



3D Domains (intrinsic)

Niloy Mitra

Iasonas Kokkinos

Federico Monti

Emanuele Rodolà

Michael Bronstein

Or Litany

Leonidas Guibas

UCL

UCL

USI Lugano

La Sapienza

Imperial College
USI Lugano

Stanford University
Facebook

Stanford University



http://geometry.cs.ucl.ac.uk/dl_for_CG/

Timetable

			Niloy	Federico	Iasonas	Emanuele
Theory/Basics	Introduction	9:00	X	X	X	X
	Machine Learning Basics	~ 9:05	X			
	Neural Network Basics	~ 9:35		X		
	Alternatives to Direct Supervision (GANs)	~11:00			X	
State of the Art	Image Domain	~11:45			X	
	3D Domains (extrinsic)	~13:30	X			
	3D Domains (intrinsic)	~ 14:15				X
	Physics and Animation	~ 16:00	X			
	Discussion	~ 16:45	X	X	X	X


Sessions: A. 9:00-10:30 (coffee) B. 11:00-12:30 [LUNCH] C. 13:30-15:00 (coffee) D. 15:30-17:00

Deep Learning on Manifolds


Shape representation




Image-based



Volumetric



Point-based



Surface-based


Extrinsic vs Intrinsic CNNs

Intrinsic


Different formulations of non-Euclidean CNNs



Spectral domain




Spatial domain



Parametric domain

Riemannian geometry in one minute

- Manifold \mathcal{X} = topological space
- No global Euclidean structure
- **Tangent plane** $T_x \mathcal{X}$ = local Euclidean representation of manifold \mathcal{X} around x




Riemannian geometry in one minute

- Manifold \mathcal{X} = topological space
- No global Euclidean structure
- **Tangent plane** $T_x \mathcal{X}$ = local Euclidean representation of manifold \mathcal{X} around x
- **Riemannian metric**

$$\langle \cdot, \cdot \rangle_{T_x \mathcal{X}} : T_x \mathcal{X} \times T_x \mathcal{X} \rightarrow \mathbb{R}$$

depending smoothly on x




Riemannian geometry in one minute

- Manifold \mathcal{X} = topological space
- No global Euclidean structure
- **Tangent plane** $T_x \mathcal{X}$ = local Euclidean representation of manifold \mathcal{X} around x
- **Riemannian metric**

$$\langle \cdot, \cdot \rangle_{T_x \mathcal{X}} : T_x \mathcal{X} \times T_x \mathcal{X} \rightarrow \mathbb{R}$$

depending smoothly on x

Isometry = metric-preserving shape deformation



Riemannian geometry in one minute


- Manifold \mathcal{X} = topological space
- No global Euclidean structure
- **Tangent plane** $T_x \mathcal{X}$ = local Euclidean representation of manifold \mathcal{X} around x
- **Riemannian metric**

$$\langle \cdot, \cdot \rangle_{T_x \mathcal{X}} : T_x \mathcal{X} \times T_x \mathcal{X} \rightarrow \mathbb{R}$$

depending smoothly on x

Isometry = metric-preserving shape deformation

Intrinsic = expressed solely in terms of the Riemannian metric



Riemannian geometry in one minute

- Manifold \mathcal{X} = topological space
- No global Euclidean structure
- **Tangent plane** $T_x \mathcal{X}$ = local Euclidean representation of manifold \mathcal{X} around x
- **Riemannian metric**


$$\langle \cdot, \cdot \rangle_{T_x \mathcal{X}} : T_x \mathcal{X} \times T_x \mathcal{X} \rightarrow \mathbb{R}$$

depending smoothly on x

Isometry = metric-preserving shape deformation


Intrinsic = expressed solely in terms of the Riemannian metric

- **Geodesic** = shortest path on \mathcal{X} between x and x'



Calculus on manifolds


- Scalar field $f : \mathcal{X} \rightarrow \mathbb{R}$



Calculus on manifolds

- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$




Calculus on manifolds

- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

- **Laplacian operator** $\Delta f = -\text{div}(\nabla f)$
“difference between $f(x)$ and average value of f around x ”



Calculus on manifolds


- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

- **Laplacian** operator $\Delta f = -\operatorname{div}(\nabla f)$

“difference between $f(x)$ and average value of f around x ”

- **Intrinsic** (expressed solely in terms of the Riemannian metric)



Calculus on manifolds


- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

- **Laplacian** operator $\Delta f = -\operatorname{div}(\nabla f)$

“difference between $f(x)$ and average value of f around x ”

- **Intrinsic** (expressed solely in terms of the Riemannian metric)
- **Isometry-invariant**



Calculus on manifolds


- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

- **Laplacian** operator $\Delta f = -\operatorname{div}(\nabla f)$

“difference between $f(x)$ and average value of f around x ”

- **Intrinsic** (expressed solely in terms of the Riemannian metric)
- **Isometry-invariant**
- **Self-adjoint** $\langle \Delta f, g \rangle_{L^2(\mathcal{X})} = \langle f, \Delta g \rangle_{L^2(\mathcal{X})}$



Calculus on manifolds


- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

- **Laplacian** operator $\Delta f = -\operatorname{div}(\nabla f)$

“difference between $f(x)$ and average value of f around x ”

- **Intrinsic** (expressed solely in terms of the Riemannian metric)
- **Isometry-invariant**
- **Self-adjoint** $\langle \Delta f, g \rangle_{L^2(\mathcal{X})} = \langle f, \Delta g \rangle_{L^2(\mathcal{X})} \Rightarrow$ orthogonal eigenfunctions



Calculus on manifolds


- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

- **Laplacian** operator $\Delta f = -\text{div}(\nabla f)$

“difference between $f(x)$ and average value of f around x ”

- **Intrinsic** (expressed solely in terms of the Riemannian metric)
- **Isometry-invariant**
- **Self-adjoint** $\langle \Delta f, g \rangle_{L^2(\mathcal{X})} = \langle f, \Delta g \rangle_{L^2(\mathcal{X})} \Rightarrow$ orthogonal eigenfunctions
- **Positive semidefinite**



Calculus on manifolds


- **Scalar field** $f : \mathcal{X} \rightarrow \mathbb{R}$
- **Hilbert space** $L^2(\mathcal{X})$ with inner product

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$


- **Laplacian** operator $\Delta f = -\text{div}(\nabla f)$

“difference between $f(x)$ and average value of f around x ”

- **Intrinsic** (expressed solely in terms of the Riemannian metric)
- **Isometry-invariant**
- **Self-adjoint** $\langle \Delta f, g \rangle_{L^2(\mathcal{X})} = \langle f, \Delta g \rangle_{L^2(\mathcal{X})} \Rightarrow$ orthogonal eigenfunctions
- **Positive semidefinite** \Rightarrow non-negative eigenvalues




Discrete Laplacian



Undirected graph $(\mathcal{V}, \mathcal{E})$

$$(\Delta f)_i \approx \sum_{(i,j) \in \mathcal{E}} w_{ij} (f_i - f_j)$$




Triangular mesh $(\mathcal{V}, \mathcal{E}, \mathcal{F})$

$$(\Delta f)_i \approx \frac{1}{a_i} \sum_{(i,j) \in \mathcal{E}} \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{2} (f_i - f_j)$$


$a_i =$ local area element

Discrete Laplacian



Undirected graph $(\mathcal{V}, \mathcal{E})$

$$(\Delta f)_i \approx \sum_{(i,j) \in \mathcal{E}} w_{ij} (f_i - f_j)$$



Triangular mesh $(\mathcal{V}, \mathcal{E}, \mathcal{F})$

$$(\Delta f)_i \approx \frac{1}{a_i} \sum_{(i,j) \in \mathcal{E}} \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{2} (f_i - f_j)$$

$a_i =$ local area element

In matrix-vector notation

$$\Delta \mathbf{f} = \mathbf{A}^{-1}(\mathbf{D} - \mathbf{W})\mathbf{f}$$

where $\mathbf{f} = (f_1, \dots, f_n)^\top$, \mathbf{W} is the **stiffness matrix**, $\mathbf{A} = \text{diag}(a_1, \dots, a_n)$ is the **mass matrix**, and $\mathbf{D} = \text{diag}(\sum_{j \neq 1} w_{1j}, \dots, \sum_{j \neq n} w_{nj})$

Laplacian eigenfunctions and eigenvalues

$$\Delta \Phi = \Phi \Lambda$$

- $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\Phi = (\phi_1, \dots, \phi_n)$ a matrix of eigenvectors

Laplacian eigenfunctions and eigenvalues

$$\mathbf{A}^{-1}(\mathbf{D} - \mathbf{W})\Phi = \Phi\Lambda$$

- $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\Phi = (\phi_1, \dots, \phi_n)$ a matrix of eigenvectors

Laplacian eigenfunctions and eigenvalues

$$(\mathbf{D} - \mathbf{W})\Phi = \mathbf{A}\Phi\Lambda$$

- $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\Phi = (\phi_1, \dots, \phi_n)$ an \mathbf{A} -orthonormal matrix of eigenvectors ($\Phi^\top \mathbf{A}\Phi = \mathbf{I}$)

Laplacian eigenfunctions and eigenvalues

$$\mathbf{A}^{-1/2}(\mathbf{D} - \mathbf{W})\mathbf{A}^{-1/2}\mathbf{A}^{1/2}\Phi = \mathbf{A}^{1/2}\Phi\Lambda$$

- $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\Phi = (\phi_1, \dots, \phi_n)$ an \mathbf{A} -orthonormal matrix of eigenvectors ($\Phi^\top \mathbf{A} \Phi = \mathbf{I}$)

