Continuous Laplacian, Functional Map, Spectral CNN

Instructor: Hao Su

CONTINUOUS LAPLACIAN-BERTRAMI OPERATOR

Planar Region

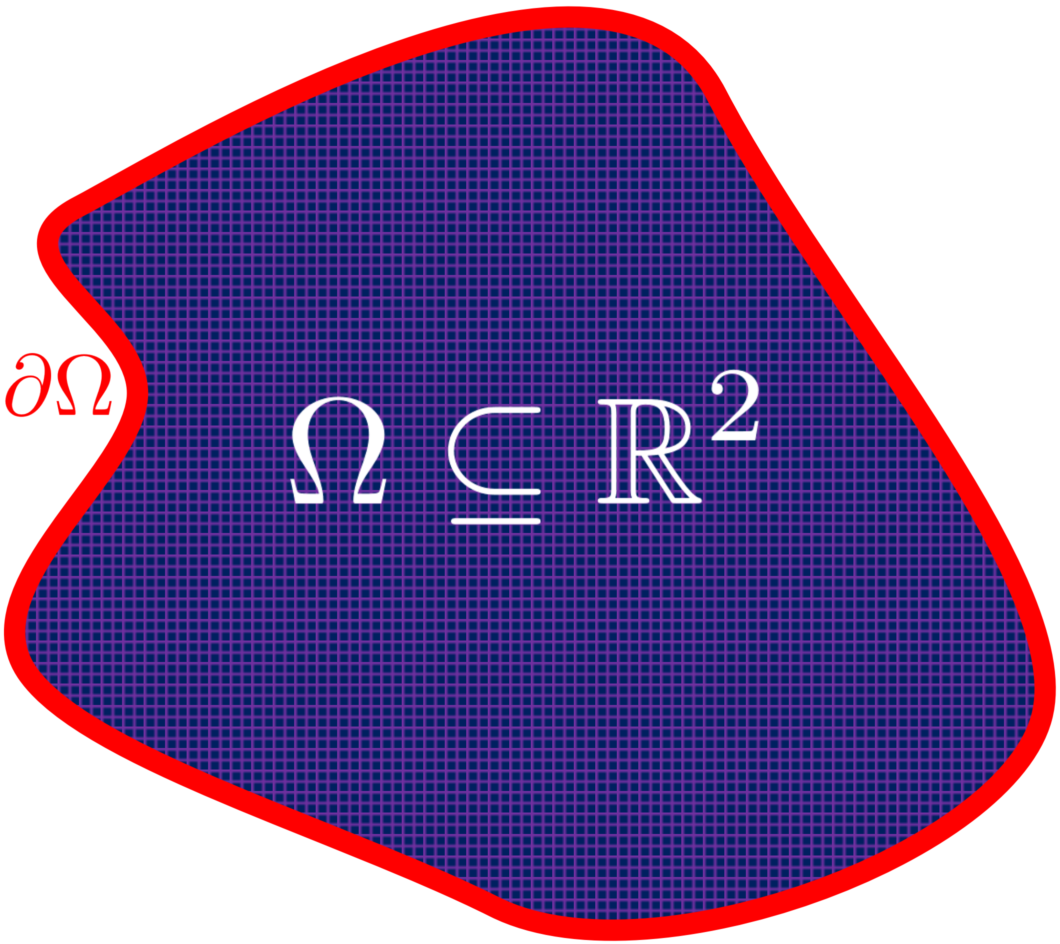

Heat equation:

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

$$\Delta := \sum_i \frac{\partial^2}{\partial x_i^2}$$

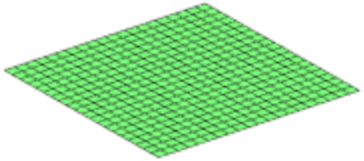

Typical Notation

$$\text{“}\Delta = \underbrace{\nabla\cdot}_{\text{divergence}}\underbrace{\nabla}_{\text{gradient}}\text{”}$$

More later...

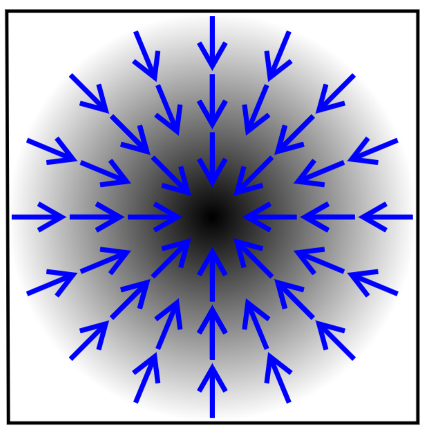

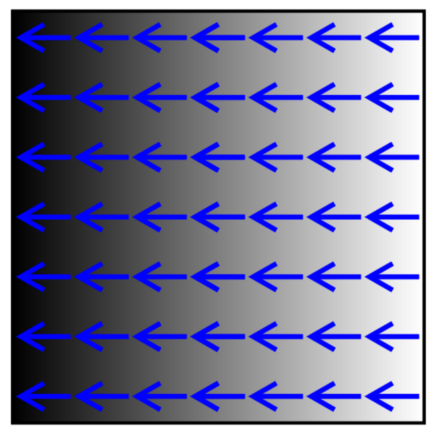

Gradient operator:

$$\nabla := \left(\frac{\partial}{\partial x_1}, \frac{\partial}{\partial x_2}, \ldots, \frac{\partial}{\partial x_n}\right)$$

http://www.gamasutra.com/db_area/images/feature/4164/figy.png, https://en.wikipedia.org/wiki/Gradient

Divergence Operator

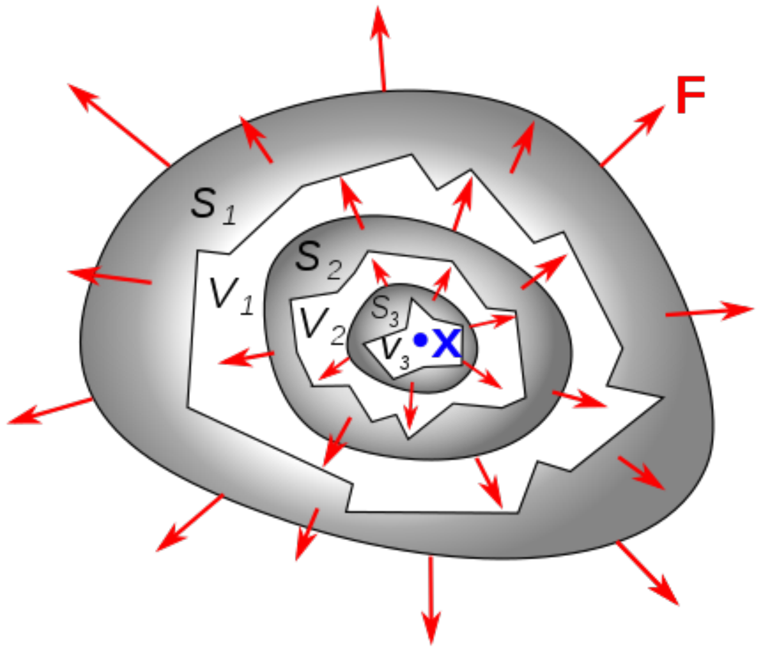

The divergence at a point **x** is the limit of the ratio of the flux through the surface S_i *(red arrows)* to the volume for any sequence of closed regions V_1, V_2, V_3... enclosing **x** that approaches zero volume: $\nabla \cdot F = \lim_{|V_i| \to 0} \frac{\Phi(S_i)}{|V_i|}$

Positivity, Self-Adjointness

$$\{f(\cdot) \in C^\infty(\Omega) : f|_{\partial\Omega} \equiv 0\}$$

"Dirichlet boundary conditions"

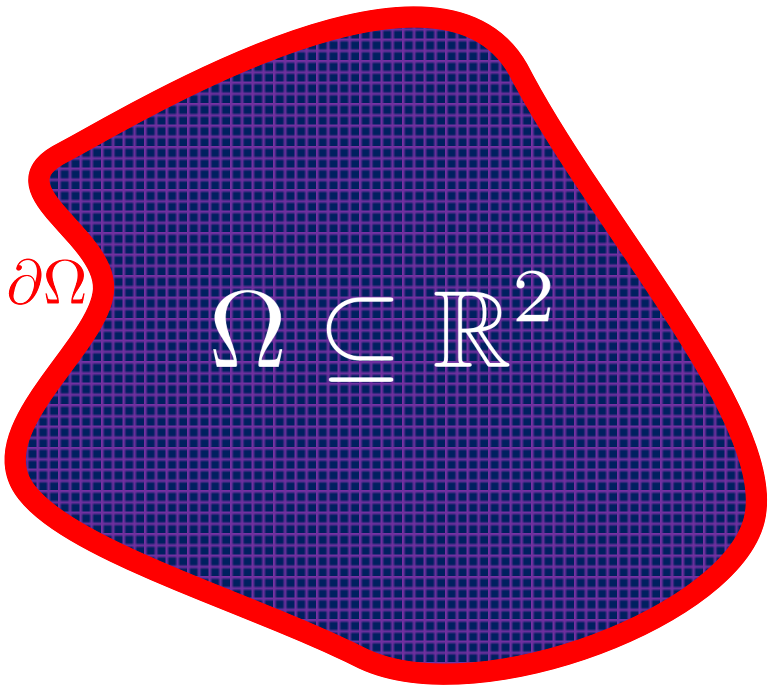

$$\mathcal{L}[f] := -\Delta f$$

$$\langle f, g \rangle := \int_\Omega f(x) g(x)\, dx$$

On board:

1. Positive: $\langle f, \mathcal{L}[f] \rangle \geq 0$
2. Self-adjoint: $\langle f, \mathcal{L}[g] \rangle = \langle \mathcal{L}[f], g \rangle$

Proof

Proof of 1

$$\langle f, \mathscr{L}[f] \rangle = \int_{\Omega} f(-\nabla \cdot \nabla f)\, dV = \int_{\partial\Omega} f(-\nabla f \cdot \overrightarrow{n})\, dS + \int_{\Omega} \nabla f \cdot \nabla f\, dV = \int_{\Omega} \nabla f \cdot \nabla f\, dV \geq 0$$

where the second equality follows from Green formula, and the third equality follows from $f|_{\partial\Omega} \equiv 0$

Proof of 2

$$\langle f, \mathscr{L}[g] \rangle = \int_{\Omega} f(-\nabla \cdot \nabla g)\, dV = \int_{\partial\Omega} f(-\nabla g \cdot \overrightarrow{n})\, dS + \int_{\Omega} \nabla f \cdot \nabla g\, dV = \int_{\Omega} \nabla f \cdot \nabla g\, dV$$

where the second equality follows from Green formula, and the third equality follows from $f|_{\partial\Omega} \equiv 0$

Similarly, $\langle \mathscr{L}[f], g \rangle = \int_{\Omega} \nabla g \cdot \nabla f\, dV$

It also shows $\langle f, \mathscr{L}[g] \rangle = \int_{\Omega} \nabla f \cdot \nabla g\, dV$

