

Machine Learning Meets Geometry

Emanuele Rodolà



USI Lugano
Switzerland

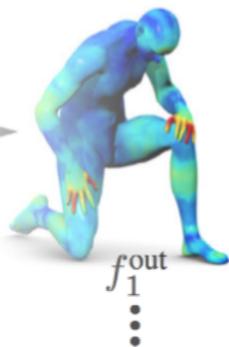


Sapienza University of Rome
Italy

SAPIENZA
UNIVERSITÀ DI ROMA



Input M -dim

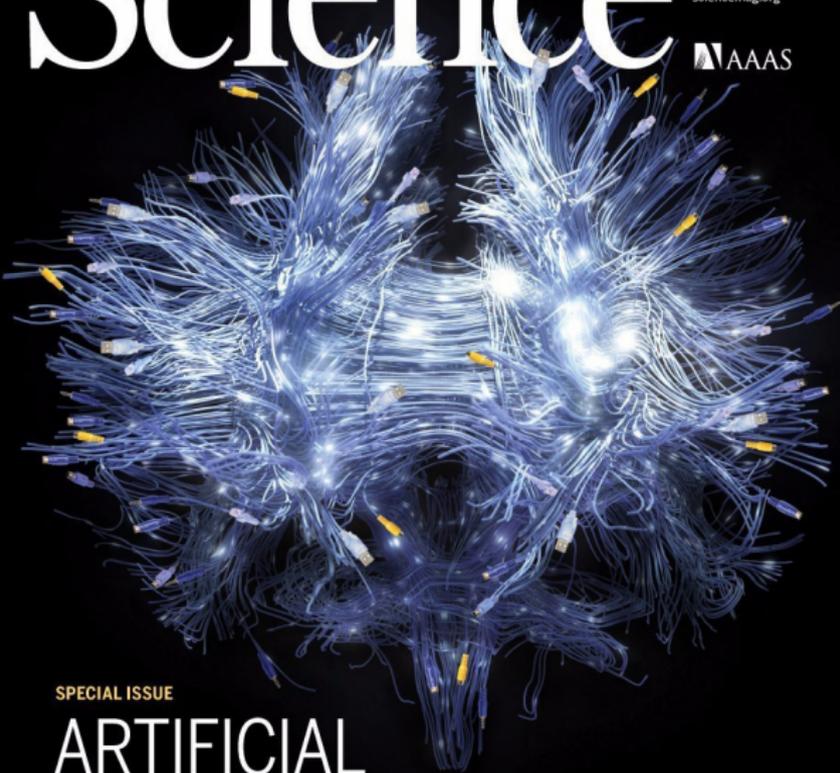


Output Q -dim

Science

\$10
17 JULY 2015
sciencemag.org

AAAS



SPECIAL ISSUE

ARTIFICIAL INTELLIGENCE

A large, light-colored wooden sign with the word "facebook" in blue, lowercase, sans-serif font. The sign is mounted on a stage with a blue and white geometric background.

facebook



Google DEEPMIND

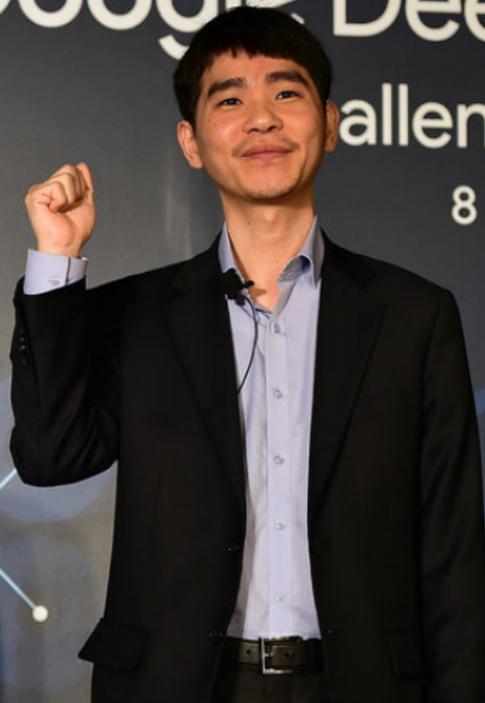


Google acquires DeepMind in 2014



Google DeepMind Challenge Match

8 - 15 March 2016





Tesla autonomous driving car

Message may be recorded to improve translations.

skype



I was wondering what you are going to do later?

Me preguntaba lo que vas a hacer después?

Going to the pub, do you want to join us? I think you met some of the team at the party in April.



Al pub, ¿quieres unirme a nosotros? Creo que conociste a algunos miembros del equipo en la fiesta en abril.

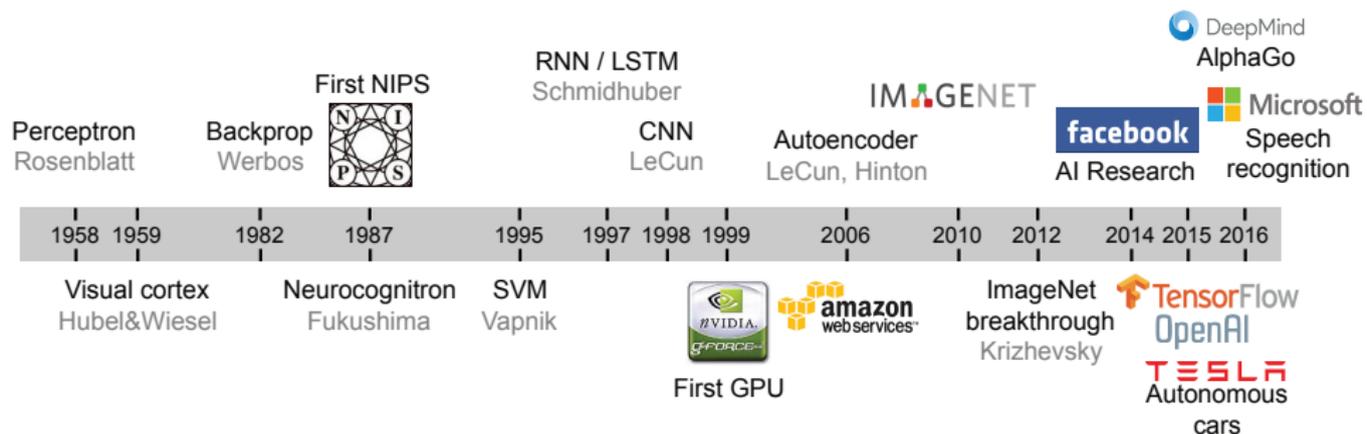


Can I join you guys after my meeting?

¿Puedo unirme a ustedes después de mi reunión?

Type a message in English here



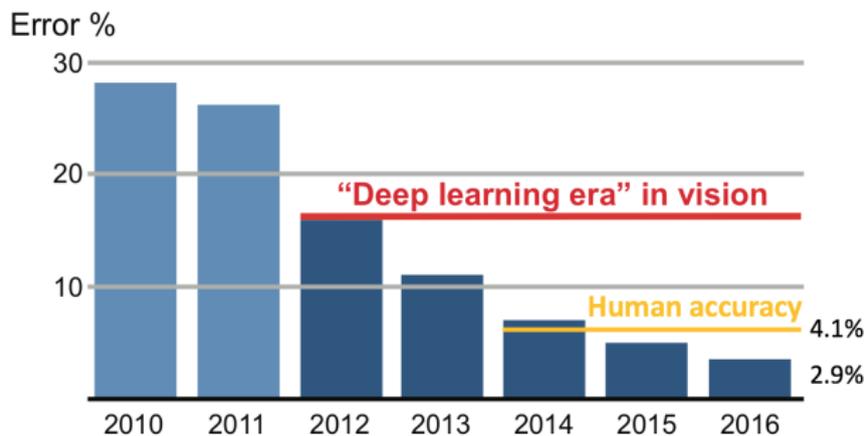


AI BIRTH

DARK AGES

RENAISSANCE

Breakthrough in image recognition



Doubt thou the stars are fire,
Doubt that the sun doth move,
Doubt truth to be a liar,
But never doubt I love...

Text



Audio signals



Images

Doubt thou the stars are fire,
Doubt that the sun doth move,
Doubt truth to be a liar,
But never doubt I love...

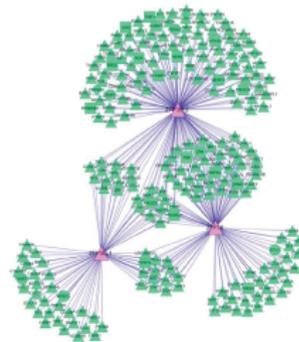
Text



Audio signals



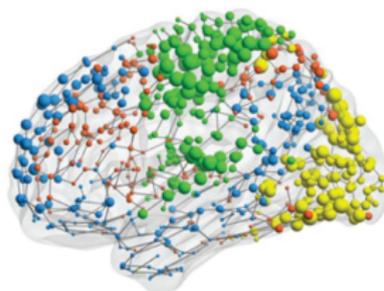
Social networks



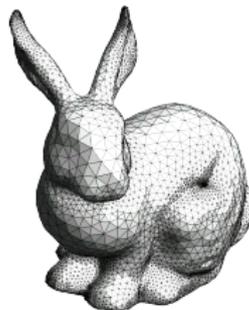
Regulatory networks



Images



Functional networks



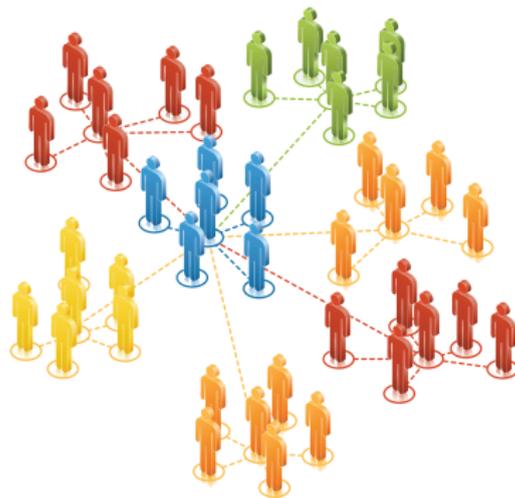
3D shapes

How to design deep nets on
non-Euclidean domains
and what to do with them?

Prototypical non-Euclidean objects



Manifolds

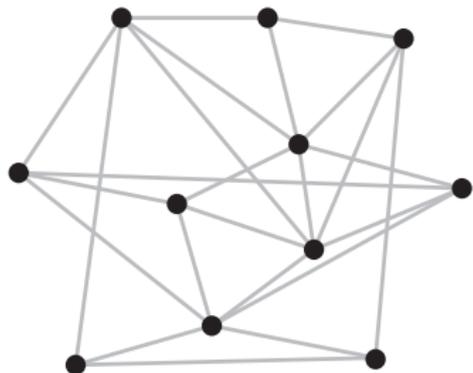


Graphs

Domain structure vs Data on a domain

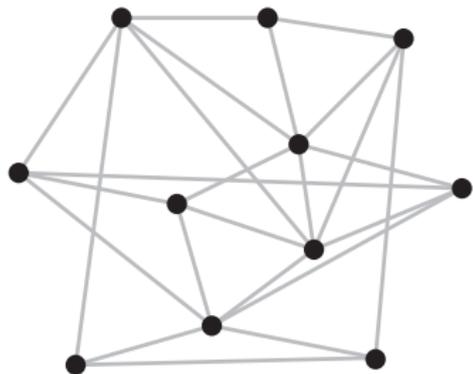


Domain structure vs Data on a domain

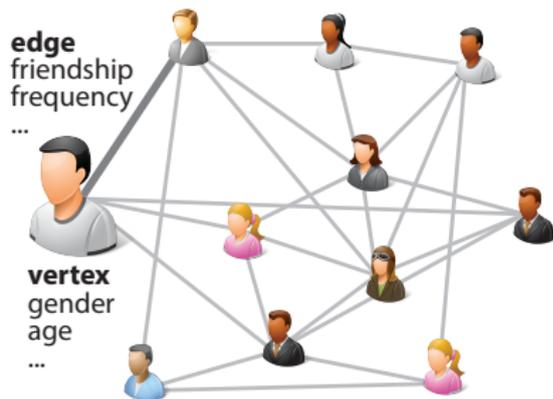


Domain structure

Domain structure vs Data on a domain

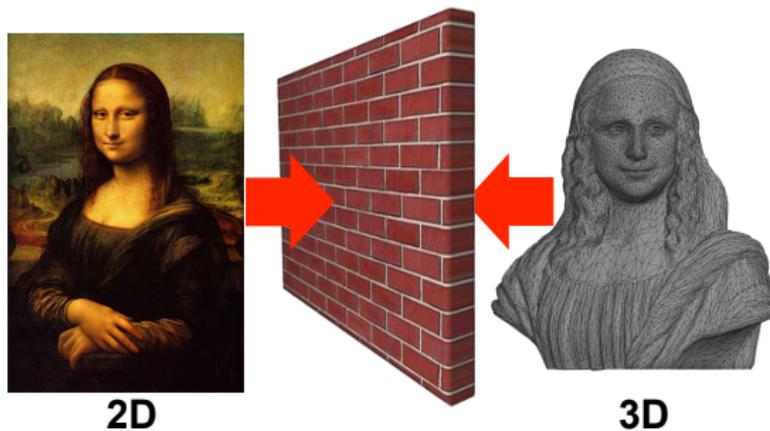


Domain structure



Data on a domain

Domain structure vs Data on a domain



Fixed vs different domain

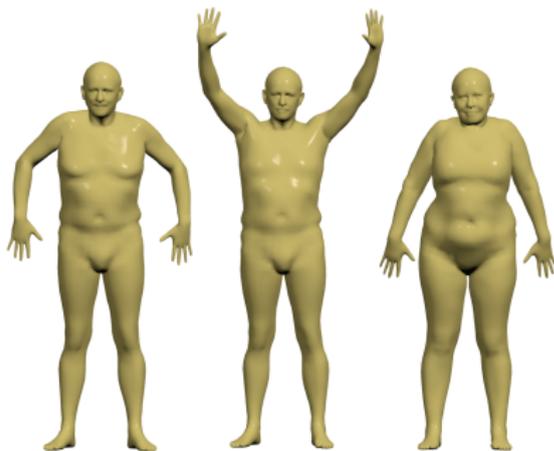


Social network
(fixed graph)

Fixed vs different domain

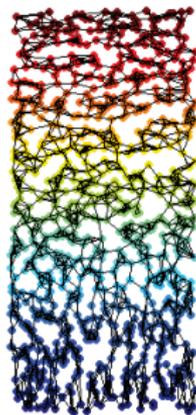
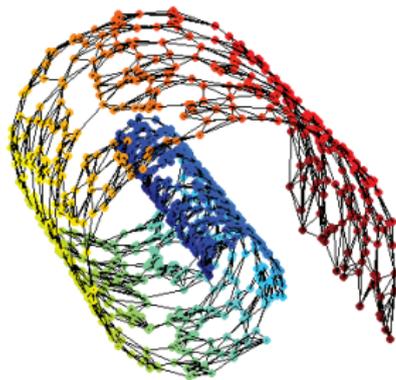


Social network
(fixed graph)



3D shapes
(different manifolds)

Geometric learning \neq Manifold learning



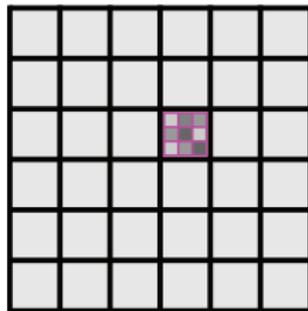
Different formulations of non-Euclidean CNNs



Spectral domain



Spatial domain

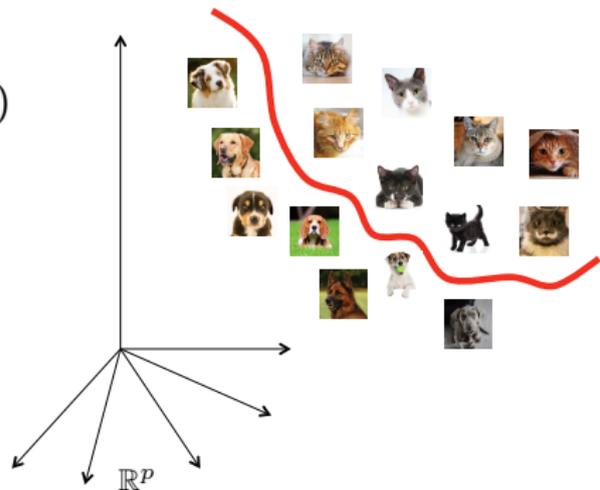


Embedding domain

Background

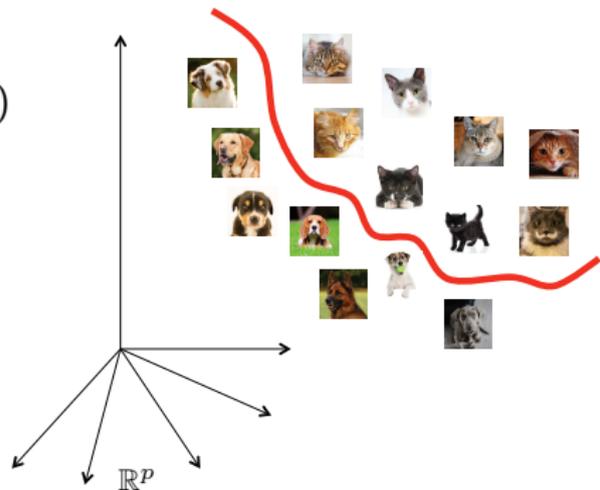
Supervised learning

- Data vectors $\mathbf{f} \in \mathbb{R}^p$
(e.g. for 512×512 images $p \approx 10^6$)
- Unknown classification functional
 $y : \mathbb{R}^p \rightarrow \{1, \dots, L\}$ in L classes
- Training set
$$S = \{(\mathbf{f}_i \in \mathbb{R}^p, y_i = y(\mathbf{f}_i))\}_{i=1}^T$$
- Parametric model y_{Θ} of y



Supervised learning

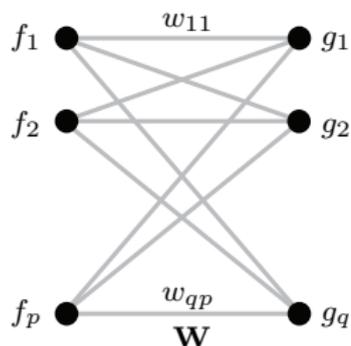
- Data vectors $\mathbf{f} \in \mathbb{R}^p$
(e.g. for 512×512 images $p \approx 10^6$)
- Unknown classification functional $y : \mathbb{R}^p \rightarrow \{1, \dots, L\}$ in L classes
- Training set
$$S = \{(\mathbf{f}_i \in \mathbb{R}^p, y_i = y(\mathbf{f}_i))\}_{i=1}^T$$
- Parametric model y_{Θ} of y



Supervised learning: find optimal model parameters by minimizing the loss ℓ on the training set

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \sum_{i=1}^T \ell(y_{\Theta}(\mathbf{f}_i), y_i)$$

Neural network (NN)