Laplacian eigenfunctions and eigenvalues


$$\mathbf{A}^{-1/2}(\mathbf{D} - \mathbf{W})\mathbf{A}^{-1/2}\mathbf{\Psi} = \mathbf{\Psi}\mathbf{\Lambda}$$

- $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\mathbf{\Psi} = (\psi_1, \dots, \psi_n)$ an orthonormal matrix of eigenvectors ($\mathbf{\Psi}^\top \mathbf{\Psi} = \mathbf{I}$)

Laplacian eigenfunctions and eigenvalues

$$\mathbf{A}^{-1/2}(\mathbf{D} - \mathbf{W})\mathbf{A}^{-1/2}\Psi = \Psi\Lambda$$

- $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\Psi = (\psi_1, \dots, \psi_n)$ an orthonormal matrix of eigenvectors ($\Psi^T \Psi = \mathbf{I}$)




First eigenfunctions of 1D Euclidean Laplacian = standard Fourier basis

Laplacian eigenfunctions and eigenvalues

$$\mathbf{A}^{-1/2}(\mathbf{D} - \mathbf{W})\mathbf{A}^{-1/2}\Psi = \Psi\Lambda$$

- $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ diagonal matrix of non-negative eigenvalues
- $\Psi = (\psi_1, \dots, \psi_n)$ an orthonormal matrix of eigenvectors ($\Psi^T \Psi = \mathbf{I}$)




First eigenfunctions of a manifold Laplacian

Fourier analysis: Euclidean

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


$$f(x) = \sum_{k \geq 0} \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x') e^{-ikx'} dx' e^{ikx}$$



Fourier analysis: Euclidean

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


$$f(x) = \sum_{k \geq 0} \frac{1}{2\pi} \underbrace{\int_{-\pi}^{\pi} f(x') e^{-ikx'} dx'}_{\hat{f}_k = \langle f, e^{ikx} \rangle_{L^2([-\pi, \pi])}} e^{ikx}$$



Fourier analysis: Euclidean

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

$$f(x) = \sum_{k \geq 0} \frac{1}{2\pi} \underbrace{\int_{-\pi}^{\pi} f(x') e^{-ikx'} dx'}_{\hat{f}_k = \langle f, e^{ikx} \rangle_{L^2([-\pi, \pi])}} e^{ikx}$$




Fourier basis = **Laplacian eigenfunctions**: $-\frac{d^2}{dx^2} e^{ikx} = k^2 e^{ikx}$

Fourier analysis: non-Euclidean

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

$$f(x) = \sum_{k \geq 1} \underbrace{\int_{\mathcal{X}} f(x') \phi_k(x') dx'}_{\hat{f}_k = \langle f, \phi_k \rangle_{L^2(\mathcal{X})}} \phi_k(x)$$



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

Convolution: Euclidean space

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

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x - x')dx'$$

Convolution: Euclidean space

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

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x - x')dx'$$

- **Shift-invariance:** $f(x - x_0) \star g(x) = (f \star g)(x - x_0)$

Convolution: Euclidean space

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

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x - x')dx'$$

- **Shift-invariance:** $f(x - x_0) \star g(x) = (f \star g)(x - x_0)$
- Convolution operator **commutes with Laplacian:** $(\Delta f) \star g = \Delta(f \star g)$

Convolution: Euclidean space

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

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x - x')dx'$$

- **Shift-invariance:** $f(x - x_0) \star g(x) = (f \star g)(x - x_0)$
- Convolution operator **commutes with Laplacian:** $(\Delta f) \star g = \Delta(f \star g)$
- **Convolution theorem:** Fourier transform diagonalizes the convolution operator

Convolution: Euclidean space

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

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x - x')dx'$$

- **Shift-invariance:** $f(x - x_0) \star g(x) = (f \star g)(x - x_0)$
- Convolution operator **commutes with Laplacian:** $(\Delta f) \star g = \Delta(f \star g)$
- **Convolution theorem:** Fourier transform diagonalizes the convolution operator \Rightarrow convolution can be computed in the Fourier domain as

$$\widehat{(f \star g)} = \hat{f} \cdot \hat{g}$$

Convolution: Euclidean space

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

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x - x')dx'$$

- **Shift-invariance:** $f(x - x_0) \star g(x) = (f \star g)(x - x_0)$
- Convolution operator **commutes with Laplacian:** $(\Delta f) \star g = \Delta(f \star g)$
- **Convolution theorem:** Fourier transform diagonalizes the convolution operator \Rightarrow convolution can be computed in the Fourier domain as

$$\widehat{(f \star g)} = \hat{f} \cdot \hat{g}$$

- **Efficient computation** using FFT

Convolution Theorem

Convolution of two vectors $\mathbf{f} = (f_1, \dots, f_n)^\top$ and $\mathbf{g} = (g_1, \dots, g_n)^\top$

$$\mathbf{f} \star \mathbf{g} = \begin{bmatrix} g_1 & g_2 & \dots & \dots & g_n \\ g_n & g_1 & g_2 & \dots & g_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ g_3 & g_4 & \dots & g_1 & g_2 \\ g_2 & g_3 & \dots & \dots & g_1 \end{bmatrix} \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix}$$

Convolution Theorem

Convolution of two vectors $\mathbf{f} = (f_1, \dots, f_n)^\top$ and $\mathbf{g} = (g_1, \dots, g_n)^\top$

$$\mathbf{f} \star \mathbf{g} = \underbrace{\begin{bmatrix} g_1 & g_2 & \dots & \dots & g_n \\ g_n & g_1 & g_2 & \dots & g_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ g_3 & g_4 & \dots & g_1 & g_2 \\ g_2 & g_3 & \dots & \dots & g_1 \end{bmatrix}}_{\text{circulant matrix}} \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix}$$

Convolution Theorem

Convolution of two vectors $\mathbf{f} = (f_1, \dots, f_n)^\top$ and $\mathbf{g} = (g_1, \dots, g_n)^\top$

$$\mathbf{f} \star \mathbf{g} = \underbrace{\begin{bmatrix} g_1 & g_2 & \dots & \dots & g_n \\ g_n & g_1 & g_2 & \dots & g_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ g_3 & g_4 & \dots & g_1 & g_2 \\ g_2 & g_3 & \dots & \dots & g_1 \end{bmatrix}}_{\text{diagonalized by Fourier basis}} \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix}$$

Convolution Theorem

Convolution of two vectors $\mathbf{f} = (f_1, \dots, f_n)^\top$ and $\mathbf{g} = (g_1, \dots, g_n)^\top$

$$\begin{aligned}\mathbf{f} \star \mathbf{g} &= \begin{bmatrix} g_1 & g_2 & \dots & \dots & g_n \\ g_n & g_1 & g_2 & \dots & g_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ g_3 & g_4 & \dots & g_1 & g_2 \\ g_2 & g_3 & \dots & \dots & g_1 \end{bmatrix} \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix} \\ &= \mathbf{\Phi} \begin{bmatrix} \hat{g}_1 & & & & \\ & \ddots & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \hat{g}_n \end{bmatrix} \mathbf{\Phi}^\top \mathbf{f}\end{aligned}$$

Convolution Theorem

Convolution of two vectors $\mathbf{f} = (f_1, \dots, f_n)^\top$ and $\mathbf{g} = (g_1, \dots, g_n)^\top$

$$\begin{aligned}\mathbf{f} \star \mathbf{g} &= \begin{bmatrix} g_1 & g_2 & \dots & \dots & g_n \\ g_n & g_1 & g_2 & \dots & g_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ g_3 & g_4 & \dots & g_1 & g_2 \\ g_2 & g_3 & \dots & \dots & g_1 \end{bmatrix} \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix} \\ &= \Phi \begin{bmatrix} \hat{g}_1 & & & \\ & \ddots & & \\ & & \hat{g}_n & \end{bmatrix} \begin{bmatrix} \hat{f}_1 \\ \vdots \\ \hat{f}_n \end{bmatrix}\end{aligned}$$

Convolution Theorem

Convolution of two vectors $\mathbf{f} = (f_1, \dots, f_n)^\top$ and $\mathbf{g} = (g_1, \dots, g_n)^\top$

$$\begin{aligned}\mathbf{f} \star \mathbf{g} &= \begin{bmatrix} g_1 & g_2 & \dots & \dots & g_n \\ g_n & g_1 & g_2 & \dots & g_{n-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ g_3 & g_4 & \dots & g_1 & g_2 \\ g_2 & g_3 & \dots & \dots & g_1 \end{bmatrix} \begin{bmatrix} f_1 \\ \vdots \\ f_n \end{bmatrix} \\ &= \Phi \begin{bmatrix} \hat{f}_1 \cdot \hat{g}_1 \\ \vdots \\ \hat{f}_n \cdot \hat{g}_n \end{bmatrix}\end{aligned}$$

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

In matrix-vector notation

$$\mathbf{f} \star \mathbf{g} = \Phi (\Phi^\top \mathbf{g}) \circ (\Phi^\top \mathbf{f})$$

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

In matrix-vector notation

$$\mathbf{f} \star \mathbf{g} = \mathbf{\Phi} \text{diag}(\hat{g}_1, \dots, \hat{g}_n) \mathbf{\Phi}^\top \mathbf{f}$$

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

In matrix-vector notation

$$\mathbf{f} \star \mathbf{g} = \underbrace{\Phi \text{diag}(\hat{g}_1, \dots, \hat{g}_n) \Phi^\top}_{\mathbf{G}} \mathbf{f}$$

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

In matrix-vector notation

$$\mathbf{f} \star \mathbf{g} = \underbrace{\Phi \text{diag}(\hat{g}_1, \dots, \hat{g}_n) \Phi^\top}_{\mathbf{G}} \mathbf{f}$$

