the plane in to the space.

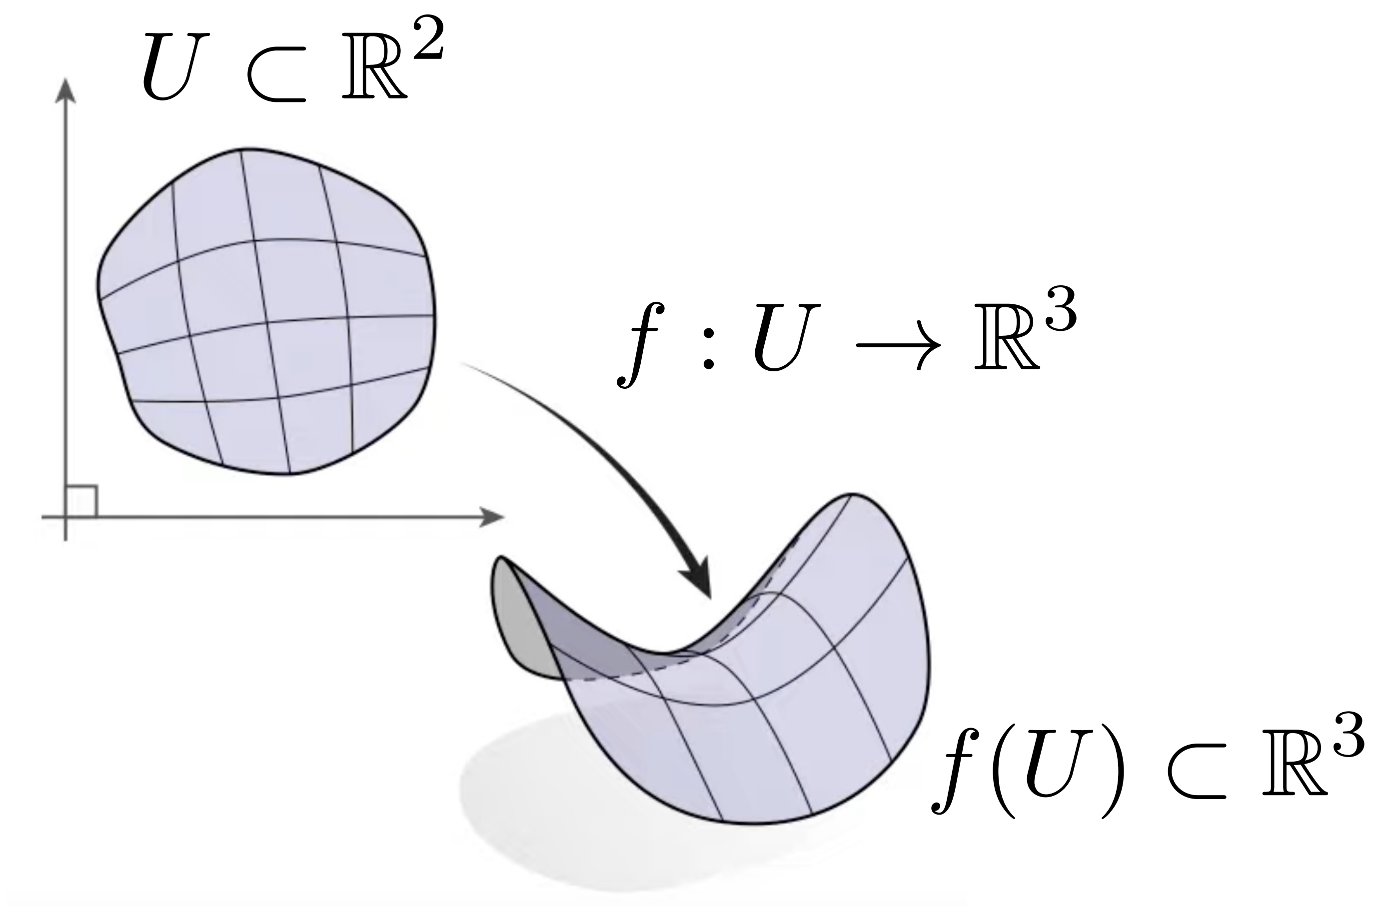

Parametrized Surfaces

A parametrized surface is a map from the plane in to the space.

$U \subset \mathbb{R}^2$

$f(U) \subset \mathbb{R}^3$

$f : U \to \mathbb{R}^3$

Parametrized Surfaces

A parametrized surface is a map from the plane in to the space.

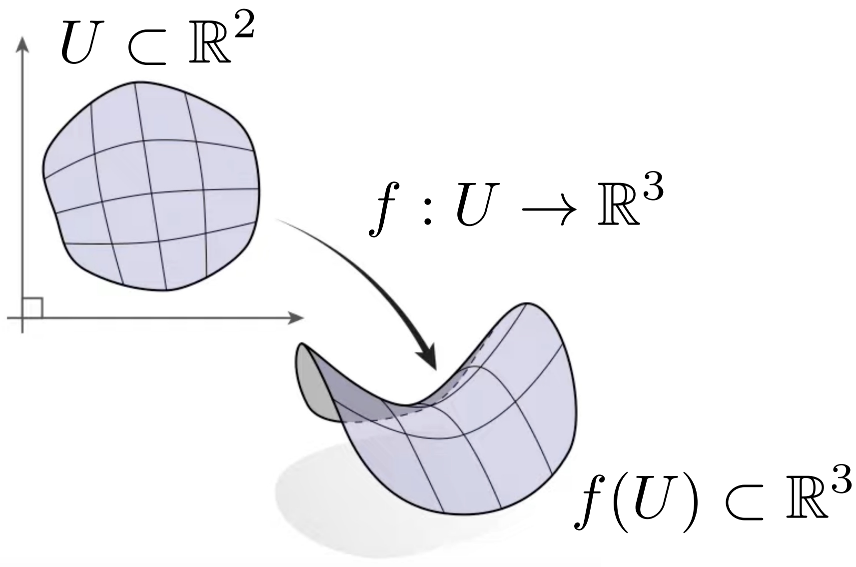

Assumption: discrete surfaces are approximation of smooth surface

Describing a Surface

What makes a surface unique?

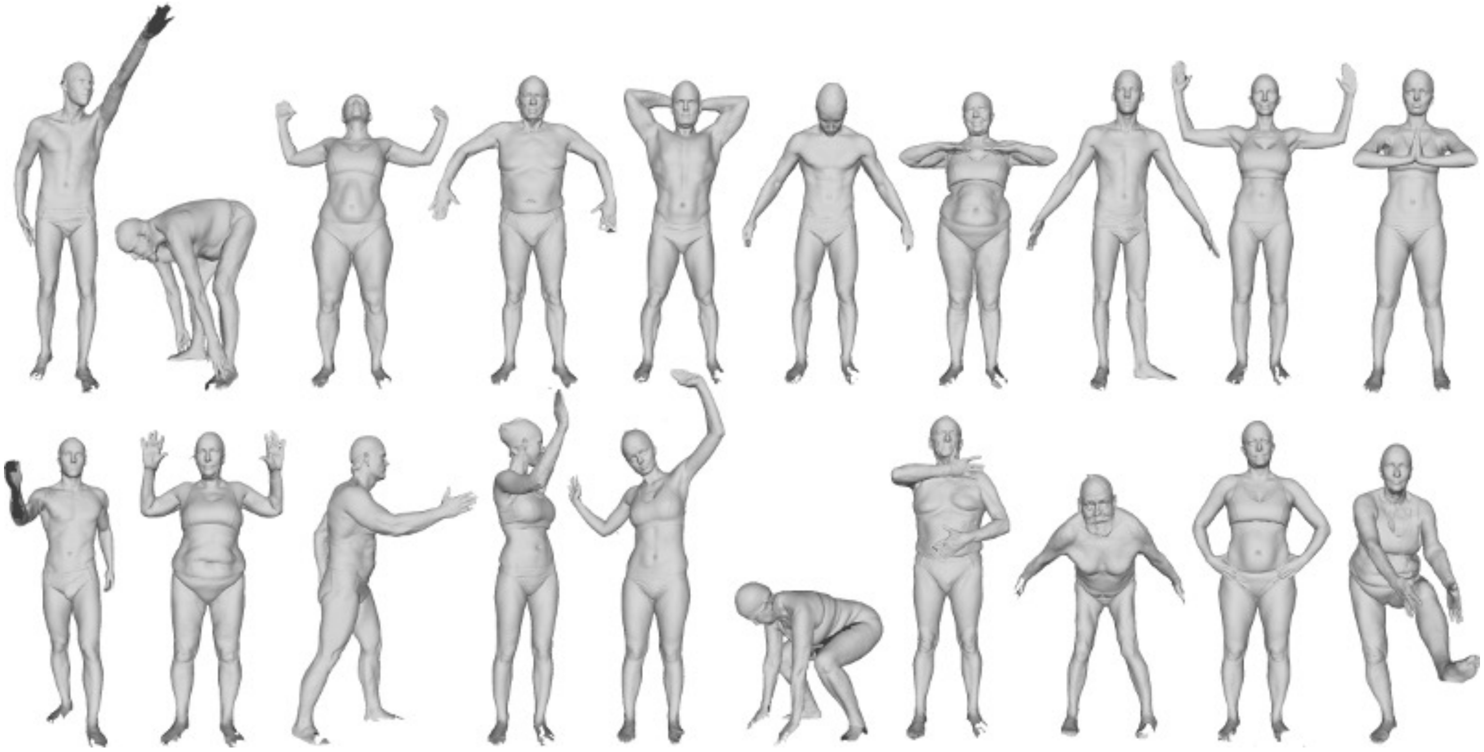

Describing a Surface

What makes a surface unique (up to rigid transformation)?

1. Geodesic distances: shortest distance between two points
2. Curvature: change in normal

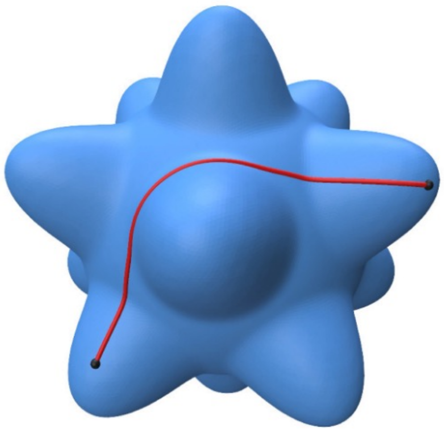

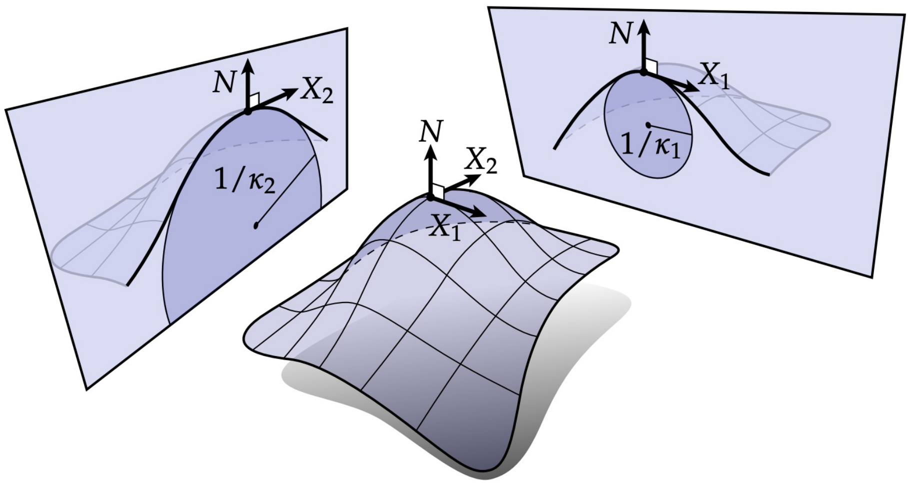

Geodesic Path

- Shortest path on a surface
  - Not always a straight line!

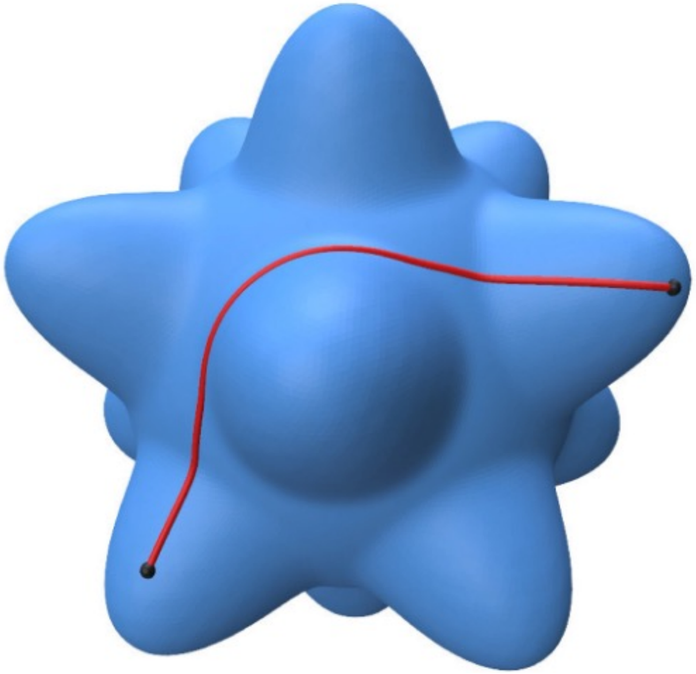

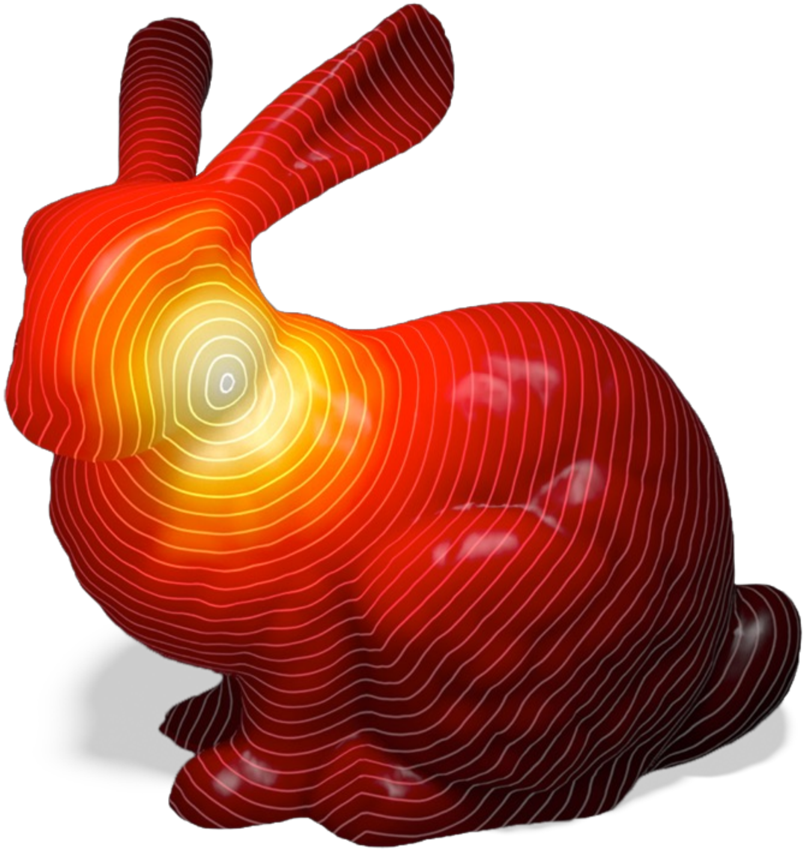

Geodesic Path

- Shortest path on a surface
  - Not always a straight line!