Single linear layer

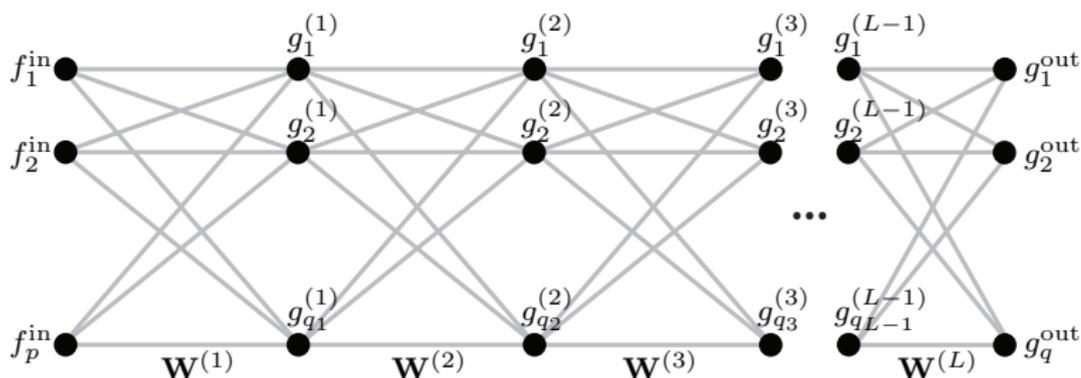
Linear layer

$$g_l = \xi \left(\sum_{l'=1}^p f_{l'} w_{l,l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters layer weights \mathbf{W} (including bias)

Neural network (NN)



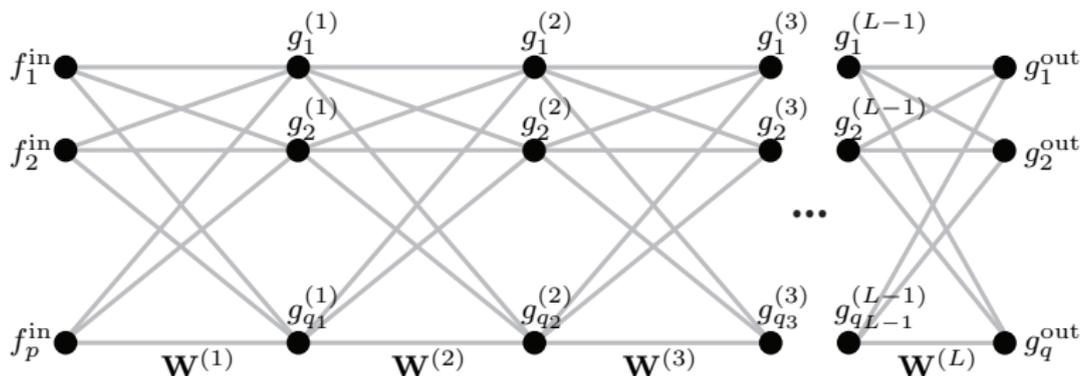
Deep neural network consisting of L layers

Linear layer
$$g_l^{(k)} = \xi \left(\sum_{l'=1}^{q_{k-1}} g_{l'}^{(k-1)} w_{l,l'}^{(k)} \right) \quad \begin{array}{l} l = 1, \dots, q_k \\ l' = 1, \dots, q_{k-1} \end{array}$$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters weights of all layers $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)}$ (including biases)

Neural network (NN)



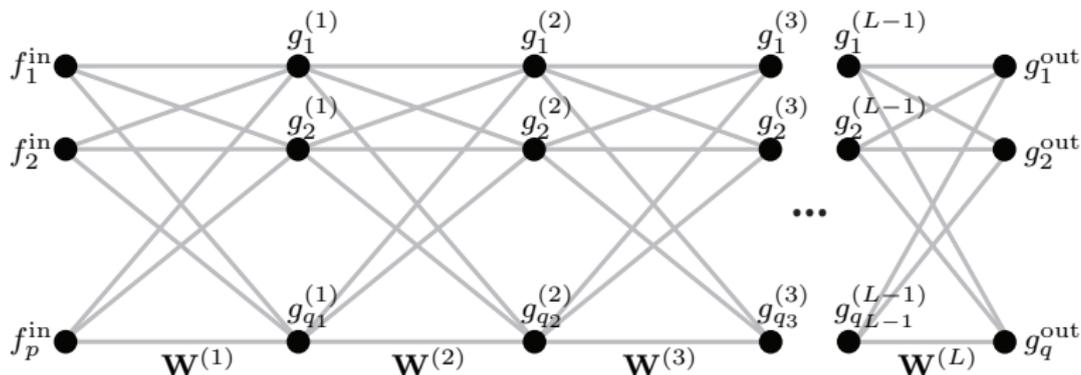
Deep neural network consisting of L layers

Linear layer $\mathbf{g}^{(k)} = \xi(\mathbf{W}^{(k)} \mathbf{g}^{(k-1)})$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters weights of all layers $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)}$ (including biases)

Neural network (NN)



Deep neural network consisting of L layers

Net output $\mathbf{g}^{\text{out}} = \xi(\dots \mathbf{W}^{(2)} \xi(\mathbf{W}^{(1)} \mathbf{f}^{\text{in}})) = y_{(\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)})}(\mathbf{f}^{\text{in}})$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters weights of all layers $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)}$ (including biases)

Neural nets as universal approximators

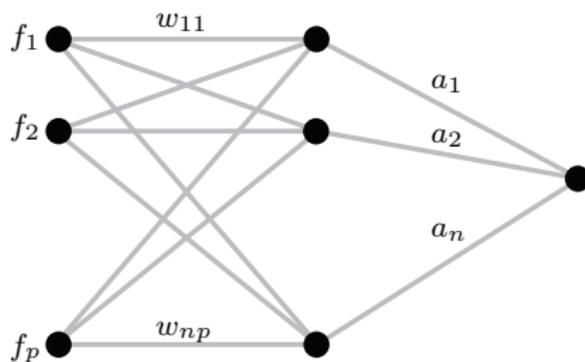
Universal Approximation Theorem Let ξ be a non-constant, bounded, and monotonically-increasing continuous activation function, $y : [0, 1]^p \rightarrow \mathbb{R}$ continuous function, and $\epsilon > 0$. Then, $\exists n$ and parameters $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times p}$ (including bias) s.t.

$$\left| \sum_{i=1}^n a_i \xi(\mathbf{w}_i^\top \mathbf{f}) - y(\mathbf{f}) \right| < \epsilon \quad \forall \mathbf{f} \in [0, 1]^p$$

Neural nets as universal approximators

Universal Approximation Theorem Let ξ be a non-constant, bounded, and monotonically-increasing continuous activation function, $y : [0, 1]^p \rightarrow \mathbb{R}$ continuous function, and $\epsilon > 0$. Then, $\exists n$ and parameters $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times p}$ (including bias) s.t.

$$\left| \sum_{i=1}^n a_i \xi(\mathbf{w}_i^\top \mathbf{f}) - y(\mathbf{f}) \right| < \epsilon \quad \forall \mathbf{f} \in [0, 1]^p$$



Neural nets as universal approximators

Universal Approximation Theorem Let ξ be a non-constant, bounded, and monotonically-increasing continuous activation function, $y : [0, 1]^p \rightarrow \mathbb{R}$ continuous function, and $\epsilon > 0$. Then, $\exists n$ and parameters $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times p}$ (including bias) s.t.

$$\left| \sum_{i=1}^n a_i \xi(\mathbf{w}_i^\top \mathbf{f}) - y(\mathbf{f}) \right| < \epsilon \quad \forall \mathbf{f} \in [0, 1]^p$$

☺ Any continuous function can be approximated arbitrarily well by a neural network with a single hidden layer

Neural nets as universal approximators

Universal Approximation Theorem Let ξ be a non-constant, bounded, and monotonically-increasing continuous activation function, $y : [0, 1]^p \rightarrow \mathbb{R}$ continuous function, and $\epsilon > 0$. Then, $\exists n$ and parameters $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times p}$ (including bias) s.t.

$$\left| \sum_{i=1}^n a_i \xi(\mathbf{w}_i^\top \mathbf{f}) - y(\mathbf{f}) \right| < \epsilon \quad \forall \mathbf{f} \in [0, 1]^p$$

- ☺ Any continuous function can be approximated arbitrarily well by a neural network with a single hidden layer
- ☹ How to find the parameters?

Neural nets as universal approximators

Universal Approximation Theorem Let ξ be a non-constant, bounded, and monotonically-increasing continuous activation function, $y : [0, 1]^p \rightarrow \mathbb{R}$ continuous function, and $\epsilon > 0$. Then, $\exists n$ and parameters $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times p}$ (including bias) s.t.

$$\left| \sum_{i=1}^n a_i \xi(\mathbf{w}_i^\top \mathbf{f}) - y(\mathbf{f}) \right| < \epsilon \quad \forall \mathbf{f} \in [0, 1]^p$$

- ☺ Any continuous function can be approximated arbitrarily well by a neural network with a single hidden layer
- ☹ How to find the parameters?
- ☹ Does it generalize well / overfit?

Neural nets as universal approximators

Universal Approximation Theorem Let ξ be a non-constant, bounded, and monotonically-increasing continuous activation function, $y : [0, 1]^p \rightarrow \mathbb{R}$ continuous function, and $\epsilon > 0$. Then, $\exists n$ and parameters $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times p}$ (including bias) s.t.

$$\left| \sum_{i=1}^n a_i \xi(\mathbf{w}_i^\top \mathbf{f}) - y(\mathbf{f}) \right| < \epsilon \quad \forall \mathbf{f} \in [0, 1]^p$$

- ☺ Any continuous function can be approximated arbitrarily well by a neural network with a single hidden layer
- ☹ How to find the parameters?
- ☹ Does it generalize well / overfit?
- ☹ Shallow nets work poorly for high-dimensional data s.a. images

Take advantage of the
structure of the data!

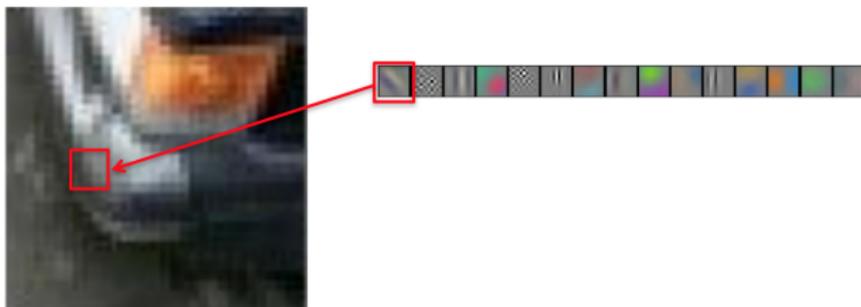
Stationarity and Self-similarity



Data is self-similar across the domain

Locality

- Local features are represented by $r \times r$ compact support kernels.
- $\mathcal{O}(1)$ parameters per filter.



Translation invariance (image classification tasks)



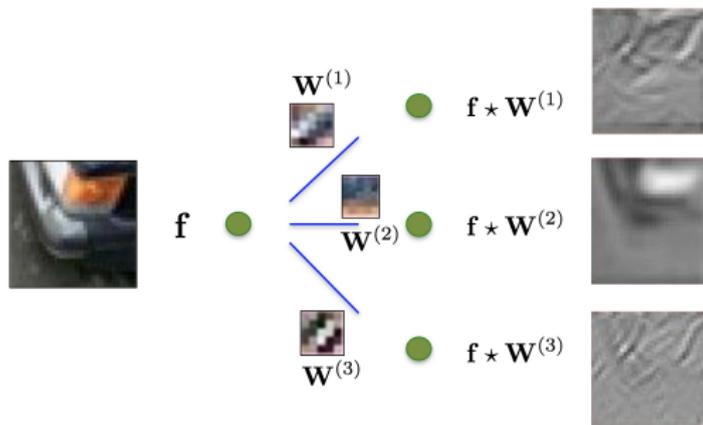
$$y(\mathcal{T}_v f) = y(f) \quad \forall f, v$$

where

- image is modeled as a function $f \in L^2([0, 1]^2)$
- $\mathcal{T}_v f(x) = f(x - v)$ is a **translation operator**
- $v \in [0, 1]^2$ is a translation vector
- $y : L^2([0, 1]^2) \rightarrow \{1, \dots, L\}$ is classification functional

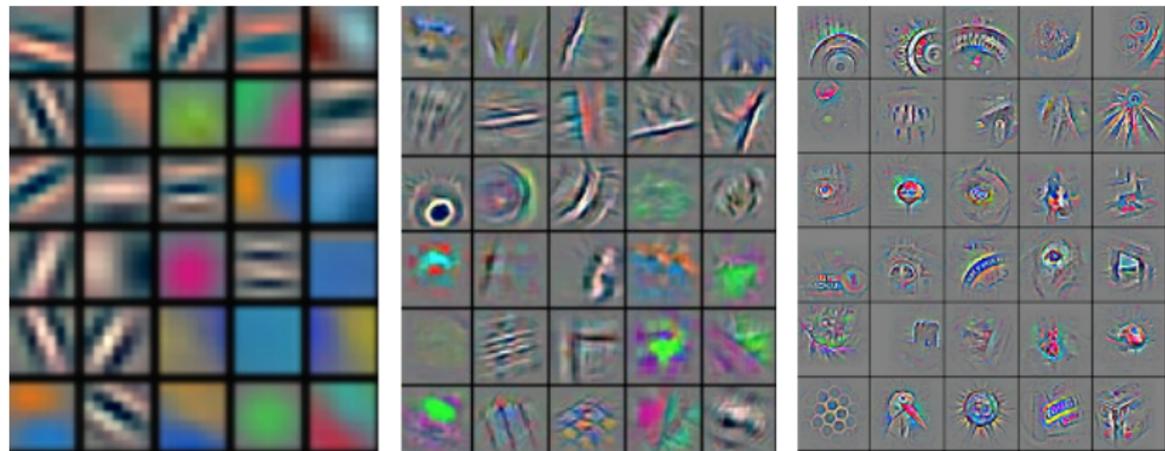
Introducing local filters

- Stationarity is implemented by shift-invariant **convolutional operators**.
- $\mathcal{O}(n \log n)$ computational complexity in general case (FFT).
- $\mathcal{O}(n)$ computational complexity for compact kernels.



Hierarchy and Compositionality

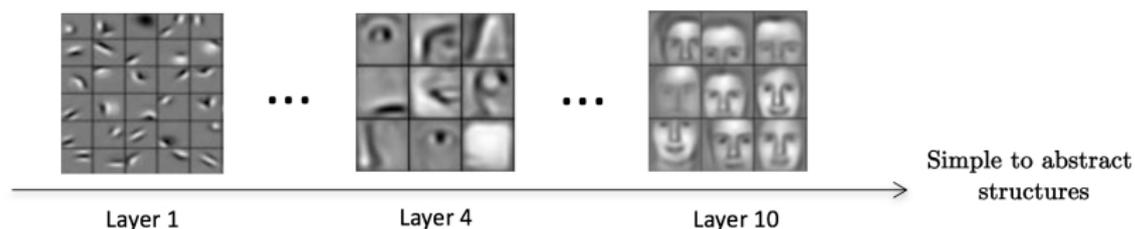
Data is **compositional**: images, video, sound are formed of **hierarchical local stationary patterns**.



Typical features learned by a CNN becoming increasingly complex at deeper layers

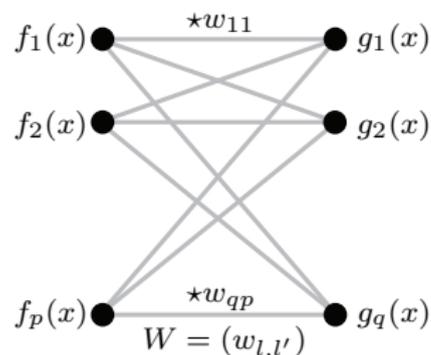
Hierarchy and Compositionality

Data is **compositional**: images, video, sound are formed of **hierarchical local stationary patterns**.



Typical features learned by a CNN becoming increasingly complex at deeper layers

Convolutional neural network (CNN)



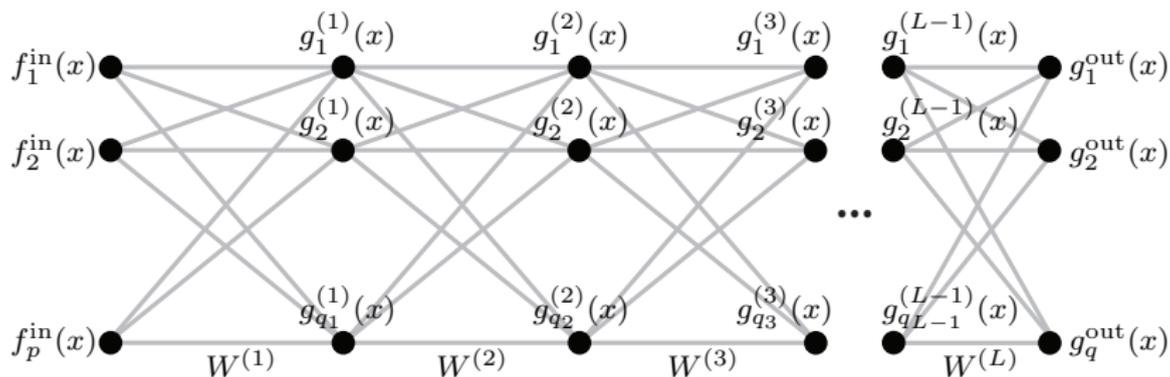
Single convolutional layer

Conv. layer $g_l(x) = \xi \left(\sum_{l'=1}^p (f_{l'} \star w_{l,l'})(x) \right) \quad \begin{matrix} l = 1, \dots, q \\ l' = 1, \dots, p \end{matrix}$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters filters W

Convolutional neural network (CNN)



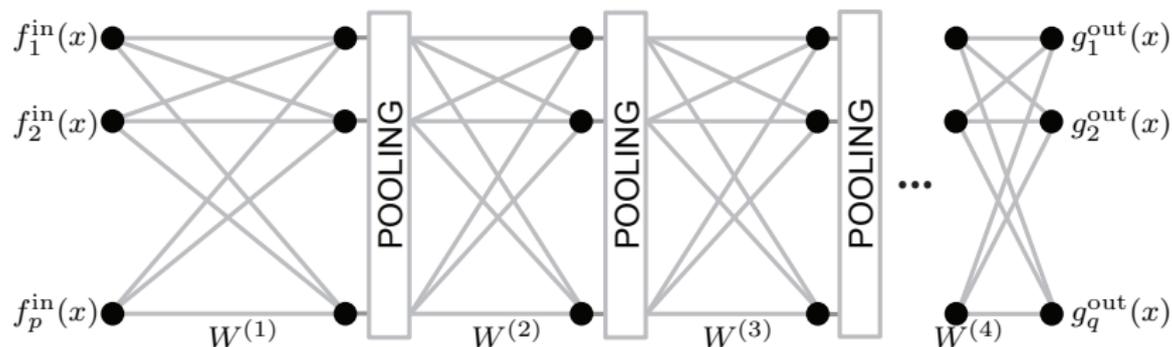
CNN consisting of L convolutional layers

Conv. layer
$$g_l^{(k)}(x) = \xi \left(\sum_{l'=1}^{q_{k-1}} (g_{l'}^{(k-1)} \star w_{l,l'}^{(k)}) (x) \right) \quad \begin{array}{l} l = 1, \dots, q_k \\ l' = 1, \dots, q_{k-1} \end{array}$$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters filters of all layers $W^{(1)}, \dots, W^{(L)}$

Convolutional neural network (CNN)