Dirichlet Energy

$$E[f] := \int_\Omega \langle \nabla f, \nabla f \rangle \, dA$$

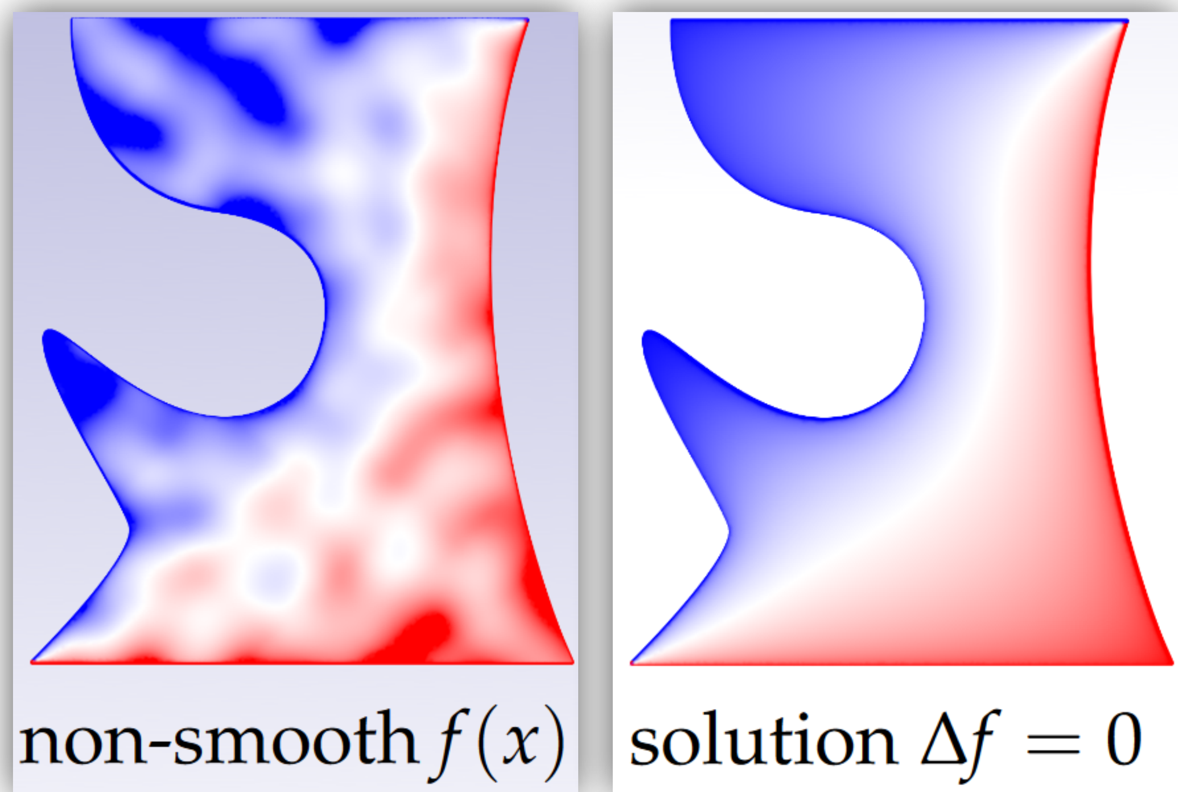

Images made by E.Vouga

On board:

$$\min_f E[f]$$
$$\text{s.t. } f|_{\partial\Omega} = g$$

$$\Delta f \equiv 0$$

"Laplace equation"
"Harmonic function"

Proof

We use variational method to derive.

Lagrangian: $$\mathbb{L}[f] = \frac{1}{2}\int \langle \nabla f, \nabla f \rangle + \int_{\partial\Omega} \lambda(x)(f(x) - g(x))$$

So

$$\delta\mathbb{L}[f] = \mathbb{L}[f+\delta h] - \mathbb{L}[f] = \int_{\Omega} \langle \nabla f, \nabla \delta h \rangle + \int_{\partial\Omega} \lambda(x)\delta h(x) = \int_{\partial\Omega} \delta h(\nabla f \cdot \overrightarrow{n}) - \int_{\Omega} \delta h(\nabla \cdot \nabla f) + \int_{\partial\Omega} \lambda(x)\delta h(x)$$

In the interior of Ω, $\Delta f \equiv 0$ so that $\delta\mathbb{L}[f] = 0$ for any δh

Harmonic Functions

$$\Delta f \equiv 0$$

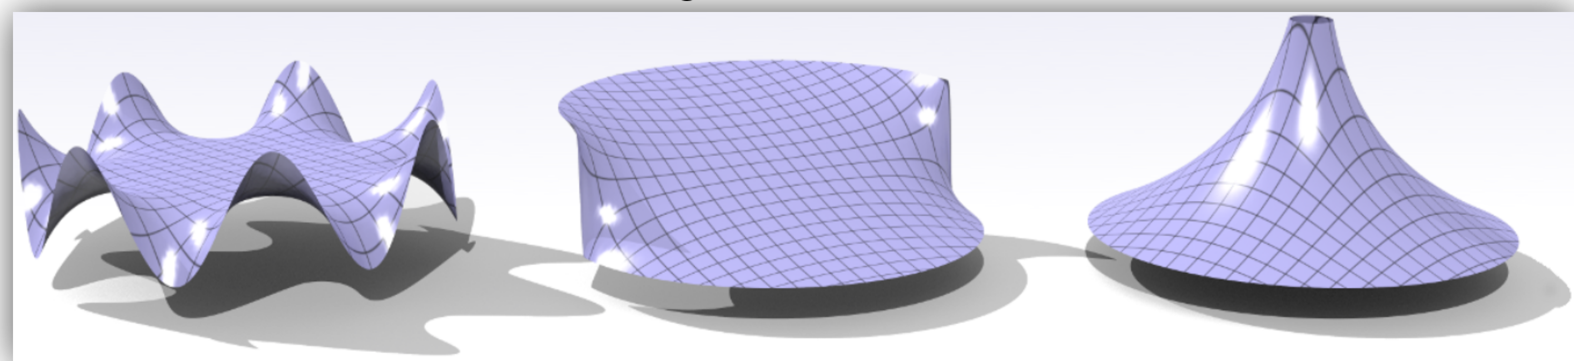

Mean value property:

$$f(x) = \frac{1}{\pi r^2} \int_{B_r(x)} f(y)\, dA$$

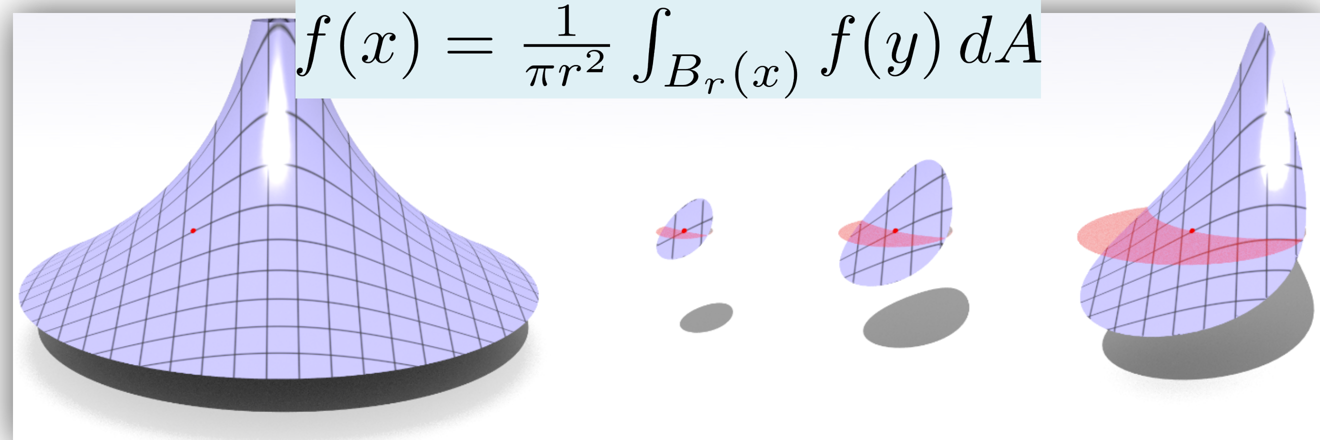

Images made by E. Vouga

Recall: Scalar Functions

$f :$ 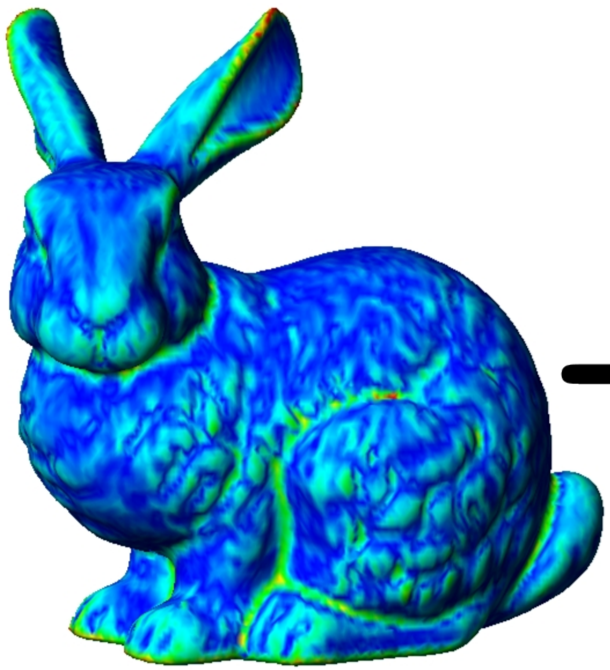 $\longrightarrow \mathbb{R}$

http://www.ieeta.pt/polymeco/Screenshots/PolyMeCo_OneView.jpg

Map points to real numbers

Gradient Vector Field

$$\nabla f : S \to \mathbb{R}^3 \text{ with}$$

$$\begin{cases} \langle (\nabla f)(p), v \rangle = (df)_p(v), v \in T_pS \\ \langle (\nabla f)(p), N(p) \rangle = 0 \end{cases}$$

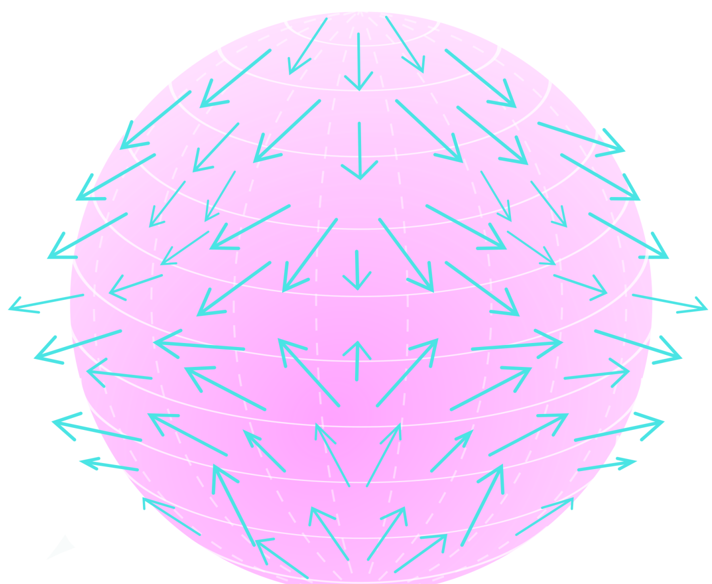

Following *Curves and Surfaces*, Montiel & Ros

Dirichlet Energy

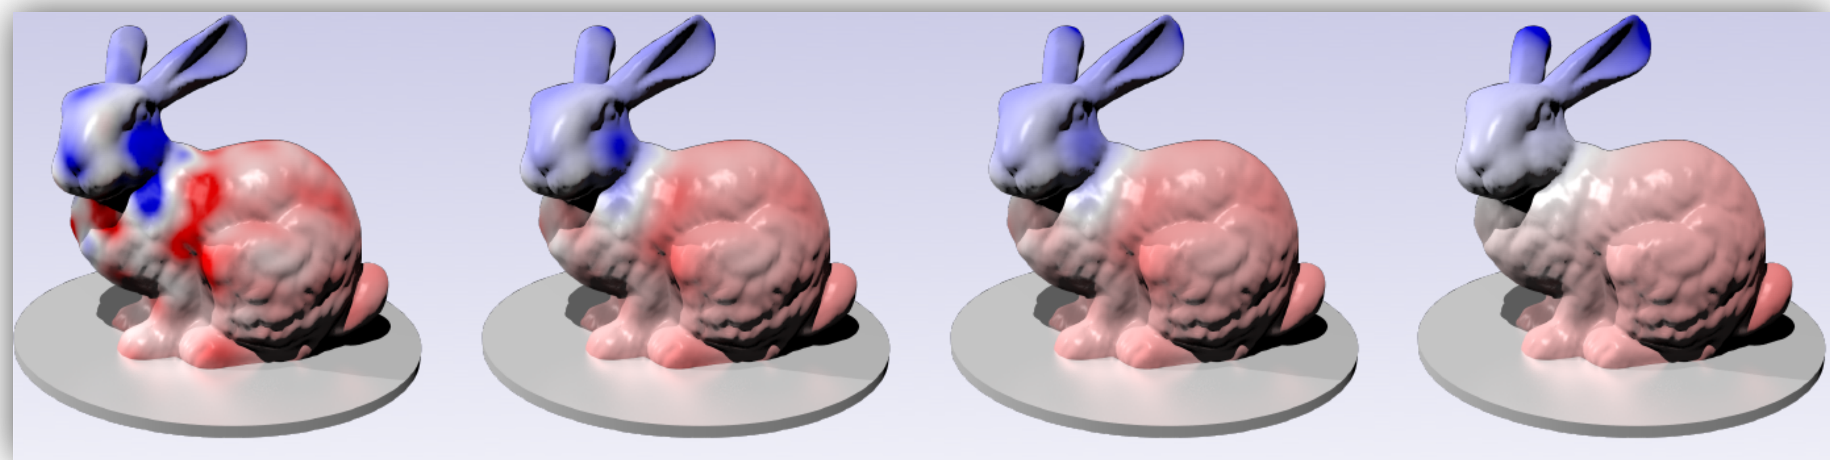

Decreasing *E*

$$E[f] := \int_S \|\nabla f\|_2^2 \, dA$$

Images made by E. Vouga

What is Divergence?

$$V : S \to \mathbb{R}^3 \text{ where } V(p) \in T_pS$$
$$dV_p : T_pS \to \mathbb{R}^3$$
$$\{e_1, e_2\} \subset T_pS \text{ orthonormal basis}$$

$$(\nabla \cdot V)_p := \sum_{i=1}^{2} \langle e_i, dV(e_i) \rangle_p$$

Things we should check (but probably won't):

- Independent of choice of basis

Eigenfunctions

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

Vibration modes of surface (not volume!)