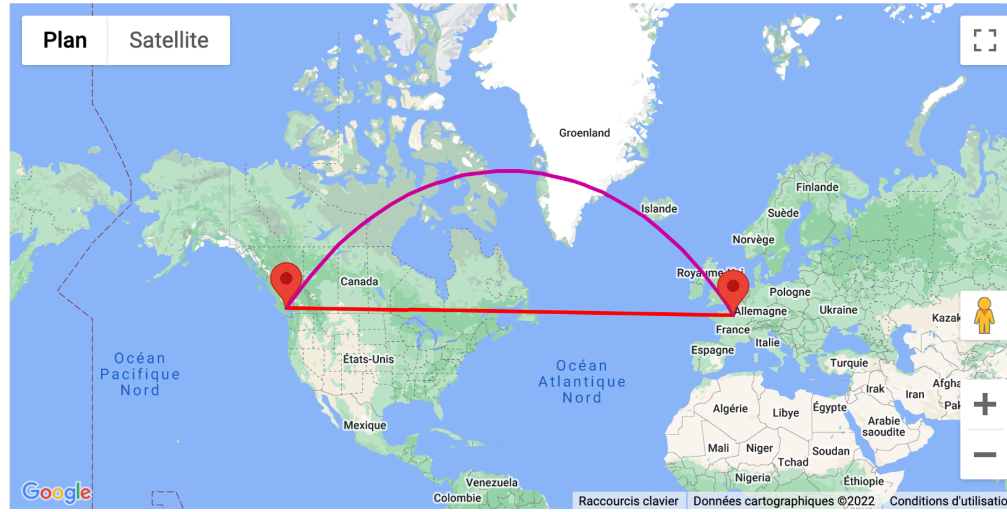

Describing a Surface

- Curvature = normal variations

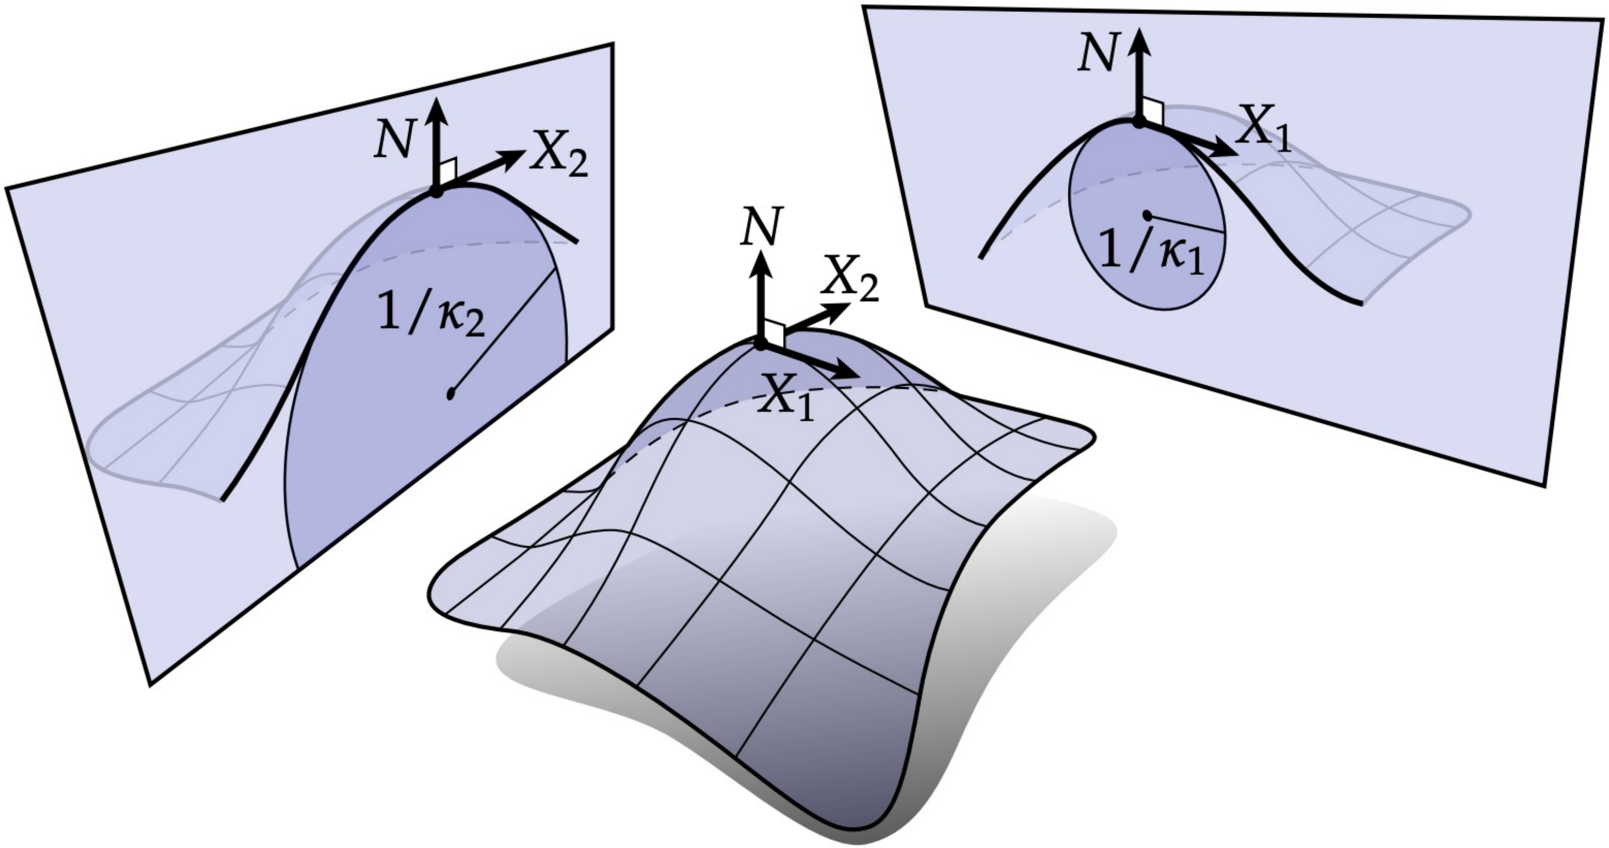

Gaussian Curvature

- Geodesic triangles
- Total Gaussian curvature: sum of inner angle minus pi

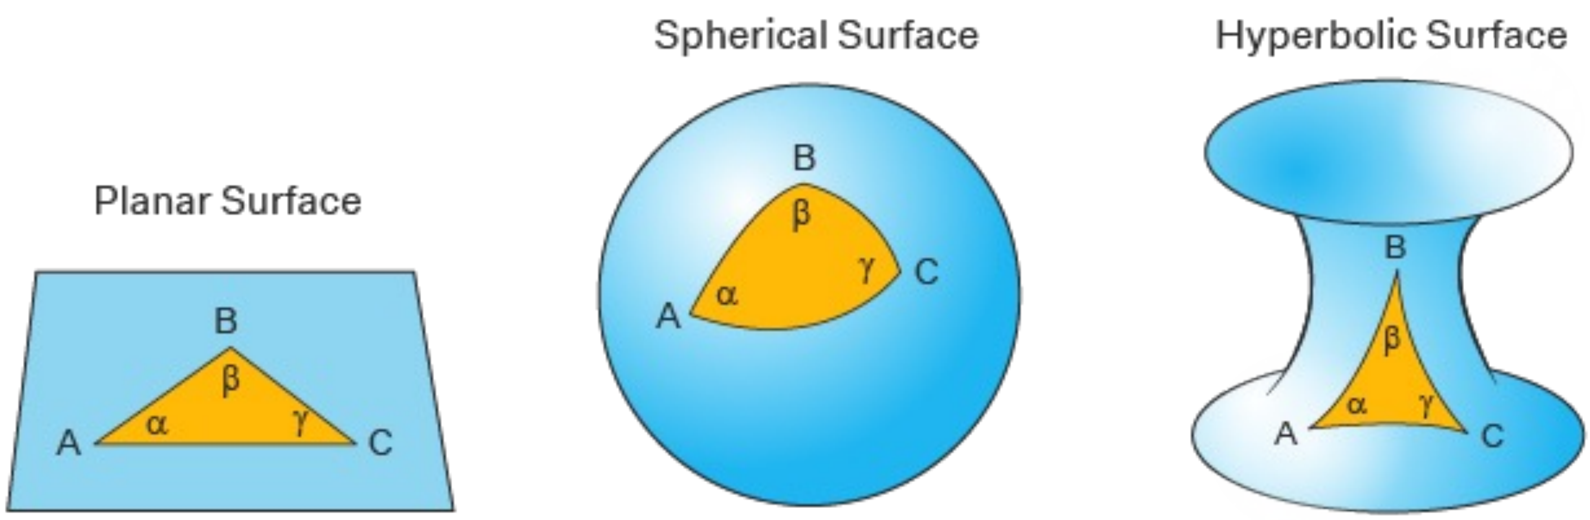

$K = \alpha + \beta + \gamma - \pi = 0$ $\quad$ $K = \alpha + \beta + \gamma - \pi > 0$ $\quad$ $K = \alpha + \beta + \gamma - \pi < 0$

Gaussian Curvature

- Defined from geodesic triangles
- Total Gaussian curvature: sum of inner angle minus pi
- Distance to a “folded” piece of paper

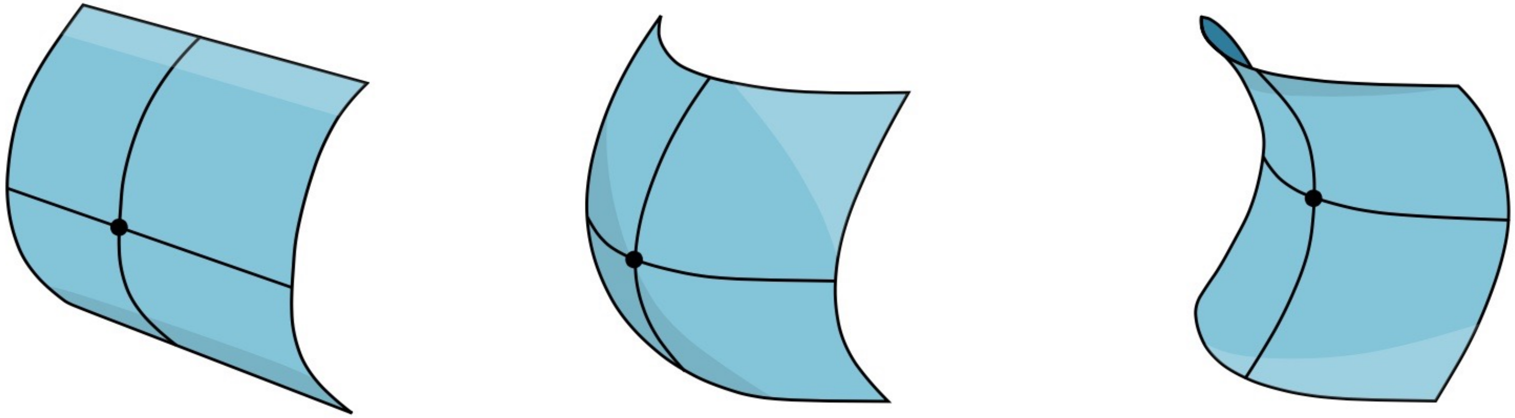

$K = 0$ “folded” paper | $K > 0$ spherical | $K < 0$ saddle

Gaussian Curvature

- Locally a surface "looks" like:
  - A sphere;
  - A saddle;
  - A piece of folded paper.

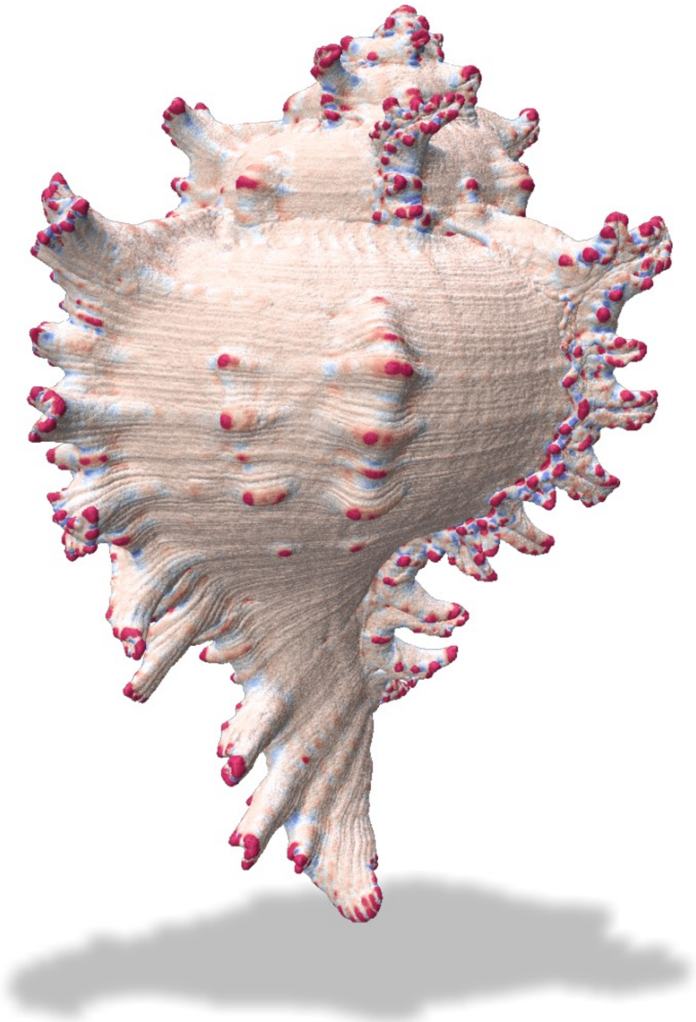

Gaussian curvature

- Not the same surface but same Gaussian curvature and geodesics
  - Intrinsic information are not enough to fully describe a surface

Feature Functions

Curvature and geodesic can difficult to compute in practice!

Shape descriptors:

- Easy to compute
- Stable under noise
- Stable under small deformation

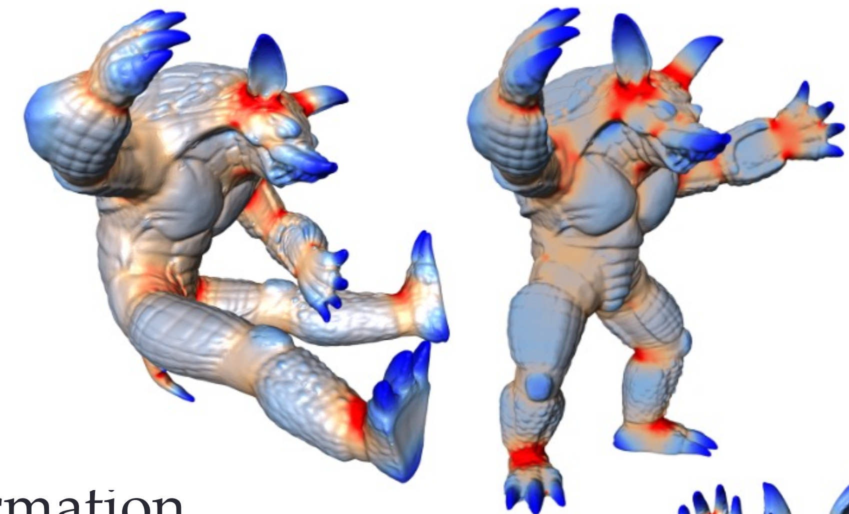

Laplacian to the rescue!