CNN consisting of L convolutional layers interleaved with pooling

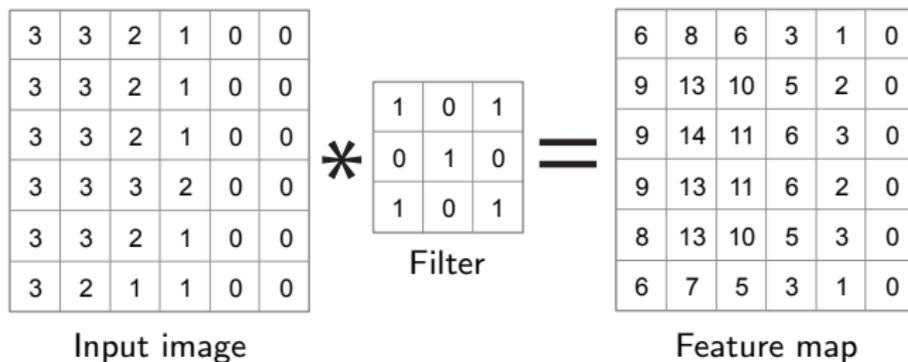
Conv. layer $g_l^{(k)}(x) = \xi \left(\sum_{l'=1}^{q_{k-1}} (g_{l'}^{(k-1)} \star w_{l,l'}^{(k)}) (x) \right) \quad \begin{matrix} l = 1, \dots, q_k \\ l' = 1, \dots, q_{k-1} \end{matrix}$

Activation, e.g. $\xi(x) = \max\{x, 0\}$ rectified linear unit (ReLU)

Parameters filters of all layers $W^{(1)}, \dots, W^{(L)}$

Pooling $g_l^{(k)}(x) = \|g_l^{(k-1)}(x') : x' \in \mathcal{N}(x)\|_p \quad p = 1, 2, \text{ or } \infty$

Pooling



Pooling

3	3	2	1	0	0
3	3	2	1	0	0
3	3	2	1	0	0
3	3	3	2	0	0
3	3	2	1	0	0
3	2	1	1	0	0

Input image

*

1	0	1
0	1	0
1	0	1

Filter

=

6	8	6	3	1	0
9	13	10	5	2	0
9	14	11	6	3	0
9	13	11	6	2	0
8	13	10	5	3	0
6	7	5	3	1	0

Feature map

13	10	2
14	11	3
13	10	3

Max pooling

Pooling

3	3	2	1	0	0
3	3	2	1	0	0
3	3	2	1	0	0
3	3	3	2	0	0
3	3	2	1	0	0
3	2	1	1	0	0

Input image

*

1	0	1
0	1	0
1	0	1

Filter

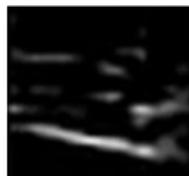
=

6	8	6	3	1	0
9	13	10	5	2	0
9	14	11	6	3	0
9	13	11	6	2	0
8	13	10	5	3	0
6	7	5	3	1	0

Feature map

13	10	2
14	11	3
13	10	3

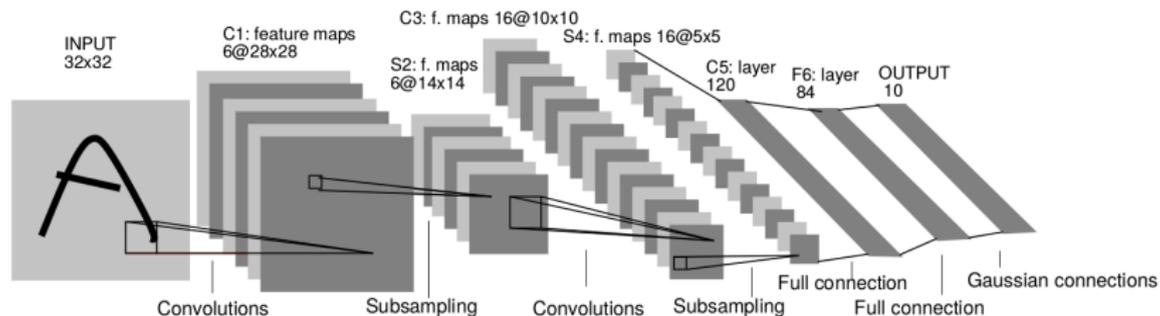
Max pooling



2x2 Max pooling
→

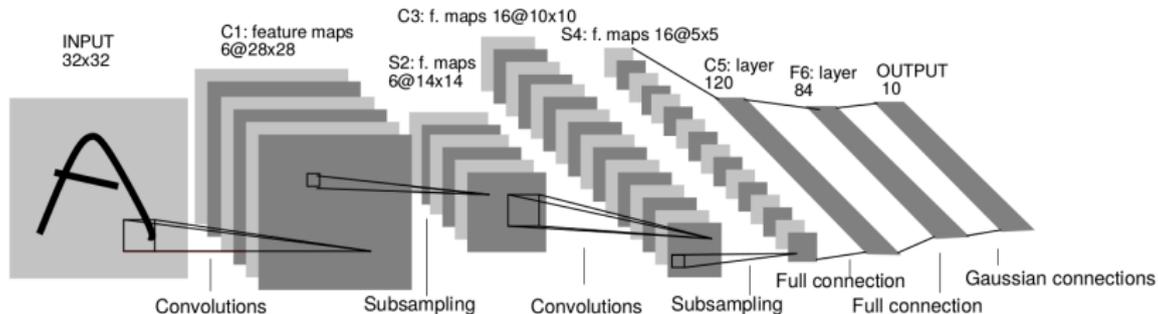


Key properties of CNNs



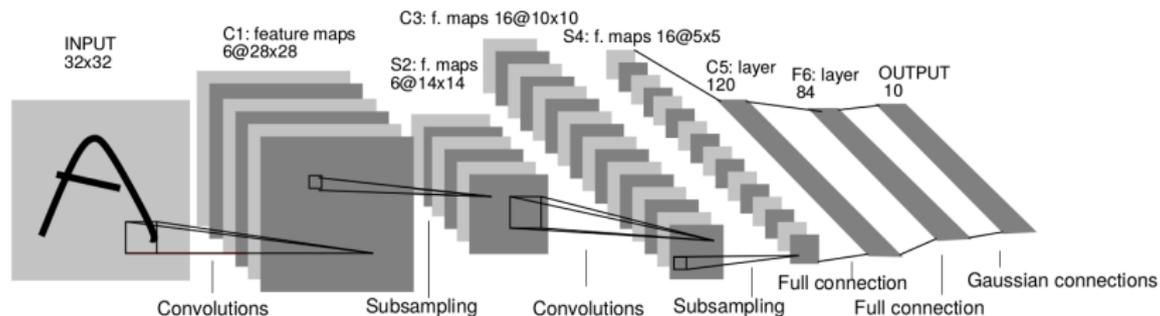
😊 Convolutional filters (**Translation invariance**)

Key properties of CNNs



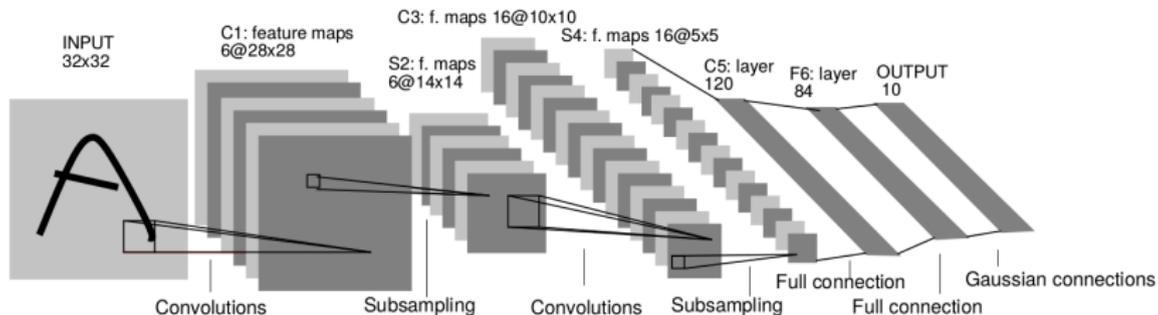
- 😊 Convolutional filters (**Translation invariance**)
- 😊 Multiple layers (**Compositionality**)

Key properties of CNNs



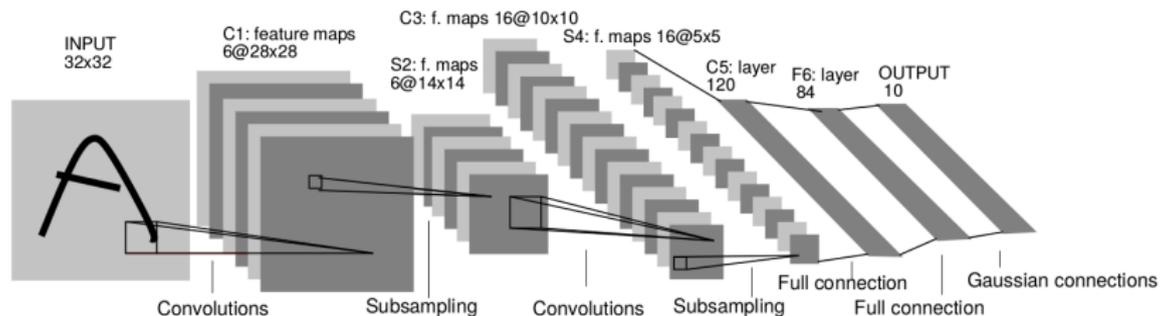
- 😊 Convolutional filters (**Translation invariance**)
- 😊 Multiple layers (**Compositionality**)
- 😊 Filters localized in space (**Locality**)

Key properties of CNNs



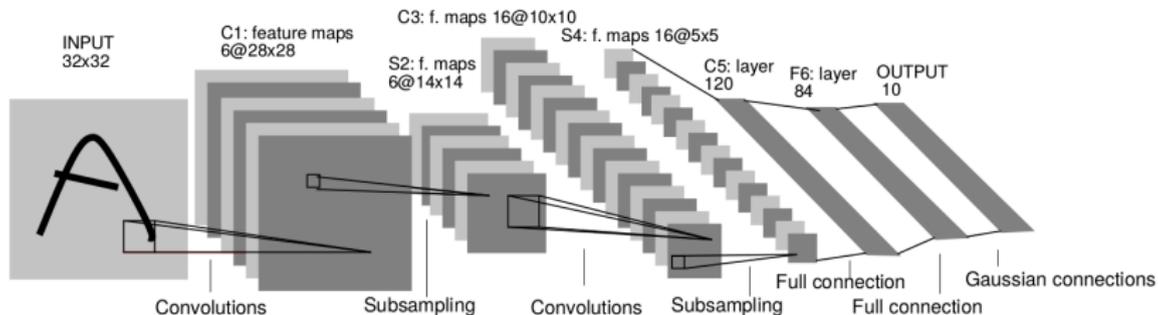
- 😊 Convolutional filters (**Translation invariance**)
- 😊 Multiple layers (**Compositionality**)
- 😊 Filters localized in space (**Locality**)
- 😊 Weight sharing (**Self-similarity**)

Key properties of CNNs



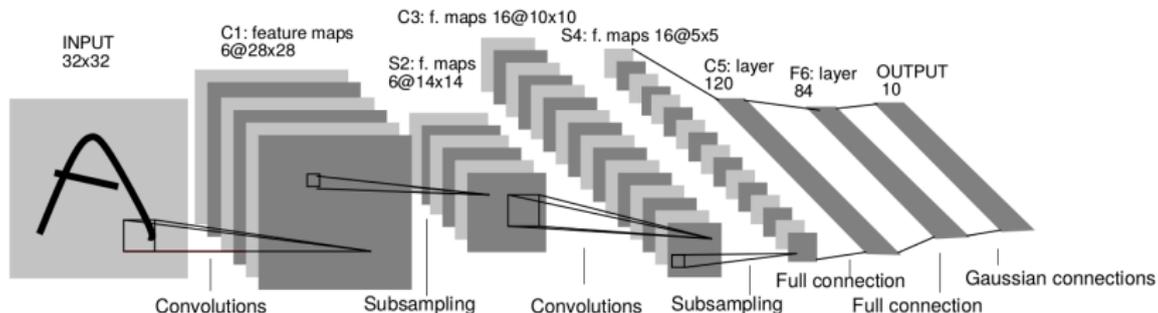
- ☺ Convolutional filters (**Translation invariance**)
- ☺ Multiple layers (**Compositionality**)
- ☺ Filters localized in space (**Locality**)
- ☺ Weight sharing (**Self-similarity**)
- ☺ $\mathcal{O}(1)$ parameters per filter (independent of input image size n)

Key properties of CNNs



- ☺ Convolutional filters (**Translation invariance**)
- ☺ Multiple layers (**Compositionality**)
- ☺ Filters localized in space (**Locality**)
- ☺ Weight sharing (**Self-similarity**)
- ☺ $\mathcal{O}(1)$ parameters per filter (independent of input image size n)
- ☺ $\mathcal{O}(n)$ complexity per layer (filtering done in the spatial domain)

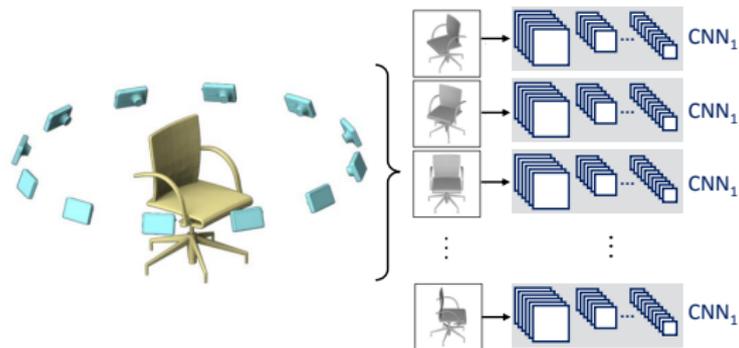
Key properties of CNNs



- ☺ Convolutional filters (**Translation invariance**)
- ☺ Multiple layers (**Compositionality**)
- ☺ Filters localized in space (**Locality**)
- ☺ Weight sharing (**Self-similarity**)
- ☺ $\mathcal{O}(1)$ parameters per filter (independent of input image size n)
- ☺ $\mathcal{O}(n)$ complexity per layer (filtering done in the spatial domain)
- ☺ $\mathcal{O}(\log n)$ layers in classification tasks

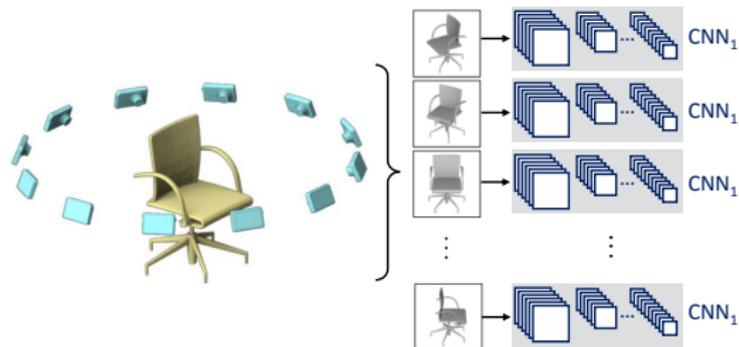
Extrinsic Methods

Multi-view CNNs



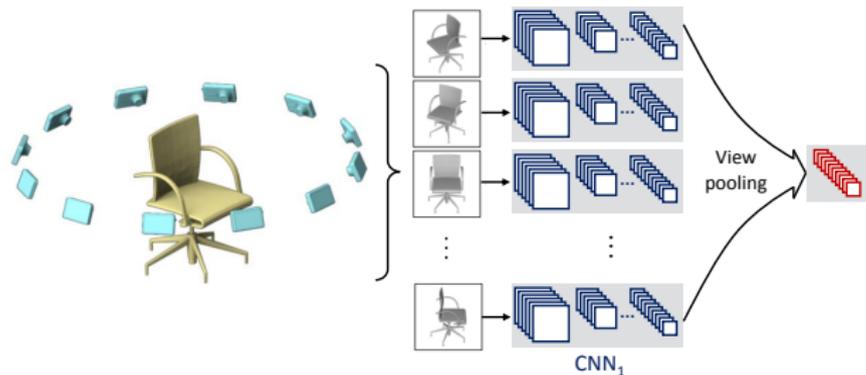
- Represent 3D object as a collection of range images from different views

Multi-view CNNs



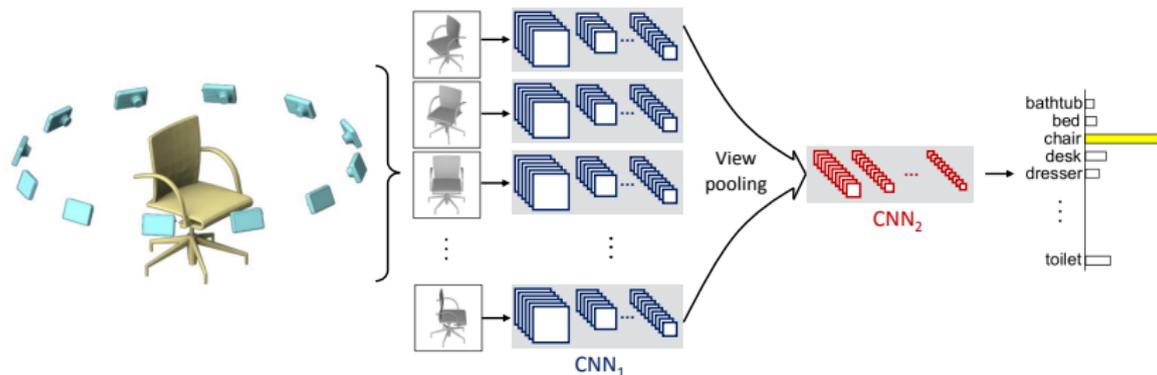
- Represent 3D object as a collection of range images from different views
- CNN_1 : Extract image features (parameters are shared across views)

Multi-view CNNs



- Represent 3D object as a collection of range images from different views
- CNN_1 : Extract image features (parameters are shared across views)
- Element-wise max pooling across all views

Multi-view CNNs



- Represent 3D object as a collection of range images from different views
- CNN₁: Extract image features (parameters are shared across views)
- Element-wise max pooling across all views
- CNN₂: Produce shape descriptors + final prediction

Applications of Multi-view CNNs

- 3D shape classification and retrieval
 - Pre-trained on ImageNet
 - Fine-tuned on 2D views



Applications of Multi-view CNNs

- 3D shape classification and retrieval

- Pre-trained on ImageNet
- Fine-tuned on 2D views



classify
⇒

“chair”

- Sketch classification

- Mimic views by jittering



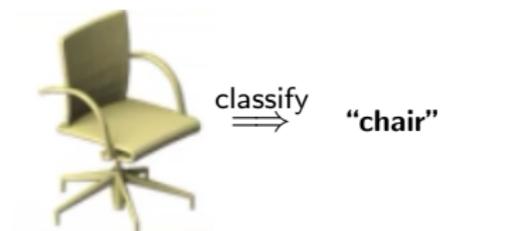
classify
⇒

“chair”