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

Nodal Domains

Theorem (Courant). The n-th eigenfunction of the Dirichlet boundary value problem has at most n nodal domains.

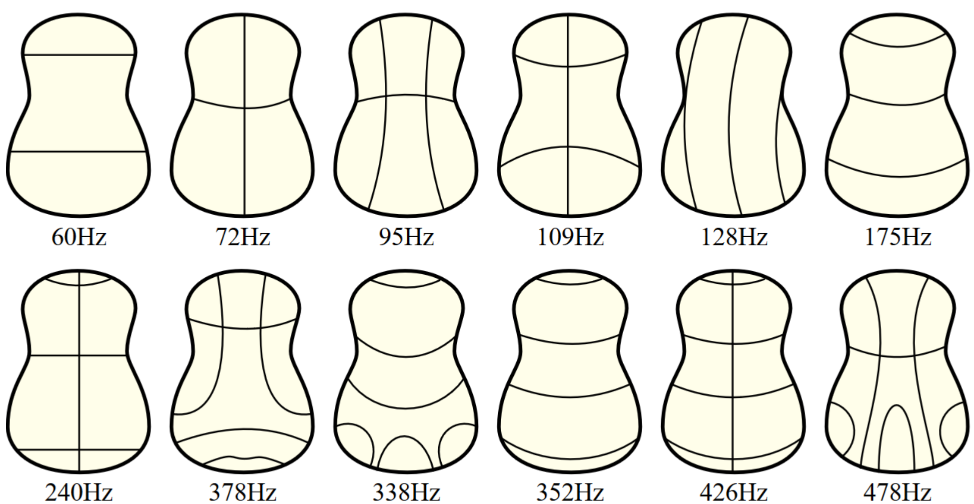

https://en.wikipedia.org/wiki/Ernst_Chladni

Practical Application

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

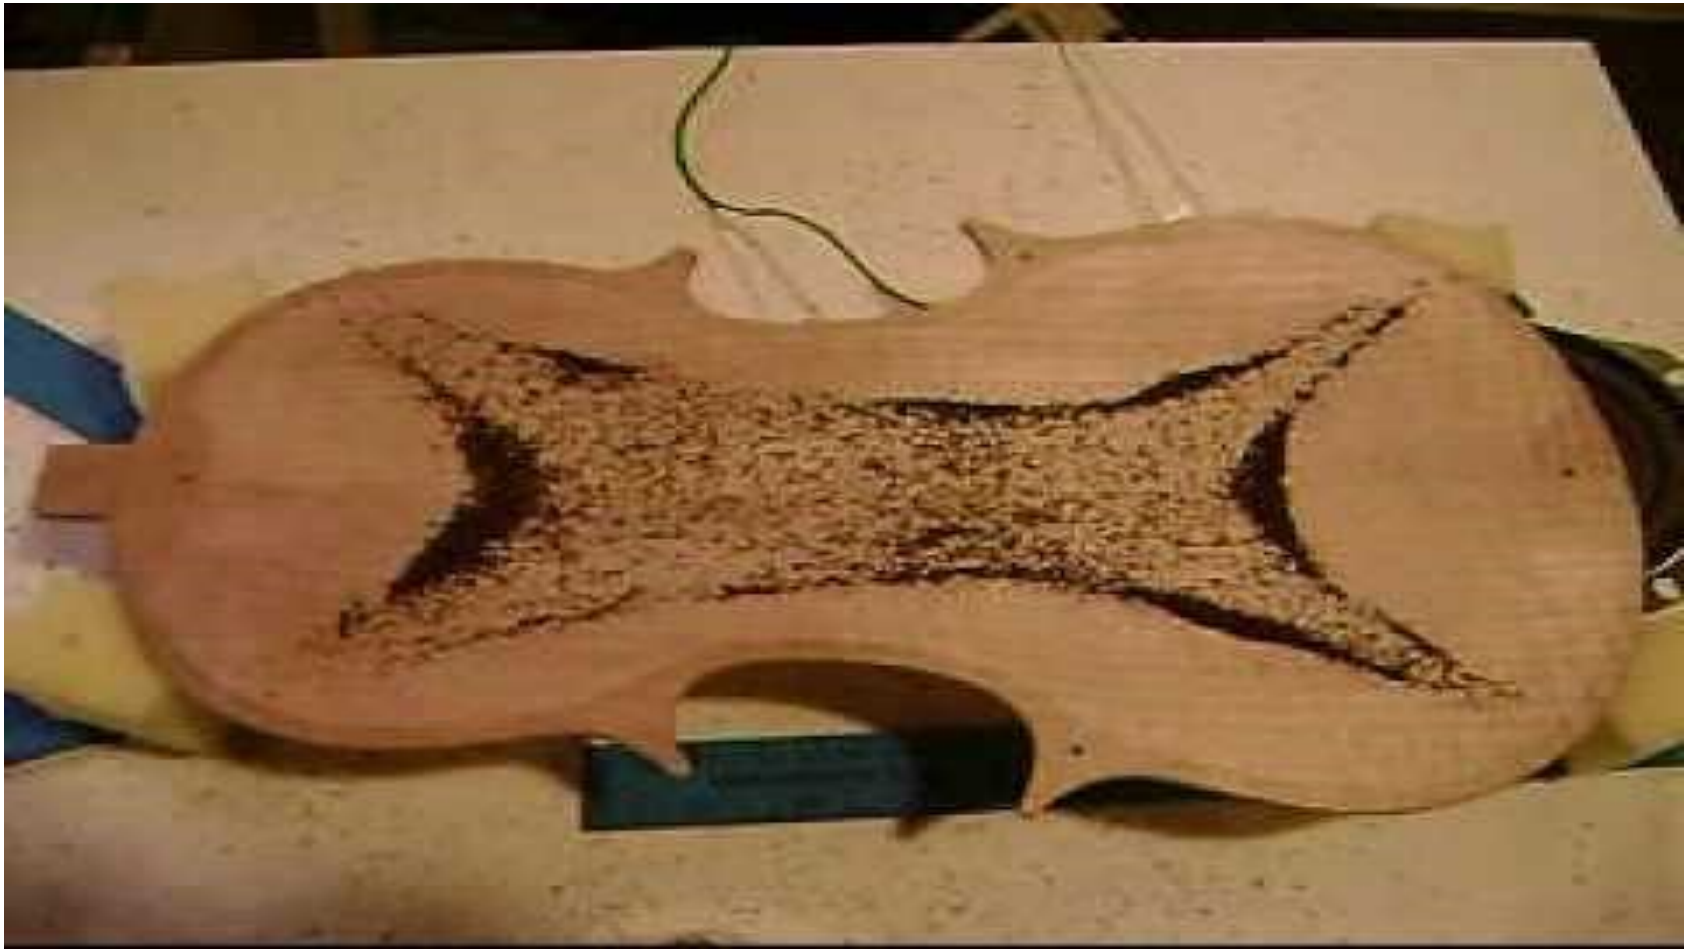

https://www.youtube.com/watch?v=3uMZzVvnSiU

FUNCTIONAL MAP

Maps

$$\phi : X \longrightarrow Y$$

Map from *X* to *Y*

Maps and Correspondences

- Multiscale mappings
  - Point/pixel level
  - part level

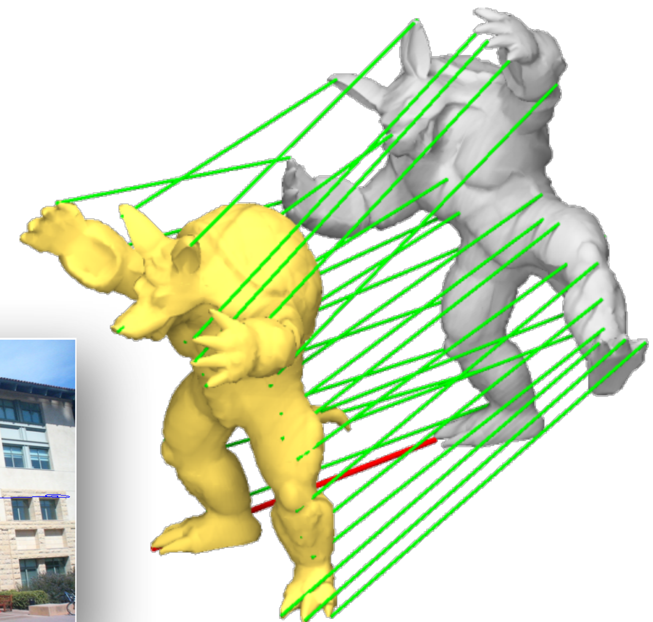

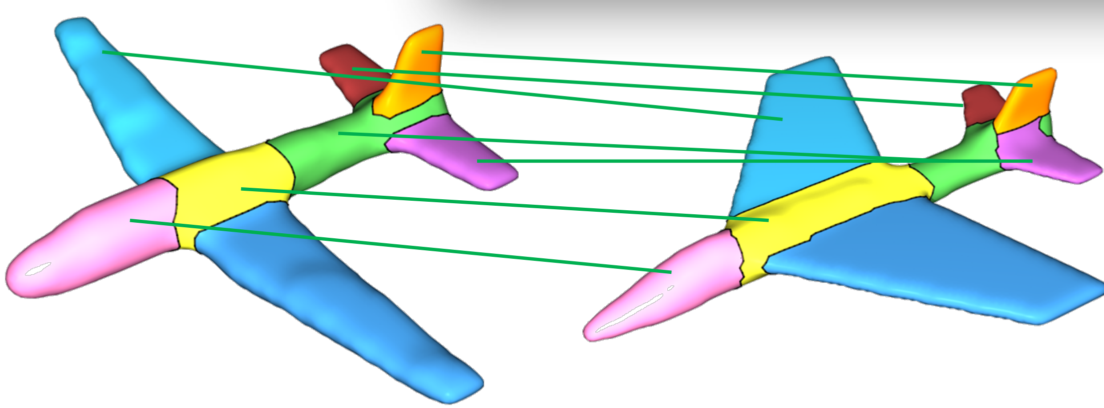

Maps capture what is the same or similar across two data sets

A Dual View: Functions and Operators

- Functions on data
  - Properties, attributes, descriptors, part indicators, etc.