Generalized Heat Kernel Signatures, Valentin Zobel, Jan Reininghaus, Ingrid Hotz

Laplacian in Physics

- Heat diffusion: $\dfrac{\partial u}{\partial t} = \Delta u$

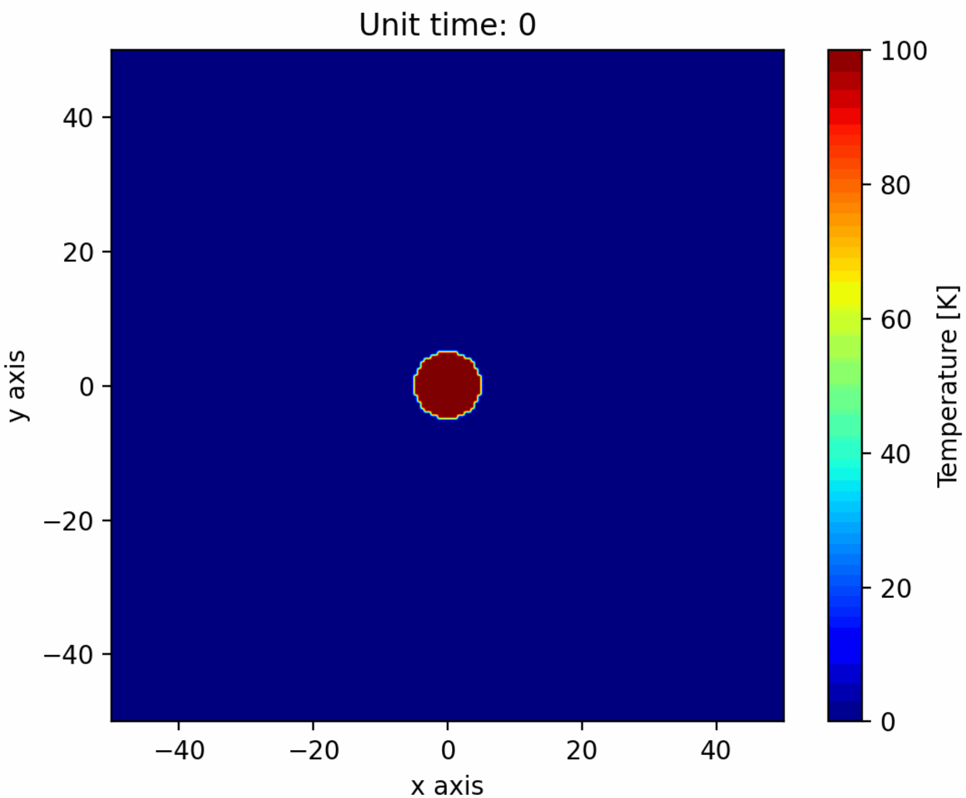

Paul Böhm

Laplacian in Physics

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

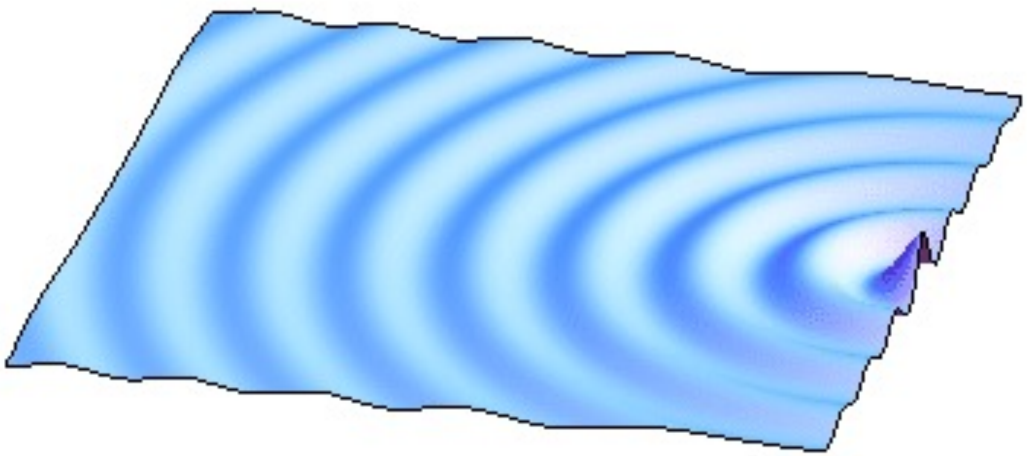

Times of diffusion is a geodesic distance!

Laplacian in Geometry

- Isometry invariance:
  - Same geodesic if and only if same Laplacian
  - Not the same su … se enough)

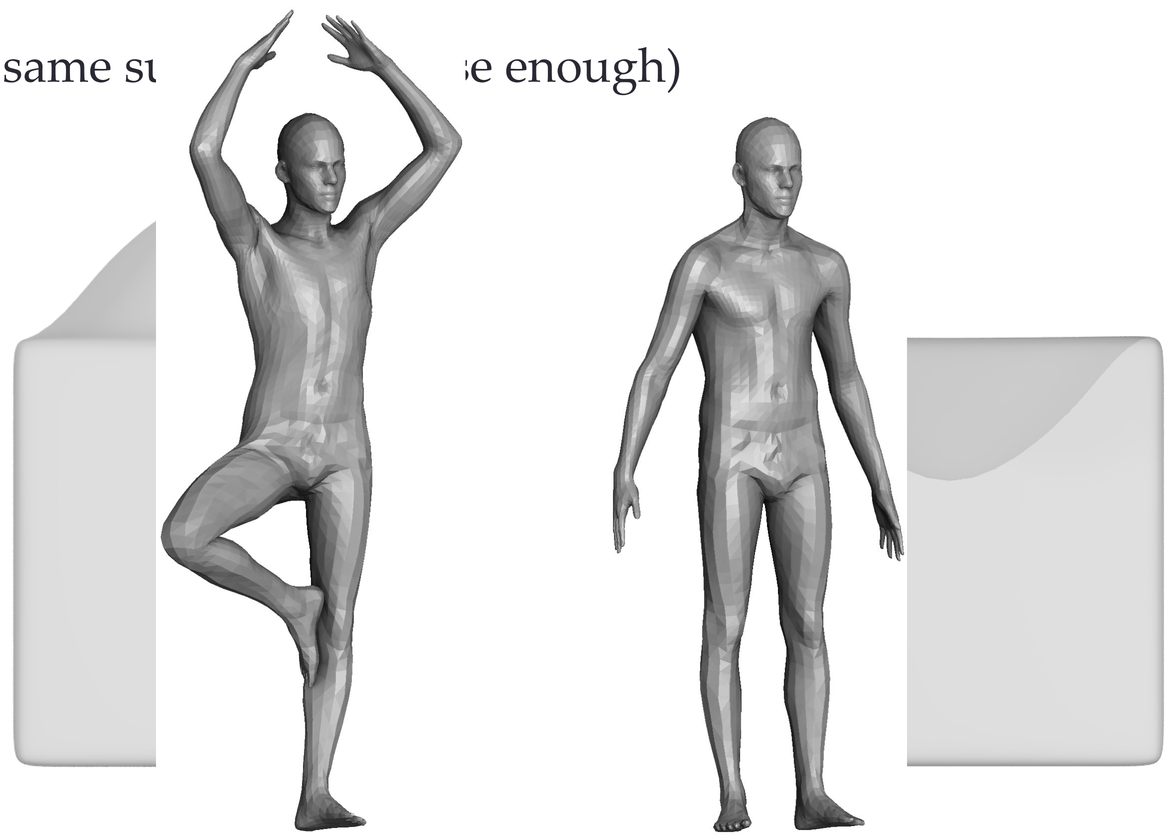

Chladni Plates

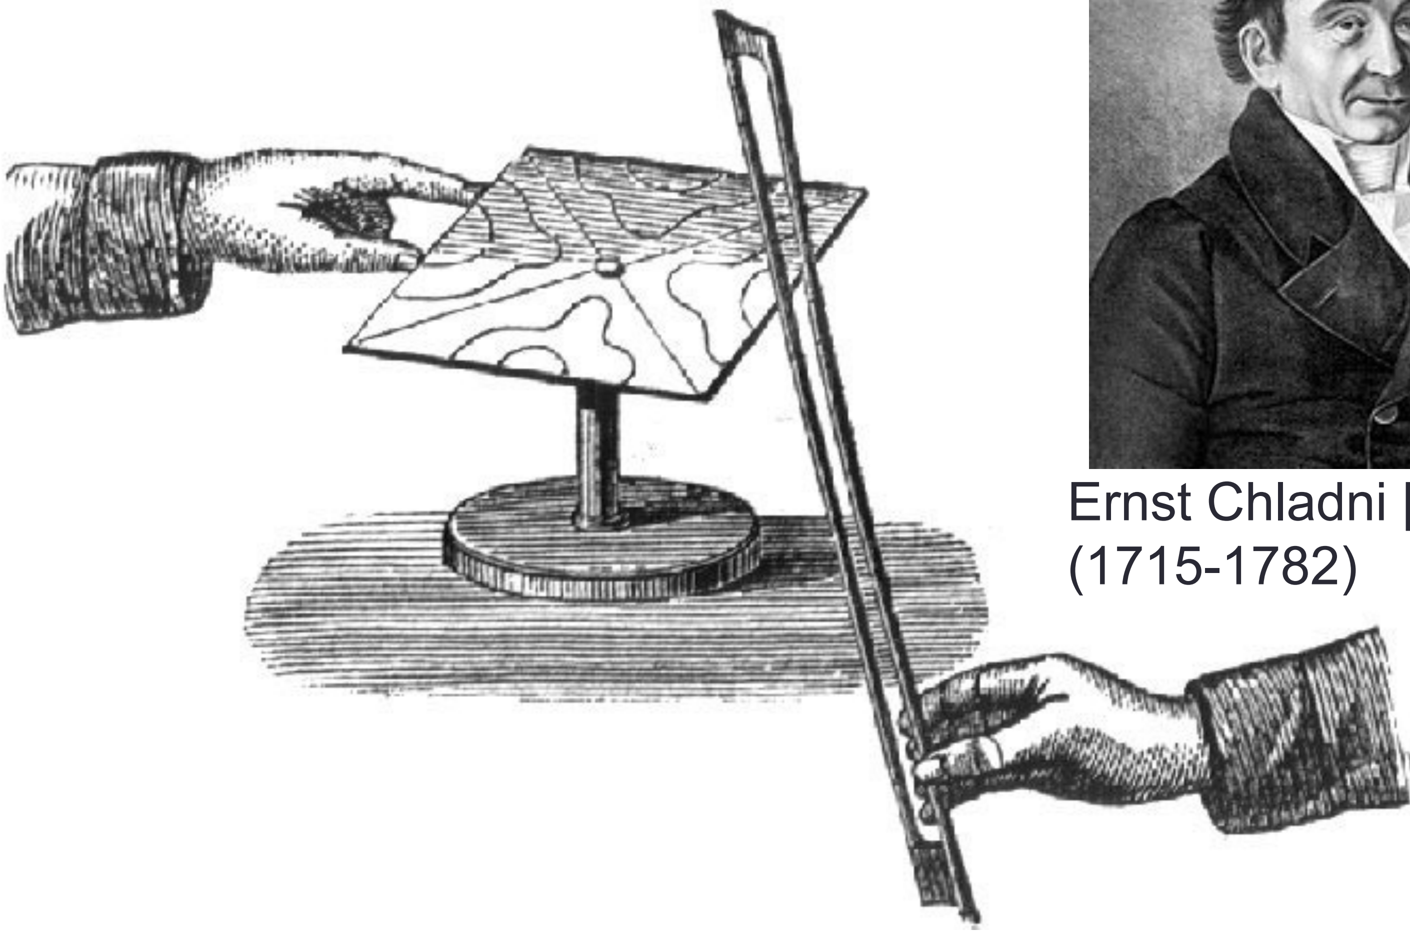

Ernst Chladni ['kladnɪ]
(1715-1782)

Chladni's experimental setup allowing to visualize acoustic waves

E. Chladni, *Entdeckungen über die Theorie des Klanges*

Laplacian in Geometry

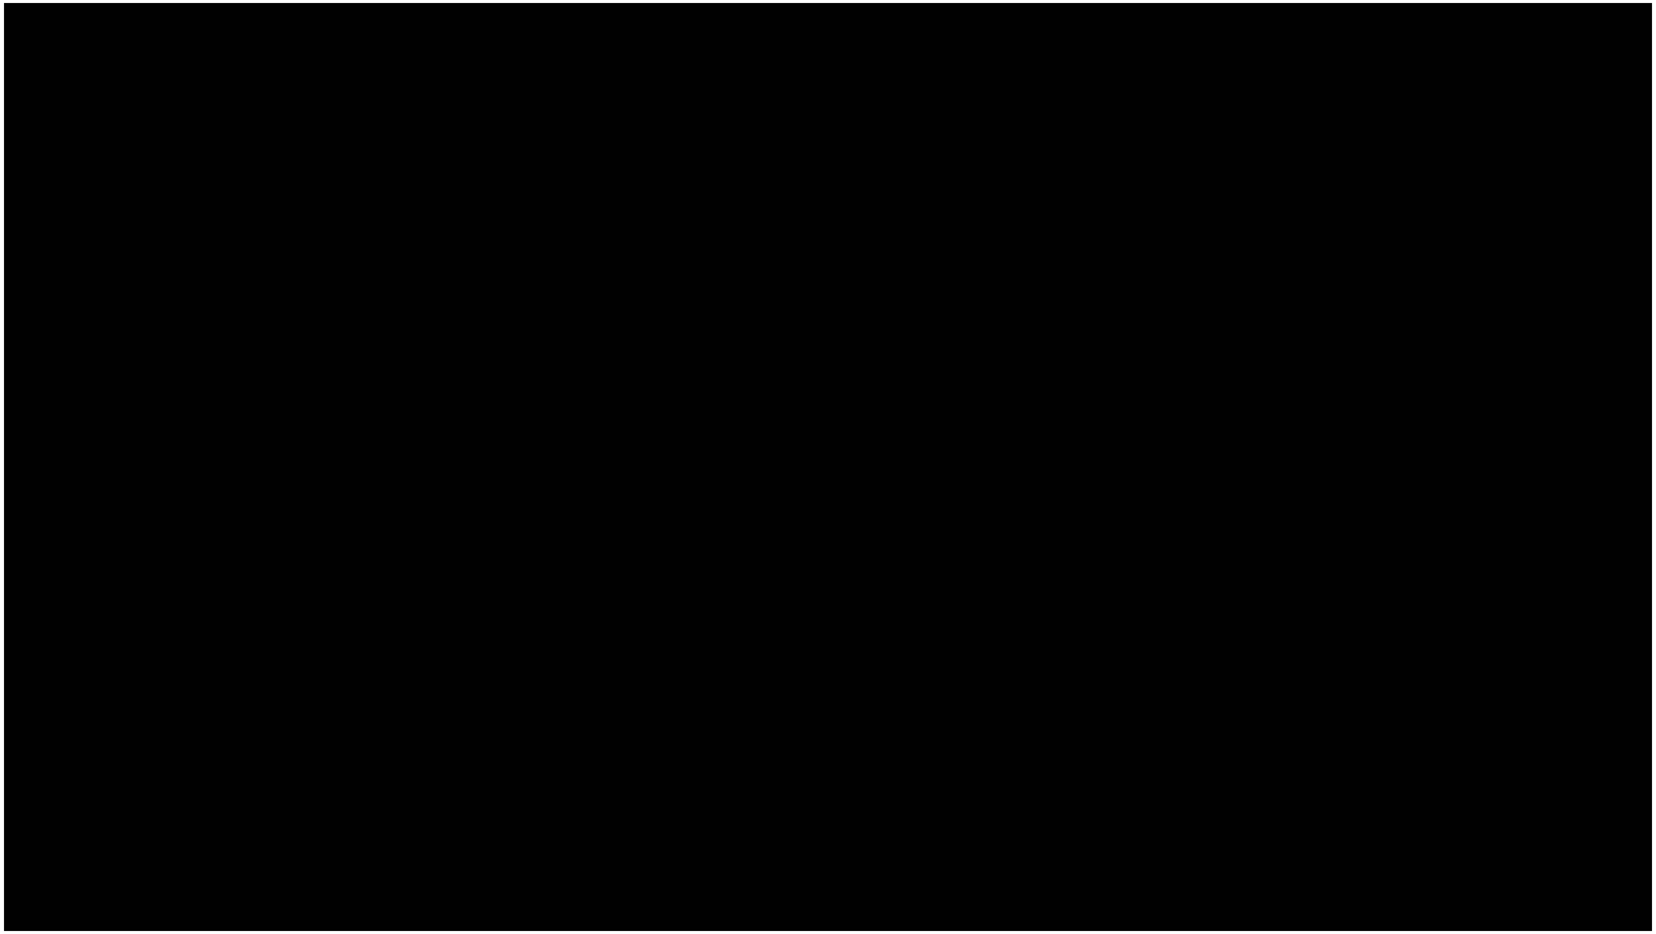

https://www.youtube.com/watch?v=wvJAgrUBF4w

Chladni Plates

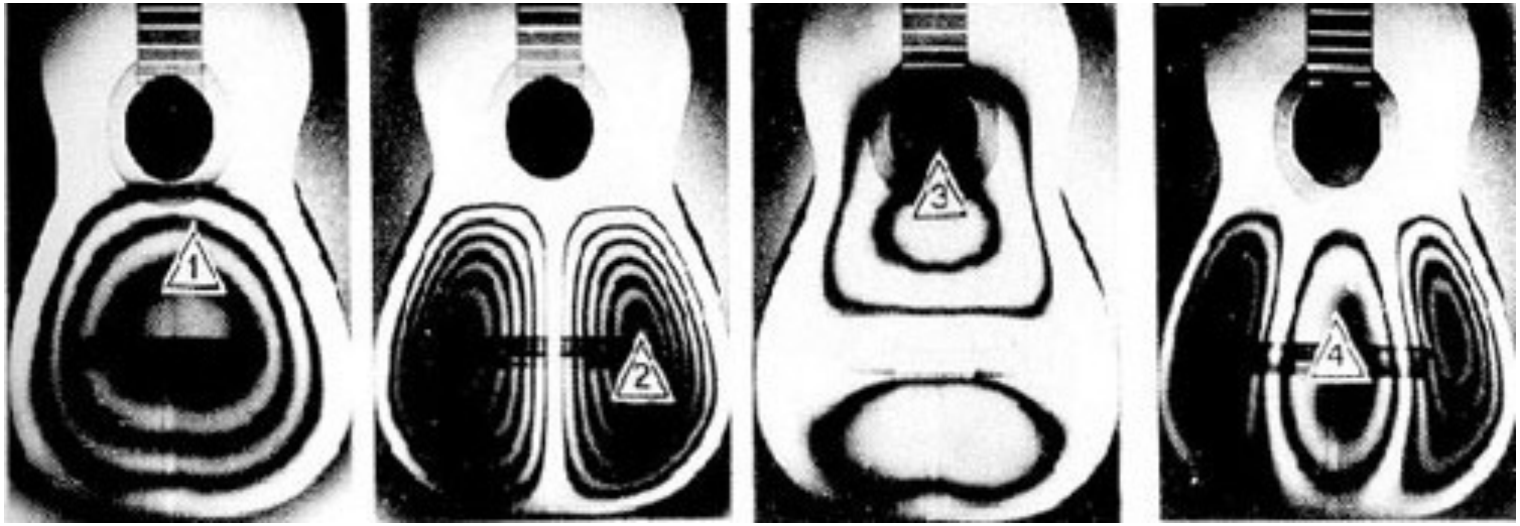

Patterns seen by Chladni are solutions to **stationary Helmholtz equation**

$$\Delta_X f = \lambda f$$

Solutions of this equation are **eigenfunction** of Laplace-Beltrami operator

"Can one hear the shape of the drum?"