- Not shift-invariant! (\mathbf{G} has no circulant structure)

Spectral convolution

Spectral convolution of $f, g \in L^2(\mathcal{X})$ can be defined by analogy

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

In matrix-vector notation

$$\mathbf{f} \star \mathbf{g} = \underbrace{\Phi \text{diag}(\hat{g}_1, \dots, \hat{g}_n) \Phi^\top}_{\mathbf{G}} \mathbf{f}$$

- Not shift-invariant! (\mathbf{G} has no circulant structure)
- Filter coefficients depend on basis ϕ_1, \dots, ϕ_n

Spectral CNN

Convolution expressed in the **spectral domain**

$$\mathbf{g} = \Phi \mathbf{W} \Phi^T \mathbf{f}$$

where \mathbf{W} is $n \times n$ diagonal matrix of learnable spectral filter coefficients

Spectral CNN

Convolution expressed in the **spectral domain**

$$\mathbf{g} = \Phi \mathbf{W} \Phi^T \mathbf{f}$$

where \mathbf{W} is $n \times n$ diagonal matrix of learnable spectral filter coefficients

☹ Filters are basis-dependent \Rightarrow do not generalize across domains

Spectral CNN

Convolution expressed in the **spectral domain**

$$\mathbf{g} = \Phi \mathbf{W} \Phi^T \mathbf{f}$$

where \mathbf{W} is $n \times n$ diagonal matrix of learnable spectral filter coefficients

- ☹ Filters are basis-dependent \Rightarrow do not generalize across domains
- ☹ $\mathcal{O}(n)$ parameters per layer

Spectral CNN

Convolution expressed in the **spectral domain**

$$\mathbf{g} = \Phi \mathbf{W} \Phi^T \mathbf{f}$$

where \mathbf{W} is $n \times n$ diagonal matrix of learnable spectral filter coefficients

- ☹ Filters are basis-dependent \Rightarrow do not generalize across domains
- ☹ $\mathcal{O}(n)$ parameters per layer
- ☹ $\mathcal{O}(n^2)$ computation of forward and inverse Fourier transforms Φ^T, Φ (no FFT on graphs)

Spectral CNN


Convolution expressed in the **spectral domain**

$$\mathbf{g} = \Phi \mathbf{W} \Phi^\top \mathbf{f}$$

where \mathbf{W} is $n \times n$ diagonal matrix of learnable spectral filter coefficients

- ☹ Filters are basis-dependent \Rightarrow do not generalize across domains
- ☹ $\mathcal{O}(n)$ parameters per layer
- ☹ $\mathcal{O}(n^2)$ computation of forward and inverse Fourier transforms Φ^\top, Φ (no FFT on graphs)
- ☹ No guarantee of spatial localization of filters

Basis dependence



Function f

Basis dependence



'Edge detecting' spectral filter $\Phi W \Phi^T \mathbf{f}$

Basis dependence

Basis dependence

Laplacian eigenbases on non-isometric domains



ϕ_2



ϕ_3



ϕ_{10}



ϕ_{15}



ϕ_{20}



ψ_2




ψ_3



ψ_{10}





ψ_{15}




ψ_{20}


Functional maps



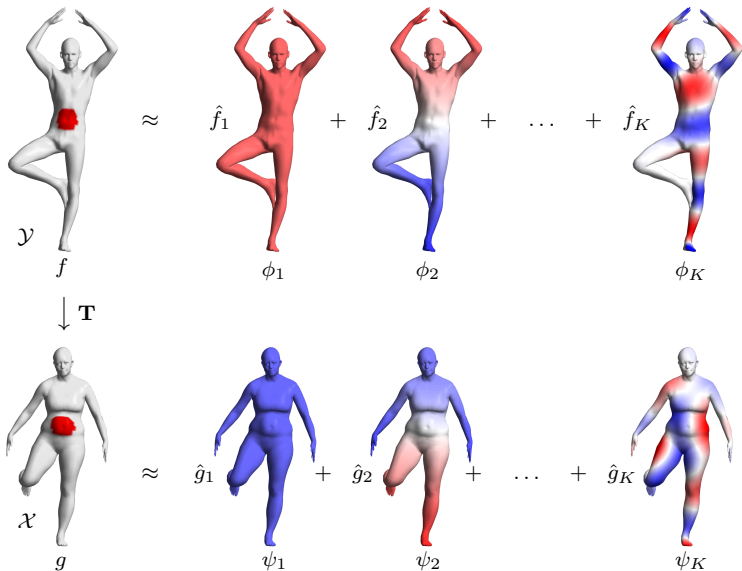
Functional maps



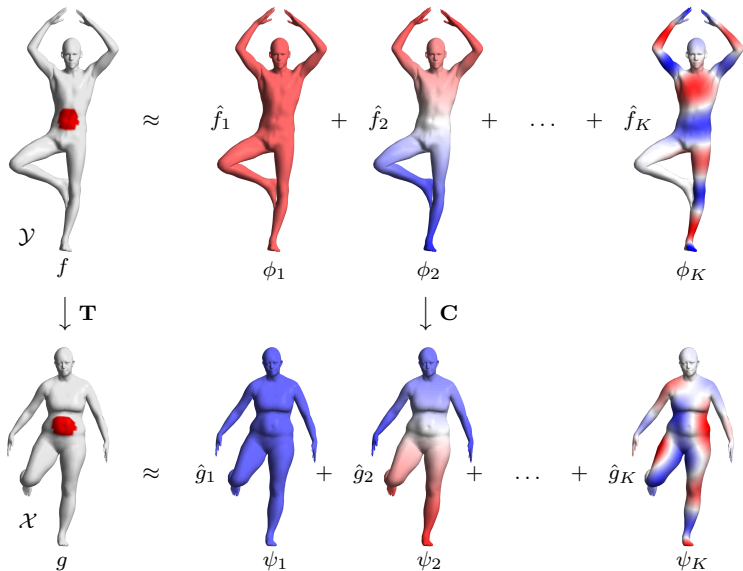
$\downarrow T : L^2(\mathcal{X}) \rightarrow L^2(\mathcal{Y})$



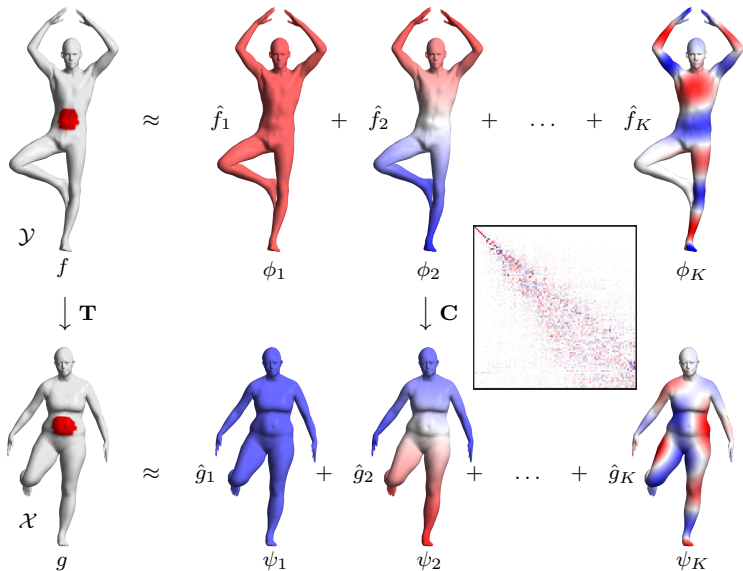
Functional maps




Functional maps




Functional maps




Basis synchronization with functional maps




$$\downarrow \mathbf{T} \approx \Psi \mathbf{C}^\top \Phi^\top$$




Basis synchronization with functional maps



$$\downarrow \mathbf{T} \approx \Psi(\Phi\mathbf{C})^\top$$





Basis synchronization with functional maps



Filtering in different bases



$$\Phi \tau(\Lambda_\Phi) \Phi^T \delta_0$$



$$\Psi \tau(\Lambda_\Psi) \Psi^T \delta_0$$

Apply spectral filter $\tau(\lambda)$ in **different bases** Φ and Ψ
 \Rightarrow **different results!**


Filtering in different bases



$$\Phi \tau(\Lambda_\Phi) \Phi^T \delta_0$$



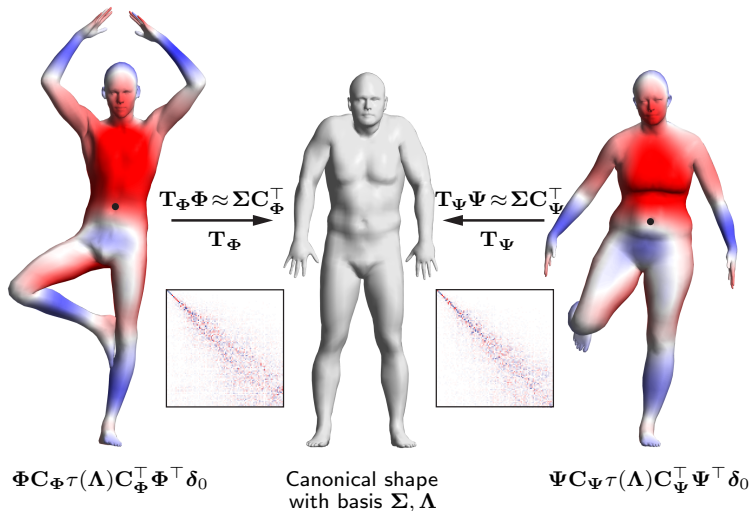
Canonical shape
with basis Σ, Λ



$$\Psi \tau(\Lambda_\Psi) \Psi^T \delta_0$$


Apply spectral filter $\tau(\lambda)$ in **different bases** Φ and Ψ
 \Rightarrow **different results!**