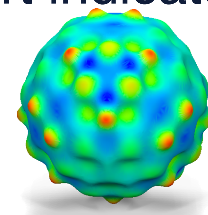

Curvature

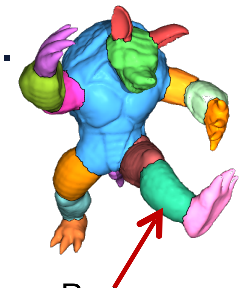

Parts

- Operators on functions
  - Maps of functions to functions
    - Laplace-Beltrami operator on a manifold M

$$\Delta : C^{\infty}(M) \rightarrow C^{\infty}(M), \Delta f = \text{div } \nabla f$$

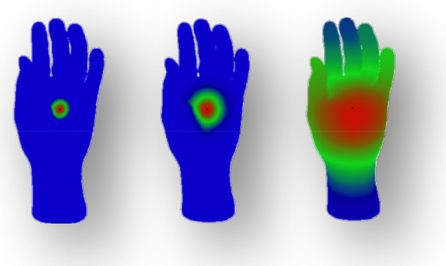

$$\frac{\partial u}{\partial t} = -\Delta u$$

heat diffusion

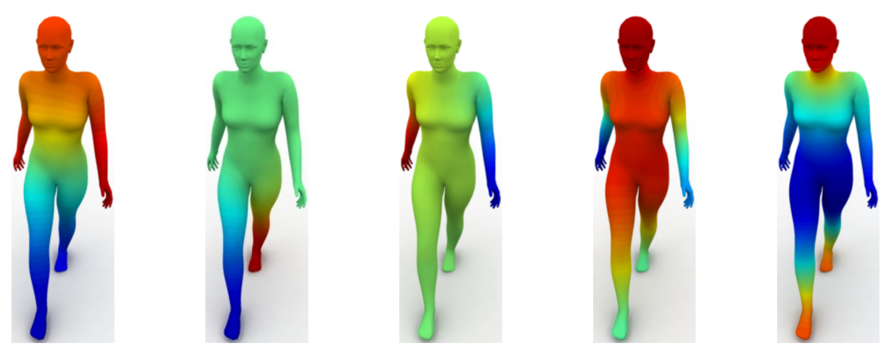

Laplace Beltrami eigenfunctions

Starting from a Regular Map ϕ

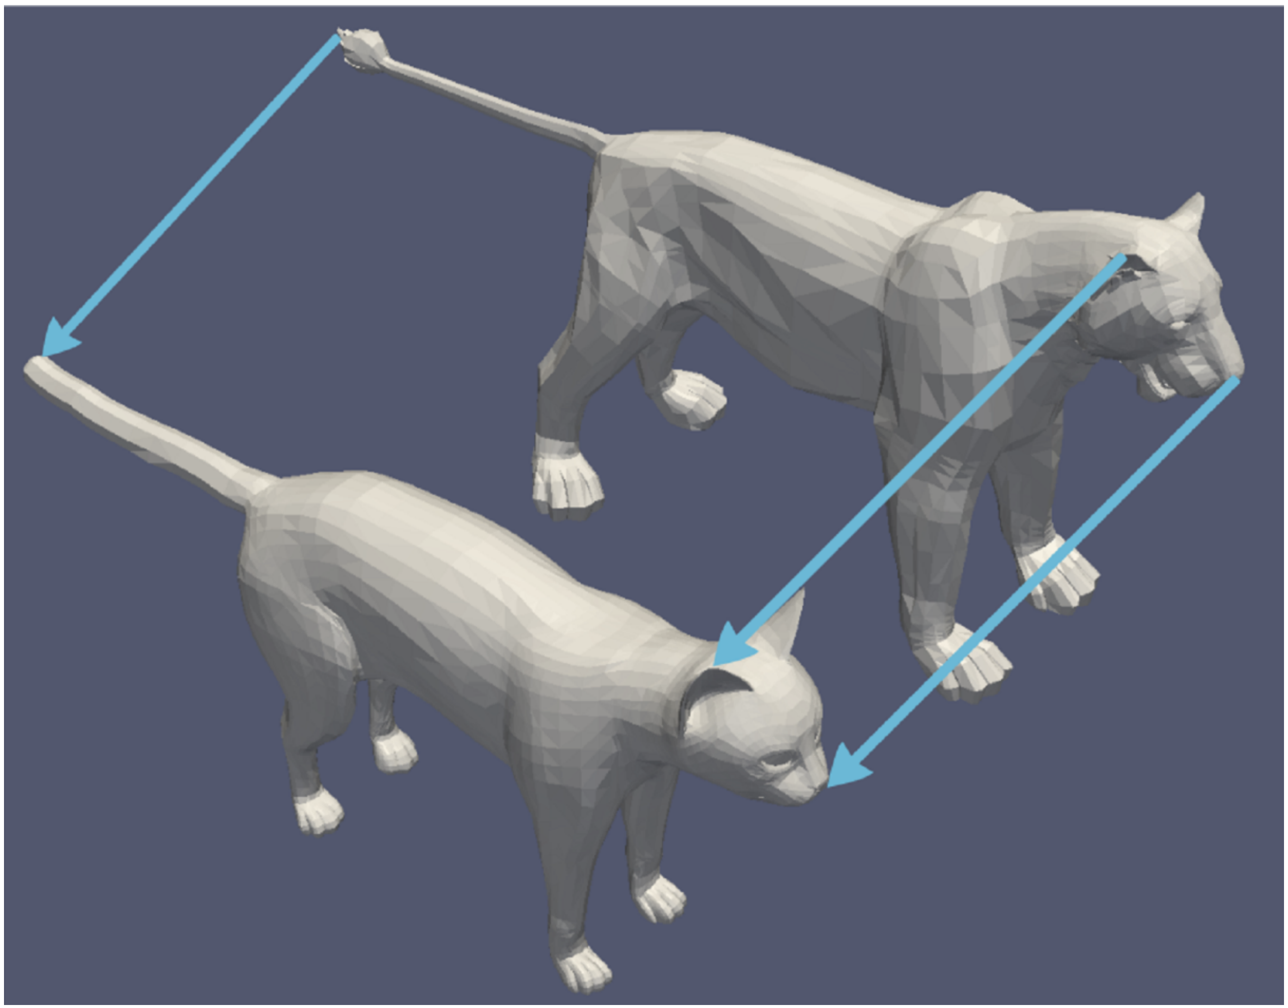

φ: lion → cat

Attribute Transfer via Pull-Back

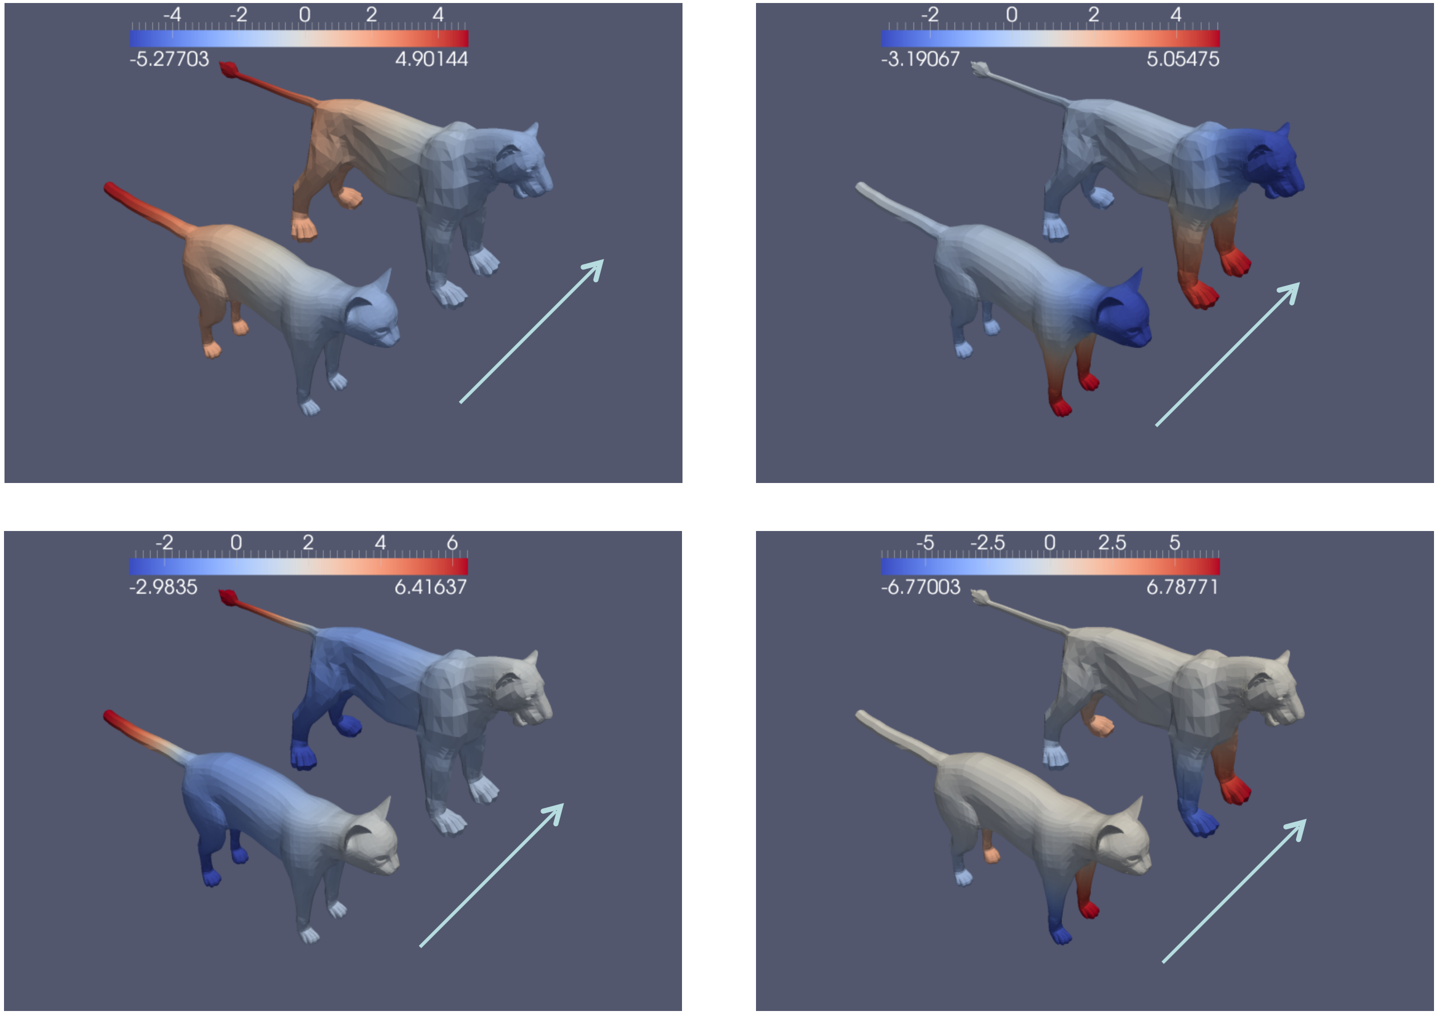

T_φ: cat → lion

from cat to lion

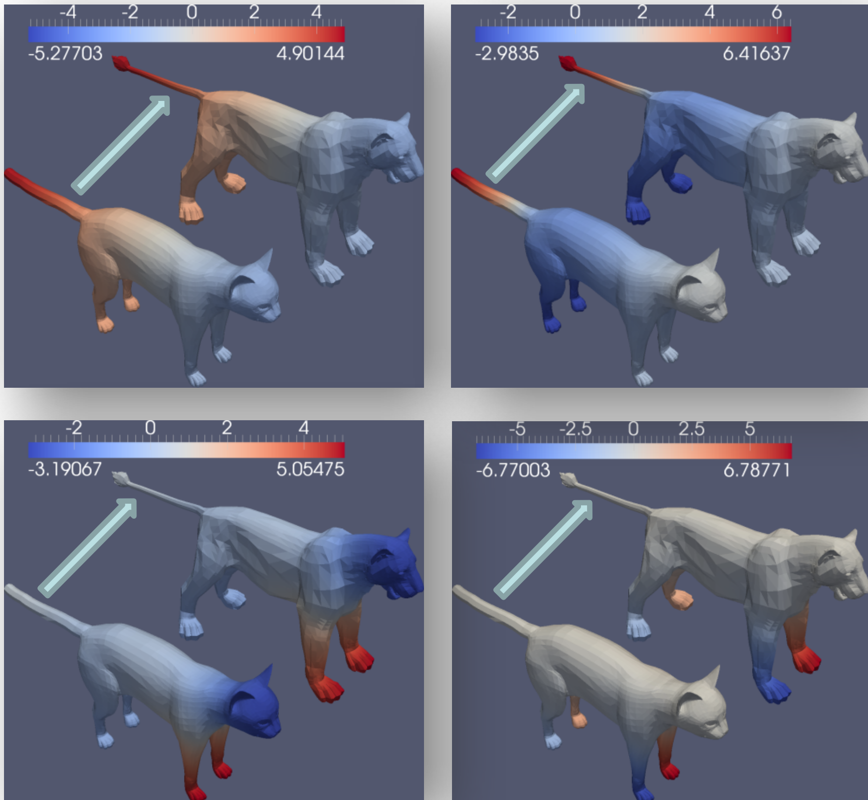

Functions on cat are transferred to lion using T_φ

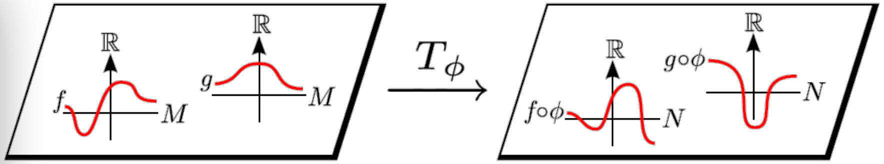

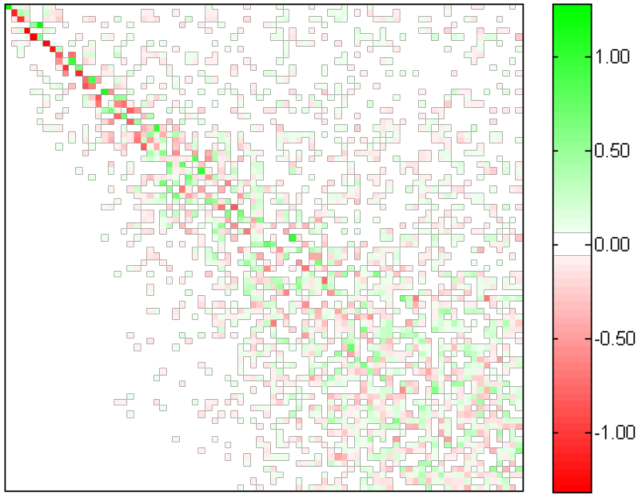

T_φ is a linear operator (matrix)