Mark Kac
(1914-1984)

More prosaically: can one reconstruct the shape (up to an isometry) from its Laplace-Beltrami spectrum?

To Hear the Shape

In Chladni's experiments, the spectrum describes acoustic characteristics of the plates ("modes" of vibrations)

What can be "heard" from the spectrum:

- Total Gaussian curvature
- Euler characteristic
- Area

Can we "hear" the geodesic distances? **NO!**

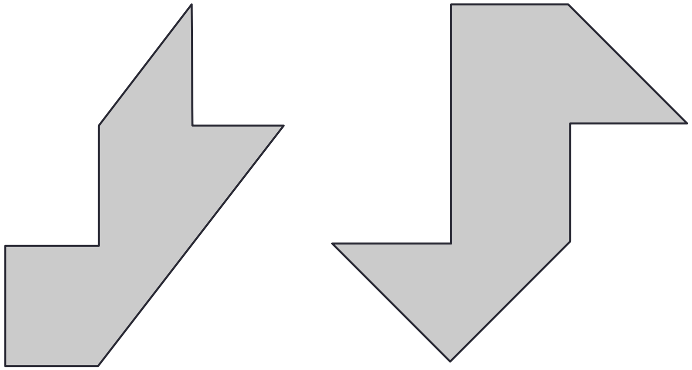

Laplacian in Geometry

- Let's build reliable descriptors on discrete surfaces with:
  - Heat diffusion
  - Eigen-decompostion

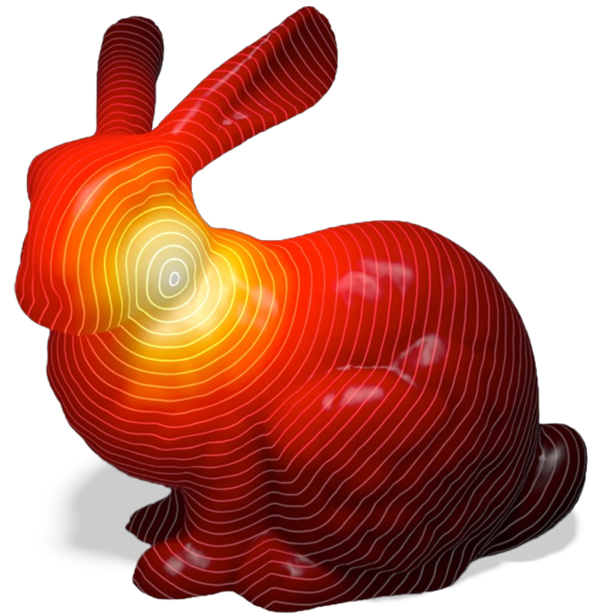

Generalized Heat Kernel Signatures, Valentin Zobel, Jan Reininghaus, Ingrid Hotz
The Heat Method for Distance Computation, Keenan Crane, Clarisse Weischedel, Max Wardetzky

Different Shape Representations

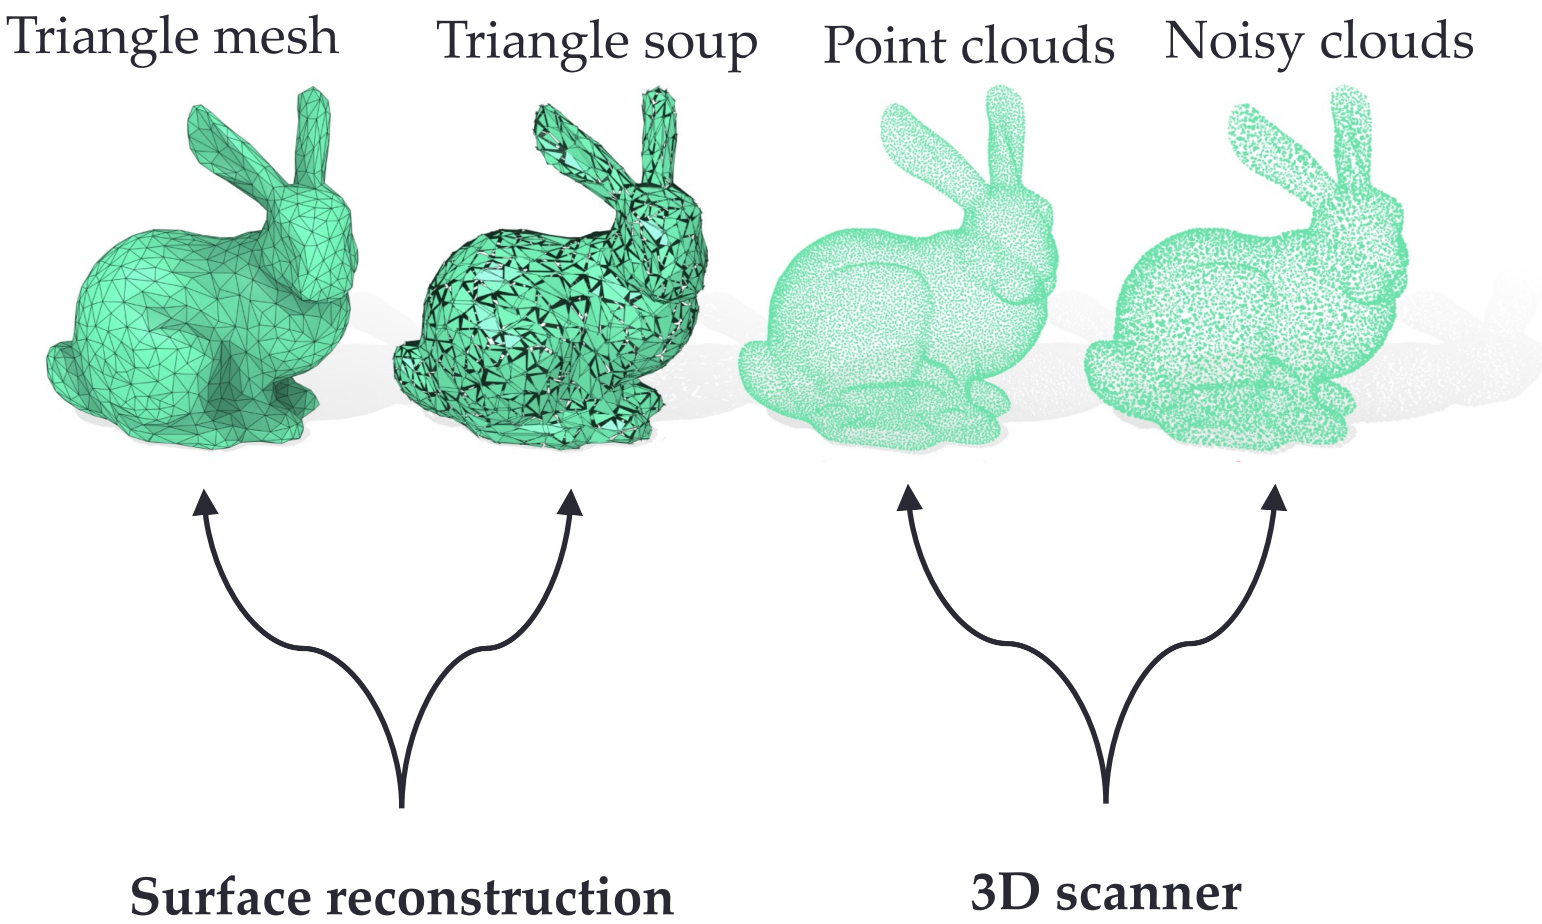

Different Shape Representations

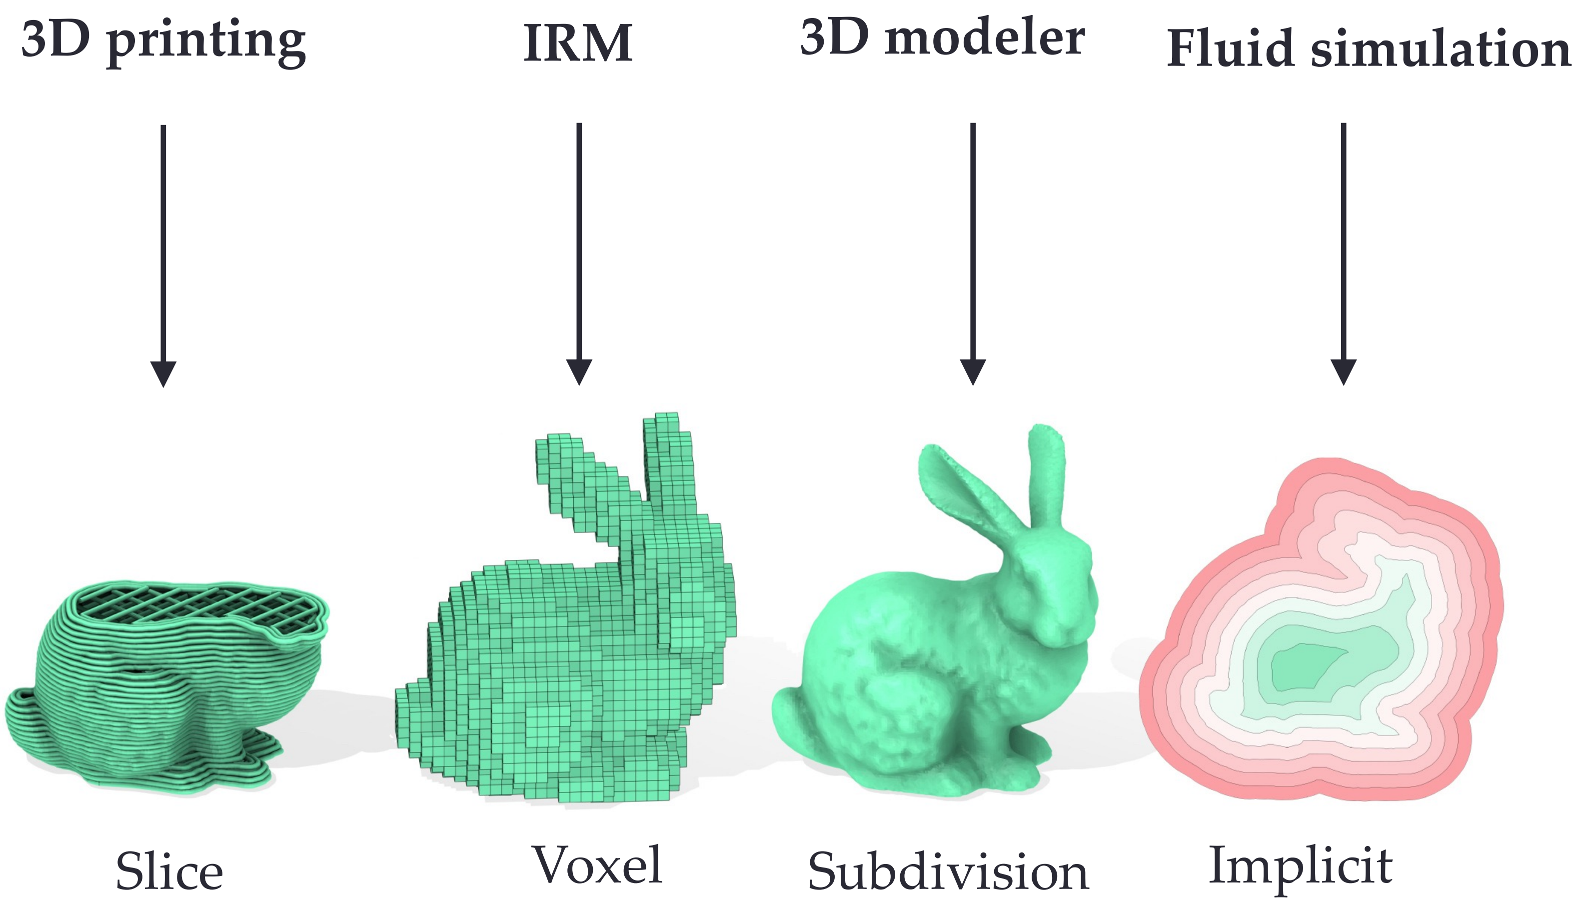

Why Different Shape Representations?

- Depends on the acquisition process
- Depends on the applications

- Which representation are we going to use?
  - Ideally, we would like a learning pipeline working on all representations
  - In this course
    - triangle meshes (today)
    - point clouds
    - Signed distance functions

Point Clouds

- Simplest shape representation
  - Only point coordinates (x,y,z) (sometimes with normal)
  - Typically results of 3D scanning
  - Need to be processed before used

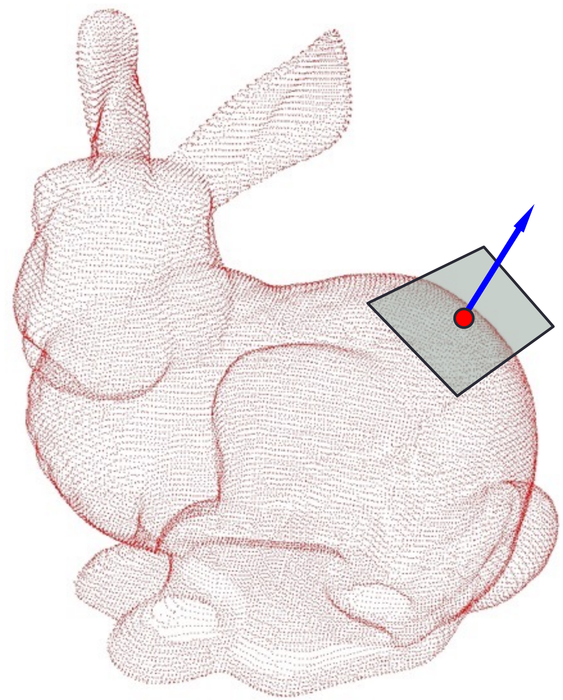

Triangle Meshes

- A very special type of graph!
- Two arrays
  - Point coordinates (x,y,z)
  - Triangle indices (i1,i2,i3)

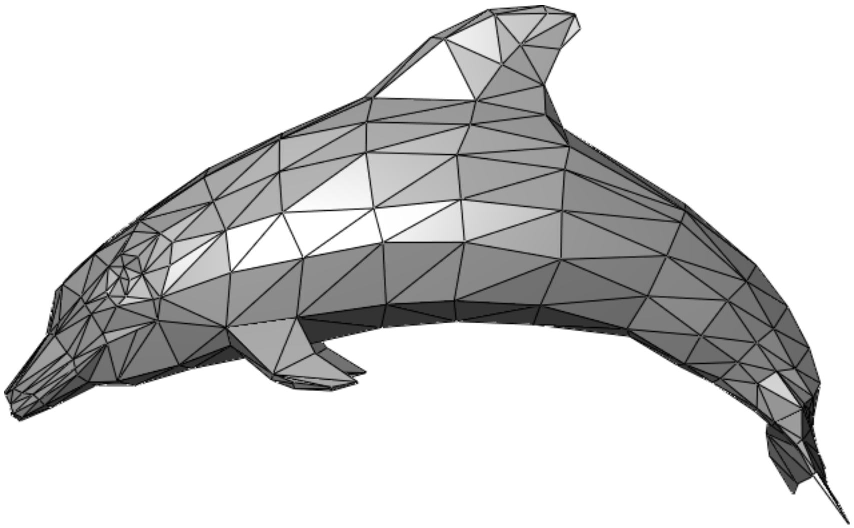

Graph Definitions

- Graph: G = {V,T}
- Vertices: V = { A,B,C,…}
- Faces: T = {(BEF), …}
- Edges: E = {(AB),(CD),…}