Applications of Multi-view CNNs

- 3D shape classification and retrieval

- Pre-trained on ImageNet
- Fine-tuned on 2D views



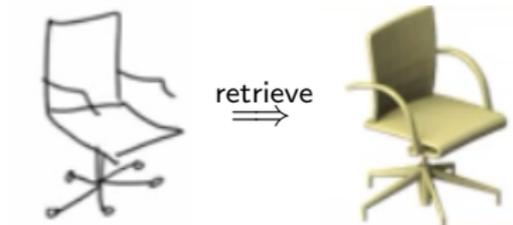
- Sketch classification

- Mimic views by jittering



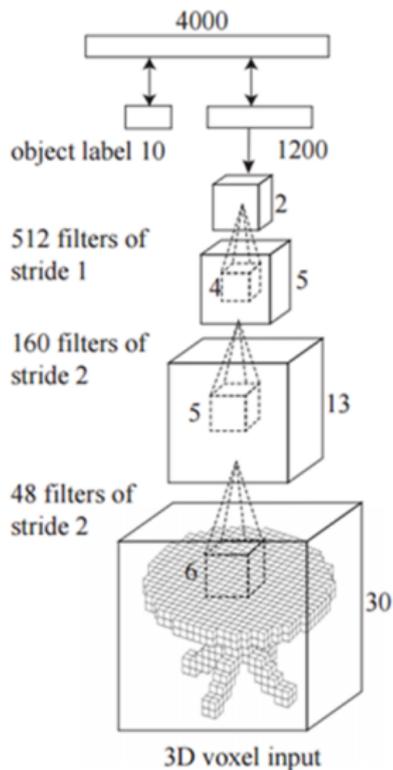
- Sketch-based shape retrieval

- Render views with hand-drawn style (edge maps)



3D ShapeNets

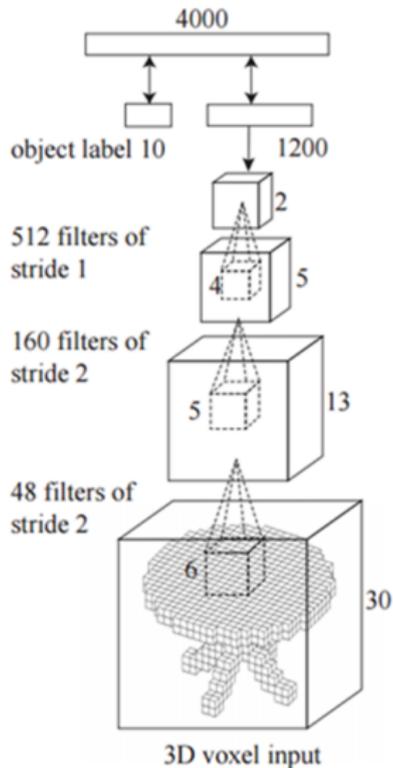
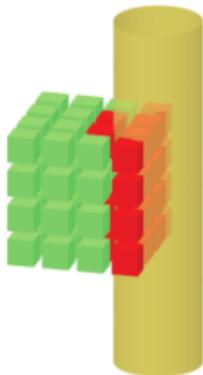
- Volumetric representation (shape = binary voxels on 3D grid)



Convolutional deep belief network

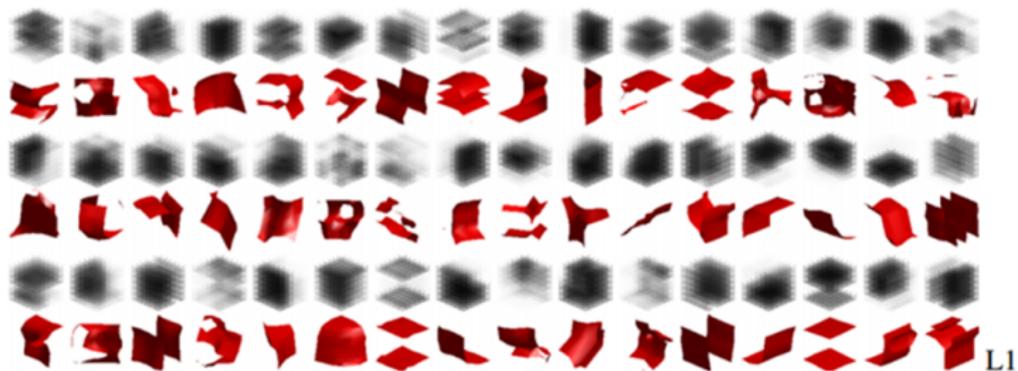
3D ShapeNets

- Volumetric representation (shape = binary voxels on 3D grid)
- 3D convolutional network

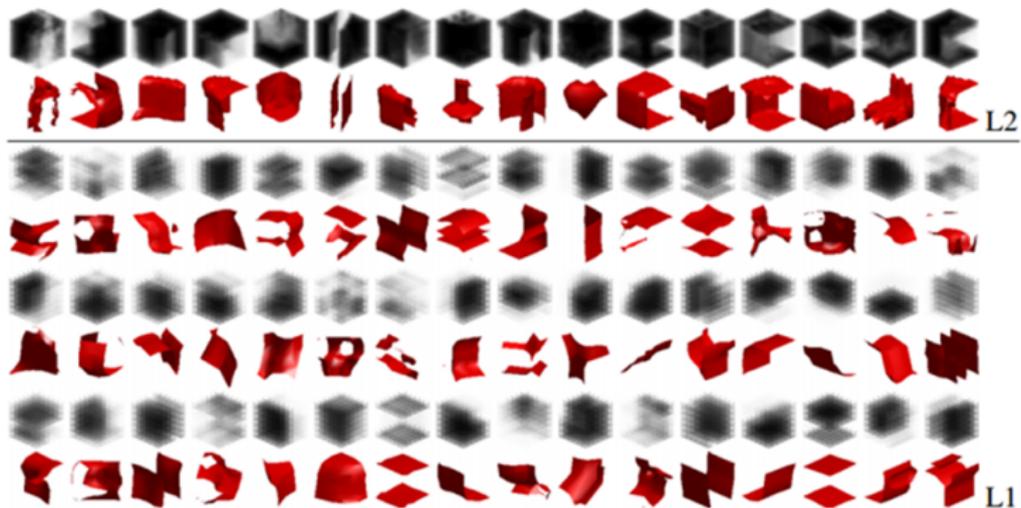


Convolutional deep belief network

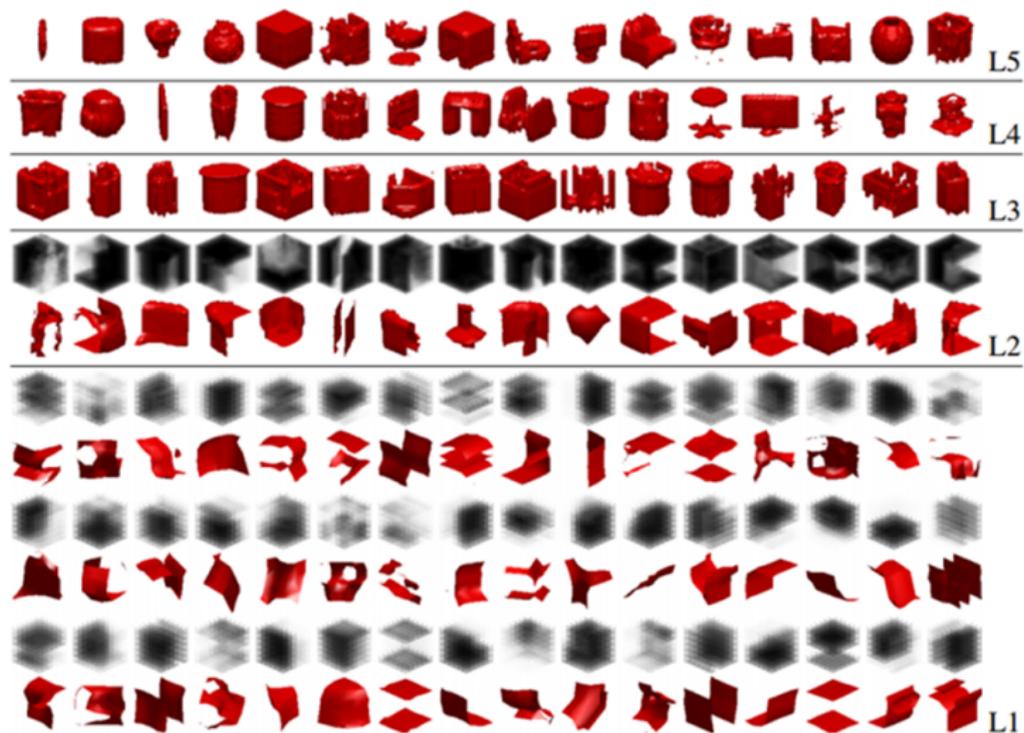
Learned features: 3D primitives



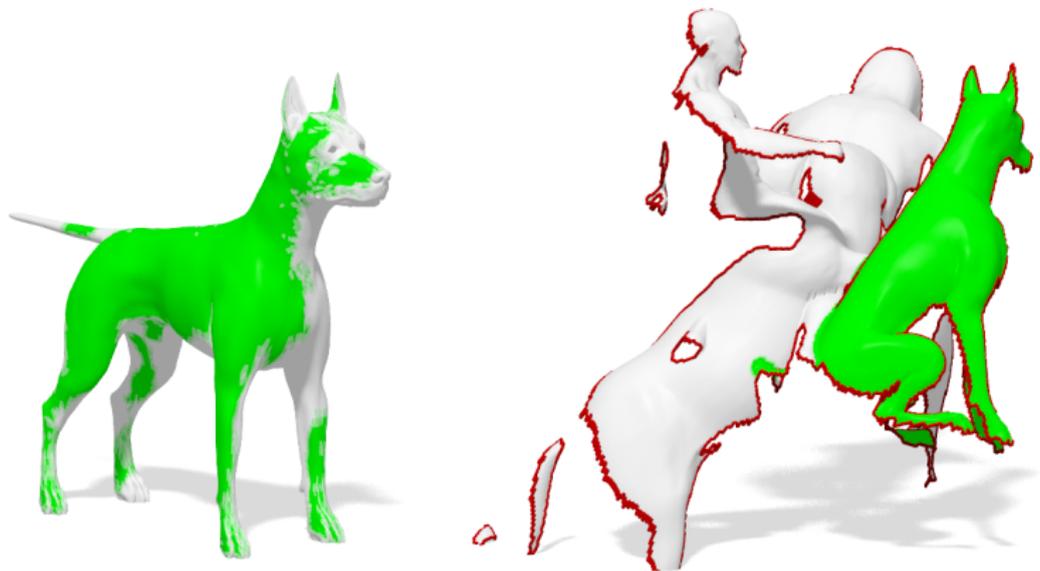
Learned features: 3D primitives



Learned features: 3D primitives

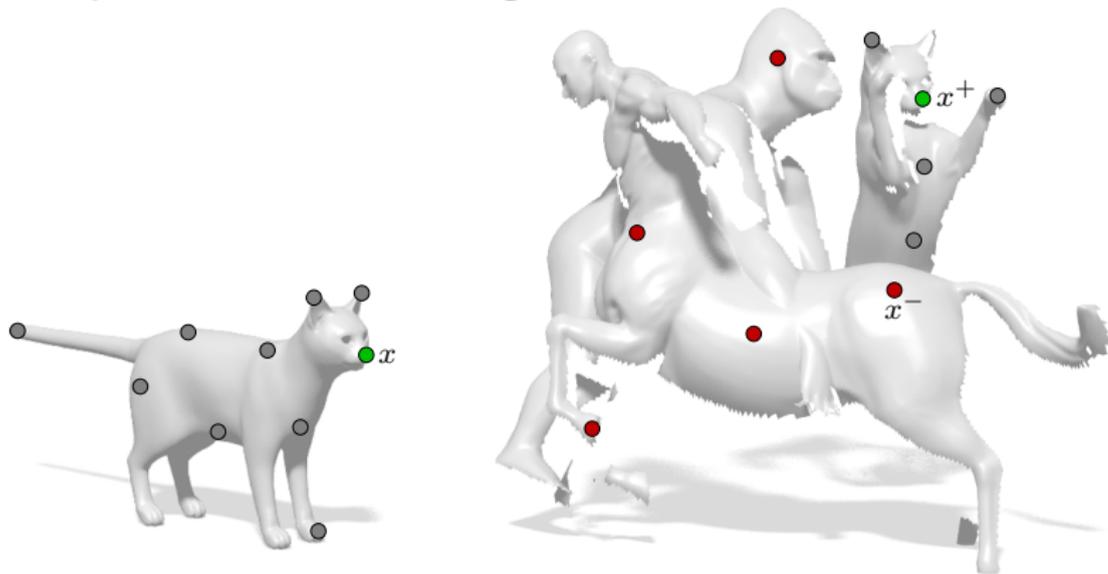


Non-rigid clutter



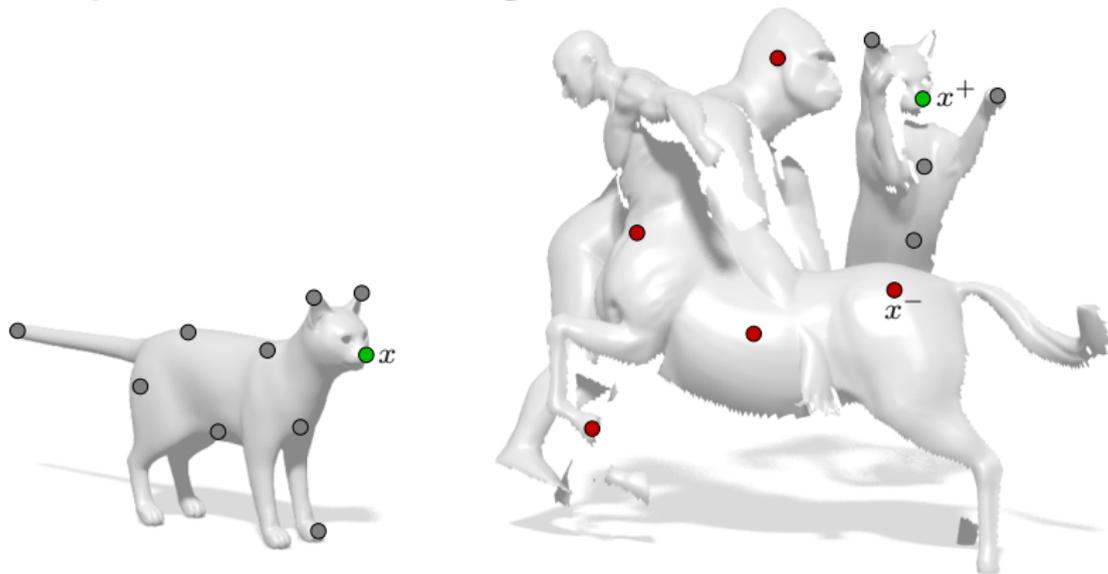
Clutter, missing parts, and non-rigid deformations

Descriptor metric learning



Training set consisting of pairs of **positives** (x, x^+) and **negatives** (x, x^-)

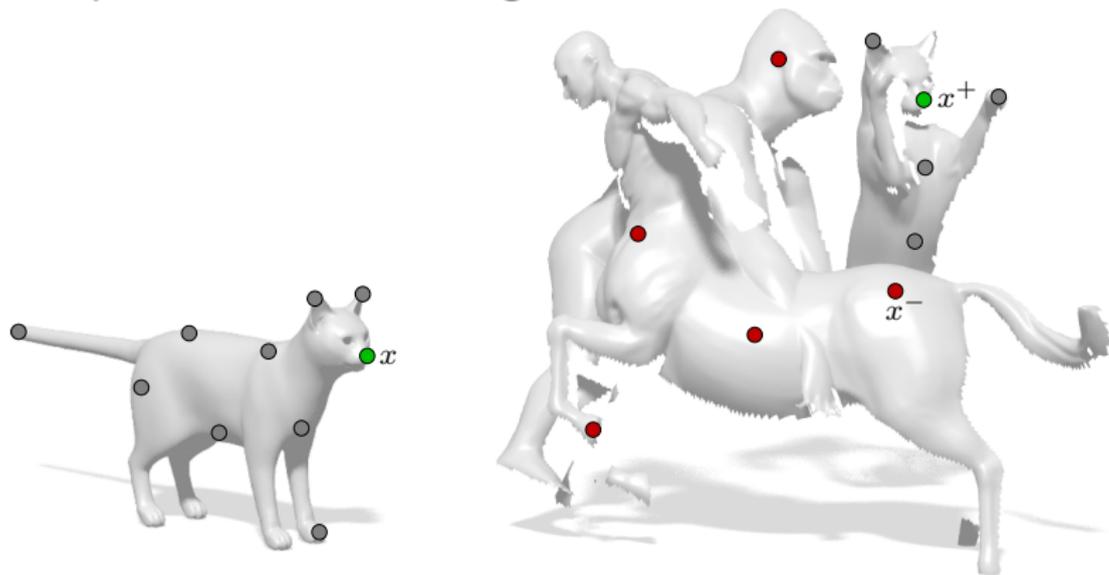
Descriptor metric learning



Training set consisting of pairs of **positives** (x, x^+) and **negatives** (x, x^-)

Each point x has an associated **local descriptor** $g(x)$

Descriptor metric learning

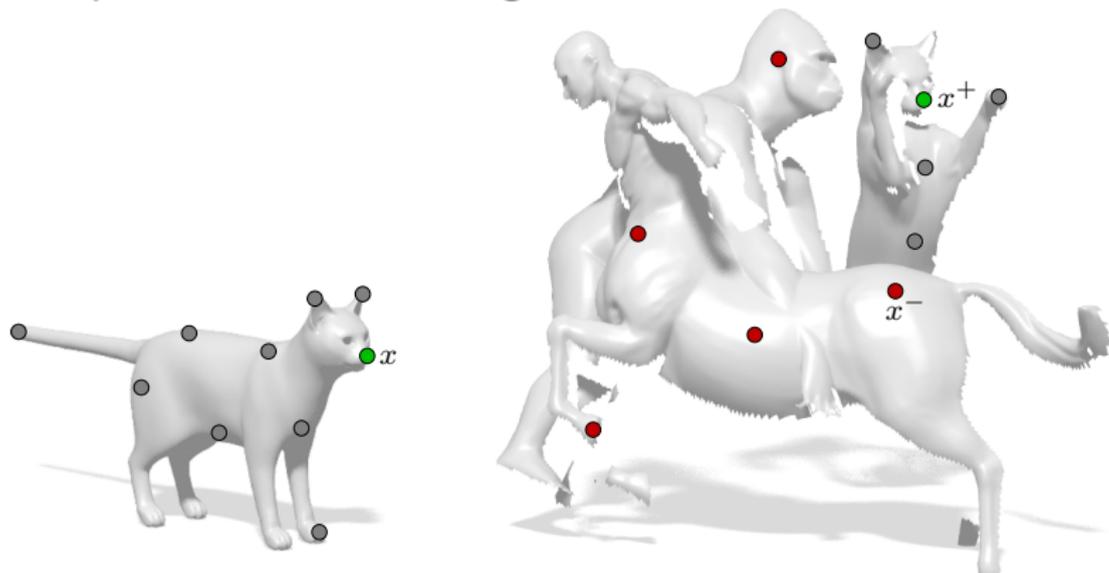


Training set consisting of pairs of **positives** (x, x^+) and **negatives** (x, x^-)

Each point x has an associated **local descriptor** $\mathbf{g}(x)$

Goal: learn a **metric** $d(\mathbf{g}(x), \mathbf{g}(x'))$ between the descriptors representing as reliably as possible the similarity/dissimilarity of positives/negatives

Descriptor metric learning



Training set consisting of pairs of **positives** (x, x^+) and **negatives** (x, x^-)

Each point x has an associated **local descriptor** $g(x)$

Goal: learn a **metric** s.t. $d(g(x), g(x^+)) \approx 0$ and $d(g(x), g(x^-)) \gg 0$

Descriptor metric learning

Parametrize the metric by applying a deep net \mathbf{f}_{Θ} to the descriptor

$$d(\mathbf{g}(x), \mathbf{g}(x')) = \|\mathbf{f}_{\Theta}(\mathbf{g}(x)) - \mathbf{f}_{\Theta}(\mathbf{g}(x'))\|_2$$