$$T_\phi : L^2(\text{cat}) \to L^2(\text{lion})$$

Functional Map

$$\phi : M \to N$$

$$T_\phi : L^2(N) \to L^2(M)$$

Dual of a point-to-point map

Bases for a Function Space

Point basis
Finite-element basis

Local bases

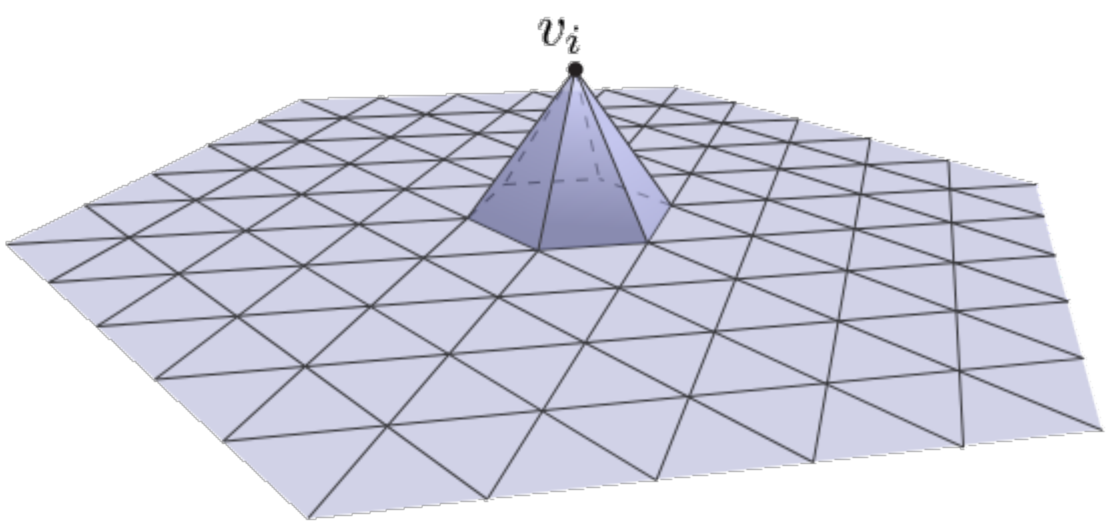

Bases for a Function Space

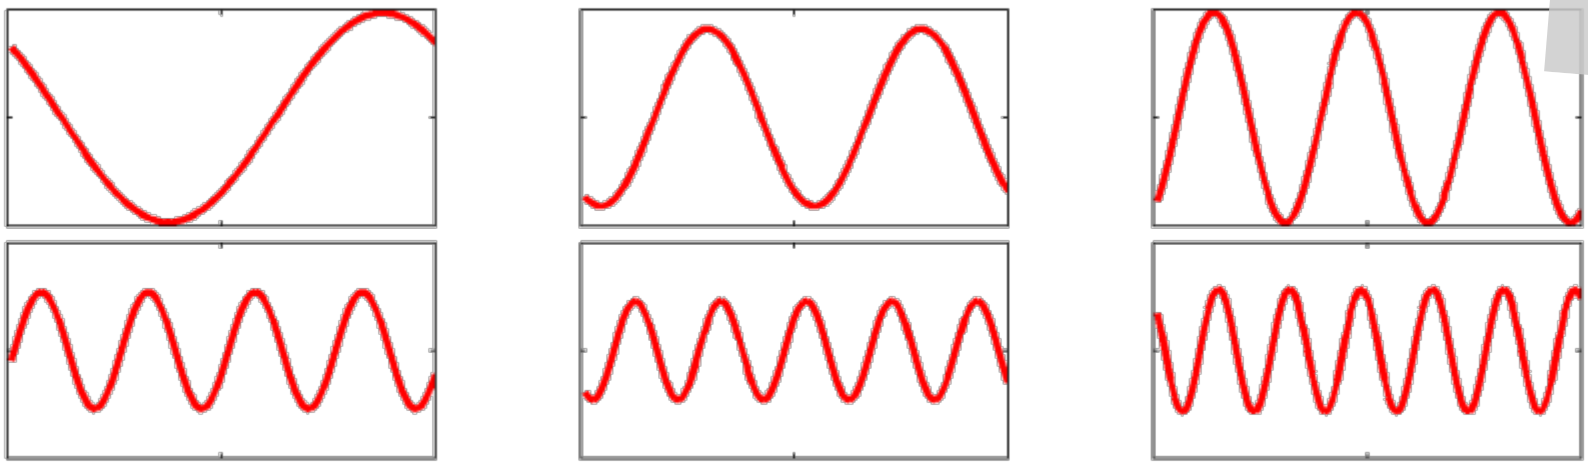

Fourier

Laplace-Beltrami

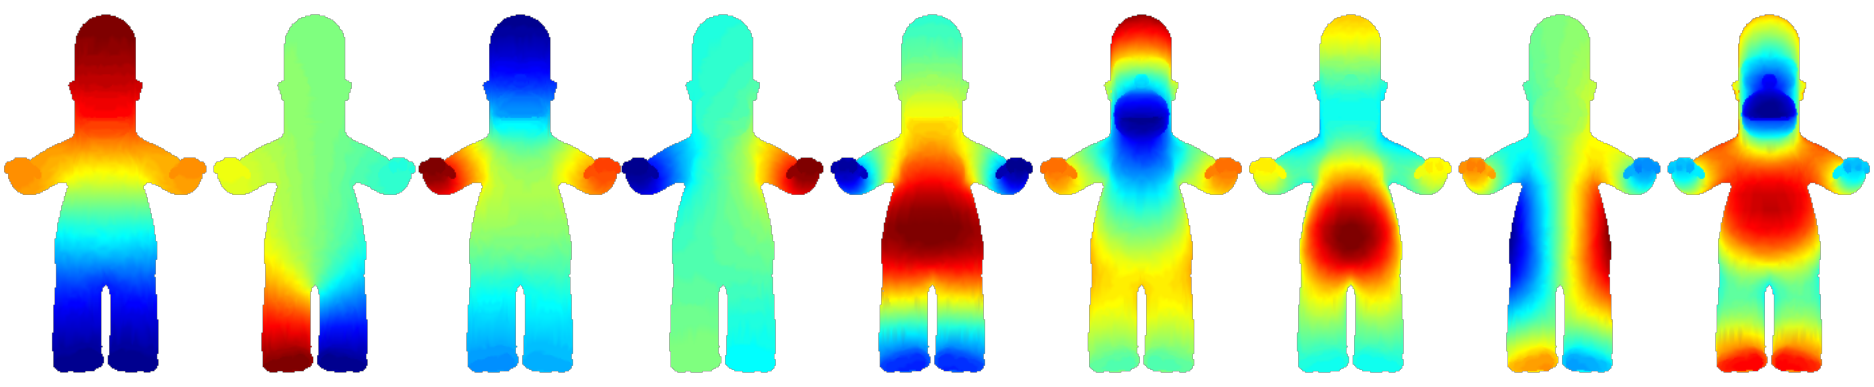

global support

More Exotic Bases Possible

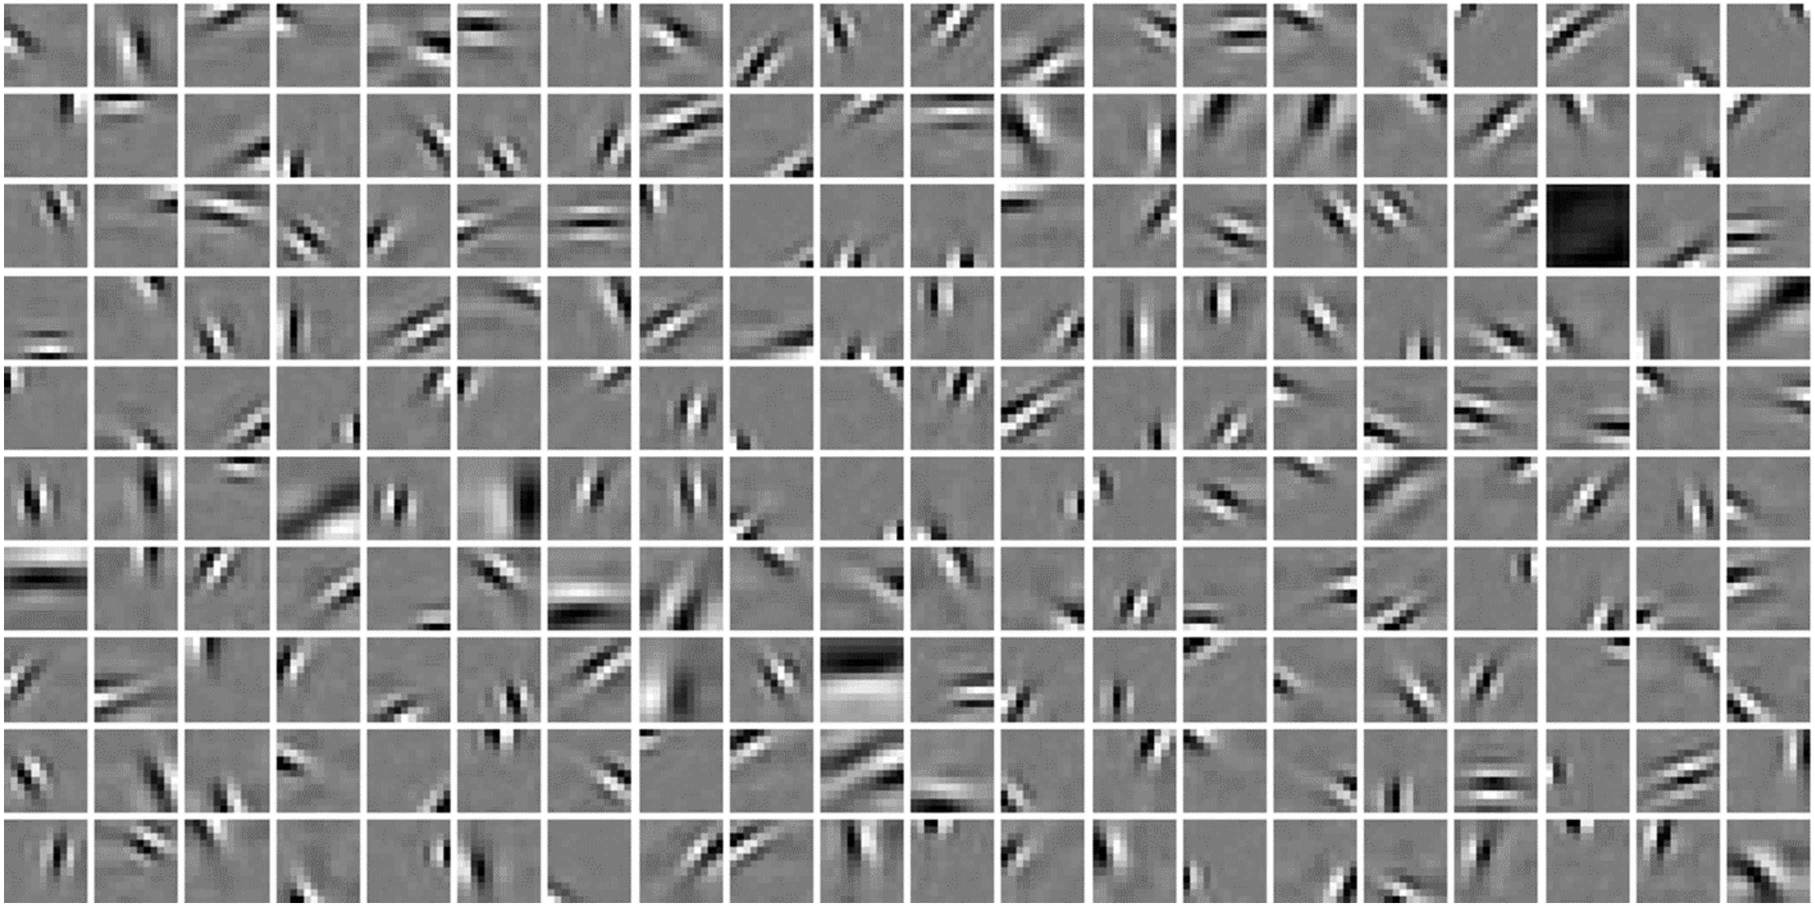

Textons, wavelets, …

Functional Map: Define Maps Across Objects by Relating Basis by a Linear Matrix (Operator)