Filtering in synchronized bases



Apply spectral filter $\tau(\lambda)$ in **synchronized bases** ΦC_Φ and ΨC_Ψ
 \Rightarrow **similar results!**

Spectral CNN

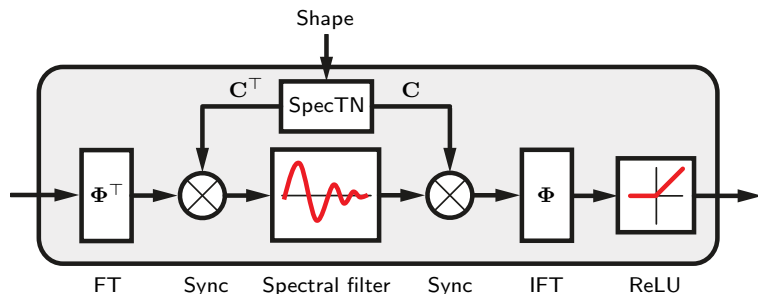


Convolutional filter of a Spectral CNN

☹ Fixed basis \Rightarrow Does not generalize across domains

☺ Possible $\mathcal{O}(n)$ complexity avoiding explicit FT and IFT

Spectral Transformer Network



Convolutional filter of a Spectral Transformer Network


😊 Basis synchronization allows generalization across domains

☹️ Explicit FT and IFT

Example: normal prediction with SpecTN




Predicted



Groundtruth

Example: shape segmentation with SpecTN



Localization and Smoothness

Vanishing moments: In the Euclidean setting

$$\int_{-\infty}^{+\infty} |x|^{2k} |f(x)|^2 dx = \int_{-\infty}^{+\infty} \left| \frac{\partial^k \hat{f}(\omega)}{\partial \omega^k} \right|^2 d\omega$$

Localization and Smoothness

Vanishing moments: In the Euclidean setting

$$\int_{-\infty}^{+\infty} |x|^{2k} |f(x)|^2 dx = \int_{-\infty}^{+\infty} \left| \frac{\partial^k \hat{f}(\omega)}{\partial \omega^k} \right|^2 d\omega$$

Localization in space = smoothness in frequency domain

Localization and Smoothness

Vanishing moments: In the Euclidean setting

$$\int_{-\infty}^{+\infty} |x|^{2k} |f(x)|^2 dx = \int_{-\infty}^{+\infty} \left| \frac{\partial^k \hat{f}(\omega)}{\partial \omega^k} \right|^2 d\omega$$

Localization in space = smoothness in frequency domain

Parametrize the filter using a **smooth spectral transfer function** $\tau(\lambda)$

Localization and Smoothness

Vanishing moments: In the Euclidean setting

$$\int_{-\infty}^{+\infty} |x|^{2k} |f(x)|^2 dx = \int_{-\infty}^{+\infty} \left| \frac{\partial^k \hat{f}(\omega)}{\partial \omega^k} \right|^2 d\omega$$

Localization in space = smoothness in frequency domain

Parametrize the filter using a **smooth spectral transfer function** $\tau(\lambda)$

Application of the filter

$$\tau(\mathbf{\Delta})\mathbf{f} = \mathbf{\Phi}\tau(\mathbf{\Lambda})\mathbf{\Phi}^T\mathbf{f}$$

Localization and Smoothness

Vanishing moments: In the Euclidean setting

$$\int_{-\infty}^{+\infty} |x|^{2k} |f(x)|^2 dx = \int_{-\infty}^{+\infty} \left| \frac{\partial^k \hat{f}(\omega)}{\partial \omega^k} \right|^2 d\omega$$

Localization in space = smoothness in frequency domain

Parametrize the filter using a **smooth spectral transfer function** $\tau(\lambda)$

Application of the filter

$$\tau(\mathbf{\Delta})\mathbf{f} = \mathbf{\Phi} \begin{pmatrix} \tau(\lambda_1) & & \\ & \ddots & \\ & & \tau(\lambda_n) \end{pmatrix} \mathbf{\Phi}^\top \mathbf{f}$$

Localization and Smoothness

Vanishing moments: In the Euclidean setting

$$\int_{-\infty}^{+\infty} |x|^{2k} |f(x)|^2 dx = \int_{-\infty}^{+\infty} \left| \frac{\partial^k \hat{f}(\omega)}{\partial \omega^k} \right|^2 d\omega$$


Localization in space = smoothness in frequency domain

Parametrize the filter using a **smooth spectral transfer function** $\tau(\lambda)$

Application of the parametric filter with learnable parameters α


$$\tau_{\alpha}(\Delta)\mathbf{f} = \Phi \begin{pmatrix} \tau_{\alpha}(\lambda_1) & & \\ & \ddots & \\ & & \tau_{\alpha}(\lambda_n) \end{pmatrix} \Phi^{\top} \mathbf{f}$$

Localization and smoothness



smooth spectral filter (delocalized in space)

Localization and smoothness



smooth spectral filter (localized in space)

Spectral CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\lambda) = \sum_{\ell=0}^r \alpha_{\ell} \lambda^{\ell}$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

Spectral CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\lambda) = \sum_{\ell=0}^r \alpha_{\ell} \lambda^{\ell}$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

😊 $\mathcal{O}(1)$ parameters per layer

Spectral CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\lambda) = \sum_{\ell=0}^r \alpha_{\ell} \lambda^{\ell}$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

☺ $\mathcal{O}(1)$ parameters per layer

☺ No explicit computation of $\Phi^{\top}, \Phi \Rightarrow \mathcal{O}(nr)$ complexity

Spectral CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\Delta) = \sum_{\ell=0}^r \alpha_{\ell} \Delta^{\ell}$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ No explicit computation of $\Phi^{\top}, \Phi \Rightarrow \mathcal{O}(nr)$ complexity
- ☺ Filters have guaranteed r -ring support

Spectral CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\Delta) = \sum_{\ell=0}^r \alpha_{\ell} \Delta^{\ell}$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters



- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ No explicit computation of $\Phi^{\top}, \Phi \Rightarrow \mathcal{O}(nr)$ complexity
- ☺ Filters have guaranteed r -ring support
- ☹ **Mesh dependent !**

Spectral CNN with polynomial filters

Represent spectral transfer function as a **polynomial** of order r

$$\tau_{\alpha}(\Delta) = \sum_{\ell=0}^r \alpha_{\ell} \Delta^{\ell}$$

where $\alpha = (\alpha_0, \dots, \alpha_r)^{\top}$ is the vector of filter parameters



Spectral CNN rational filters (CayleyNet)

Represent spectral transfer function as a **Cayley polynomial** of order r

$$\tau_{\mathbf{c},h}(\lambda) = c_0 + 2\text{Re}\left\{ \sum_{\ell=1}^r c_{\ell}(h\lambda - i)^{\ell}(h\lambda + i)^{-\ell} \right\}$$

where the filter parameters are the vector of real/complex coefficients $\mathbf{c} = (c_0, \dots, c_r)^{\top}$ and the **spectral zoom** h

Spectral CNN rational filters (CayleyNet)

Represent spectral transfer function as a **Cayley polynomial** of order r

$$\tau_{\mathbf{c},h}(\lambda) = c_0 + 2\text{Re}\left\{\sum_{\ell=1}^r c_\ell (h\lambda - i)^\ell (h\lambda + i)^{-\ell}\right\}$$

where the filter parameters are the vector of real/complex coefficients $\mathbf{c} = (c_0, \dots, c_r)^\top$ and the **spectral zoom** h

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed exponential spatial decay
- ☹ $\mathcal{O}(n^3)$ computational complexity with direct matrix inversion

Spectral CNN rational filters (CayleyNet)

Represent spectral transfer function as a **Cayley polynomial** of order r

$$\tau_{\mathbf{c},h}(\lambda) = c_0 + 2\text{Re}\left\{\sum_{\ell=1}^r c_\ell (h\lambda - i)^\ell (h\lambda + i)^{-\ell}\right\}$$

where the filter parameters are the vector of real/complex coefficients $\mathbf{c} = (c_0, \dots, c_r)^\top$ and the **spectral zoom** h

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed exponential spatial decay
- ☺ $\mathcal{O}(n)$ computational complexity with **Jacobi approximate matrix inversion**

Spectral CNN rational filters (CayleyNet)

Represent spectral transfer function as a **Cayley polynomial** of order r

$$\tau_{\mathbf{c},h}(\lambda) = c_0 + 2\text{Re}\left\{\sum_{\ell=1}^r c_{\ell}(h\lambda - i)^{\ell}(h\lambda + i)^{-\ell}\right\}$$

where the filter parameters are the vector of real/complex coefficients $\mathbf{c} = (c_0, \dots, c_r)^{\top}$ and the **spectral zoom** h

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed exponential spatial decay
- ☺ $\mathcal{O}(n)$ computational complexity with **Jacobi approximate matrix inversion**
- ☺ Spectral zoom property allowing to better localize in frequency

Spectral CNN rational filters (CayleyNet)

Represent spectral transfer function as a **Cayley polynomial** of order r

$$\tau_{\mathbf{c},h}(\lambda) = c_0 + 2\text{Re}\left\{ \sum_{\ell=1}^r c_\ell (h\lambda - i)^\ell (h\lambda + i)^{-\ell} \right\}$$

where the filter parameters are the vector of real/complex coefficients $\mathbf{c} = (c_0, \dots, c_r)^\top$ and the **spectral zoom** h