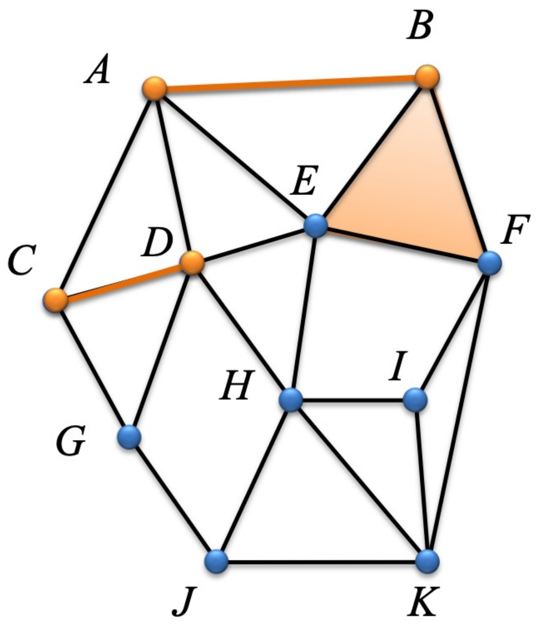

Graph Embedding

Embedding: G is embedded in $\mathbb{R}^d$ if every vertices is assigned a position in $\mathbb{R}^d$

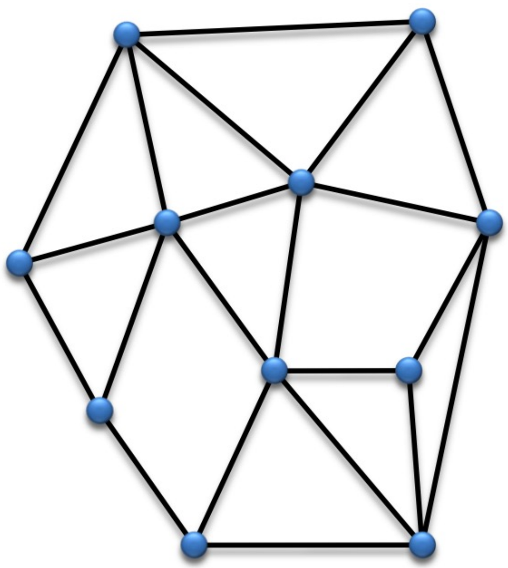

Embedded in $\mathbb{R}^2$

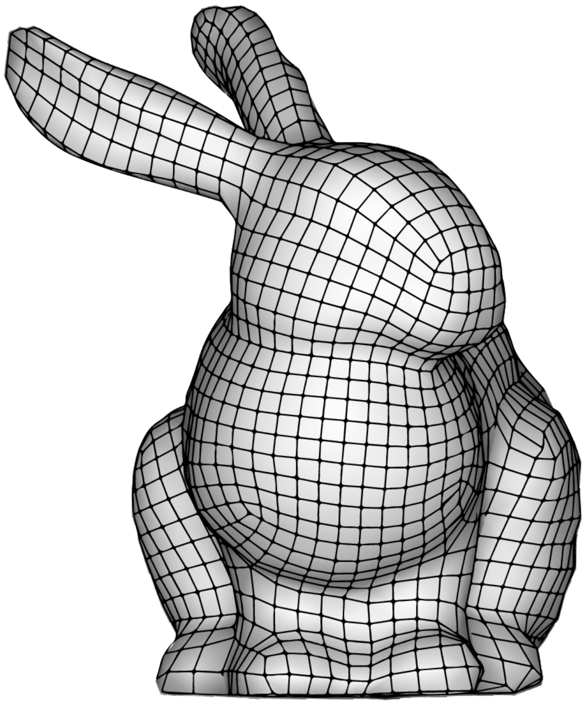

Embedded in $\mathbb{R}^3$

Triangle Mesh

Triangulation: every face is a triangle

- Connectivity: triangle list
- Embedding: vertex list

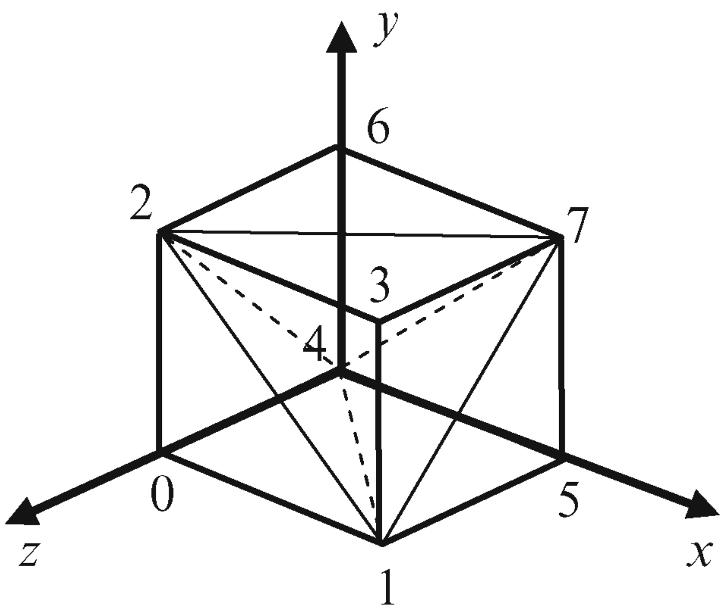

Vertex List		
x	y	z
0.0	0.0	1.0
1.0	0.0	1.0
0.0	1.0	1.0
1.0	1.0	1.0
0.0	0.0	0.0
1.0	0.0	0.0
0.0	1.0	0.0
1.0	1.0	0.0

Triangle List		
i	j	k
0	1	2
1	3	2
2	3	7
2	7	6
1	7	3
1	5	7
6	7	4
7	5	4
0	4	1
1	4	5
2	6	4
0	2	4

Data Structure

- Needs a mesh data structure to ”walk” on the mesh
  - For a triangle find incident edges, vertices
  - For a vertex visit 1-ring vertex
  - Iterates on vertices/faces/edges

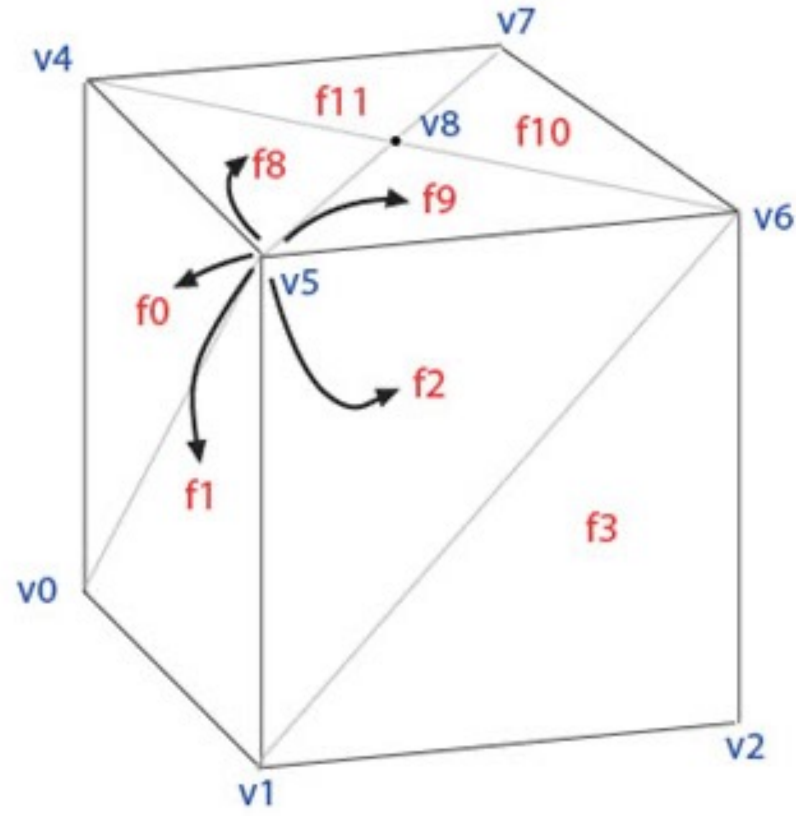

Manifold (aka “Nice”) Meshes

Disk-like neighborhood:

- Edges adjacent to at most two faces
- Triangles incident to a vertex can be sorted

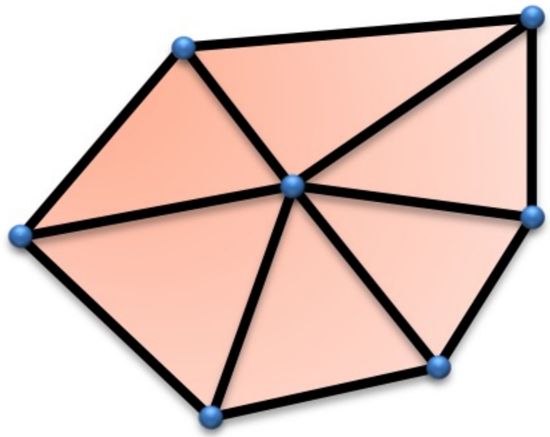

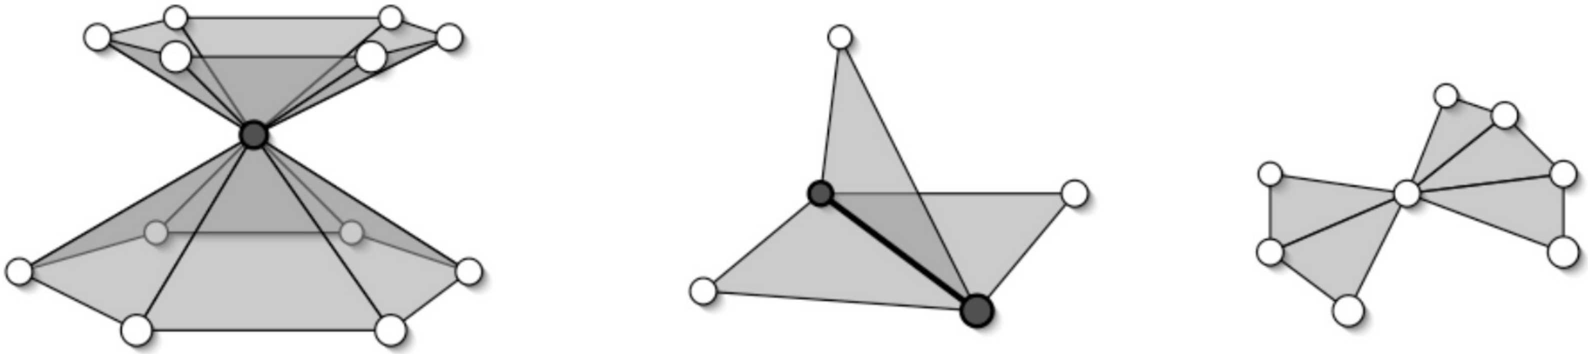

Non-manifold triangle meshes

Mesh Orientation

- Face orientation is defined by vertex order or normal direction
- A mesh is **orientable** if all faces can be orientated consistently

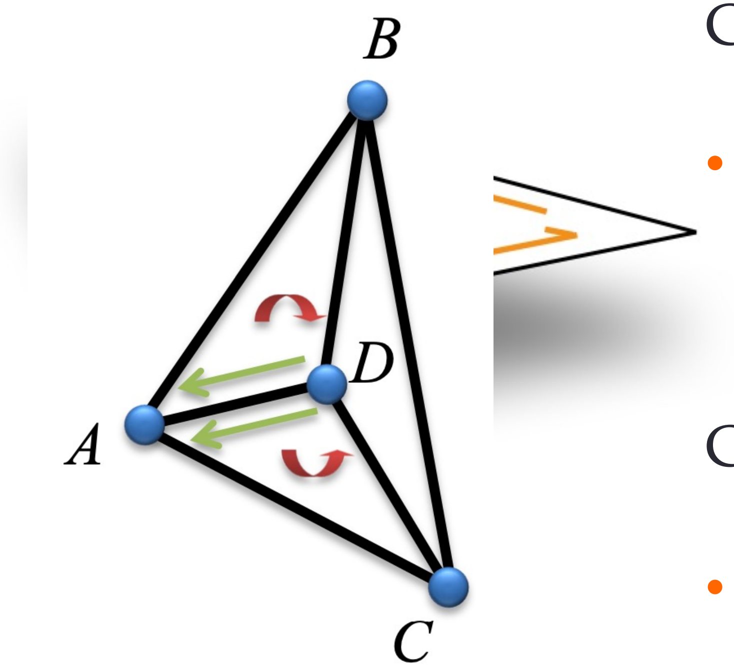

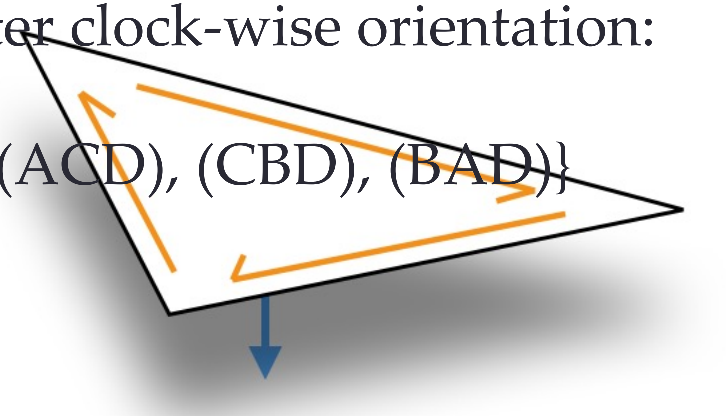

Counter clock-wise orientation:

- T = {(ACD), (CBD), (BAD)}

Clock-wise orientation:

- T = {(ADC), (CDB), (BDA)}

Non-Orientable Meshes

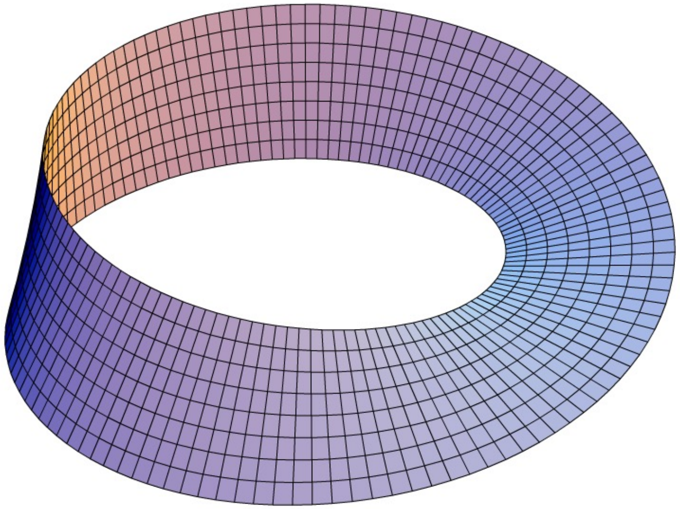

Moebius strip

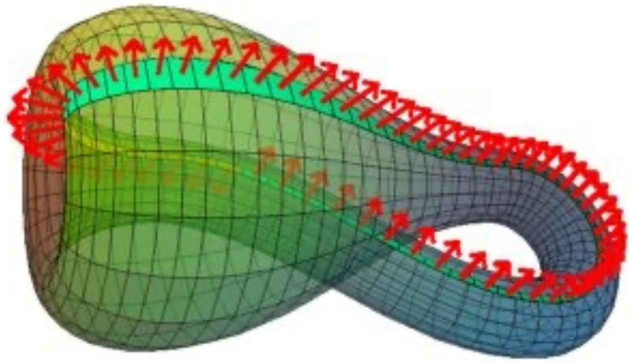

Klein bottle

Gradient

$$\nabla f = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z} \right)$$

Input: scalar function: $f : \mathbb{R}^n \to \mathbb{R}$

Output: vector field: $\nabla f : \mathbb{R}^n \to \mathbb{R}^n$

Intuition: steepest ascent direction

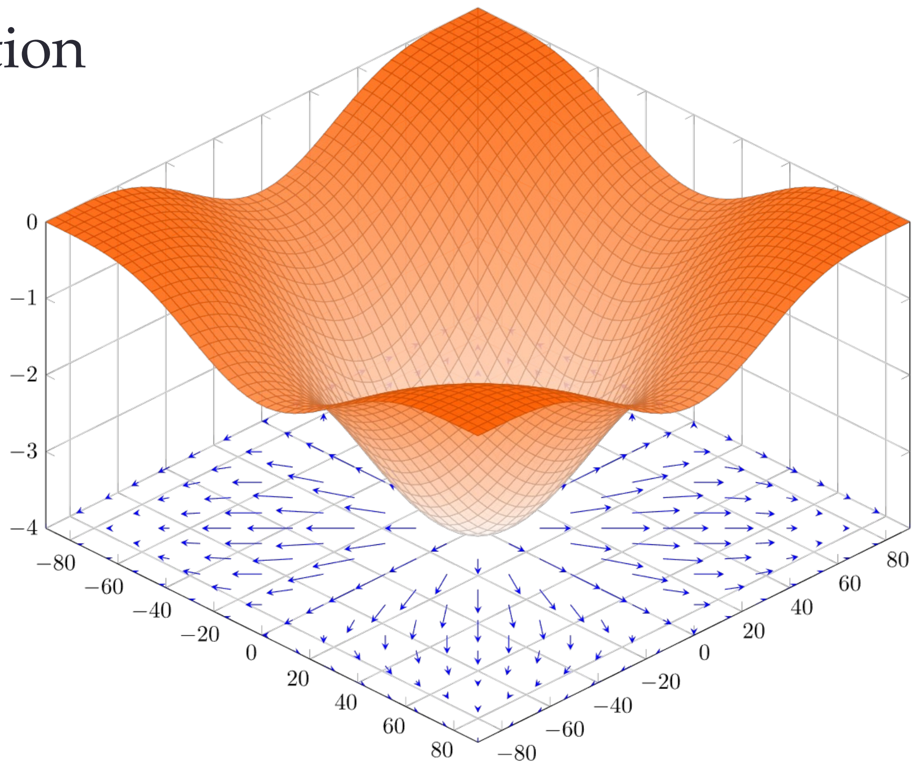

Divergence

$$\mathrm{div}V = \frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y}$$

Input: vector field: $V : \mathbb{R}^2 \to \mathbb{R}^2$

Output: scalar function: $\mathrm{div}V : \mathbb{R}^2 \to \mathbb{R}$

Intuition: source/sinks

Divergence theorem:

$$\int_U \mathrm{div}V = \int_{\partial U} V.n$$

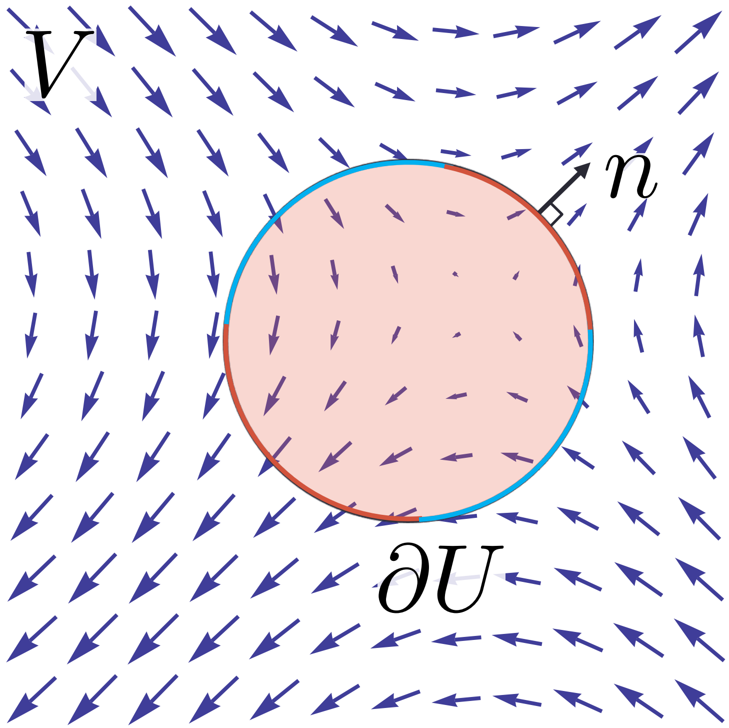

Divergence

$$\mathrm{div} V = \frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y}$$

Input: vector field: $V : \mathbb{R}^2 \to \mathbb{R}^2$

Output: scalar function: $\mathrm{div} V : \mathbb{R}^2 \to \mathbb{R}$

Intuition: source/sinks

$$\mathrm{div} V = 0$$

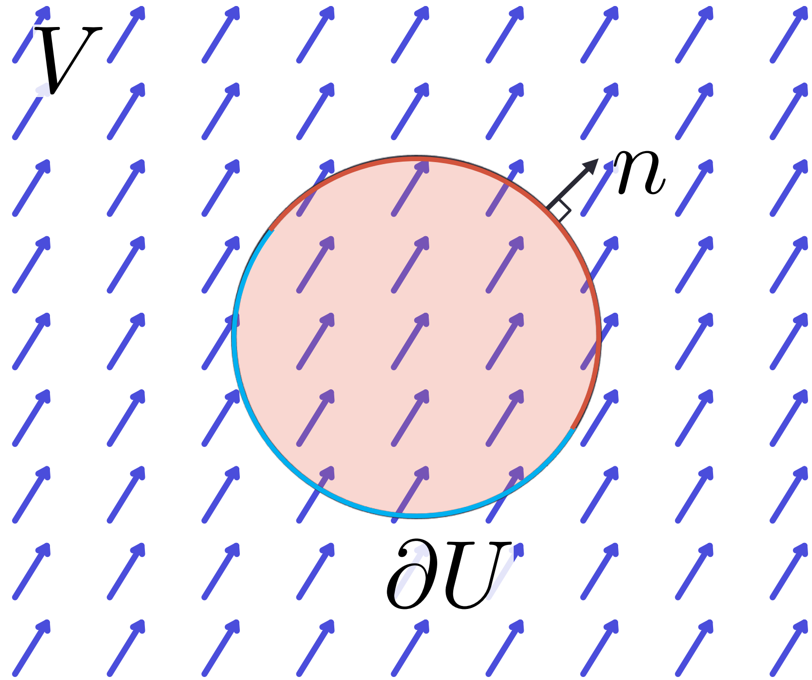

Divergence

$$\text{div}V = \frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y}$$

Input: vector field: $V : \mathbb{R}^2 \to \mathbb{R}^2$

Output: scalar function: $\text{div}V : \mathbb{R}^2 \to \mathbb{R}$

Intuition: source/sinks

$$\text{div}V < 0$$

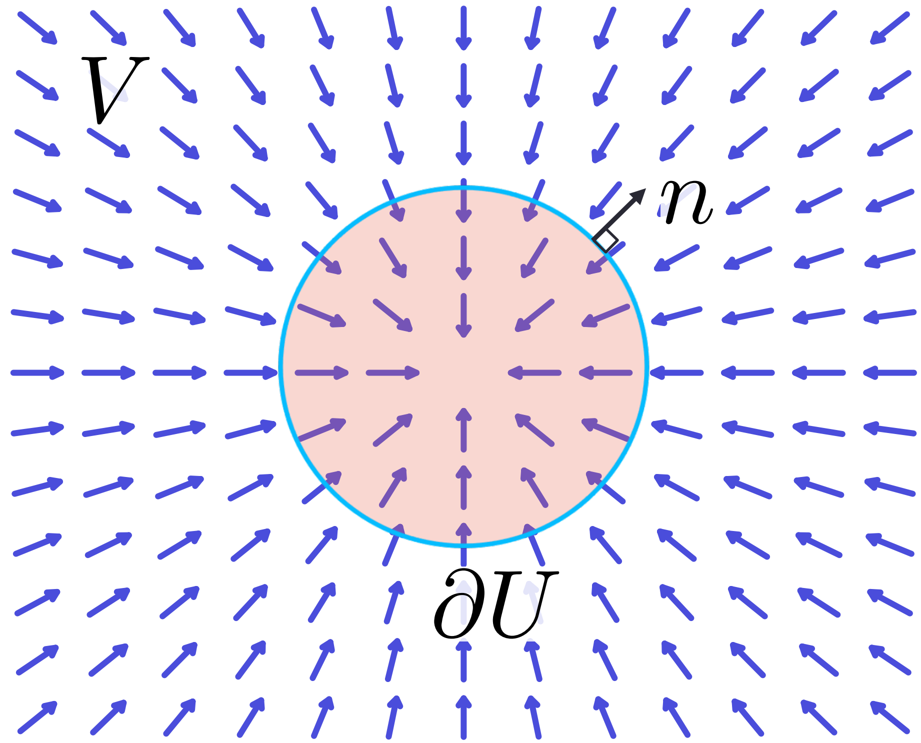

Laplacian

$$\Delta f = \mathrm{div} \nabla f = \sum_{i=1}^{3} \frac{\partial f}{\partial x_i^2}$$

Input: vector field: $f : \mathbb{R}^n \to \mathbb{R}$

Output: scalar function: $\Delta f : \mathbb{R}^n \to \mathbb{R}$

Intuition: smoothness, deviation from average

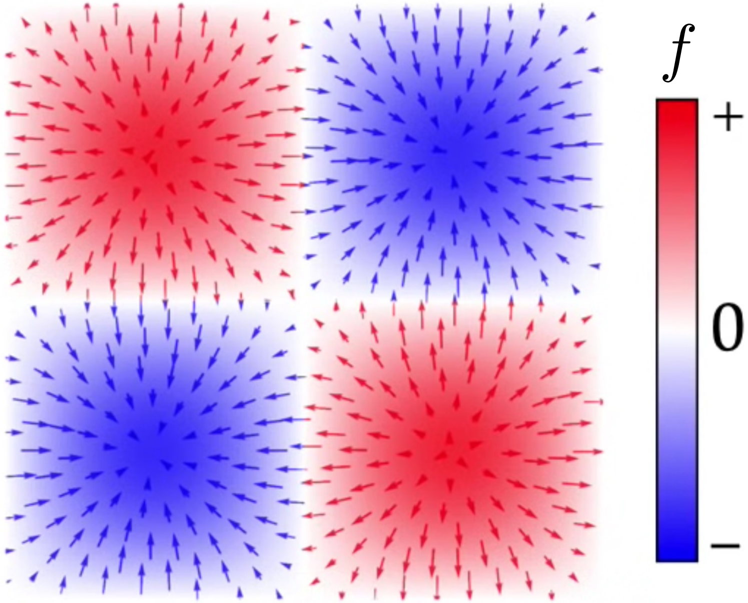

Functions on Meshes

- Assignment of a number per vertex: $f(x_i) = f_i$
- Linearly interpolated inside triangles:

$$f(x) = f_i B_i(x) + f_j B_j(x) + f_k B_k(x)$$

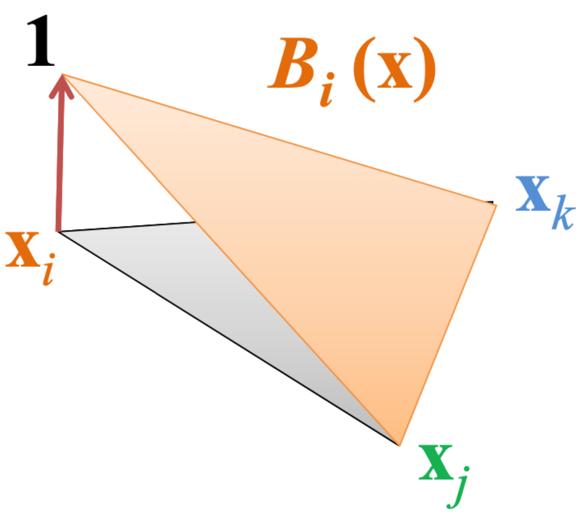

Gradient of a Function

Inside a single triangle, use piecewise-linear interpolation:

$$\nabla f(x) = f_i \nabla B_i(x) + f_j \nabla B_j(x) + f_k \nabla B_k(x)$$

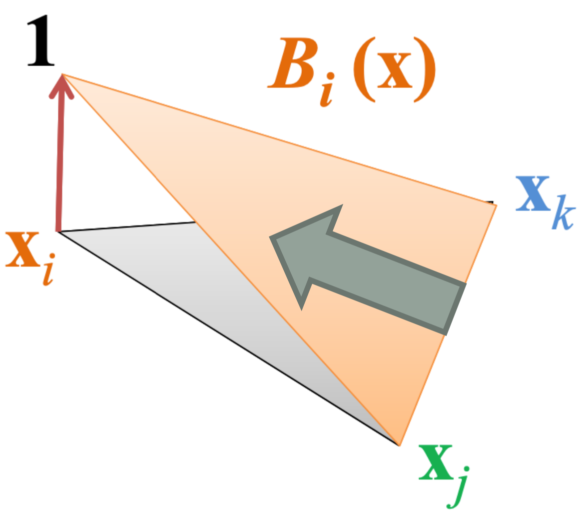

Steepest ascent direction perpendicular to opposite edge

$$\nabla B_i(x) = \nabla B_i = \frac{(x_k - x_j)^{\perp}}{2A_T}$$

Gradient is constant on a triangle.

Gradient of a Function

Inside a single triangle, use piecewise-linear interpolation:

$$\nabla f(x) = f_i \nabla B_i(x) + f_j \nabla B_j(x) + f_k \nabla B_k(x)$$
$$= \frac{f_i}{2A_T}(x_k - x_j)^\perp + \frac{f_j}{2A_T}(x_i - x_k)^\perp + \frac{f_k}{2A_T}(x_j - x_i)^\perp$$

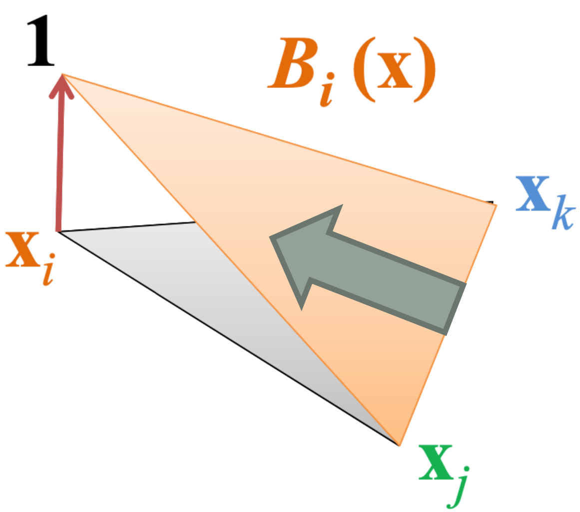

Steepest ascent direction perpendicular to opposite edge

$$\nabla B_i(x) = \nabla B_i = \frac{(x_k - x_j)^\perp}{2A_T}$$

Gradient is constant on a triangle.