Descriptor metric learning

Parametrize the metric by applying a deep net \mathbf{f}_Θ to the descriptor

$$d(\mathbf{g}(x), \mathbf{g}(x')) = \|\mathbf{f}_\Theta(\mathbf{g}(x)) - \mathbf{f}_\Theta(\mathbf{g}(x'))\|_2$$

Learn network parameters Θ by minimizing the [siamese loss](#)

$$\begin{aligned} \mathcal{L}_s(\Theta) = & \sum_{x, x^+ \in \mathcal{T}^+} \gamma \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^+)\|_2^2 \\ & + \sum_{x, x^- \in \mathcal{T}^-} (1 - \gamma) (m_s - \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^-)\|_2)_+^2 \end{aligned}$$

Descriptor metric learning

Parametrize the metric by applying a deep net \mathbf{f}_Θ to the descriptor

$$d(\mathbf{g}(x), \mathbf{g}(x')) = \|\mathbf{f}_\Theta(\mathbf{g}(x)) - \mathbf{f}_\Theta(\mathbf{g}(x'))\|_2$$

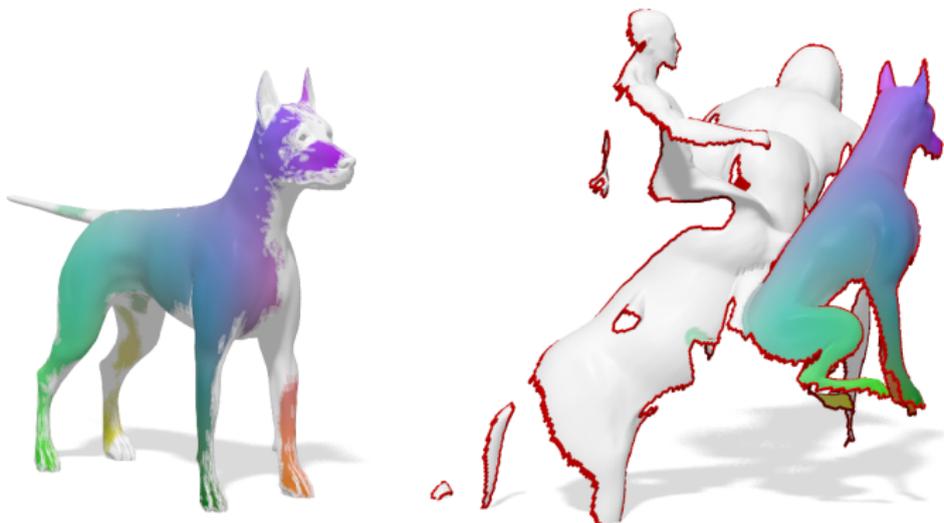
Learn network parameters Θ by minimizing the [siamese loss](#)

$$\begin{aligned} \mathcal{L}_s(\Theta) = & \sum_{x, x^+ \in \mathcal{T}^+} \gamma \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^+)\|_2^2 \\ & + \sum_{x, x^- \in \mathcal{T}^-} (1 - \gamma) (m_s - \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^-)\|_2)_+^2 \end{aligned}$$

Regularize with [global distribution](#) penalty

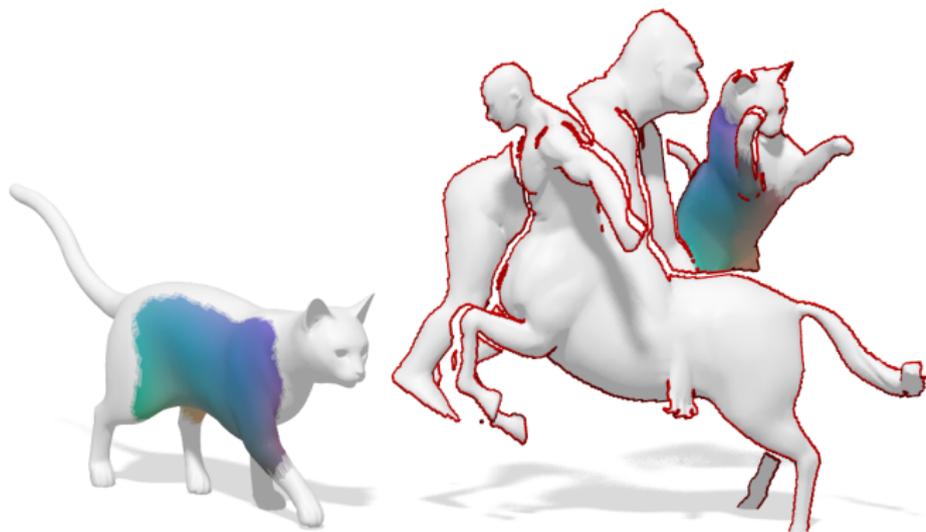
$$\mathcal{L}_g(\Theta) = \sigma_{\mathbf{g}}^+ + \sigma_{\mathbf{g}}^- + (m_g + \mu_{\Theta}^+ - \mu_{\Theta}^-)_+$$

Cluttered matching examples



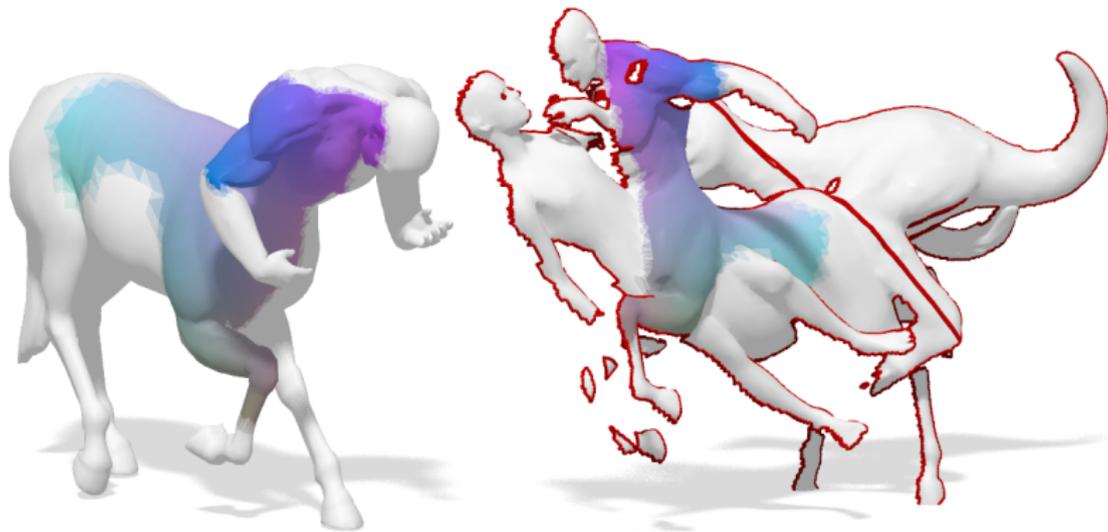
Correspondence examples (corresponding points are marked with same color)

Cluttered matching examples



Correspondence examples (corresponding points are marked with same color)

Cluttered matching examples



Correspondence examples (corresponding points are marked with same color)

Going non-Euclidean

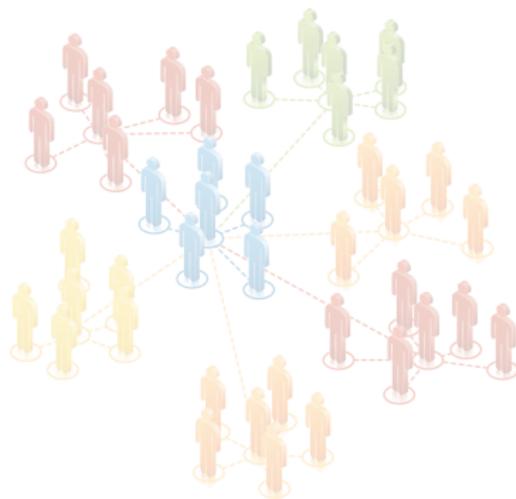
Challenges of geometric deep learning

- Non-Euclidean data often have local, self-similar, and hierarchical structure
- How to define convolution?
- How to do pooling?
- How to work fast?

Prototypical non-Euclidean objects

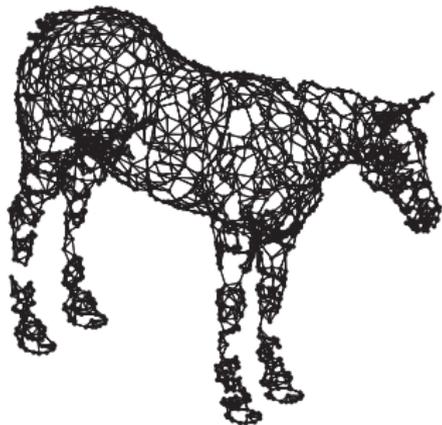


Manifolds



Graphs

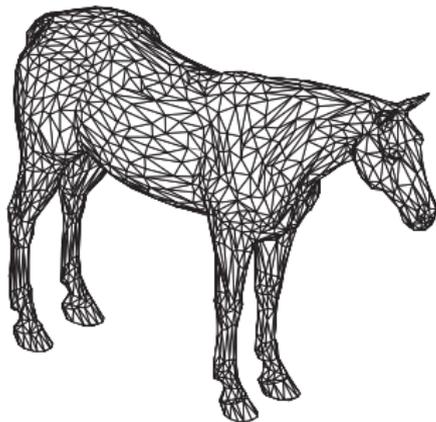
Discrete manifolds



Nearest neighbor graph

Vertices $\mathcal{V} = \{1, \dots, n\}$

Edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$



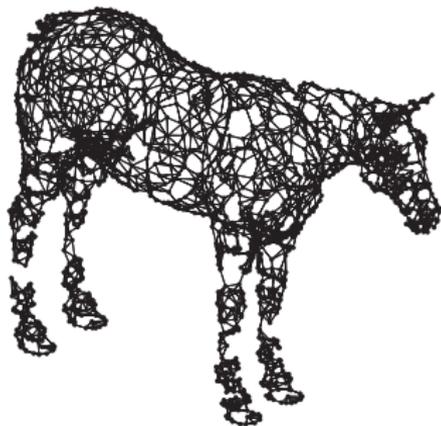
Triangular mesh

Vertices $\mathcal{V} = \{1, \dots, n\}$

Edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$

Faces $\mathcal{F} = \{(i, j, k) \in \mathcal{V} \times \mathcal{V} \times \mathcal{V} : (i, j), (j, k), (k, i) \in \mathcal{E}\}$

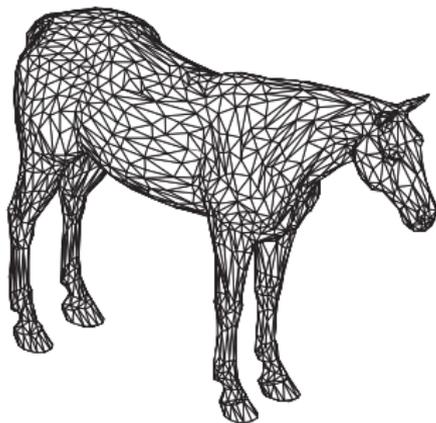
Discrete manifolds



Nearest neighbor graph

Vertices $\mathcal{V} = \{1, \dots, n\}$

Edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$



Triangular mesh

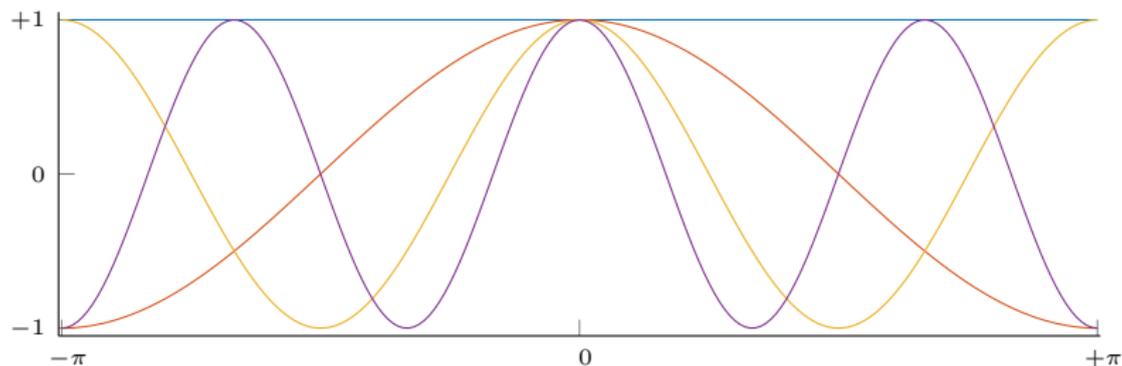
Vertices $\mathcal{V} = \{1, \dots, n\}$

Edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$

Faces $\mathcal{F} = \{(i, j, k) \in \mathcal{V} \times \mathcal{V} \times \mathcal{V} : (i, j), (j, k), (k, i) \in \mathcal{E}\}$

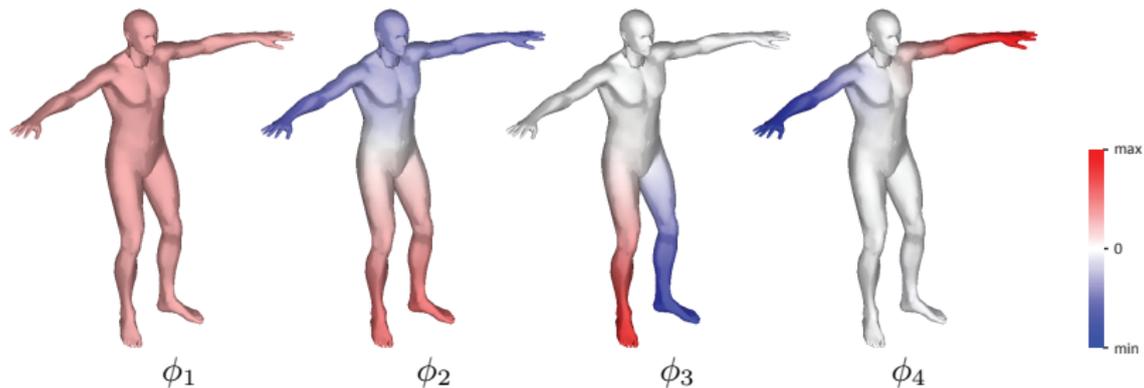
Manifold mesh = each edge is shared by 2 faces + each vertex has 1 loop

Laplacian eigenfunctions: Euclidean



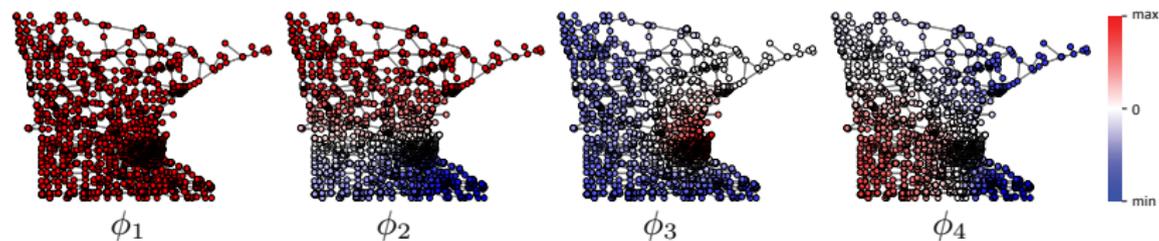
First eigenfunctions of 1D Euclidean Laplacian = standard Fourier basis

Laplacian eigenfunctions: manifold



First eigenfunctions of a manifold Laplacian

Laplacian eigenfunctions: graph

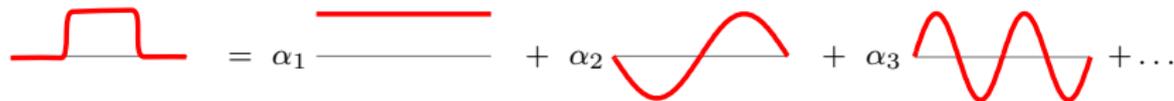


First eigenfunctions of a graph Laplacian

Fourier analysis: Euclidean space

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

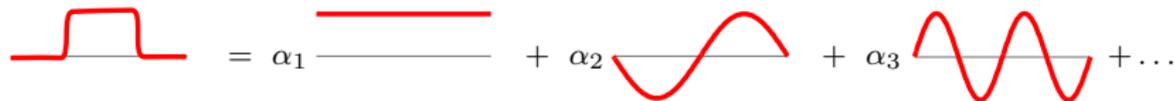
$$f(x) = \sum_{k \geq 0} \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x') e^{-ikx'} dx' e^{ikx}$$



Fourier analysis: Euclidean space

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

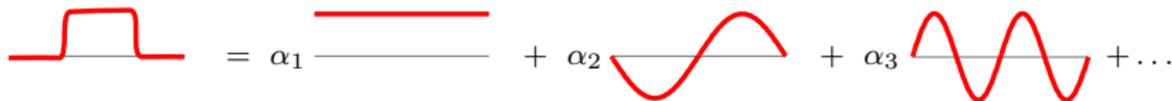
$$f(x) = \sum_{k \geq 0} \frac{1}{2\pi} \underbrace{\int_{-\pi}^{\pi} f(x') e^{-ikx'} dx'}_{\hat{f}_k = \langle f, e^{ikx} \rangle_{L^2([-\pi, \pi])}} e^{ikx}$$



Fourier analysis: Euclidean space

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

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

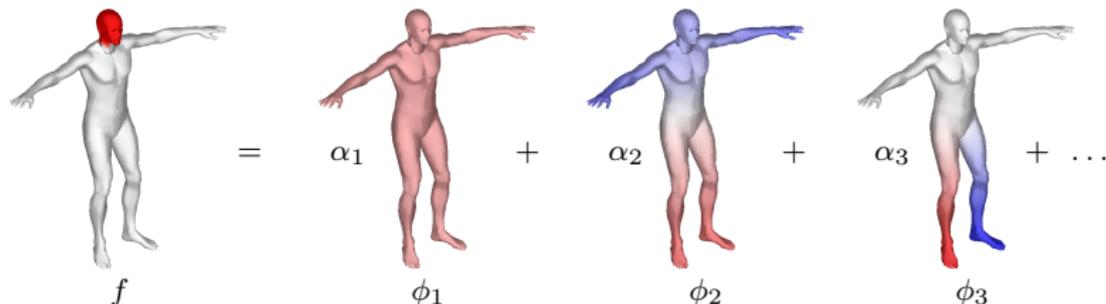


Fourier basis = **Laplacian eigenfunctions**: $-\frac{d^2}{dx^2} e^{ikx} = k^2 e^{ikx}$

Fourier analysis: non-Euclidean space

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: discrete case

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

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

Convolution Theorem: discrete case

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

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

Convolution Theorem: discrete case

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

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

Convolution Theorem: discrete case

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

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

Convolution Theorem: discrete case

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

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

Convolution Theorem: discrete case

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

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

Spectral convolution

Generalized 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

Generalized 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

Generalized 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

Generalized 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

Generalized 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

Generalized 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

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

Spectral convolution

Generalized 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

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

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

Spectral convolution

Generalized 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)
- Commutes with Laplacian: $\mathbf{G}\Delta\mathbf{f} = \Delta\mathbf{G}\mathbf{f}$

Spectral convolution

Generalized 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

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

- Not shift-invariant! (\mathbf{G} has no circulant structure)
- Commutes with Laplacian: $\mathbf{G}\Delta\mathbf{f} = \Delta\mathbf{G}\mathbf{f}$
- Filter coefficients depend on basis ϕ_1, \dots, ϕ_n

Spectral domain deep learning methods

Spectral graph CNN

Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \mathbf{\Phi} \mathbf{W}_{l,l'} \mathbf{\Phi}^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where $\mathbf{W}_{l,l} = n \times n$ diagonal matrix of filter coefficients

Spectral graph CNN

Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \Phi \mathbf{W}_{l,l'} \Phi^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where $\mathbf{W}_{l,l} = n \times n$ diagonal matrix of filter coefficients

☹ Filters are basis-dependent \Rightarrow does not generalize across graphs!