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed exponential spatial decay
- ☺ $\mathcal{O}(n)$ computational complexity with **Jacobi approximate matrix inversion**
- ☺ Spectral zoom property allowing to better localize in frequency
- ☺ Richer class of filters than polynomials for the same order

Spectral CNN rational filters (CayleyNet)

Represent spectral transfer function as a **Cayley polynomial** of order r

$$\tau_{\mathbf{c},h}(\lambda) = c_0 + 2\text{Re}\left\{ \sum_{\ell=1}^r c_\ell (h\lambda - i)^\ell (h\lambda + i)^{-\ell} \right\}$$

where the filter parameters are the vector of real/complex coefficients $\mathbf{c} = (c_0, \dots, c_r)^\top$ and the **spectral zoom** h

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Filters have guaranteed exponential spatial decay
- ☺ $\mathcal{O}(n)$ computational complexity with **Jacobi approximate matrix inversion**
- ☺ Spectral zoom property allowing to better localize in frequency
- ☺ Richer class of filters than polynomials for the same order
- ☺ Scale invariance


Different formulations of non-Euclidean CNNs



Spectral domain



Spatial domain



Parametric domain

Convolution

Euclidean

Spatial domain

$$(f \star g)(x) = \int_{-\pi}^{\pi} f(x')g(x-x')dx'$$

Spectral domain

$$\widehat{(f \star g)}(\omega) = \hat{f}(\omega) \cdot \hat{g}(\omega)$$


'Convolution Theorem'

Non-Euclidean


?

$$\widehat{(f \star g)}_k = \langle f, \phi_k \rangle_{L^2(\mathcal{X})} \langle g, \phi_k \rangle_{L^2(\mathcal{X})}$$

Spatial convolution




Euclidean




Non-Euclidean

Spatial convolution




Euclidean



Non-Euclidean


Convolution on meshes

- Local system of coordinates \mathbf{u}_{ij} around i (e.g. geodesic polar)




Convolution on meshes

- **Local system of coordinates** \mathbf{u}_{ij} around i (e.g. geodesic polar)
- **Local weights** $w_1(\mathbf{u}), \dots, w_L(\mathbf{u})$ w.r.t. \mathbf{u}



Convolution on meshes

- **Local system of coordinates** \mathbf{u}_{ij} around i (e.g. geodesic polar)
- **Local weights** $w_1(\mathbf{u}), \dots, w_L(\mathbf{u})$ w.r.t. \mathbf{u} , e.g. Gaussians
$$w_\ell = \exp\left(-(\mathbf{u} - \boldsymbol{\mu}_\ell)^\top \boldsymbol{\Sigma}_\ell^{-1}(\mathbf{u} - \boldsymbol{\mu}_\ell)\right)$$




Convolution on meshes

- **Local system of coordinates** \mathbf{u}_{ij} around i (e.g. geodesic polar)
- **Local weights** $w_1(\mathbf{u}), \dots, w_L(\mathbf{u})$ w.r.t. \mathbf{u} , e.g. Gaussians
 $w_\ell = \exp\left(-(\mathbf{u} - \boldsymbol{\mu}_\ell)^\top \boldsymbol{\Sigma}_\ell^{-1}(\mathbf{u} - \boldsymbol{\mu}_\ell)\right)$
- **Spatial convolution with filter g**

$$\mathbf{x}'_i \propto \sum_{\ell=1}^L g_\ell \sum_{j=1}^n w_\ell(\mathbf{u}_{ij}) \mathbf{x}_j$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i




Convolution on meshes

- **Local system of coordinates** \mathbf{u}_{ij} around i (e.g. geodesic polar)
- **Local weights** $w_1(\mathbf{u}), \dots, w_L(\mathbf{u})$ w.r.t. \mathbf{u} , e.g. Gaussians
 $w_\ell = \exp(-(\mathbf{u} - \boldsymbol{\mu}_\ell)^\top \boldsymbol{\Sigma}_\ell^{-1}(\mathbf{u} - \boldsymbol{\mu}_\ell))$
- **Spatial convolution with filter** g

$$\mathbf{x}'_i \propto \sum_{\ell=1}^L g_\ell \underbrace{\sum_{j=1}^n w_\ell(\mathbf{u}_{ij}) \mathbf{x}_j}_{\text{patch operator}}$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i




Geodesic polar coordinates

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

ρ_{ij} = geodesic distance from i to j

θ_{ij} = direction of geodesic from i to j



Geodesic polar coordinates


- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

ρ_{ij} = geodesic distance from i to j

θ_{ij} = direction of geodesic from i to j

- Orientation ambiguity!



Geodesic polar coordinates

- Geodesic polar coordinates


$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

ρ_{ij} = geodesic distance from i to j

θ_{ij} = direction of geodesic from i to j

- Orientation ambiguity!

- Canonical direction (e.g. intrinsic vector field, max curvature direction)



Geodesic polar coordinates

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

ρ_{ij} = geodesic distance from i to j


θ_{ij} = direction of geodesic from i to j

- Orientation ambiguity!

- Canonical direction (e.g. intrinsic vector field, max curvature direction)

- Angular max pooling: apply a rotating filter

$$\mathbf{x}'_i \propto \max_{\Delta\theta \in [0, 2\pi)} \sum_{\ell=1}^L g_{\ell} \sum_{j=1}^n w_{\ell}(\rho_{ij}, \theta_{ij} + \Delta\theta) \mathbf{x}_j$$



Geodesic polar coordinates

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

ρ_{ij} = geodesic distance from i to j

θ_{ij} = direction of geodesic from i to j


- Orientation ambiguity!

- Canonical direction (e.g. intrinsic vector field, max curvature direction)


- Angular max pooling: apply a rotating filter

$$\mathbf{x}'_i \propto \max_{\Delta\theta \in [0, 2\pi)} \sum_{\ell=1}^L g_{\ell} \sum_{j=1}^n w_{\ell}(\rho_{ij}, \theta_{ij} + \Delta\theta) \mathbf{x}_j$$


- Fourier transform magnitude w.r.t. θ




Patch operator weight functions



GCNN




ACNN



MoNet

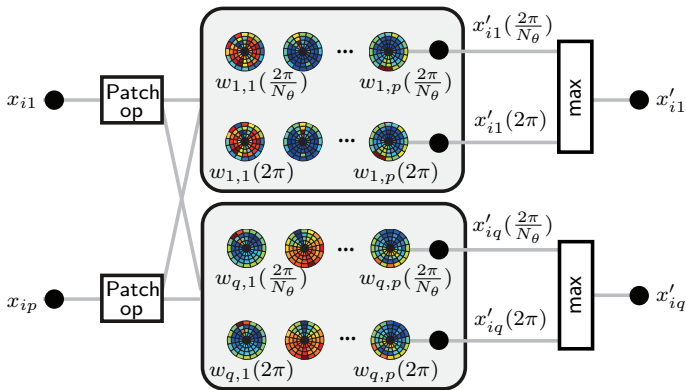
Geodesic convolution layer



Conv. layer
$$x'_{il}(\Delta\theta) \propto \sum_{l'=1}^p \sum_{\ell=1}^L g_\ell \sum_{j=1}^n w_{l,l',\ell}(\rho_{ij}, \theta_{ij} + \Delta\theta) x_{jl'}$$

 $\underbrace{\hspace{15em}}_{\text{Conv. with filter rotated by } \Delta\theta}$

Geodesic convolution layer




Conv. layer

$$x'_{il}(\Delta\theta) \propto \sum_{l'=1}^p \sum_{\ell=1}^L g_{\ell} \underbrace{\sum_{j=1}^n w_{l,l',\ell}(\rho_{ij}, \theta_{ij} + \Delta\theta)}_{\text{Conv. with filter rotated by } \Delta\theta} x_{jl}$$

Angular
max pooling

$$x'_{il} = \max_{\Delta\theta} x'_{il}(\Delta\theta)$$

Learning local descriptors with GCNN



Training set

positive (i, i^+) and negative (i, i^-) pairs of points

Siamese net

two net instances with shared parameters Θ

Poitwise feature cost

$$\begin{aligned} \ell_S(\Theta) = & \gamma \sum_{i, i^+} \|\mathbf{f}_\Theta(\mathbf{x}_i) - \mathbf{f}_\Theta(\mathbf{x}_{i^+})\|_2^2 \\ & + (1 - \gamma) \sum_{i, i^-} [\mu - \|\mathbf{f}_\Theta(\mathbf{x}_i) - \mathbf{f}_\Theta(\mathbf{x}_{i^-})\|_2]_+ \end{aligned}$$

HKS descriptor



Distance in the space of local Heat Kernel Signature (HKS) features
(shown is distance from a point on the shoulder marked in white)

Descriptor: Sun, Ovsjanikov, Guibas 2009 (HKS); data: B et al. 2008 (TOSCA);
Anguelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)

WKS descriptor



Distance in the space of local Wave Kernel Signature (WKS) features
(shown is distance from a point on the shoulder marked in white)

Descriptor: Aubry, Schlickewei, Cremers 2011 (WKS); data: B et al. 2008 (TOSCA);
Anguelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)


Descriptor learned with GCNN



Distance in the space of local GCNN features
(shown is distance from a point on the shoulder marked in white)

Descriptor: Masci et al. 2015 (GCNN); data: B et al. 2008 (TOSCA); Angelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)

Descriptor quality comparison



Descriptor performance using symmetric Princeton benchmark
(training and testing: disjoint subsets of FAUST)

Methods: Sun et al. 2009 (HKS); Aubry et al. 2011 (WKS); Litman, B 2014 (OSD); Masci et al. 2015 (GCNN); data: Bogo et al. 2014 (FAUST); benchmark: Kim et al.