Divergence of Vector Field

A vector field is piecewise constant inside a triangle:

Divergence theorem:

$$\int_U \mathrm{div} V = \int_{\partial U} V.n$$

$$\mathrm{div}(V)_i A(i) = \sum_{ijk} V_{ijk}.(x_j - x_k)^{\perp}$$

Vertex area: $A(i) = \frac{1}{3} \sum_{i \in T} A_T$

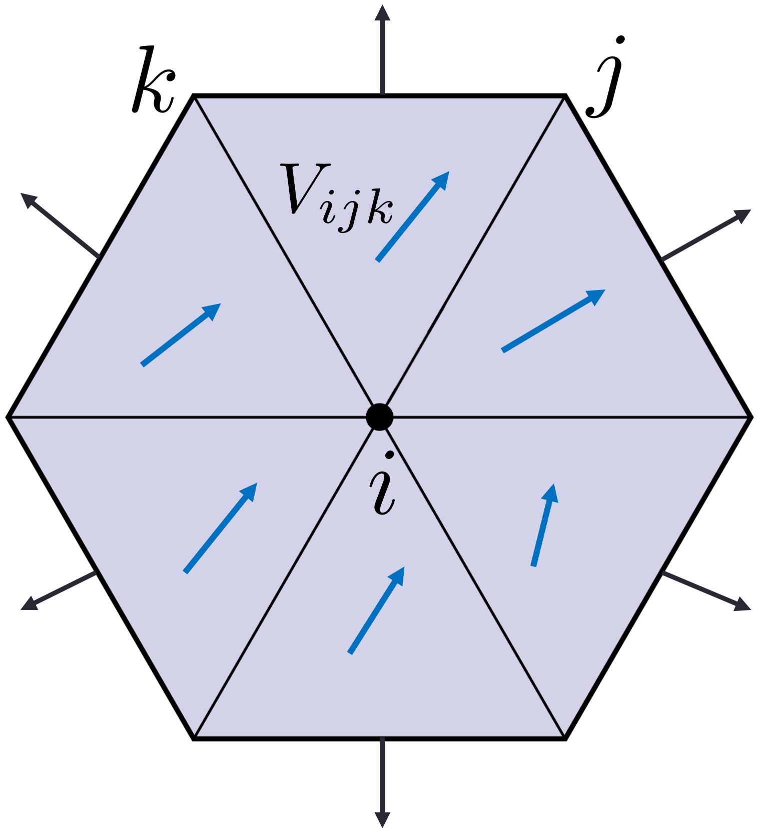

Laplacian on Meshes

Simply compose the divergence and gradient:

$$(\nabla f)_{ijk} = \frac{f_i}{2A_{ijk}}(x_k - x_j)^{\perp} + \frac{f_j}{2A_{ijk}}(x_i - x_k)^{\perp} + \frac{f_k}{2A_{ijk}}(x_j - x_i)^{\perp}$$

$$\begin{aligned}(Lf)_i A(i) &= \sum_{ijk} (\nabla f)_{ijk} . (x_j - x_k)^{\perp} \\ &= \sum_{ij} \frac{1}{2}(\cot \alpha_{ij} + \cot \beta_{ij})(f_i - f_j)\end{aligned}$$

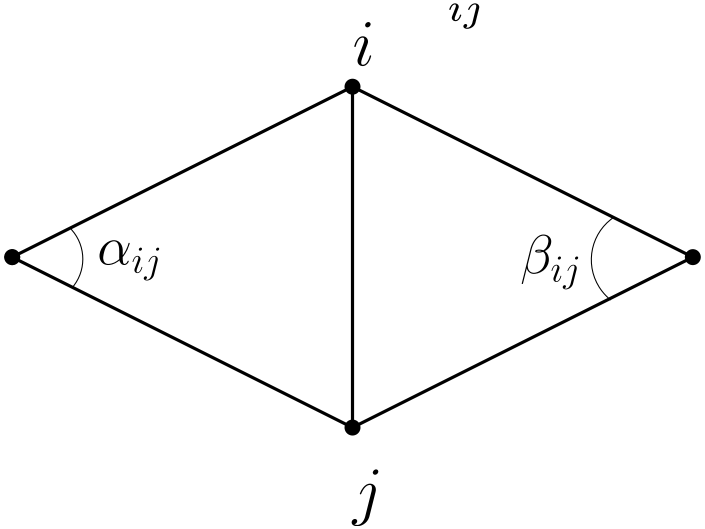

For a constant function f:

$$Lf = 0$$

Laplacian on Meshes

L is a matrix of size n x n where n is the number of vertices:

$$L_{ij} A(j) = \frac{1}{2} \cot(\alpha) + \frac{1}{2} \cot(\beta)$$

In matrix notation: $AL = -W$

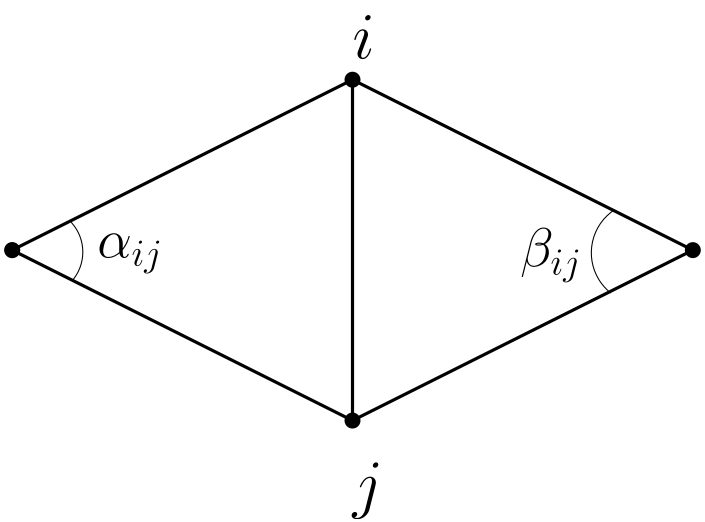

$$W_{ij} = \begin{cases} -\frac{1}{2}\left(\cot(\alpha) + \cot(\beta)\right) & \text{if } i \sim j \\ -\sum_j W_{ij} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

$$A_{ij} = \begin{cases} A(j) & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

Laplace-Beltrami – Applications

Define a multiscale signature for every point
Compare points by comparing their signatures
Compute geodesic distances

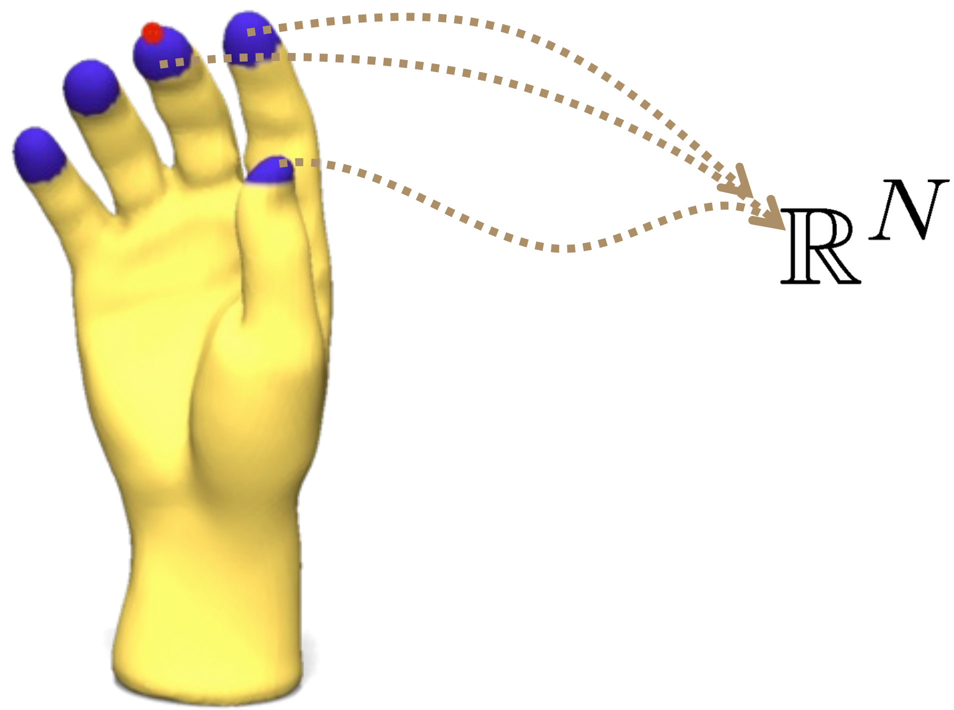

Many Signatures are derived from the LB operator

Laplacian Eigen-Decomposition

The matrix W is

- symmetric: $W_{ij} = W_{ji}$
- positive: $f^\top W f \geq 0$

There exists positive eigenvalues and eigenfunctions solutions of:

$$W\phi_i = \lambda_i A \phi_i \qquad \phi_i^\top A \phi_j = \delta_{ij}$$

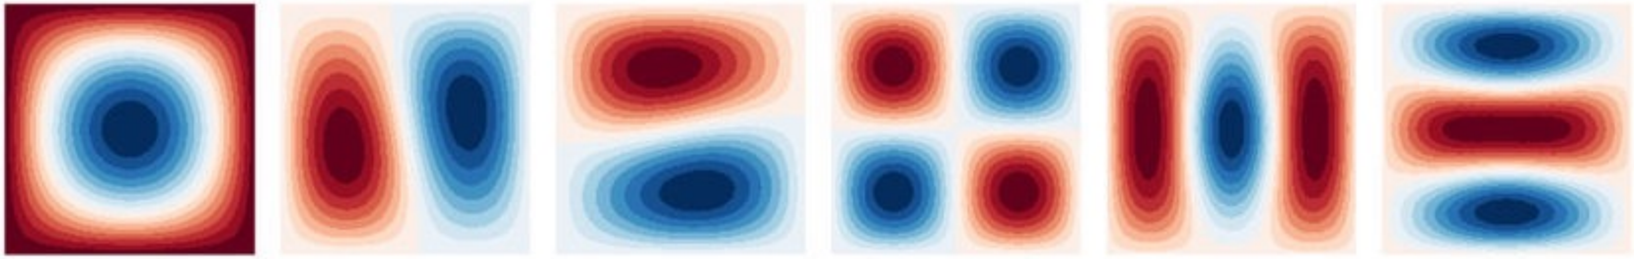

Laplacian Eigen-Decomposition

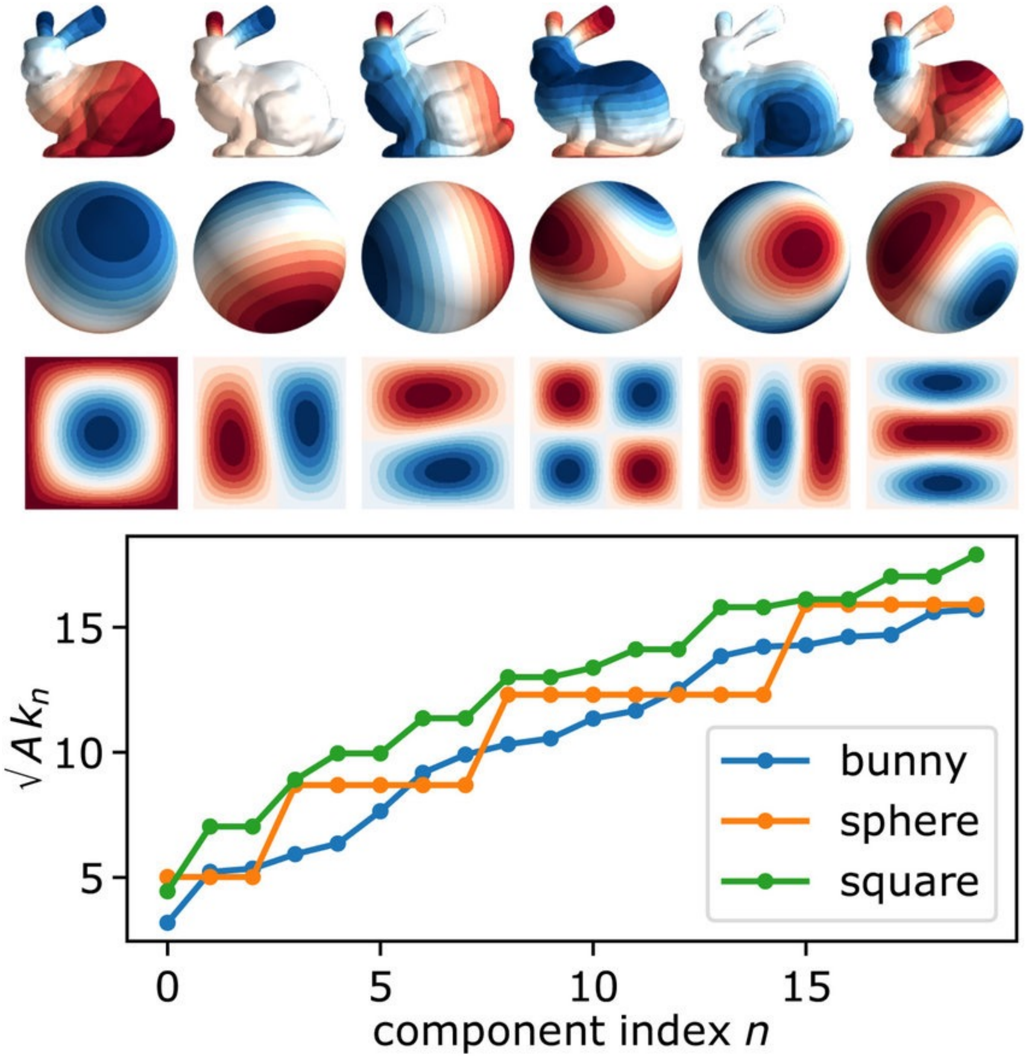

Eigenfunctions as Basis

Signal Processing on a manifold (generalizing Fourier analysis):

Given a function $f : \mathcal{M} \to \mathbb{R}$.

$$f(x) = \sum_{i=0}^{\infty} \phi_i(x) \langle \phi_i, f \rangle$$

Filter out high frequency “noise”, by truncating the series early:

$$f'(x) = \sum_{i=0}^{N} \phi_i(x) \langle \phi_i, f \rangle$$

New function will preserve the “global” properties of f.

Frequency Analysis

Multiscale nature of the spectrum:

Intuitively, eigenfunctions corresponding to larger eigenvalues, capture *smaller details* (higher frequency) of the geometry.

$\lambda_0 = 0$ $\lambda_1 = 2.6$ $\lambda_2 = 3.4$ $\lambda_3 = 5.1$ $\lambda_4 = 7.6$

- *n*-th eigenfunction has at most *n* nodal domains.
- Integral of the gradient increases.

$$\lambda_i = \int_{\mathcal{M}} \phi_i \Delta \phi_i d\mu = \int_{\mathcal{M}} \|\nabla \phi_i\|^2 d\mu$$

Heat Equation on a Surface

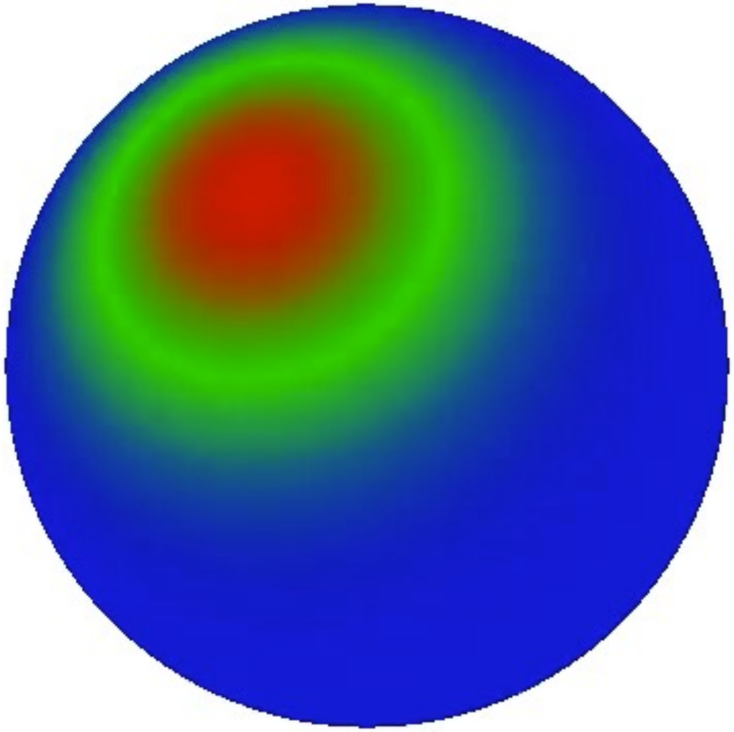

Heat Equation on a Surface

Given a compact surface without the evolution of heat is given by: $\frac{\partial f}{\partial t} = \Delta f$

Discretization in time: $\frac{f_{t+1} - f_t}{\mathrm{d}t} = \Delta f_{t+1}$

Discretization in space: $\frac{f_{t+1} - f_t}{\mathrm{d}t} = -A^{-1} W f_{t+1}$

New heat distribution solution of:

$$(A + \mathrm{d}tW) f_{t+1} = A f_t$$

Heat Equation on a Surface

Heat kernel $k_t(x,y) : \mathbb{R}^+ \times \mathcal{M} \times \mathcal{M} \to \mathbb{R}$

$$f(x,t) = \int_{\mathcal{M}} k_t(x,y) f(y,0) dy$$

$k_t(x,y)$: amount of heat transferred from x to y in time t.