$f(x) = a_1 \cdot$  $+ a_2 \cdot$  $+ a_3 \cdot$  $+ \cdots$

$\phi :$ 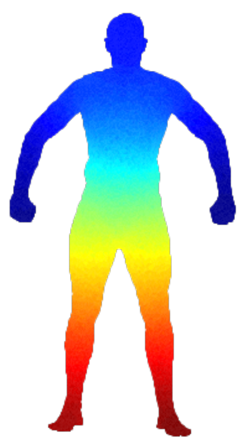 $\longrightarrow$ 

Functional Map: Define Maps Across Objects by Relating Basis by a Linear Matrix (Operator)

$T_\phi[f](x) = T_\phi[a_1 \cdot$  $+ a_2 \cdot$  $+ a_3 \cdot$  $+ \cdots]$

$= a_1 T_\phi[$  $] + a_2 T_\phi[$  $] + a_3 T_\phi[$  $] + \cdots$

Enough to know these

Functional Map Matrix

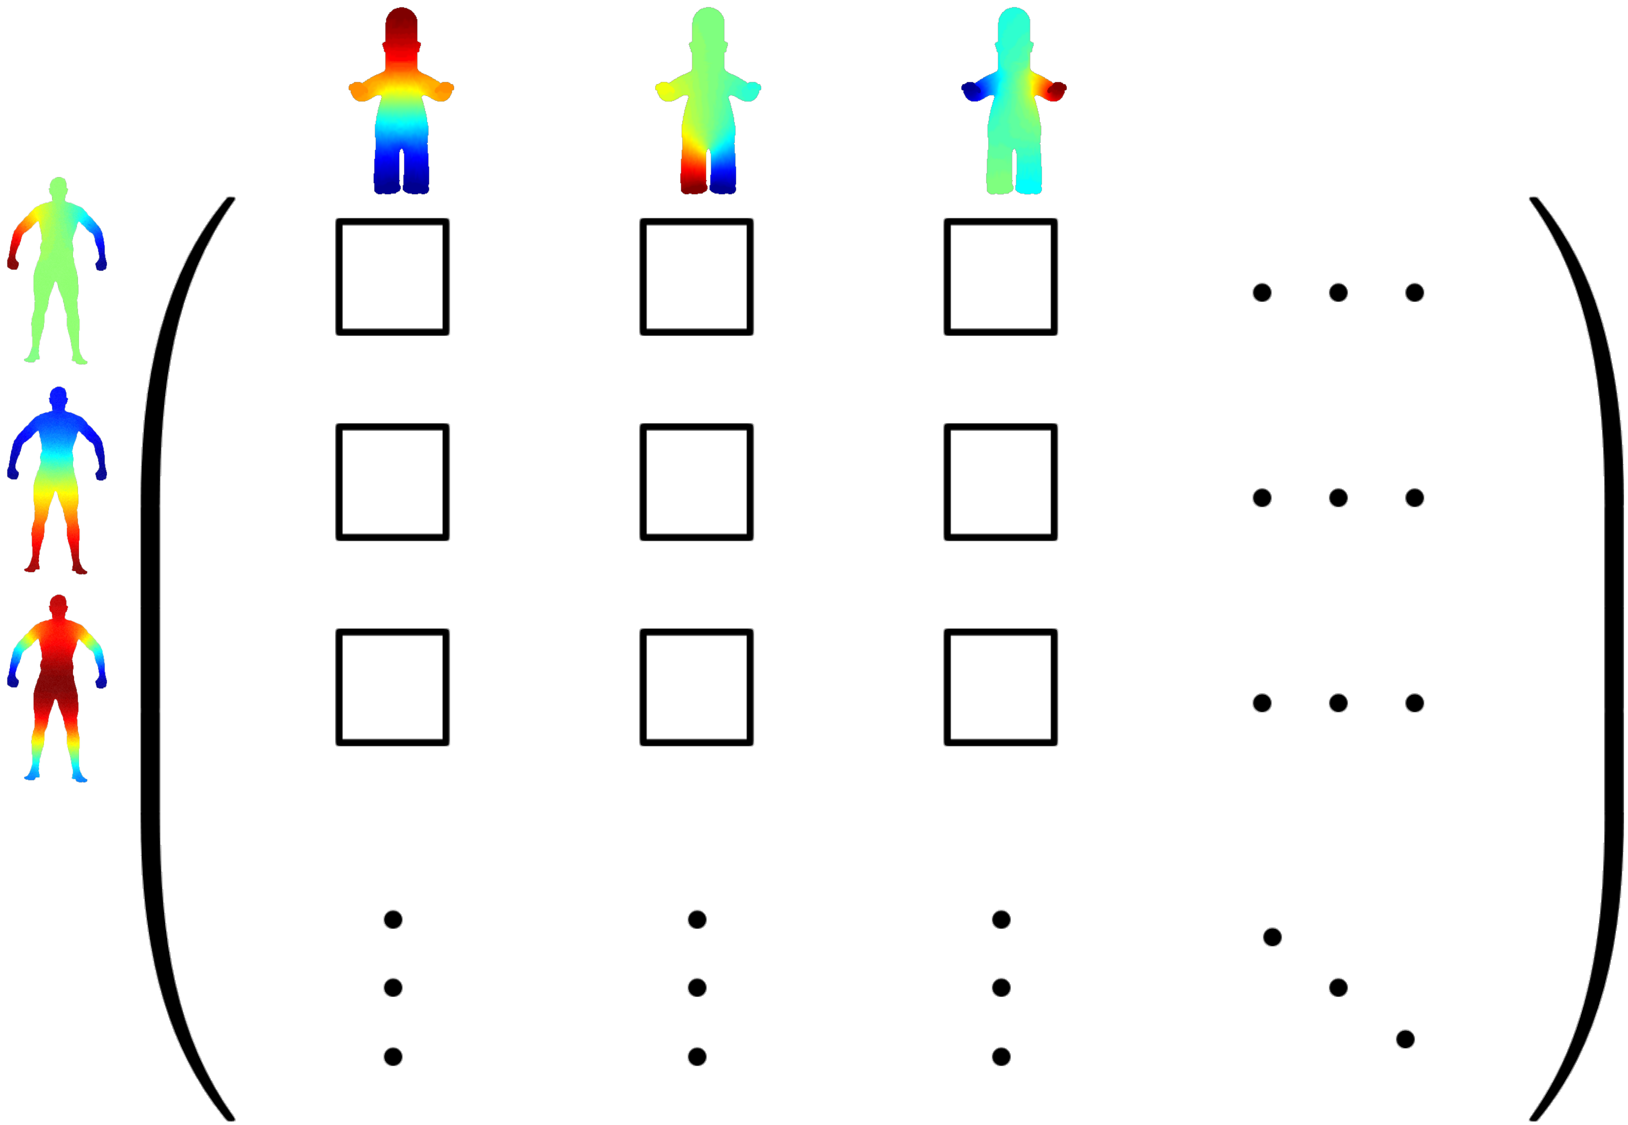

Maps as Linear Operators

- An ordinary shape map lifts to a linear operator mapping the function spaces
- With a truncated hierarchical basis, compact representations of functional maps are possible as ordinary matrices
- Map composition becomes ordinary matrix multiplication
- Functional maps can express many-to-many associations, generalizing classical 1-1 maps

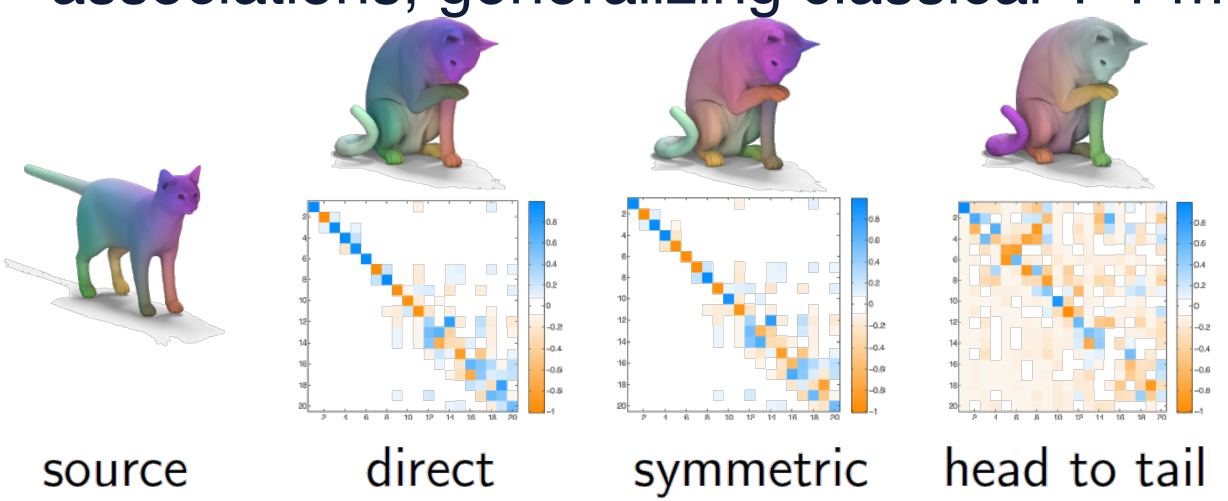

Using truncated Laplace-Beltrami basis

Estimating the Mapping Matrix

Suppose we don't know *C*. However, we expect a pair of functions $f : M \to \mathbb{R}$ and $g : N \to \mathbb{R}$ to correspond. Then, *C* must be s.t.

$$C\mathbf{a} \approx \mathbf{b}$$

where $f = \sum_i \mathbf{a_i} \phi_i^M, \quad g = \sum_i \mathbf{b}_i \phi_i^N$

Given enough $\{\mathbf{a_i}, \mathbf{b}_i\}$ pairs in correspondence, we can recover C through a linear least squares system.