Spectral graph CNN

Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \Phi \mathbf{W}_{l,l'} \Phi^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where $\mathbf{W}_{l,l} = n \times n$ diagonal matrix of filter coefficients

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

Spectral graph CNN

Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \mathbf{\Phi} \mathbf{W}_{l,l'} \mathbf{\Phi}^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where $\mathbf{W}_{l,l} = n \times n$ diagonal matrix of filter coefficients

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

Spectral graph CNN

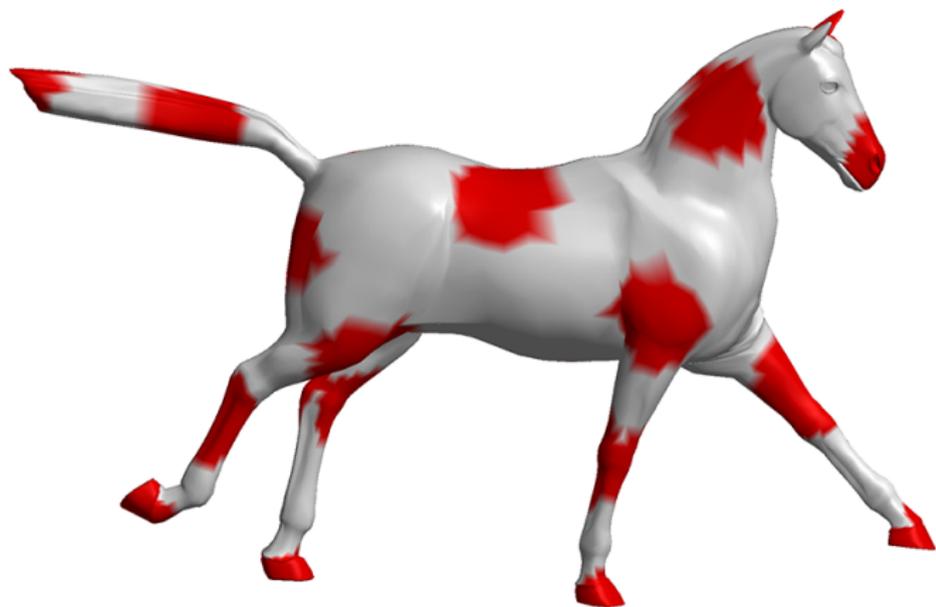
Convolutional layer expressed in the **spectral domain**

$$\mathbf{g}_l = \xi \left(\sum_{l'=1}^p \mathbf{\Phi} \mathbf{W}_{l,l'} \mathbf{\Phi}^\top \mathbf{f}_{l'} \right) \quad \begin{array}{l} l = 1, \dots, q \\ l' = 1, \dots, p \end{array}$$

where $\mathbf{W}_{l,l} = n \times n$ diagonal matrix of filter coefficients

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

Basis dependence



Function f

Basis dependence



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

Basis dependence

Same spectral filter, different basis $\Psi \mathbf{W} \Psi^T \mathbf{f}$

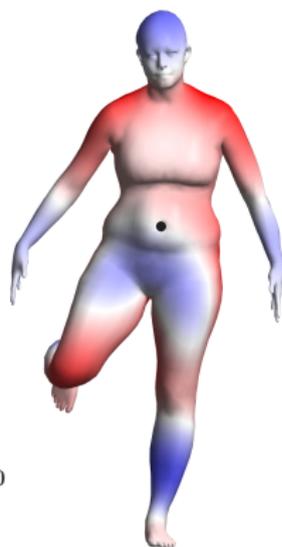
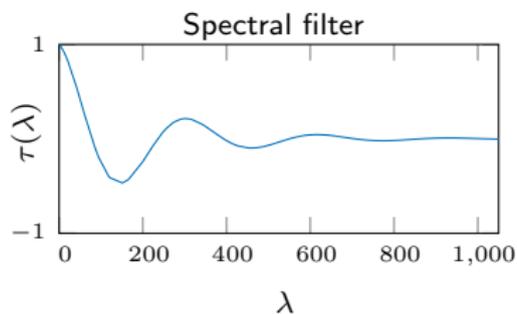
Basis dependence

High-frequency Laplacian eigenvector ϕ_{50}

Filtering in different bases



$$\Phi_1 \tau(\Lambda_1) \Phi_1^T \delta_0$$



$$\Phi_2 \tau(\Lambda_2) \Phi_2^T \delta_0$$

Apply spectral filter $\tau(\lambda)$ in different bases Φ_1 and Φ_2
 \Rightarrow different results!

Filtering in different bases



$$\Phi_1 \tau(\Lambda_1) \Phi_1^T \delta_0$$



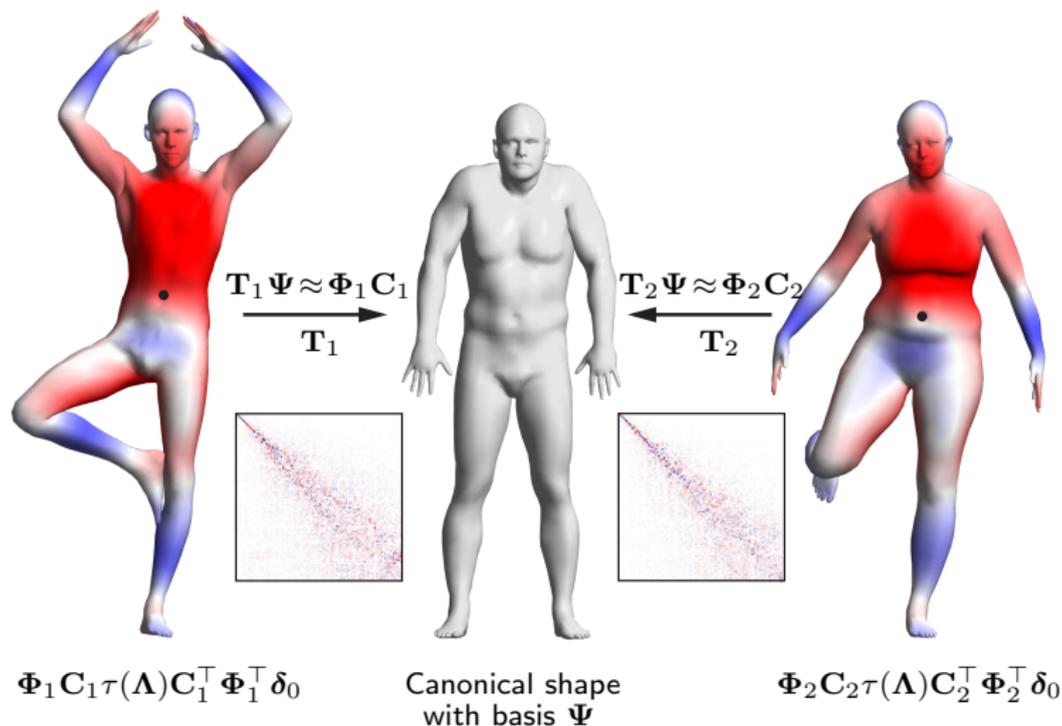
Canonical shape
with basis Ψ



$$\Phi_2 \tau(\Lambda_2) \Phi_2^T \delta_0$$

Apply spectral filter $\tau(\lambda)$ in different bases Φ_1 and Φ_2
 \Rightarrow different results!

Filtering in synchronized bases



Apply spectral filter $\tau(\lambda)$ in synchronized bases $\Phi_1 C_1$ and $\Phi_2 C_2$
 \Rightarrow similar results!

Localization and Smoothness

In the Euclidean setting (by Parseval's identity)

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

Localization and Smoothness

In the Euclidean setting (by Parseval's identity)

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

⇒ Localization in space = smoothness in frequency domain

Localization and Smoothness

In the Euclidean setting (by Parseval's identity)

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

⇒ Localization in space = smoothness in frequency domain

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

Localization and Smoothness

In the Euclidean setting (by Parseval's identity)

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

⇒ Localization in space = smoothness in frequency domain

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

Application of the filter

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

Localization and Smoothness

In the Euclidean setting (by Parseval's identity)

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

⇒ 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

In the Euclidean setting (by Parseval's identity)

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

⇒ 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

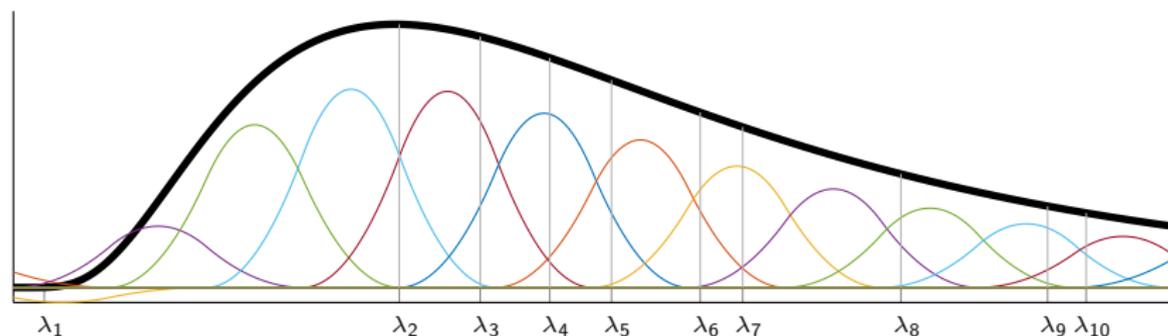
$$\tau_{\alpha}(\Delta)\mathbf{f} = \Phi \begin{pmatrix} \tau_{\alpha}(\lambda_1) & & \\ & \ddots & \\ & & \tau_{\alpha}(\lambda_n) \end{pmatrix} \Phi^{\top} \mathbf{f}$$

Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

$$\tau_{\alpha}(\lambda) = \sum_{j=1}^r \alpha_j \beta_j(\lambda)$$

where $\alpha = (\alpha_1, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

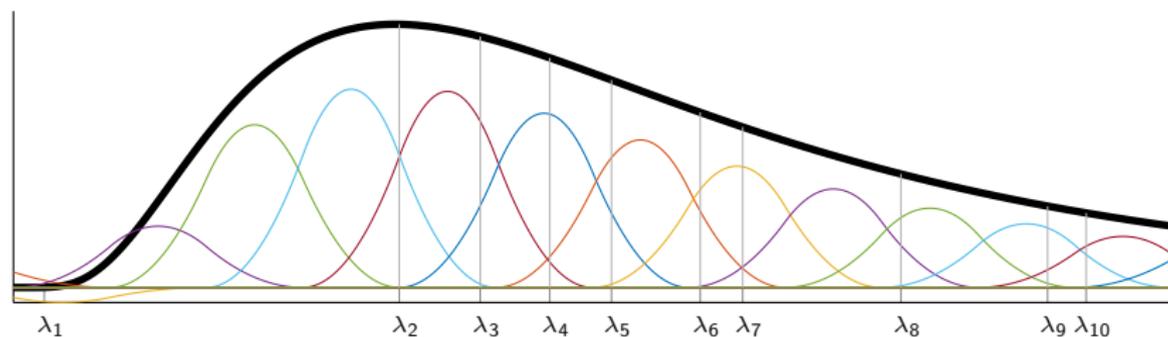


Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

$$\tau_{\alpha}(\lambda_k) = \sum_{j=1}^r \alpha_j \beta_j(\lambda_k)$$

where $\alpha = (\alpha_1, \dots, \alpha_r)^{\top}$ is the vector of filter parameters

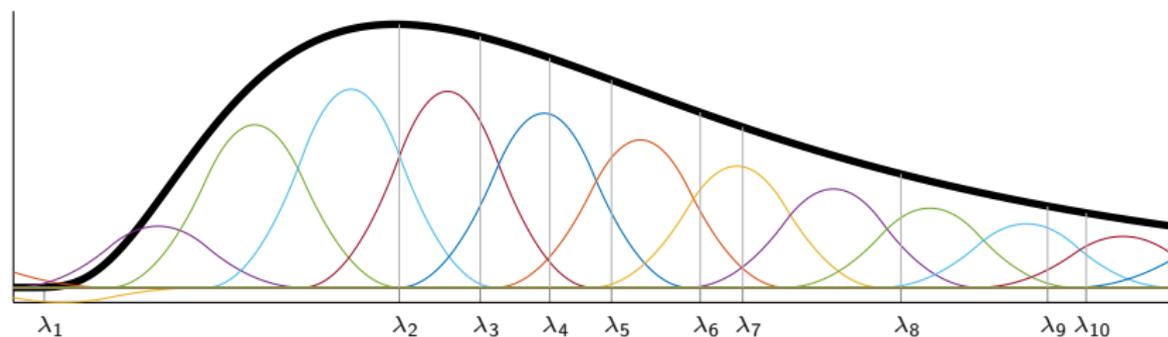


Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

$$\tau_{\alpha}(\lambda_k) = \sum_{j=1}^r \alpha_j \beta_j(\lambda_k) = (\mathbf{B}\alpha)_k$$

where $\alpha = (\alpha_1, \dots, \alpha_r)^\top$ is the vector of filter parameters

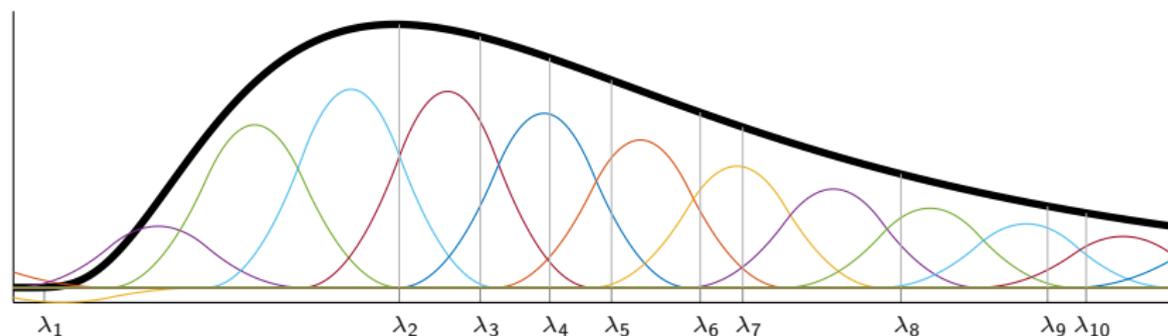


Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

$$\mathbf{W} = \text{Diag}(\mathbf{B}\boldsymbol{\alpha})$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r)^\top$ is the vector of filter parameters



Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

$$\mathbf{W} = \text{Diag}(\mathbf{B}\boldsymbol{\alpha})$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r)^\top$ is the vector of filter parameters

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

Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

$$\mathbf{W} = \text{Diag}(\mathbf{B}\boldsymbol{\alpha})$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r)^\top$ is the vector of filter parameters

- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Fast-decaying filters in space

Spectral graph CNN with smooth spectral filters

Represent spectral transfer function as a **linear combination of smooth kernel functions** $\beta_1(\lambda), \dots, \beta_r(\lambda)$

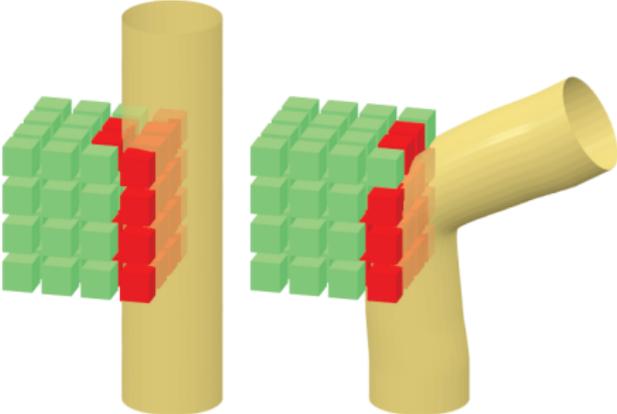
$$\mathbf{W} = \text{Diag}(\mathbf{B}\boldsymbol{\alpha})$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r)^\top$ is the vector of filter parameters

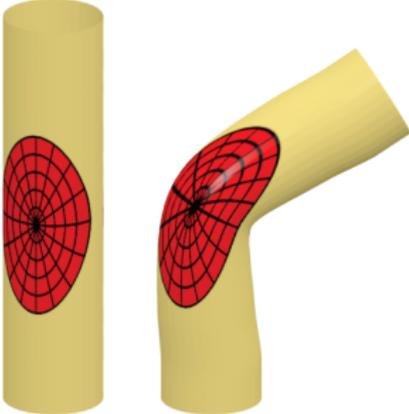
- ☺ $\mathcal{O}(1)$ parameters per layer
- ☺ Fast-decaying filters in space
- ☹ $\mathcal{O}(n^2)$ computation of forward and inverse Fourier transforms Φ^\top, Φ (no FFT on graphs)

Spatial domain deep learning methods

Extrinsic vs Intrinsic



Extrinsic



Intrinsic

Extrinsic vs Intrinsic

Extrinsic

Intrinsic

What is convolution on manifolds?