Homogeneous diffusion

$$f_t(x) = -\operatorname{div}(c\nabla f(x))$$

c = **thermal diffusivity constant** describing heat conduction properties of the material (diffusion speed is equal everywhere)

Anisotropic diffusion


$$f_t(x) = -\operatorname{div}(\mathbf{A}(x)\nabla f(x))$$

$\mathbf{A}(x)$ = heat conductivity tensor describing heat conduction properties of the material (diffusion speed is position + direction dependent)

Anisotropic diffusion


Anisotropic

Anisotropic diffusion on manifolds



$$f_t(x) = -\text{div} \left(\mathbf{R}_\theta \begin{pmatrix} \alpha & \\ & 1 \end{pmatrix} \mathbf{R}_\theta^\top \nabla f(x) \right)$$

Anisotropic diffusion on manifolds




$$f_t(x) = -\operatorname{div} \left(\underbrace{\mathbf{R}_\theta \begin{pmatrix} \alpha & \\ & 1 \end{pmatrix} \mathbf{R}_\theta^\top}_{\mathbf{D}_{\alpha\theta}(x)} \nabla f(x) \right)$$

- **Anisotropic Laplacian** $\Delta_{\alpha\theta} f(x) = \operatorname{div} (D_{\alpha\theta}(x) \nabla f(x))$
- θ = orientation w.r.t. max curvature direction
- α = 'elongation'


Anisotropic heat kernels

$$h_{\alpha\theta t}(x, x') = \sum_{k \geq 0} e^{-t\lambda_{\alpha\theta k}} \phi_{\alpha\theta k}(x) \phi_{\alpha\theta k}(x')$$


Patch operator weight functions



GCNN



ACNN




MoNet

Learnable patches on manifolds

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$



Learnable patches on manifolds


- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

- Gaussian weighting functions

$$w_{\mu, \Sigma}(\mathbf{u}) = \exp\left(-\frac{1}{2}(\mathbf{u} - \mu)^\top \Sigma^{-1}(\mathbf{u} - \mu)\right)$$

with learnable covariance Σ and mean μ



Learnable patches on manifolds

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$


- Gaussian weighting functions

$$w_{\mu, \Sigma}(\mathbf{u}) = \exp\left(-\frac{1}{2}(\mathbf{u} - \mu)^\top \Sigma^{-1}(\mathbf{u} - \mu)\right)$$

with learnable covariance Σ and mean μ

- Spatial convolution

$$\mathbf{x}'_i \propto \sum_{\ell=1}^L g_\ell \sum_{j=1}^n w_{\mu_\ell, \Sigma_\ell}(\mathbf{u}_{ij}) \mathbf{x}_j$$



Learnable patches on manifolds

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$


- Gaussian weighting functions

$$w_{\mu, \Sigma}(\mathbf{u}) = \exp\left(-\frac{1}{2}(\mathbf{u} - \mu)^\top \Sigma^{-1}(\mathbf{u} - \mu)\right)$$

with learnable covariance Σ and mean μ

- Spatial convolution

$$\mathbf{x}'_i \propto \sum_{j=1}^n \sum_{\ell=1}^L g_\ell w_{\mu_\ell, \Sigma_\ell}(\mathbf{u}_{ij}) \mathbf{x}_j$$



Learnable patches on manifolds (MoNet)

- Geodesic polar coordinates

$$\mathbf{u}_{ij} = (\rho_{ij}, \theta_{ij})$$

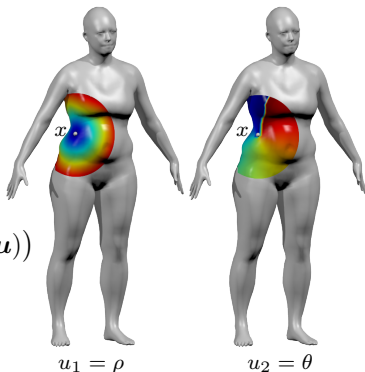
- Gaussian weighting functions

$$w_{\mu, \Sigma}(\mathbf{u}) = \exp\left(-\frac{1}{2}(\mathbf{u} - \mu)^\top \Sigma^{-1}(\mathbf{u} - \mu)\right)$$

with learnable covariance Σ and mean μ

- Spatial convolution

$$\mathbf{x}'_i \propto \sum_{j=1}^n \underbrace{\sum_{\ell=1}^L g_\ell w_{\mu_\ell, \Sigma_\ell}(\mathbf{u}_{ij})}_{\text{Gaussian mixture}} \mathbf{x}_j$$



MoNet as generalization of previous methods


Method	Coordinates \mathbf{u}_{ij}	Weight function $w_{\Theta}(\mathbf{u})$
CNN ¹	$\mathbf{u}_j - \mathbf{u}_i$	$\delta(\mathbf{u} - \mathbf{v})$
GCNN ²	ρ_{ij}, θ_{ij}	$\exp\left(-\frac{1}{2}(\mathbf{u} - \mathbf{v})^{\top} \begin{pmatrix} \sigma_{\rho}^2 & \\ & \sigma_{\theta}^2 \end{pmatrix}^{-1} (\mathbf{u} - \mathbf{v})\right)$
ACNN ³	ρ_{ij}, θ_{ij}	$\exp\left(-t\mathbf{u}^{\top} \mathbf{R}_{\varphi} \begin{pmatrix} \alpha & \\ & 1 \end{pmatrix} \mathbf{R}_{\varphi}^{\top} \mathbf{u}\right)$
MoNet ⁴	ρ_{ij}, θ_{ij}	$\exp\left(-\frac{1}{2}(\mathbf{u} - \boldsymbol{\mu})^{\top} \boldsymbol{\Sigma}^{-1}(\mathbf{u} - \boldsymbol{\mu})\right)$ learnable parameters $\Theta = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$

Some CNN models can be considered as particular settings of MoNet with weighting functions of different form

Methods: ¹LeCun et al. 1998; ²Masci et al. 2015; ³Boscaini et al. 2016;


⁴Monti et al. 2016

Application of MoNet: Protein-Protein Interaction



Designing protein binder for the PD-L1 protein target


Application of MoNet: Protein-Protein Interaction




PDB
2JAB



PDB
3S0D




PDB
3ONJ



Experimentally confirmed computational design of PD-L1 binders

Learning deformation-invariant correspondence

- Groundtruth correspondence
 $\pi^* : \mathcal{X} \rightarrow \mathcal{Y}$ from query shape \mathcal{X}
to some reference shape \mathcal{Y}
- Correspondence = label each query
vertex $i \in \{1, \dots, n\}$ as reference
vertex $\pi_i \in \{1, \dots, m\}$

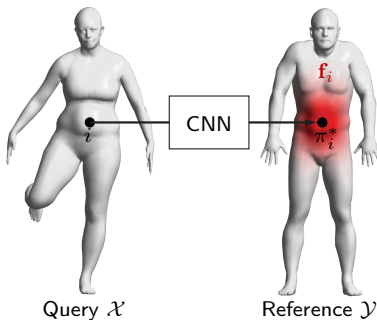


Learning deformation-invariant correspondence

- Groundtruth correspondence
 $\pi^* : \mathcal{X} \rightarrow \mathcal{Y}$ from query shape \mathcal{X}
to some reference shape \mathcal{Y}
- Correspondence = label each query
vertex $i \in \{1, \dots, n\}$ as reference
vertex $\pi_i \in \{1, \dots, m\}$
- Net output at i after softmax layer

$$\mathbf{f}_{\Theta}(\mathbf{x}_i) = (f_{i1}, \dots, f_{im})$$

= probability distribution on \mathcal{Y}




Learning deformation-invariant correspondence

- Groundtruth correspondence
 $\pi^* : \mathcal{X} \rightarrow \mathcal{Y}$ from query shape \mathcal{X}
to some reference shape \mathcal{Y}
- Correspondence = label each query
vertex $i \in \{1, \dots, n\}$ as reference
vertex $\pi_i \in \{1, \dots, m\}$
- Net output at i after softmax layer

$$\mathbf{f}_{\Theta}(\mathbf{x}_i) = (f_{i1}, \dots, f_{im})$$


= probability distribution on \mathcal{Y}



Minimize on training set the cross entropy between groundtruth correspondence and output probability distribution w.r.t. net parameters Θ

$$\min_{\Theta} \sum_{i=1}^n H(\delta_{\pi_i^*}, \mathbf{f}_{\Theta}(\mathbf{x}_i))$$


Correspondence evaluation: Princeton benchmark



Pointwise correspondence error = geodesic distance from the groundtruth

$$\epsilon_i = d_{\mathcal{Y}}(\pi_i^*, \pi_i)$$

Correspondence quality comparison



Correspondence evaluated using asymmetric Princeton benchmark
(training and testing: disjoint subsets of FAUST)


Methods: Kim et al. 2011 (BIM); Rodolà et al. 2014 (RF); Boscaini et al. 2015 (ADD); Masci et al. 2015 (GCNN); Boscaini et al. 2016 (ACNN); Monti et al. 2016 (MoNet); data: Bogo et al. 2014 (FAUST); benchmark: Kim et al. 2011

Shape correspondence error: Blended Intrinsic Map



Pointwise correspondence error (geodesic distance from groundtruth)

Shape correspondence error: Geodesic CNN




Pointwise correspondence error (geodesic distance from groundtruth)

Shape correspondence error: Anisotropic CNN




Pointwise correspondence error (geodesic distance from groundtruth)

Shape correspondence error: MoNet



Pointwise correspondence error (geodesic distance from groundtruth)