SPECTRAL NEURAL NETWORKS

UC San Diego

Fourier analysis

A function $f : [-\pi, \pi] \to \mathbb{R}$ can be written as Fourier series

$$f(x) = \sum_{\omega} \underbrace{\frac{1}{2\pi} \int_{-\pi}^{\pi} f(x') e^{i\omega x'} dx'}_{\hat{f}(\omega) = \langle f, e^{-i\omega x} \rangle_{L^2([-\pi,\pi])}} \; e^{-i\omega x}$$

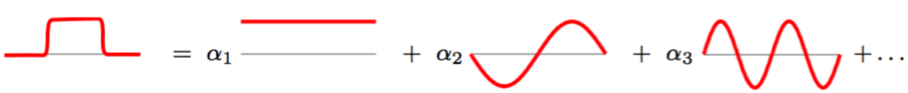

Fourier basis = Laplacian eigenfunctions: $\Delta e^{-i\omega x} = \omega^2 e^{-i\omega x}$

Euclidean domain

A function $f : \mathcal{X} \to \mathbb{R}$ can be written as Fourier series

$$f(x) = \sum_{k \geq 0} \underbrace{\int_{\mathcal{X}} f(x') \phi_k(x') dx'}_{\hat{f}_k = \langle f, \phi_k \rangle_{L^2(\mathcal{X})}} \; \phi_k(x)$$

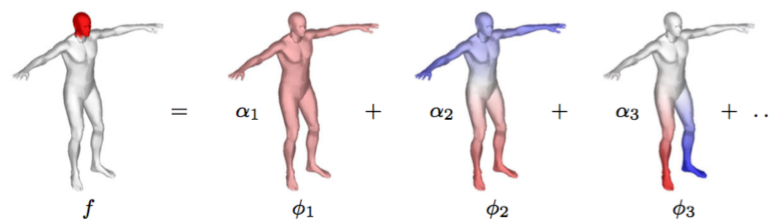

Fourier basis = Laplacian eigenfunctions: $\Delta \phi_k(x) = \lambda_k \phi_k(x)$

non Euclidean domain

from Jonathan Masci et al

Convolution Theorem in Euclidean domain

Given two functions $f, g : [-\pi, \pi] \to \mathbb{R}$ their convolution is a function

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(\xi) g(x - \xi) d\xi$$

Convolution Theorem: Fourier transform diagonalizes the convolution operator $\Rightarrow$ convolution can be computed in the Fourier domain as

$$f \star g = \mathcal{F}^{-1}(\mathcal{F}f \cdot \mathcal{F}g)$$

from Jonathan Masci et al

Convolution Theorem in Euclidean domain

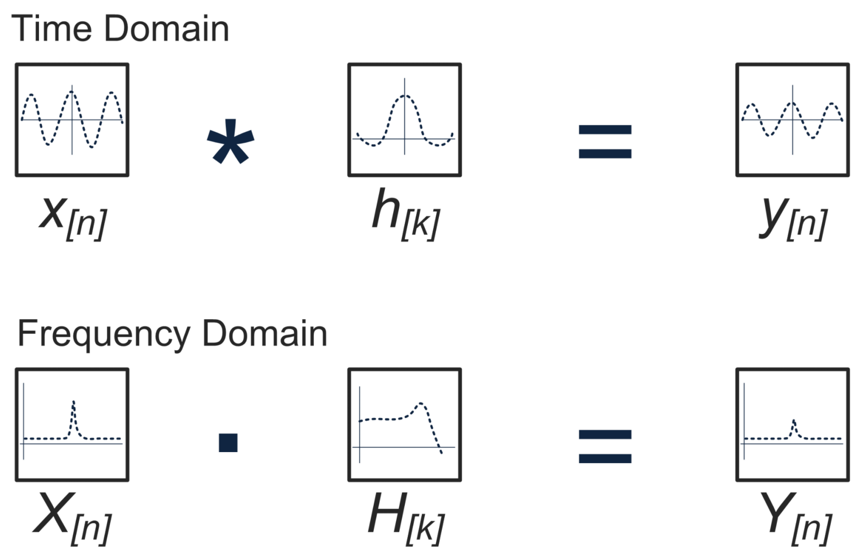

Convolution Theorem in non Euclidean domain

Generalized convolution of $f, g \in L^2(X)$ can be defined by analogy

$$(f \star g)(x) = \sum_{k \geq 1} \langle f, \phi_k \rangle_{L^2(X)} \langle g, \phi_k \rangle_{L^2(X)} \, \phi_k(x)$$

from Jonathan Masci et al

Convolution Theorem in non Euclidean domain

Generalized convolution of $f, g \in L^2(X)$ can be defined by analogy

$$(f \star g)(x) \quad = \quad \underbrace{\sum_{k \geq 1} \underbrace{\langle f, \phi_k \rangle_{L^2(X)} \langle g, \phi_k \rangle_{L^2(X)}}_{\text{product in the Fourier domain}} \phi_k(x)}_{\text{inverse Fourier transform}}$$

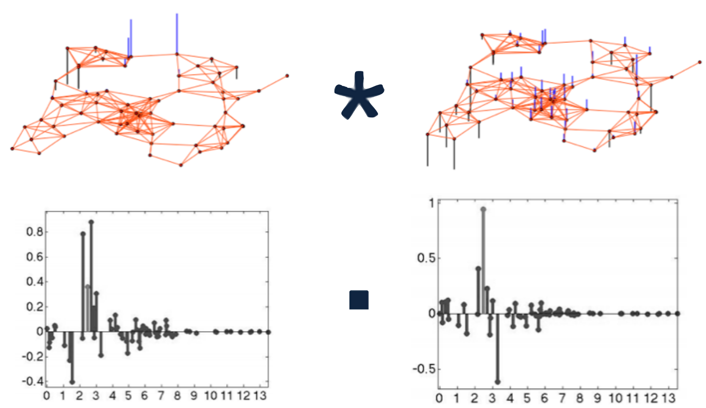

modified from Jonathan Masci et al

Convolution Theorem in non Euclidean domain

Generalized convolution of $f, g \in L^2(X)$ can be defined by analogy

$$(f \star g)(x) = \underbrace{\sum_{k \geq 1} \underbrace{\langle f, \phi_k \rangle_{L^2(X)} \langle g, \phi_k \rangle_{L^2(X)}}_{\text{product in the Fourier domain}} \phi_k(x)}_{\text{inverse Fourier transform}}$$

directly design convolution kernel in the spectral domain

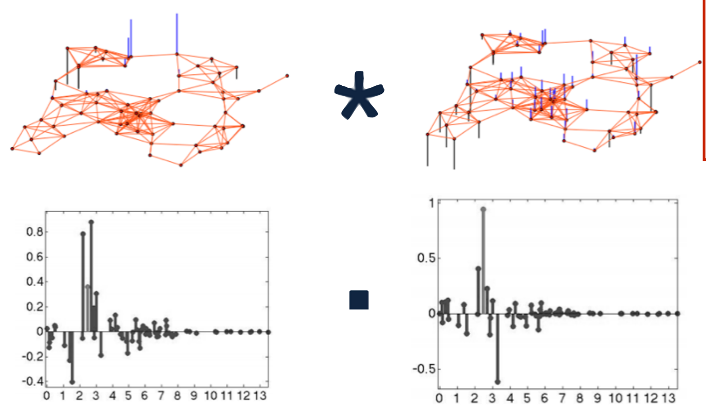

modified from Jonathan Masci et a

Spectral CNN

- Observation:

  In Fourier analysis, smoothness and sparsity are dual notions

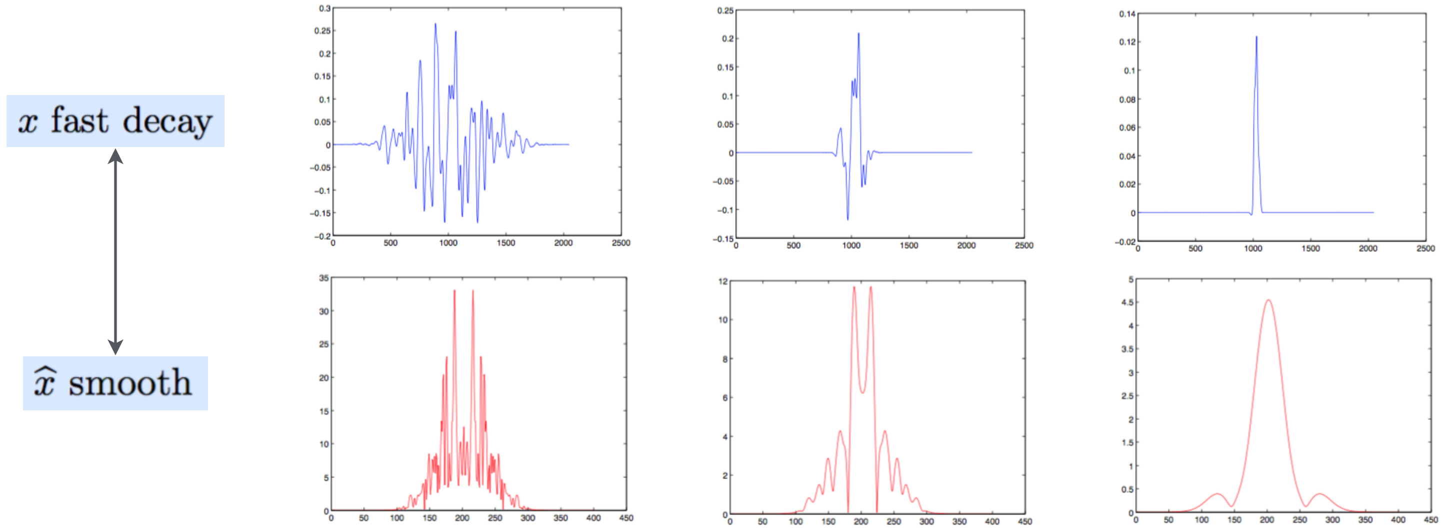

Joan Bruna et al. 2013

Spectral CNN

- Use smooth interpolation kernels (spline, polynomial, heat kernel, etc.) to parameterize the filters

spatially locally concentrated

Joan Bruna et al. 2013

Spectral CNN

- Use smooth interpolation kernels (spline, polynomial, heat kernel, etc.) to parameterize the filters

spatially locally concentrated

control #parameter

Joan Bruna et al. 2013

Spectral Dilated Convolution

- Parameterize filters with interpolation kernels.
- Shrink kernel bandwidth to increase spatial support of filters

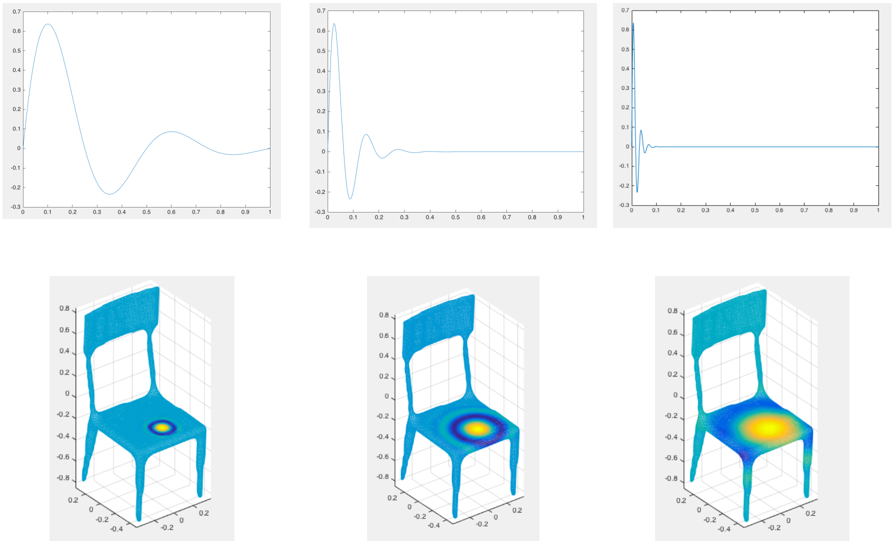

Li Yi et al. 2017

Cross Domain Discrepancy

- Parameterize filters with interpolation kernels.
- Shrink kernel bandwidth to increase spatial support of filters

Spectral Domain 1

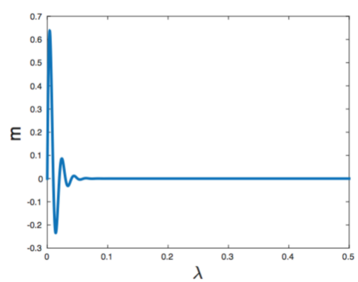

Spectral Domain 2

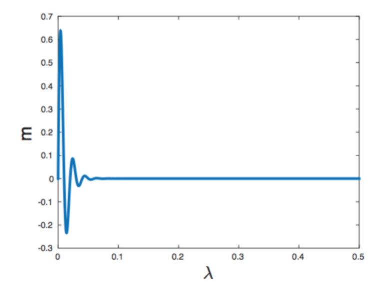

Spectral domain is independently defined for each shape graph

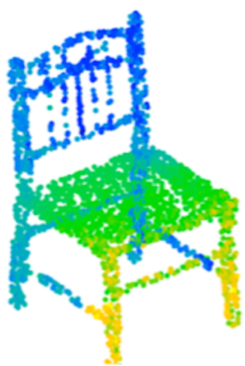

The same spectral function would induce very different spatial functions on different graphs