Convolution

Euclidean

Spatial domain

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

Spectral domain

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

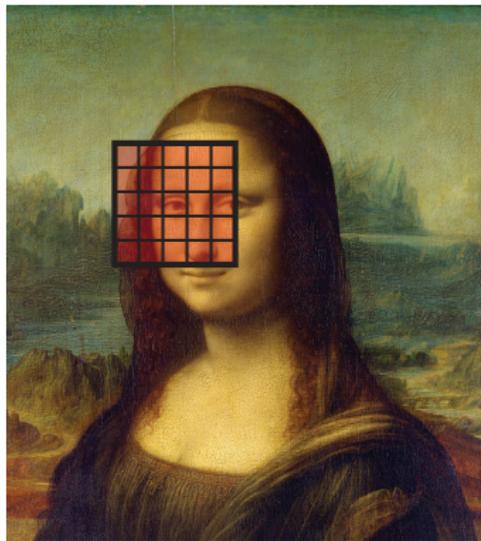
'Convolution Theorem'

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

Convolution in the spatial domain

A	B	C	D	E	A	B	C	D	E
F	G	H	I	J	F	G	H	I	J
K	L	M	N	O	K	L	M	N	O
P	R	S	T	U	P	R	S	T	U
V	W	X	Y	Z	V	W	X	Y	Z
A	B	C	D	E	A	B	C	D	E
F	G	H	I	J	F	G	H	I	J
K	L	M	N	O	K	L	M	N	O
P	R	S	T	U	P	R	S	T	U
V	W	X	Y	Z	V	W	X	Y	Z

Euclidean



Non-Euclidean

- No canonical global system of coordinates

Convolution in the spatial domain

A	B	C	D	E	A	B	C	D	E
F	G	H	I	J	F	G	H	I	J
K	L	M	N	O	K	L	M	N	O
P	R	S	T	U	P	R	S	T	U
V	W	X	Y	Z	V	W	X	Y	Z
A	B	C	D	E	A	B	C	D	E
F	G	H	I	J	F	G	H	I	J
K	L	M	N	O	K	L	M	N	O
P	R	S	T	U	P	R	S	T	U
V	W	X	Y	Z	V	W	X	Y	Z

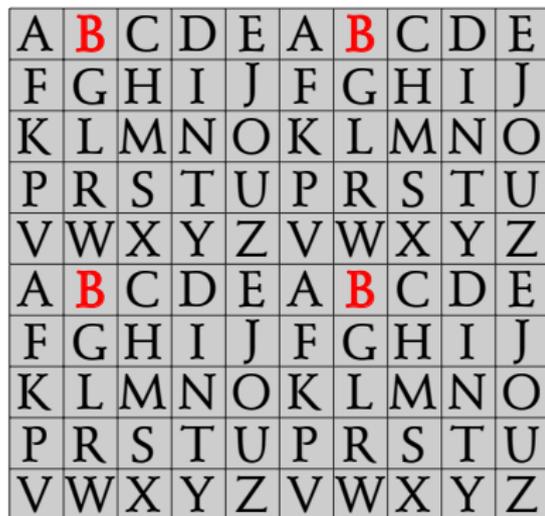
Euclidean



Non-Euclidean

- No canonical global system of coordinates
- No grid structure (no regular memory access)

Convolution in the spatial domain



Euclidean



Non-Euclidean

- No canonical global system of coordinates
- No grid structure (no regular memory access)
- No shift-invariance (patch is position-dependent)

Patch operator

$$(f \star g)(x) = \int \begin{matrix} \text{[Color Map]} \\ (D(x)f)(\mathbf{u}) \end{matrix} \times \begin{matrix} \text{[Color Map]} \\ g(\mathbf{u}) \end{matrix} d\mathbf{u}$$



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'

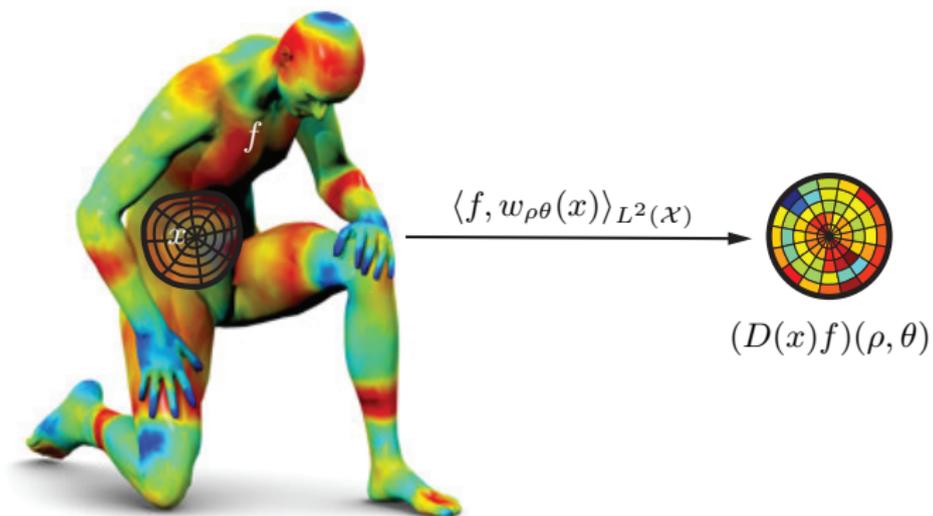
Non-Euclidean

$$(f \star g)(x) = \iint (D(x)f)(\rho, \theta)g(\rho, \theta)d\rho d\theta$$

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

Geodesic CNNs

Patch operator

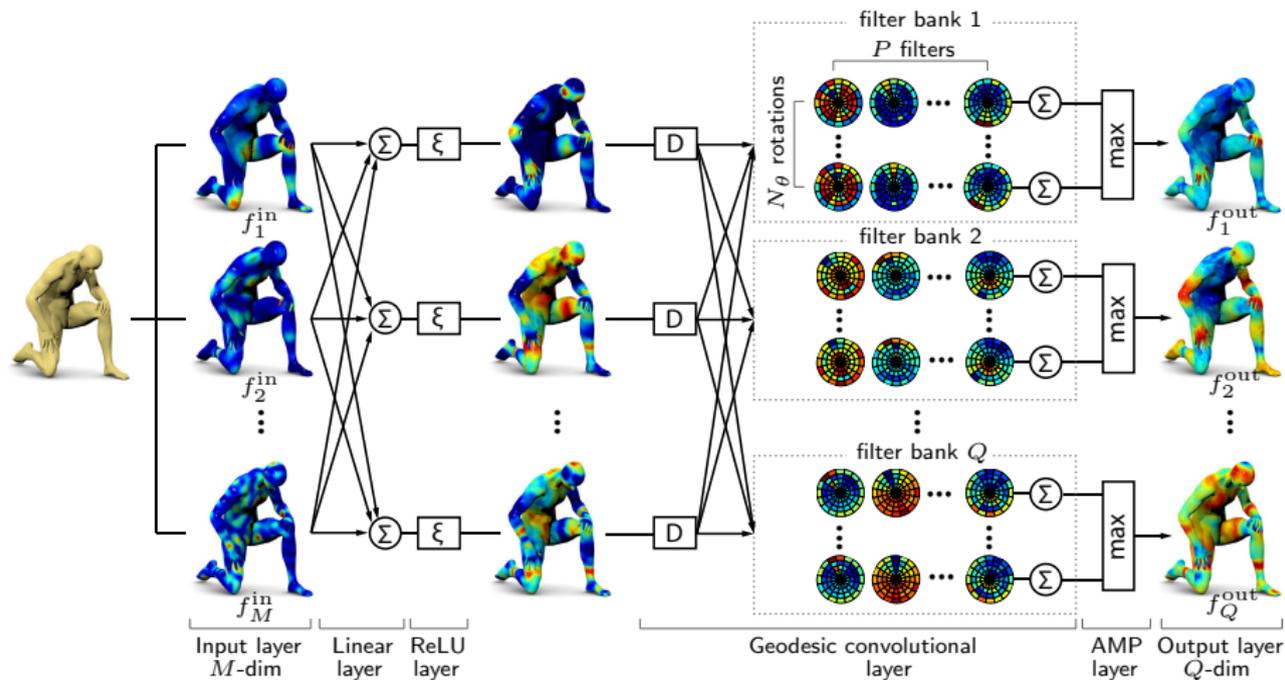


Patch operator

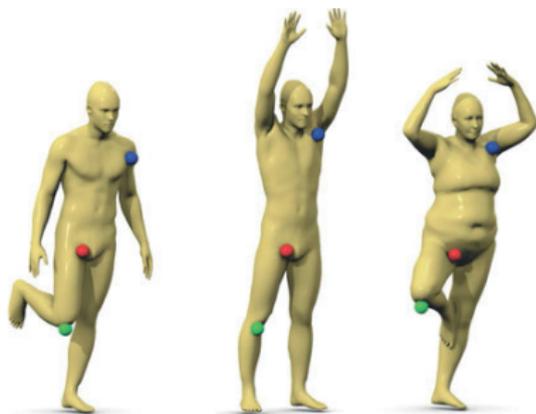


$$(f \star g)(x) = \iint (D(x)f)(\rho, \theta) \cdot g(\rho, \theta) d\rho d\theta$$

Toy Geodesic CNN (GCNN) architecture

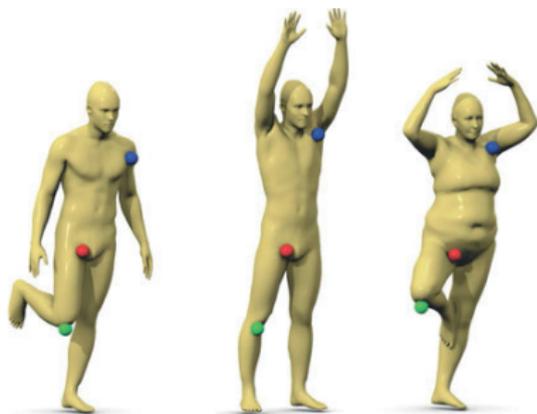


Learning local descriptors with GCNN



- As similar as possible on positives \mathcal{T}^+
- As dissimilar as possible on negatives \mathcal{T}^-

Learning local descriptors with GCNN



- As similar as possible on **positives** \mathcal{T}^+
- As dissimilar as possible on **negatives** \mathcal{T}^-
- Minimize **siamese loss** w.r.t. GCNN parameters Θ

$$\begin{aligned} \ell(\Theta) = & (1 - \gamma) \sum_{(x, x^+) \in \mathcal{T}^+} \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^+)\| \\ & + \gamma \sum_{(x, x^-) \in \mathcal{T}^-} (\mu - \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^-)\|)_+ \end{aligned}$$

Descriptor robustness



HKS descriptor distance

Descriptors: Sun, Ovsjanikov, Guibas 2009 (HKS); Aubry, Schlickewei, Cremers 2011 (WKS); Masci, Boscaini, Bronstein, Vandergheynst 2015 (GCNN); data: Bronstein et al. 2008 (TOSCA); Angelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)

Descriptor robustness



WKS descriptor distance

Descriptors: Sun, Ovsjanikov, Guibas 2009 (HKS); Aubry, Schlickewei, Cremers 2011 (WKS); Masci, Boscaini, Bronstein, Vandergheynst 2015 (GCNN); data: Bronstein et al. 2008 (TOSCA); Angelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)

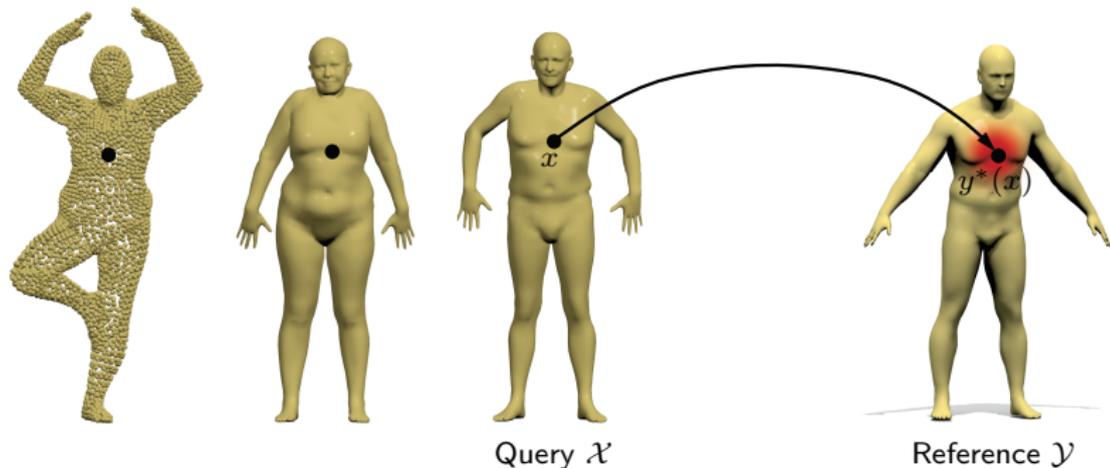
Descriptor robustness



GCNN descriptor distance

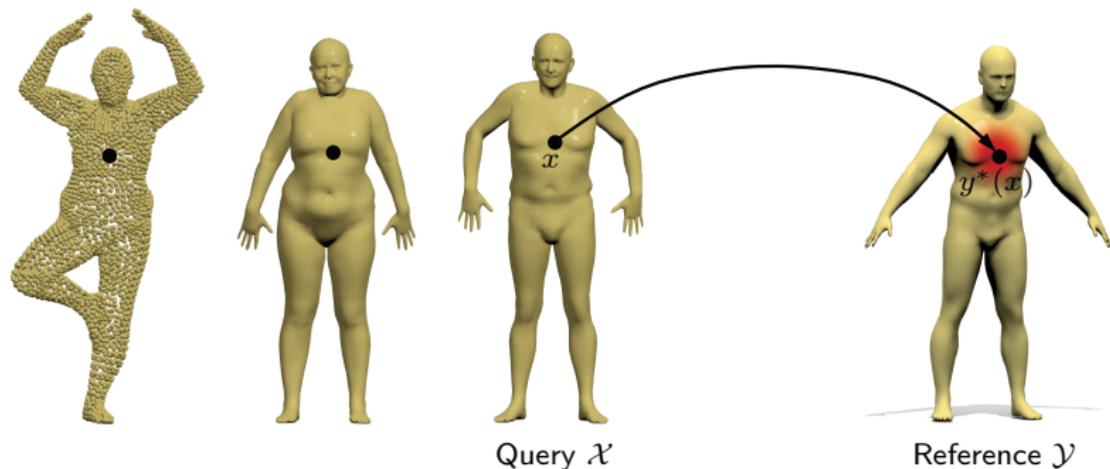
Descriptors: Sun, Ovsjanikov, Guibas 2009 (HKS); Aubry, Schlickewei, Cremers 2011 (WKS); Masci, Boscaini, Bronstein, Vandergheynst 2015 (GCNN); data: Bronstein et al. 2008 (TOSCA); Angelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)

Learning shape correspondence with GCNN



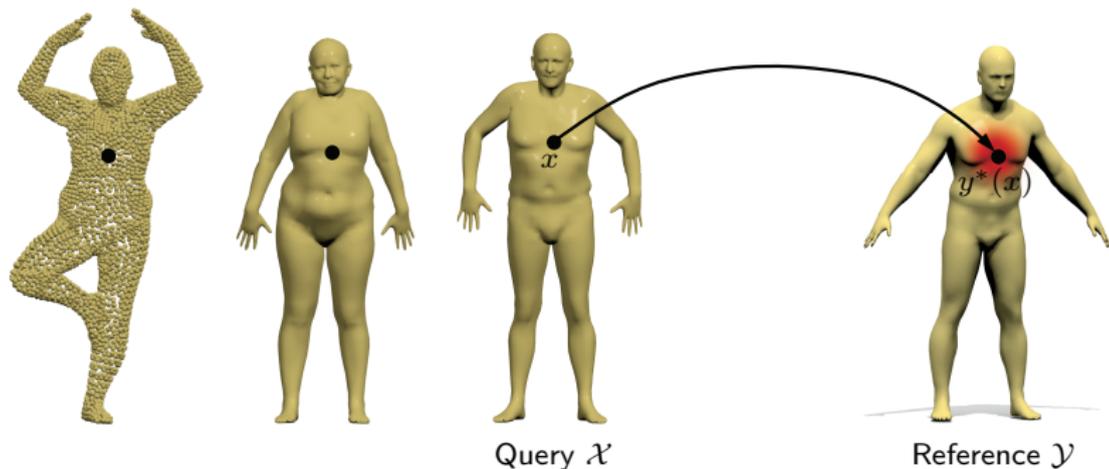
- Correspondence = labeling problem

Learning shape correspondence with GCNN



- Correspondence = labeling problem
- GCNN output $\mathbf{f}_{\Theta}(x)$ = probability distribution on reference \mathcal{Y}

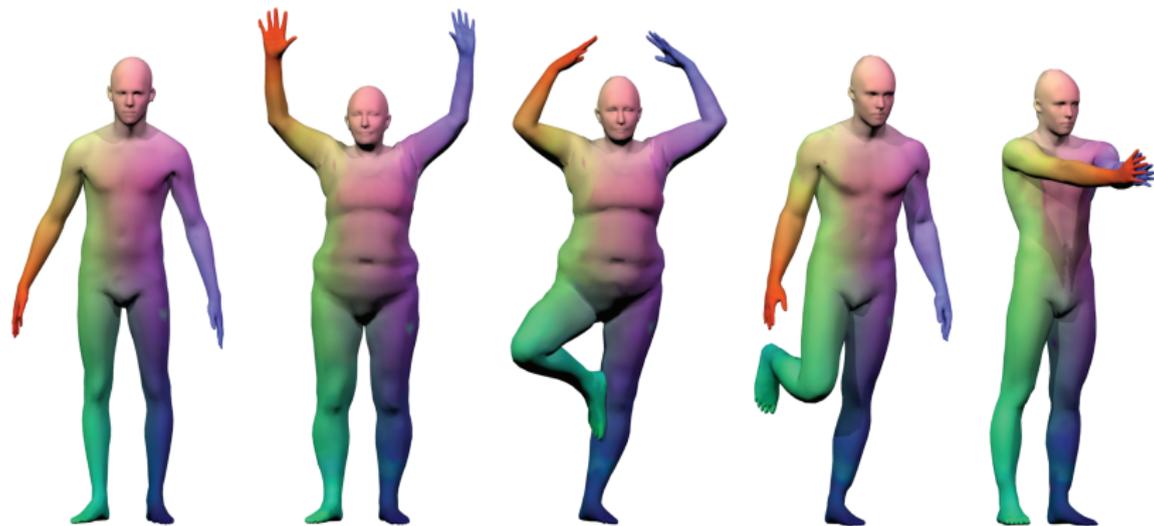
Learning shape correspondence with GCNN



- Correspondence = **labeling problem**
- GCNN output $\mathbf{f}_{\Theta}(x)$ = probability distribution on reference \mathcal{Y}
- Minimize **logistic regression** cost w.r.t. GCNN parameters Θ

$$\ell(\Theta) = - \sum_{(x, y^*(x)) \in \mathcal{T}} \langle \delta_{y^*(x)}, \log \mathbf{f}_{\Theta}(x) \rangle_{L^2(\mathcal{Y})}$$

Correspondence example



Correspondence found using GCNN
(similar colors encode corresponding points)

Anisotropic CNNs

Anisotropic heat kernels

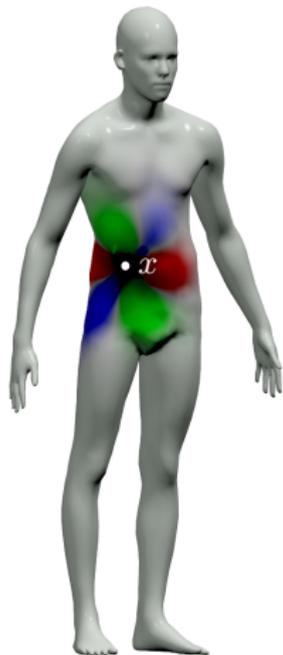
$$h_{\alpha\theta t}(x, x') = \sum_{k \geq 0} e^{-t\lambda_{\alpha\theta k}} \phi_{\alpha\theta k}(x) \phi_{\alpha\theta k}(x')$$

Scale t

Orientation θ

Elongation α

Anisotropic CNN (ACNN)



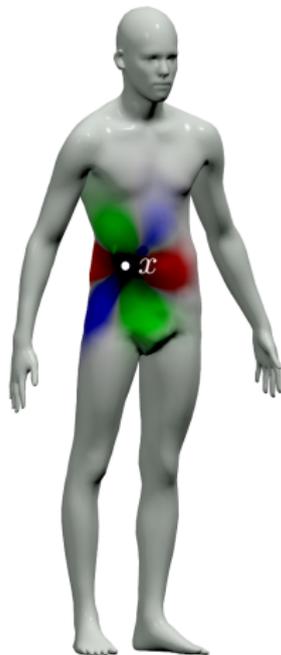
Given a function $f \in L^2(\mathcal{X})$, the patch operator

$$(D(x)f)(\theta, t) = \langle f, h_{\alpha\theta t}(x, \cdot) \rangle_{L^2(\mathcal{X})}$$

produces a local representation of f around point x

- θ = 'angular coordinate'
- t = 'radial coordinate'

Anisotropic CNN (ACNN)



Given a function $f \in L^2(\mathcal{X})$, the patch operator

$$(D(x)f)(\theta, t) = \langle f, h_{\alpha\theta t}(x, \cdot) \rangle_{L^2(\mathcal{X})}$$

produces a local representation of f around point x

- θ = 'angular coordinate'
- t = 'radial coordinate'

Anisotropic convolution

$$(f \star a)(x) = \sum_{\theta, t} (D(x)f)(\theta, t) a(\theta, t)$$

Learned correspondence visualization



Texture transferred from reference to query shapes
using correspondence learned by ACNN