Shape correspondence visualization: MoNet



Reference



Texture transferred from reference to query shapes

Correspondence on range images: MoNet



Pointwise correspondence error (geodesic distance from groundtruth)

Correspondence with MoNet: Range images



Reference



Correspondence visualization (similar colors encode corresponding points)

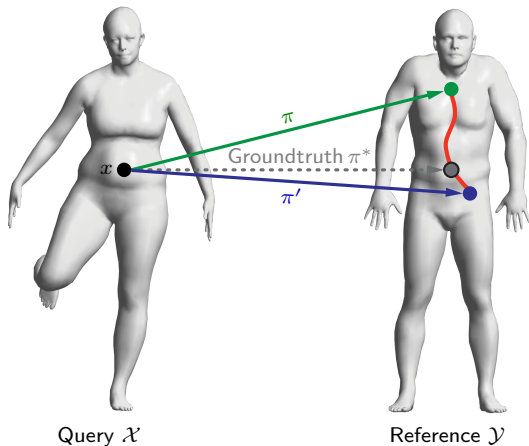
Correspondence with MoNet: Range images



Reference

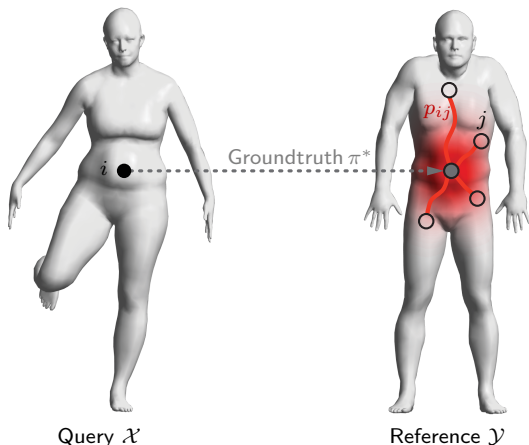
Correspondence visualization (similar colors encode corresponding points)

Correspondence as classification problem, revisited



Classification cost considers equally correspondences that deviate from the groundtruth (no matter how far)


Soft correspondence error



Soft correspondence error = probability-weighted geodesic distance from the groundtruth


$$\bar{\epsilon}_i = \sum_{j=1}^m p_{ij} d_{\mathcal{Y}}(\pi_i^*, j)$$

Pointwise vs Structured learning



Nearby points i, i' on query shape are **not guaranteed** to map to nearby points j, j' on reference shape at **test time**


Functional maps: spectral domain



Functional correspondence $T =$ linear map \mathbf{C} between Fourier coefficients

$$\hat{\mathbf{g}}^T = \hat{\mathbf{f}}^T \mathbf{C}$$


Functional maps: spectral domain



Recover correspondence from $q \geq k$ dimensional pointwise features

$$\begin{pmatrix} \hat{g}_{11} & \hat{g}_{12} & \dots & \hat{g}_{1K} \\ \vdots & \vdots & & \vdots \\ \hat{g}_{q1} & \hat{g}_{q2} & \dots & \hat{g}_{qK} \end{pmatrix} = \begin{pmatrix} \hat{f}_{11} & \hat{f}_{12} & \dots & \hat{f}_{1K} \\ \vdots & \vdots & & \vdots \\ \hat{f}_{q1} & \hat{f}_{q2} & \dots & \hat{f}_{qK} \end{pmatrix} \mathbf{C}$$


Functional maps: spectral domain



Recover correspondence from $q \geq k$ dimensional pointwise features

$$\hat{G} = \hat{F}C$$


Functional maps: spectral domain



Recover correspondence from $q \geq k$ dimensional pointwise features

$$\mathbf{C}^* = \underset{\mathbf{C}}{\operatorname{argmin}} \|\hat{\mathbf{F}}\mathbf{C} - \hat{\mathbf{G}}\|_{\mathbb{F}}^2$$


Functional maps: spatial domain



Rank- K approximation of spatial correspondence

$$\mathbf{T} \approx \Psi \mathbf{C} \Phi^T$$


Functional maps: spatial domain



Probability p_{ij} of point j mapping to i

$$\mathbf{P} \approx \|\Psi\mathbf{C}\Phi^T\|_{\|\cdot\|}$$

Siamese metric learning



Training set

positive (i, i^+) and negative (i, i^-) pairs of points

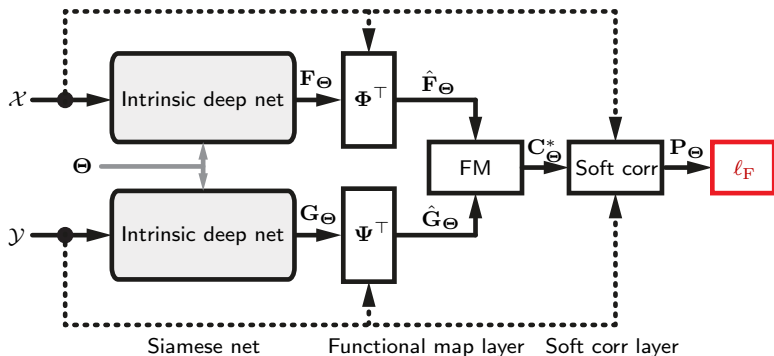
Siamese net

two net instances with shared parameters Θ

Pointwise feature cost

$$l_S(\Theta) = \gamma \sum_{i, i^+} \|f_{\Theta}(x_i) - f_{\Theta}(x_{i^+})\|_2^2 \\ + (1 - \gamma) \sum_{i, i^-} [\mu - \|f_{\Theta}(x_i) - f_{\Theta}(x_{i^-})\|_2]_+^2$$

Structured correspondence with FMNet



Siamese net

two net instances with shared parameters Θ

Functional map layer

$$C_{\Theta}^* = \underset{C}{\operatorname{argmin}} \|\hat{F}_{\Theta} C - \hat{G}_{\Theta}\|_F^2$$

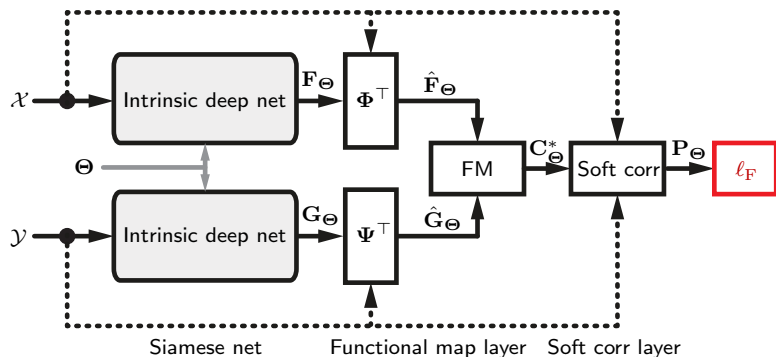
Soft correspondence layer

$$P_{\Theta} = |\Psi C_{\Theta} \Phi^T|_{\|\cdot\|}$$

Soft error cost

$$\ell_F(\Theta) = \sum_{i=1}^n \sum_{j=1}^m p_{\Theta, ij} d_{\mathcal{Y}}(\pi_i^*, j)$$

Structured correspondence with FMNet



Siamese net

two net instances with shared parameters Θ

Functional map layer

$$C_\Theta^* = \hat{F}_\Theta^\dagger \hat{G}_\Theta$$


Soft correspondence layer

$$P_\Theta = |\Psi C_\Theta \Phi^\top|_{\|\cdot\|}$$

Soft error cost

$$l_F(\Theta) = \|P_\Theta \circ D_y\|$$

Correspondence quality comparison




Correspondence evaluated using asymmetric Princeton benchmark
(training and testing: disjoint subsets of FAUST)

Methods: Kim et al. 2011 (BIM); Rodolà et al. 2014 (RF); Boscaini et al. 2015 (ADD); Masci et al. 2015 (GCNN); Boscaini et al. 2016 (ACNN); Monti et al. 2016 (MoNet); Litany et al. 2017 (FMNet); data: Bogo et al. 2014 (FAUST); benchmark: Kim et al. 2011


Shape representation




Image-based




Volumetric



Point-based



Surface-based



PointNet: learning on sets




PointNet: learning on sets

- **Permutation-invariant** function

$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = f(\mathbf{x}_{\pi_1}, \dots, \mathbf{x}_{\pi_n})$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i




PointNet: learning on sets

- **Permutation-invariant** function

$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = f(\mathbf{x}_{\pi_1}, \dots, \mathbf{x}_{\pi_n})$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i




PointNet: learning on sets

- **Permutation-invariant** function

$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = f(\mathbf{x}_{\pi_1}, \dots, \mathbf{x}_{\pi_n})$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i

- **Shared function** $h_{\Theta}(\cdot)$ applied to each point + permutation-invariant aggregation (max or \sum)



(Vinyals, Bengio, Kudlur 2015); Qi et al. 2017


PointNet: learning on sets

- **Permutation-invariant** function

$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = f(\mathbf{x}_{\pi_1}, \dots, \mathbf{x}_{\pi_n})$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i

- **Shared function** $h_{\Theta}(\cdot)$ applied to each point + permutation-invariant aggregation (max or \sum)
- **Spatial transformer units**



(Vinyals, Bengio, Kudlur 2015); Qi et al. 2017

PointNet: learning on sets

- **Permutation-invariant** function


$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = f(\mathbf{x}_{\pi_1}, \dots, \mathbf{x}_{\pi_n})$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is feature at vertex i


- **Shared function** $h_{\Theta}(\cdot)$ applied to each point + permutation-invariant aggregation (max or \sum)