Cross domain parameter sharing is not valid

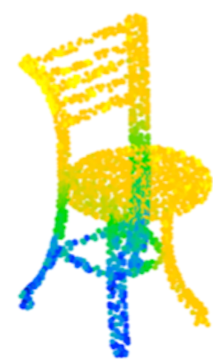

Li Yi et al. 2017

Functional Map for Domain Synchronization

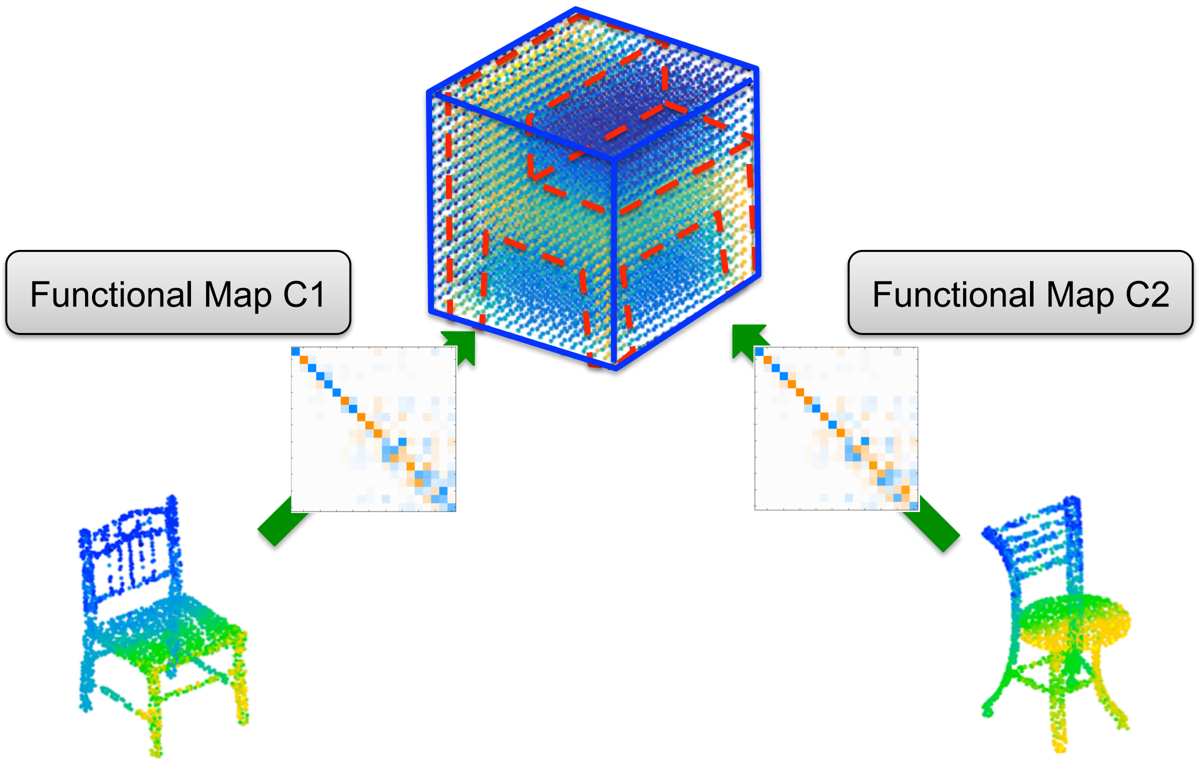

Li Yi et al. 2017

Synchronization Visualization

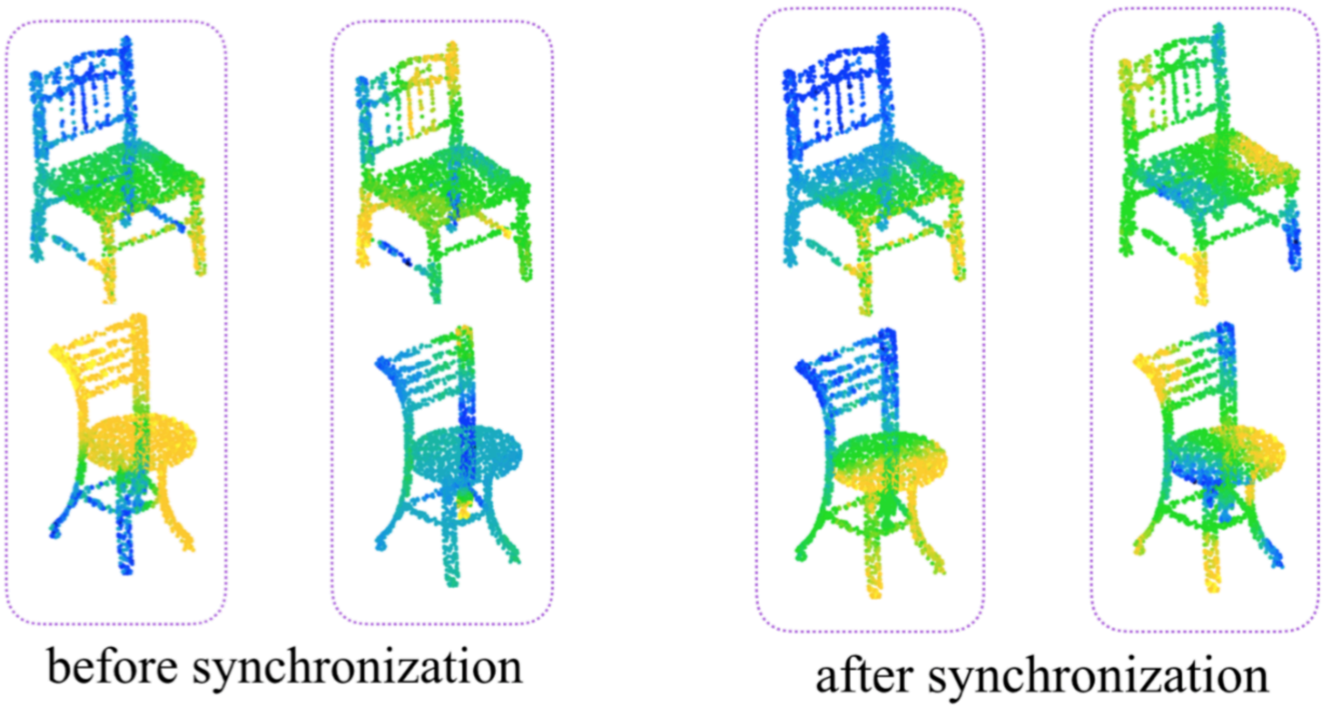

Li Yi et al. 2017

SyncSpecCNN

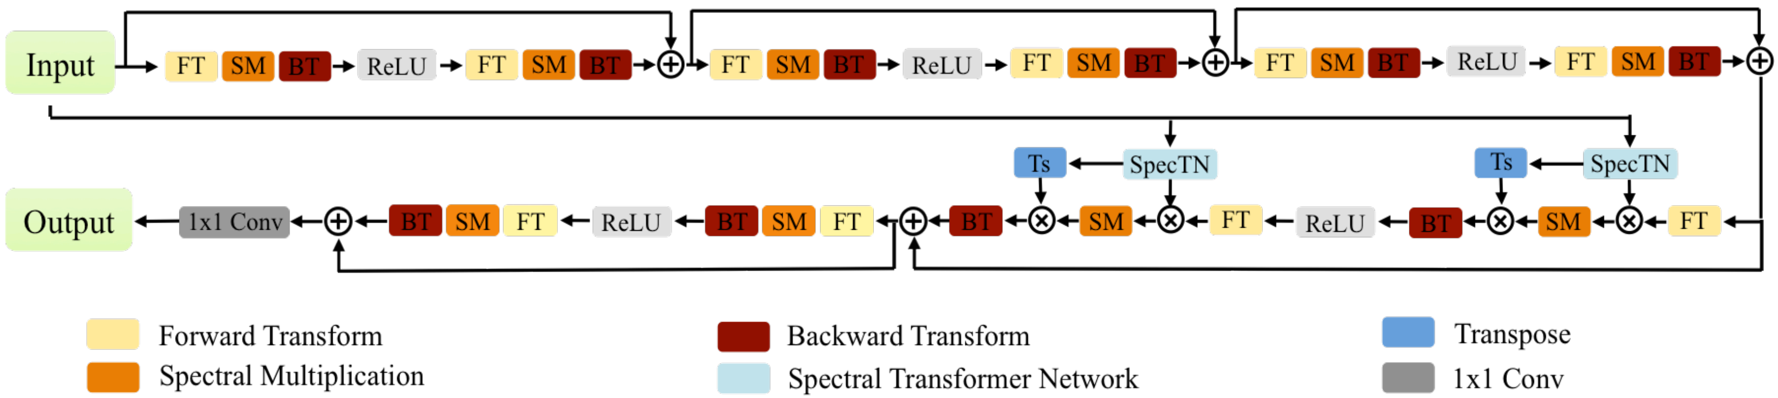

Li Yi et al. 2017

SyncSpecCNN

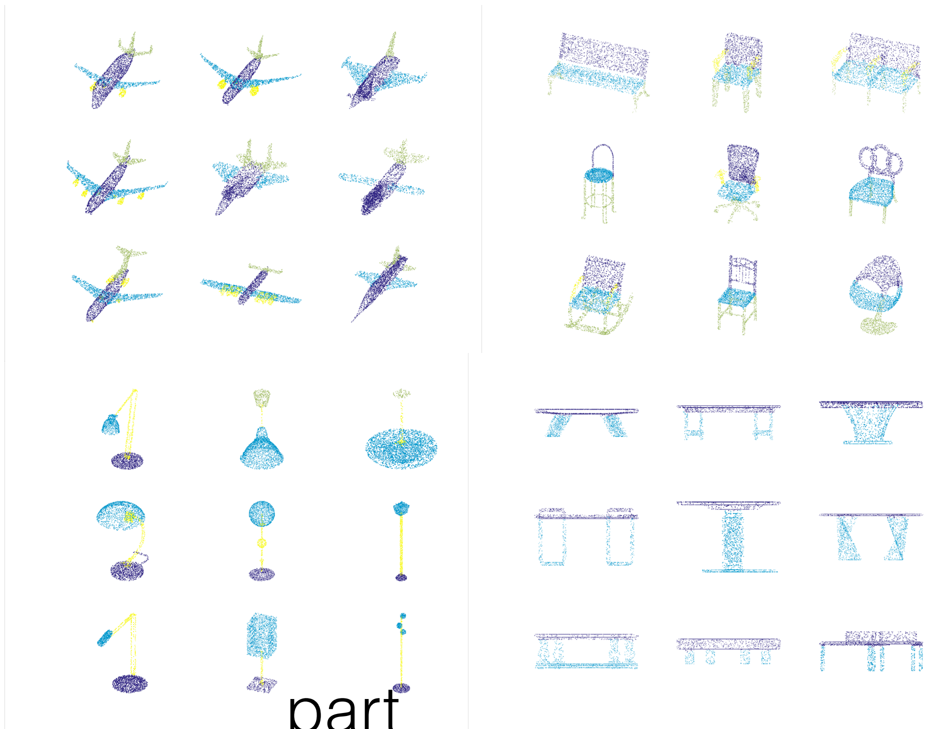

part
segmentation

key point
prediction

Li Yi et al. 2017

SPHERICAL CNN (A SPECTIAL CASE OF SPECTRAL CNN)

A Special Case: Spherical CNN

- If the surface is always a SPHERE, no worry about the functional space alignment anymore

- Generate a spherical representation

input

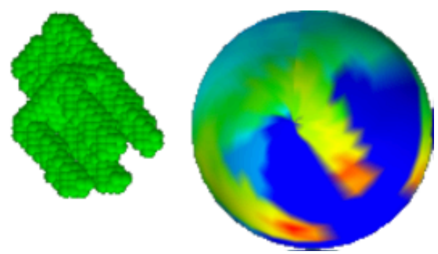

spherical representation

- Do Spectral CNN
    - Has numerical tricks exploiting the symmetry of sphere

Cohen et al., "**Spherical CNN**", *ICLR 2018*

Esteves et al., "**Learning SO(3) Equivariant Representations with Spherical CNNs**", *ECCV 2018*