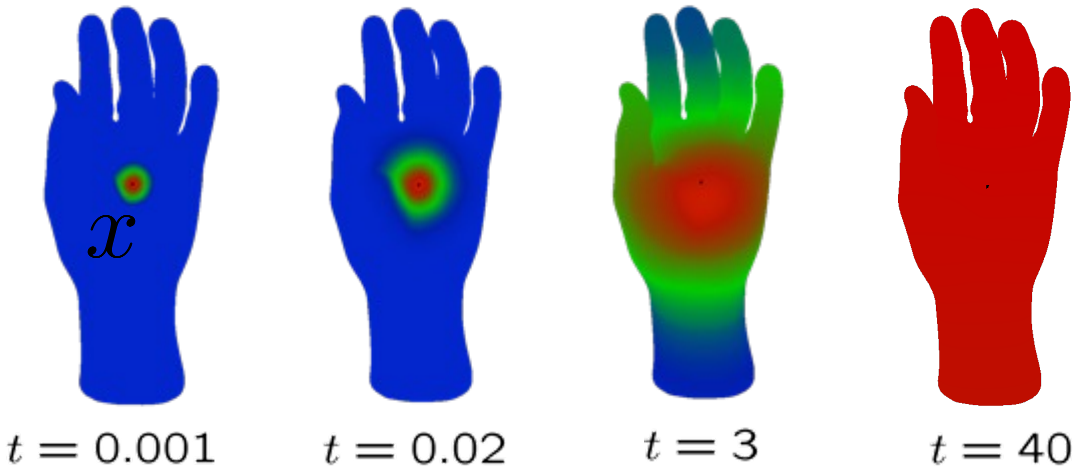

Heat Kernel

Heat kernel $k_t(x,y) : \mathbb{R}^+ \times \mathcal{M} \times \mathcal{M} \to \mathbb{R}$

$$f(x,t) = \int_{\mathcal{M}} k_t(x,y) f(y,0) dy$$

$k_t(x,y)$: amount of heat transferred from x to y in time t.

$$k_t(x,y) = \sum_i \exp(-t\lambda_i)\phi_i(x)\phi_i(y)$$

λ_i, ϕ_i eigenvalues/eigenfunctions of the LB operator.

Can be computed on a mesh using the eigenfunctions of the discrete LB operator.

Discrete Heat Kernel

Heat kernel k_t is a matrix:

$$k_t = \sum_i \exp(-t\lambda_i)\phi_i\phi_i^\top$$

Heat diffusion for time t from an initial heat distribution f_0:

$$f_t = \sum_i \exp(-t\lambda_i)\phi_i\phi_i^\top A f_0$$

λ_i, ϕ_i eigenvalues/eigenfunctions of the LB operator.

Remember: $f_0 = \sum_i \phi_i \left(\phi_i^\top A f_0\right)$

Heat Kernel Signature

$$\mathrm{HKS}(x) = k_t(x,x) = \sum_i \exp(-t\lambda_i)\phi_i(x)^2$$

λ_i, ϕ_i eigenvalues/eigenfunctions of the LB operator.

$k_t(x,x)$: amount of heat **remaining at** x after time t

A Concise and Provably Informative Multi-scale Signature Based on Heat Diffusion, Sun et al. 2009

Multiscale Matching

Comparing points through their HKS

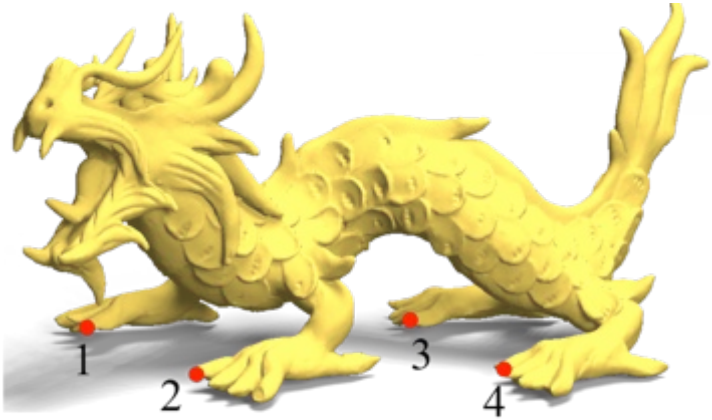

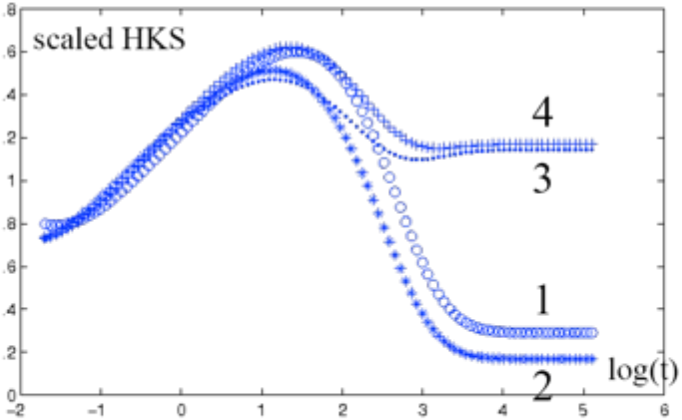

Multiscale Matching

Finding similar points across multiple shapes:

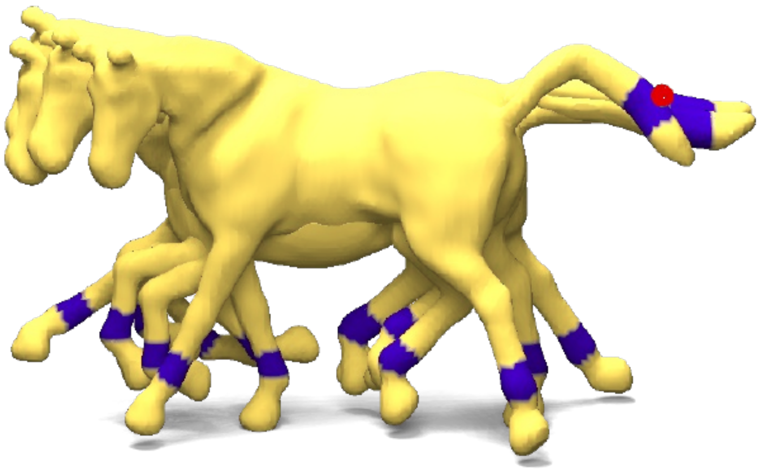

Medium scale

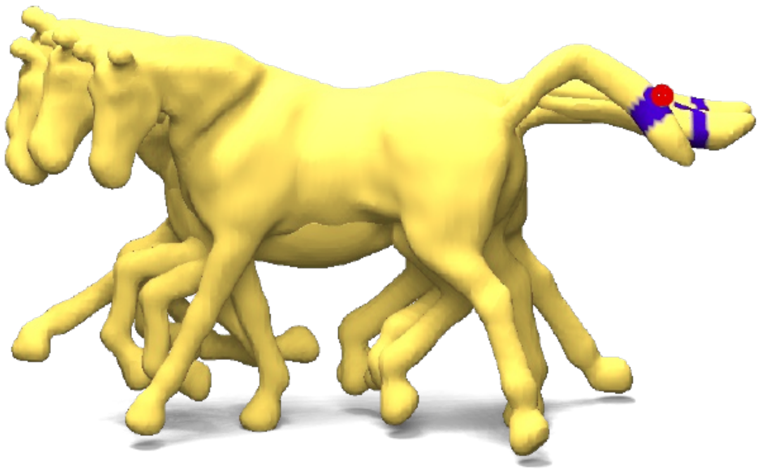

Full scale

A Concise and Provably Informative Multi-scale Signature Based on Heat Diffusion, Sun et al. 2009

Multiscale Matching

HKS is stable under mild deformations

Generalized Heat Kernel Signatures, Valentin Zobel, Jan Reininghaus, Ingrid Hotz

Heat Kernel Signature

Relation to scalar curvature for small t :

$$k_t(x,x) = \frac{1}{4\pi t}\sum_{i=0}^{\infty} a_i t^i \quad a_0 = 1, a_1 = \frac{1}{6}K(x)$$

$K(x)$: Gaussian Curvature

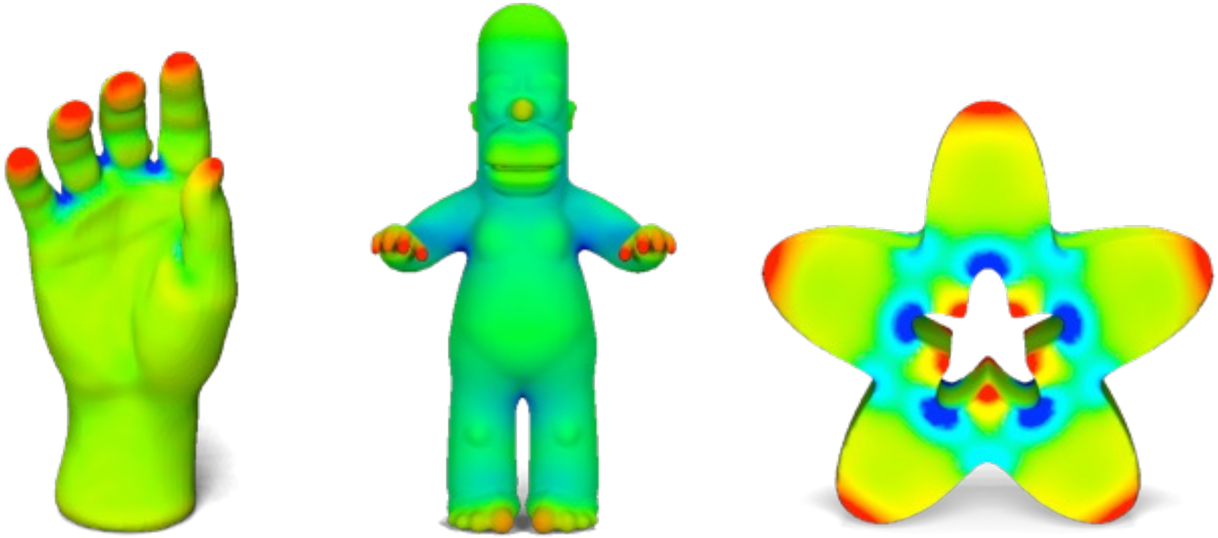

A Concise and Provably Informative Multi-scale Signature Based on Heat Diffusion, Sun et al. 2009

Heat Kernel Signature

Can be interpreted as multi-scale intrinsic curvature.

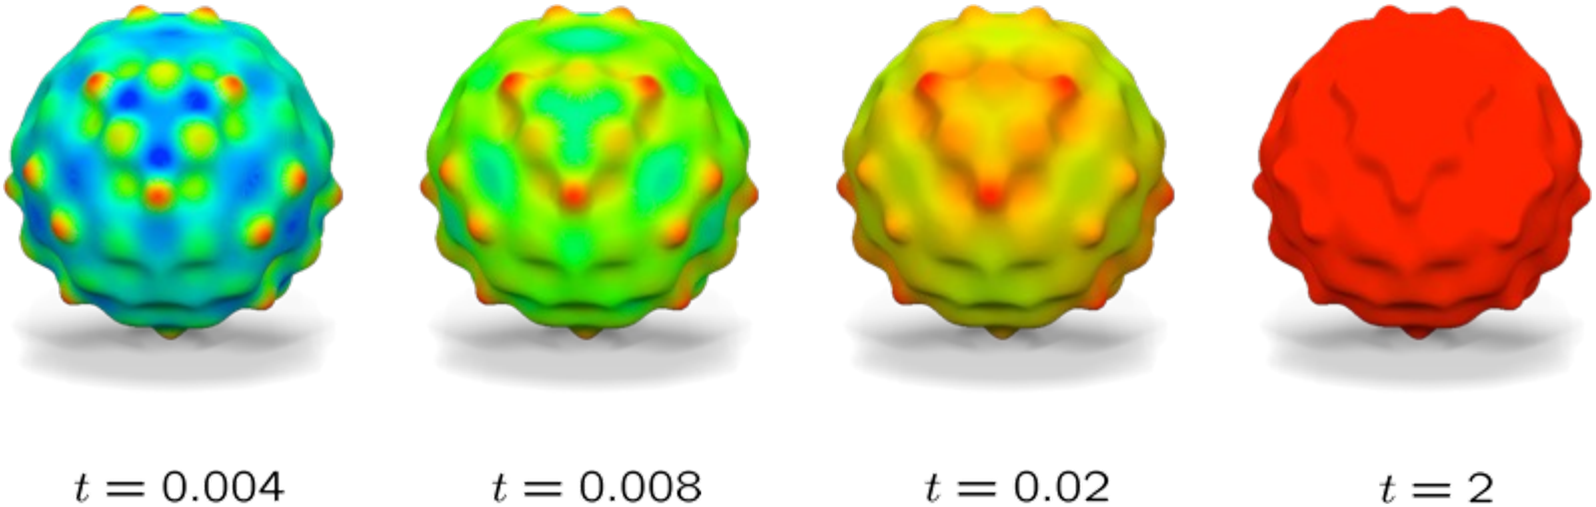

A Concise and Provably Informative Multi-scale Signature Based on Heat Diffusion, Sun et al. 2009

Wave Kernel Signature

$$\mathrm{WKS}(x,e) = \sum_{i=0}^{N} \exp\left(-\frac{(e-\log\lambda_i)^2}{2\sigma^2}\right)\phi_i(x)^2$$

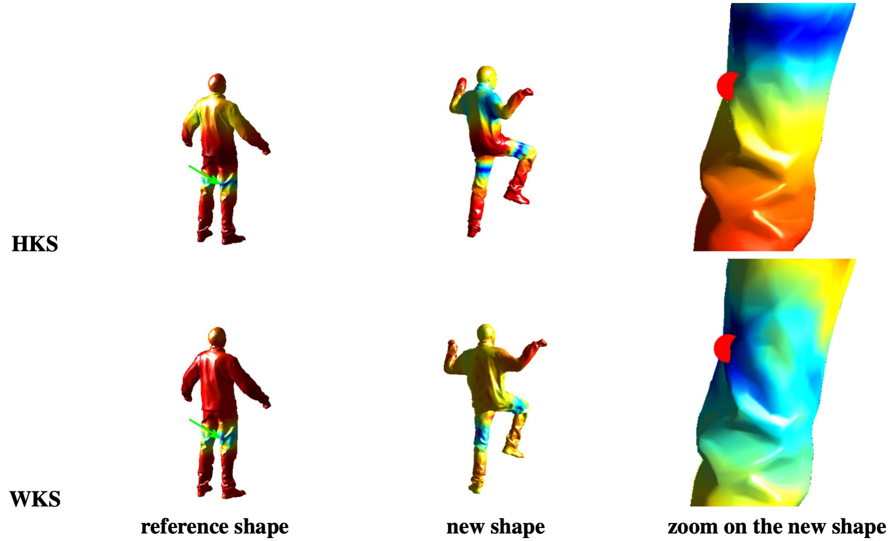

Gives more prominence to medium frequencies.
Can result in more accurate predictions.

The wave kernel signature: A quantum mechanical approach to shape analysis, Aubry et al. 2009

Generalization

Learning-based Spectral Descriptors

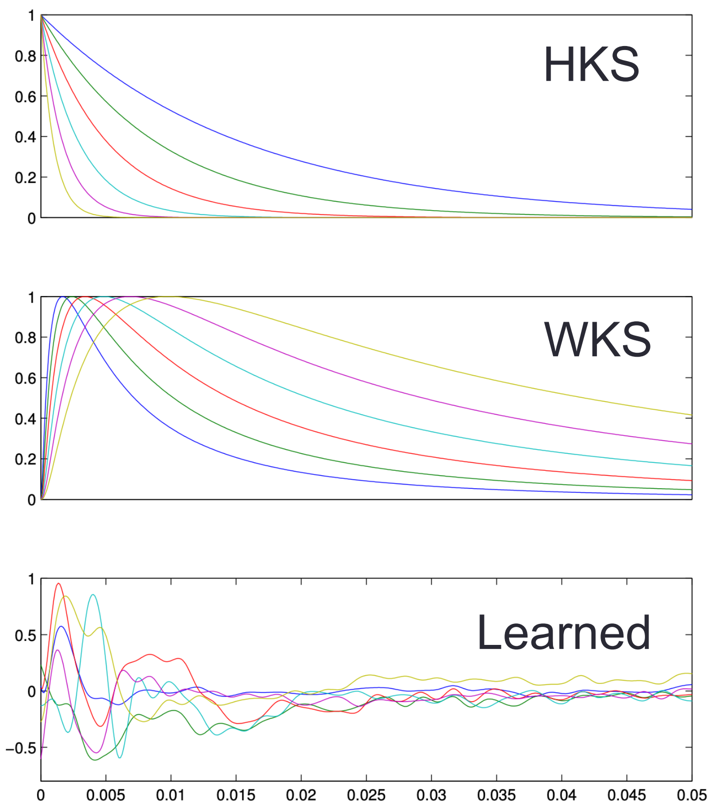

$$\mathrm{LKS}(x,t) = \sum_{i=0}^{N} f_t(\lambda_i)\phi_i(x)^2$$

Learn the optimal kernel from data

Learning spectral descriptors for deformable shape correspondence, Litman et al. 2014

Conclusion

- Spectral Methods in Shape Analysis
  - Discrete (graph) Laplacians
  - Laplace-Beltrami operator and its properties
  - Some applications
- Key message:
  - Laplacian matrices allow to organize shape information in a multi-scale, easy to manipulate way.