- Spatial transformer units

- Local grouping (PointNet++, PCPNet)




PointNet applications



Object recognition




Part segmentation



Semantic segmentation



Curvature




Normals



Point cloud generation


Graph-based edge convolution

- Local neighborhood structure modeled as a **graph**




Graph-based edge convolution

- Local neighborhood structure modeled as a **graph**
- **Edge feature** function $h_{\Theta}(\cdot, \cdot)$ parametrized by Θ



Graph-based edge convolution


- Local neighborhood structure modeled as a **graph**
- **Edge feature** function $h_{\Theta}(\cdot, \cdot)$ parametrized by Θ
- Permutation-invariant **aggregation operator** \square (e.g. \sum or \max) on the neighborhood of i



Graph-based edge convolution

- Local neighborhood structure modeled as a **graph**
- **Edge feature** function $h_{\Theta}(\cdot, \cdot)$ parametrized by Θ
- Permutation-invariant **aggregation operator** \square (e.g. \sum or \max) on the neighborhood of i
- **Edge convolution** (EdgeConv)


$$\mathbf{x}'_i = \square_j h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j)$$



Graph-based edge convolution

- Local neighborhood structure modeled as a **graph**
- **Edge feature** function $h_{\Theta}(\cdot, \cdot)$ parametrized by Θ
- Permutation-invariant **aggregation operator** \square (e.g. \sum or \max) on the neighborhood of i
- **Edge convolution** (EdgeConv)

$$\mathbf{x}'_i = \square_j h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j)$$



Learnable local (nonlinear) operator

Particular cases


Method	Aggregation \square	Edge feature $h(\mathbf{x}_i, \mathbf{x}_j)$
Laplacian	\sum	$w_{ij}(\mathbf{x}_j - \mathbf{x}_i)$
PointNet ¹	-	$h(\mathbf{x}_i)$
PointNet++ ²	max	$h(\mathbf{x}_i)$
MoNet ³	\sum	$\sum_{\ell} g_{\ell} w_{\ell}(\mathbf{u}_{ij}) \mathbf{x}_j$
PCNN ⁴	\sum	$\sum_{\ell m} c(\mathbf{x}_i \cdot \mathbf{k}_{\ell m}) w_{i\ell} q_{\ell}(\mathbf{x}_i, \mathbf{x}_j)$

Wang et al. 2018; ¹Qi et al. 2017; ²Qi, Su et al. 2017; ³Monti et al. 2017; ⁴Atzmon

et al. 2018

Dynamic Graph CNN (DynGCNN)

Construct k -NN graph in feature space and **update it after each layer**




Layer l

Features $\mathbf{x}_1^{(l)}, \dots, \mathbf{x}_n^{(l)} \in \mathbb{R}^{d_l}$

k -NN graph $\mathcal{G}^{(l)}$

$h^{(l)} : \mathbb{R}^{d_l} \times \mathbb{R}^{d_l} \rightarrow \mathbb{R}^{d_{l+1}}$




Layer $l + 1$

Features $\mathbf{x}_1^{(l+1)}, \dots, \mathbf{x}_n^{(l+1)} \in \mathbb{R}^{d_{l+1}}$

k -NN graph $\mathcal{G}^{(l+1)}$

$h^{(l+1)} : \mathbb{R}^{d_{l+1}} \times \mathbb{R}^{d_{l+1}} \rightarrow \mathbb{R}^{d_{l+2}}$

Learning semantic features



Left: Distance from red point in the feature space of different DynGCNN layers
Right: semantic segmentation results


Shape classification (ModelNet40)

Method	Mean class accuracy	Overall accuracy
3DShapeNet ¹	77.3%	84.7%
VoxNet ²	83.0%	85.9%
Subvolume ³	86.0%	89.2%
ECC ⁴	83.2%	87.4%
PointNet ⁵	86.0%	89.2%
PointNet++ ⁶	–	90.7%
Kd-Net ⁷	–	91.8%
DynGCNN (baseline) ⁸	88.8%	91.2%
DynGCNN⁸	90.2%	92.2%

Classification accuracy of different methods on ModelNet40

Methods: ¹Wu et al. 2015; ²Maturana et al. 2015; Qi et al. 2016; ⁴Simonovsky, Komodakis 2017; ⁵Qi et al. 2017; ⁶Qi, Su et al. 2017; ⁷Klokov, Lempitsky 2017; ⁸Wang et al. 2018; data: Wu et al. 2015 (ModelNet)


Semantic segmentation: synthetic (ShapeNet)




Methods: Qi et al. 2017 (PointNet); Wang et al. 2018 (DynGCNN); data: Yi et al.

2016 (ShapeNet)

Semantic segmentation: indoor scans (S3DIS)




PointNet




DynGCNN



Groundtruth



Scene



Results of semantic segmentation of point cloud+RGB data using different architectures

Methods: Qi et al. 2017 (PointNet); Wang et al. 2018 (DynGCNN); data: Armeni et al. 2016 (S3DIS)


Shape segmentation: indoor scans (S3DIS)

Method	Mean IoU	Overall accuracy
PointNet (Baseline) ¹	20.1%	53.2%
PointNet ¹	47.6%	78.5%
MS + CU(2) ²	47.8%	79.2%
G + RCU ²	49.7%	81.1%
DynGCNN³	56.1%	84.1%

S3DIS indoor scene semantic segmentation accuracy

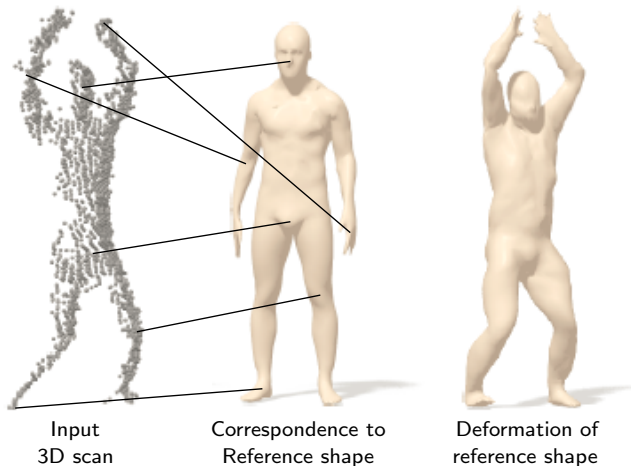
Methods: ¹Qi et al. 2017; ²Engelmann et al. 2017 ³Wang et al. 2018; data: Armeni et al, 2016 (S3DIS)

Surface normal prediction



Surface normal predicted using DynGCNN (odd columns) and groundtruth (even columns). Normal direction is color-coded

3D shape analysis and synthesis




Intrinsic Variational Autoencoder (VAE)

Minimize

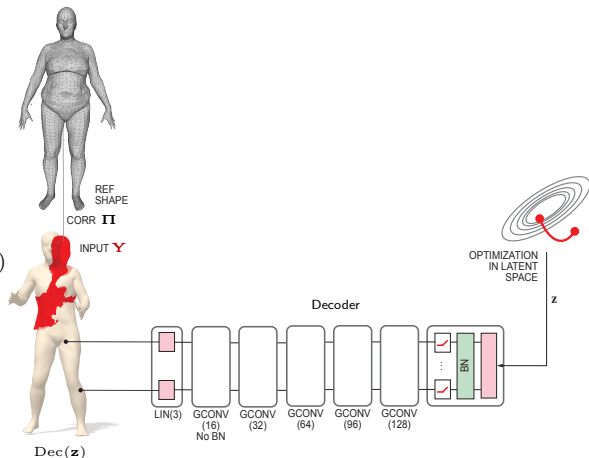
$$\| \text{Dec}(\text{Enc}(\mathbf{X})) - \mathbf{X} \|_F + \lambda D_{\text{KL}}(q(\mathbf{z}|\mathbf{X}) \| p(\mathbf{z}))$$

w.r.t. net parameters




Shape completion

Minimize
 $\|\text{Dec}(\mathbf{z})\mathbf{\Pi} - \mathbf{T}\mathbf{Y}\|_F$
in alternating manner
w.r.t. \mathbf{z} and $\mathbf{T} \in \text{SO}(3)$



Shape completion comparison




Methods: Litany et al. 2017; Dai et al. 2016 (3D-EPN); Kazhdan et al. 2013

(Poisson)

Shape completion examples




Shape completion examples




Generative models of faces




X



$0.5X + 0.5Y$




Y




$-0.5X + 1.5Y$

Generative models of faces



Face from DNA



Summary: intrinsic deep learning

- Intrinsic deep learning allows architectures that are deformation invariant by construction
- Vastly less parameters / training data
- Part of a bigger trend of Geometric deep learning on non-Euclidean domains such as graphs
- Several ways of defining intrinsic convolution - each has its own advantages / disadvantages
- Intrinsic shape synthesis (especially with different topology) is a big open question - part of a broader problem of graph generating networks

Summary: deep learning on 3D data




Image-based

- 😊 Simple
- 😊 Efficient*
- 😞 Not geometric
- 😞 No invariance***




Volumetric

- 😊 Simple
- 😞 High complexity**
- 😞 Coarse
- 😞 No invariance***



Point-based

- 😊 Efficient
- 😊 Simple
- 😊 Accurate
- 😞 No invariance***




Surface-based

- 😊 Efficient
- 😊 Accurate
- 😊 Deform. invariant
- 😞 Non-standard operations

*Rendering can be slow and memory-heavy. **Can be remedied to some extent by hierarchical data structures. ***No invariance to deformations. Rigid transformations can be remedied to some extent by transformer units.

Course Information (slides/code/comments)



http://geometry.cs.ucl.ac.uk/dl_for_CG/