Learned correspondence on point clouds



Colors transferred from reference to query shapes using correspondence learned with ACNN (similar colors encode corresponding points)

Partial correspondence with ACNN



Correspondence



Correspondence error

Partial correspondence with ACNN



Correspondence

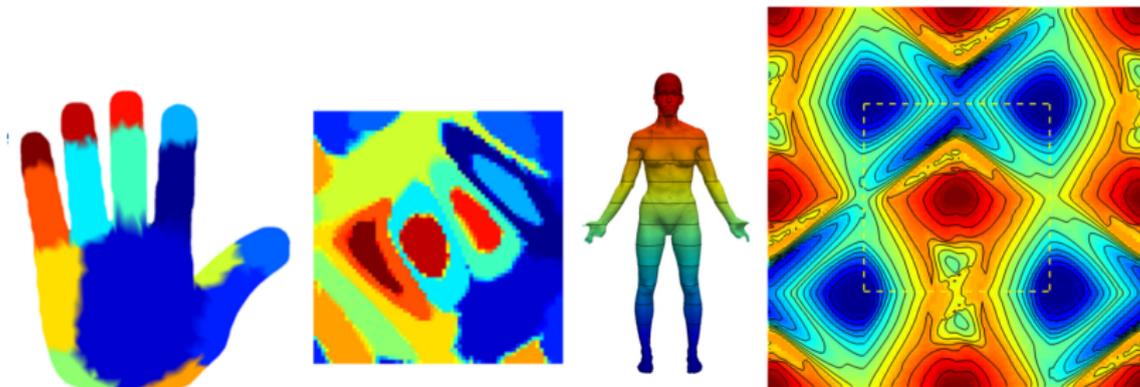


Correspondence error

Embedding domain deep learning methods

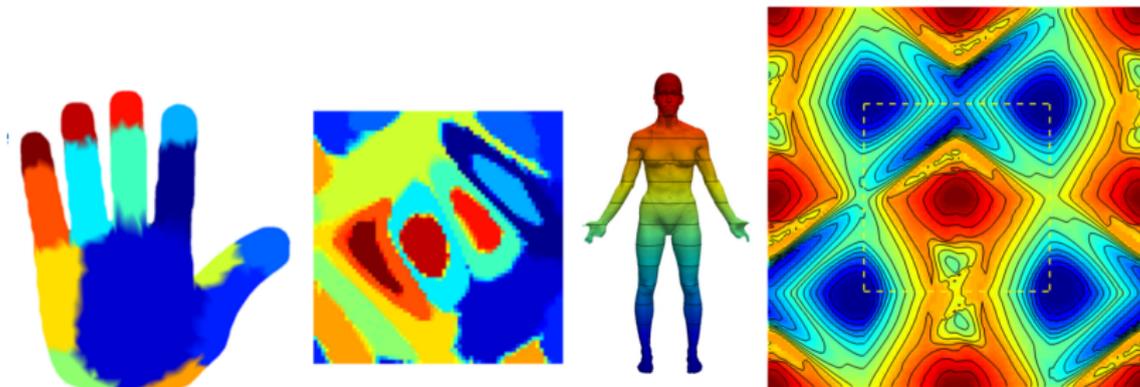
Global parametrization

Key idea: map the input surface to some **parametric domain** (e.g. 2D plane) where operations can be defined more easily.



Global parametrization

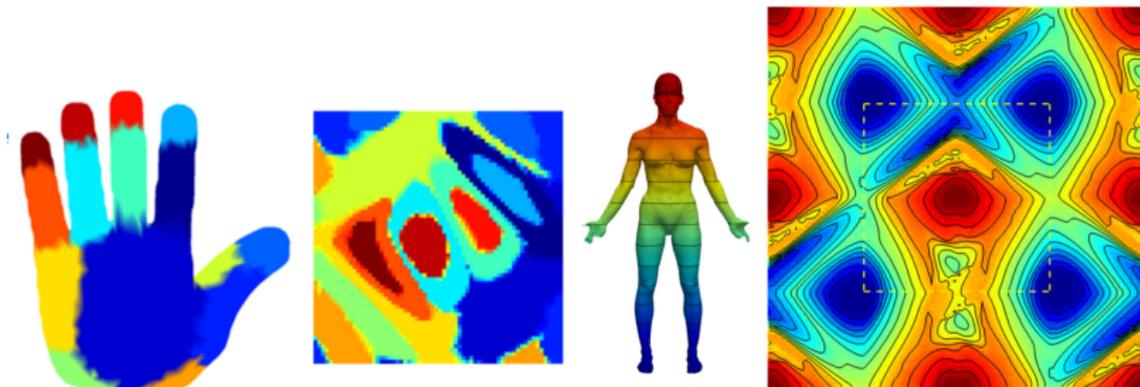
Key idea: map the input surface to some **parametric domain** (e.g. 2D plane) where operations can be defined more easily.



- Enables adoption of Euclidean techniques in the embedding space

Global parametrization

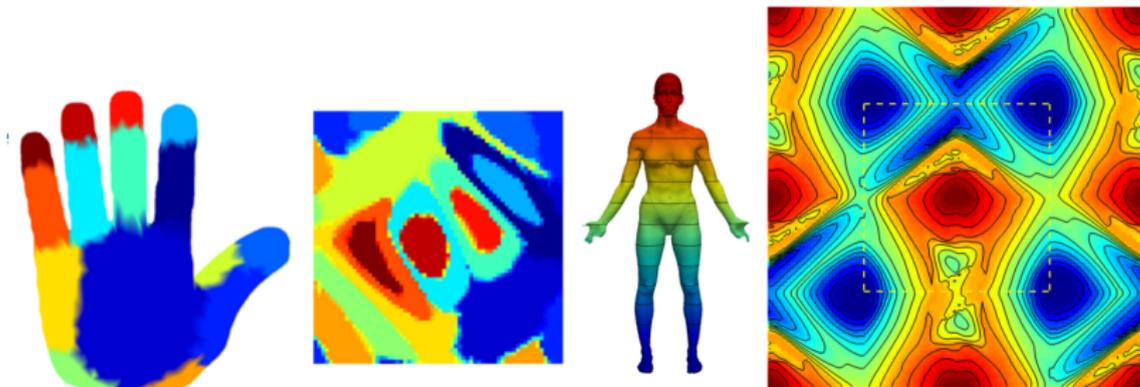
Key idea: map the input surface to some **parametric domain** (e.g. 2D plane) where operations can be defined more easily.



- Enables adoption of Euclidean techniques in the embedding space
- Provides invariance to certain operations

Global parametrization

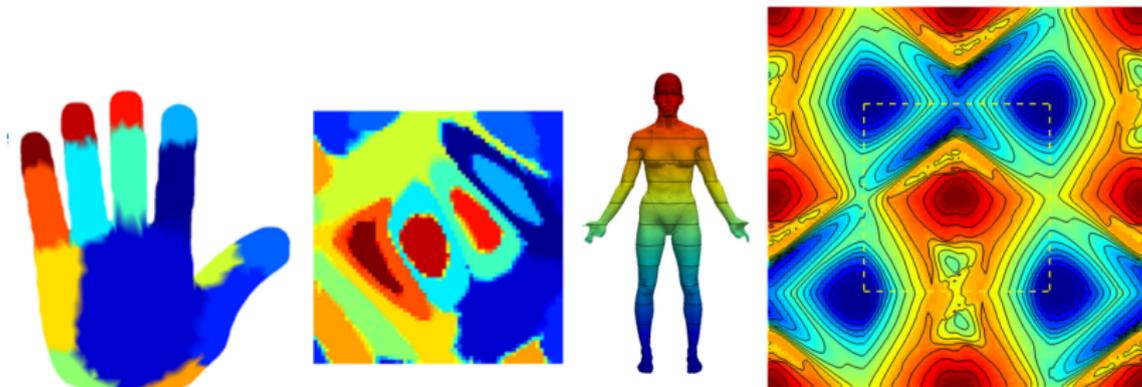
Key idea: map the input surface to some **parametric domain** (e.g. 2D plane) where operations can be defined more easily.



- Enables adoption of Euclidean techniques in the embedding space
- Provides invariance to certain operations
- Parametrization may be non-unique

Global parametrization

Key idea: map the input surface to some **parametric domain** (e.g. 2D plane) where operations can be defined more easily.



- Enables adoption of Euclidean techniques in the embedding space
- Provides invariance to certain operations
- Parametrization may be non-unique
- The map can introduce distortions

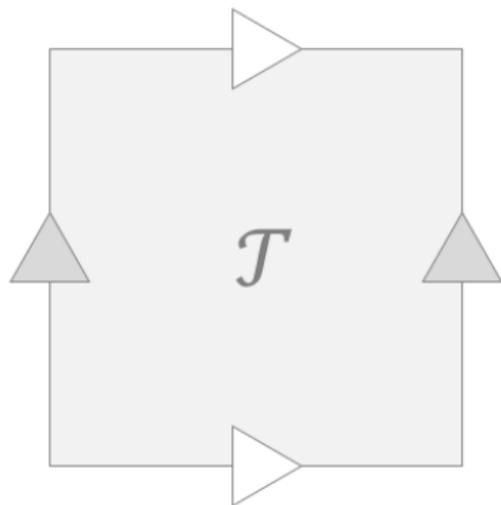
Convolution on surfaces

Is **translation-invariant** convolution on surfaces possible?

Convolution on surfaces

Is [translation-invariant](#) convolution on surfaces possible?

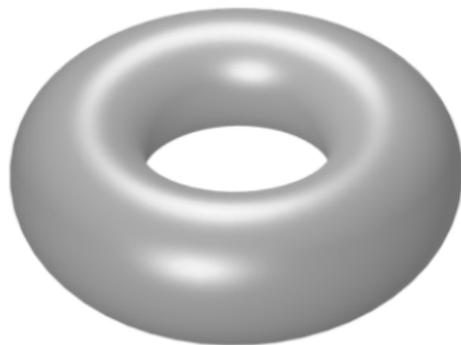
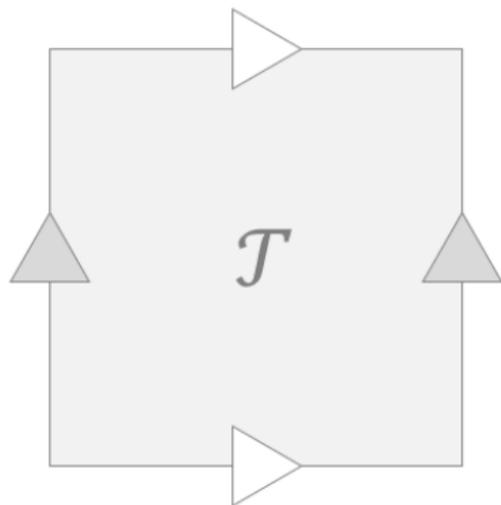
Yes! The [torus](#) is the only closed orientable surface admitting a translational group.



Convolution on surfaces

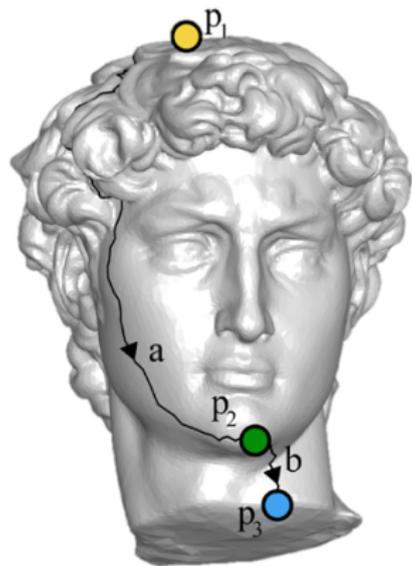
Is **translation-invariant** convolution on surfaces possible?

Yes! The **torus** is the only closed orientable surface admitting a translational group.



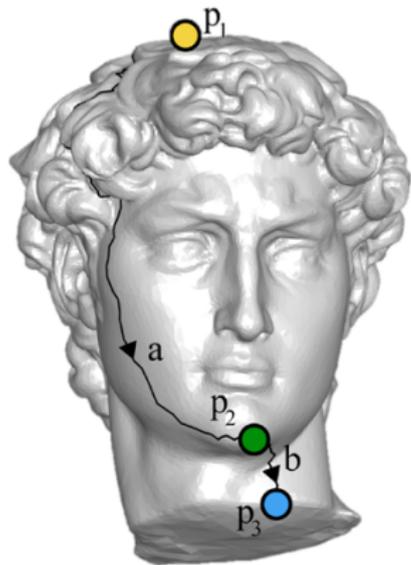
CNNs can be well-defined over the flat-torus!

Torus 4-cover

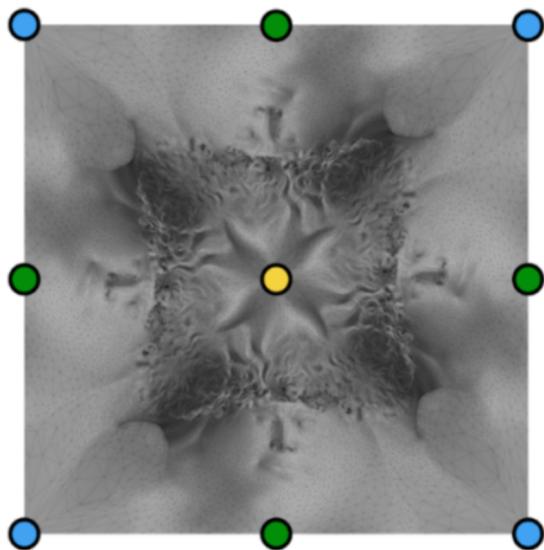


Surface \mathcal{S} with sphere topology

Torus 4-cover

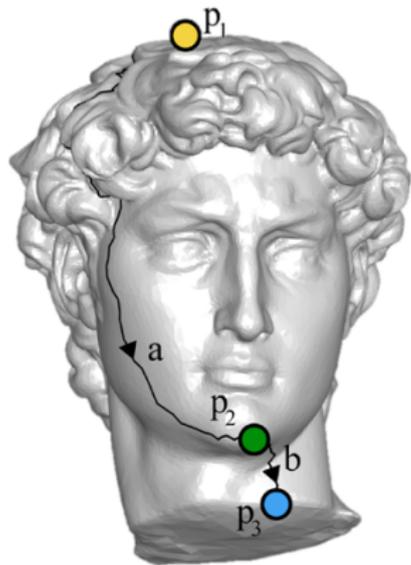


Surface \mathcal{S} with sphere topology

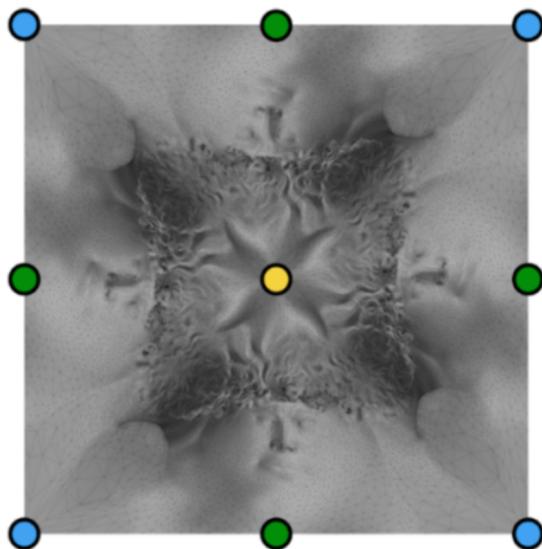


Flat-torus \mathcal{T} with 4 replicas of \mathcal{S}

Torus 4-cover



Surface \mathcal{S} with sphere topology

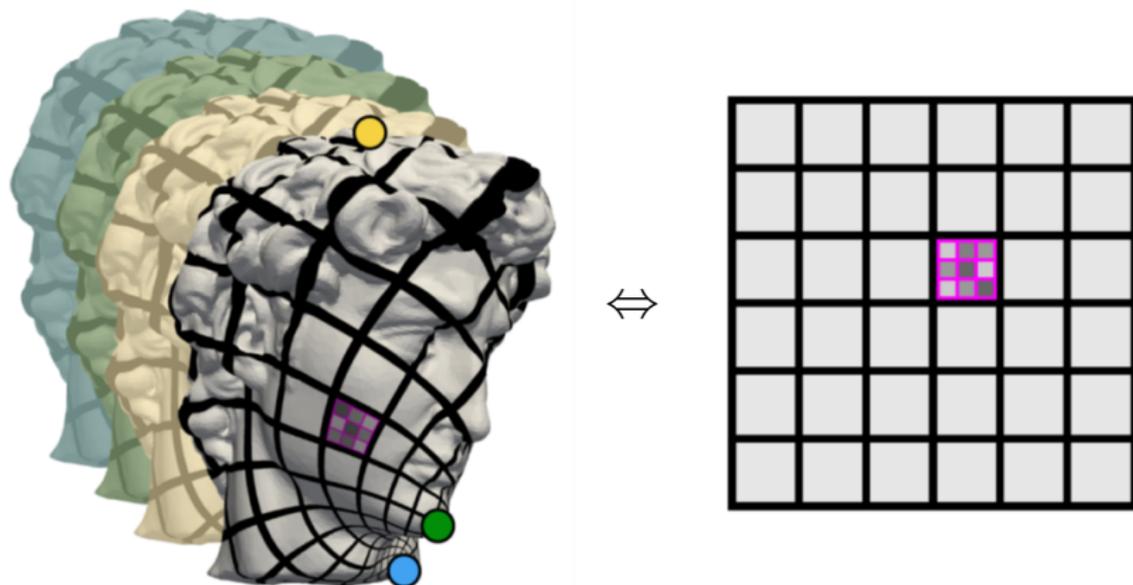


Flat-torus \mathcal{T} with 4 replicas of \mathcal{S}

Standard [Euclidean 2D CNN](#) architectures can now be used on \mathcal{T} .

Torus 4-cover

For each triplet $\{p_1, p_2, p_3\} \in \mathcal{S}$, use [orbifold-Tutte](#) to map S^4 to \mathcal{T} .



The mapping from S^4 to \mathcal{T} is a [conformal homeomorphism](#).

Conformal zoom

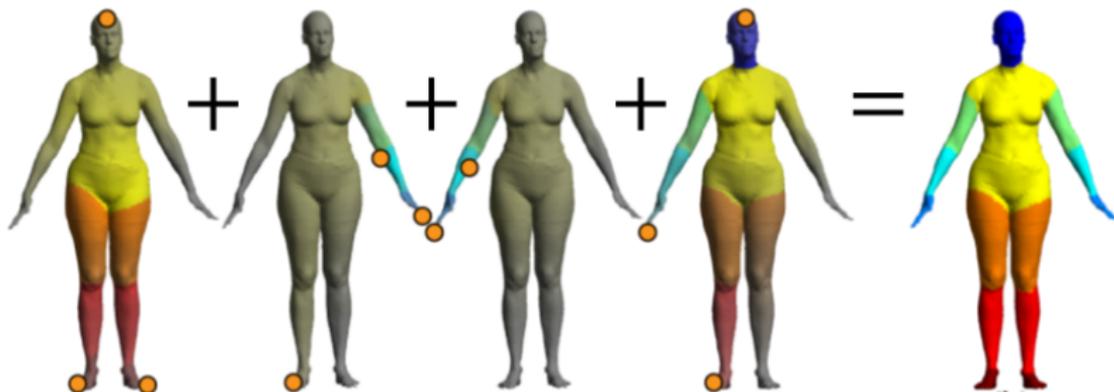
Conformality introduces **scale changes**.

- “Magnifying glass” effect

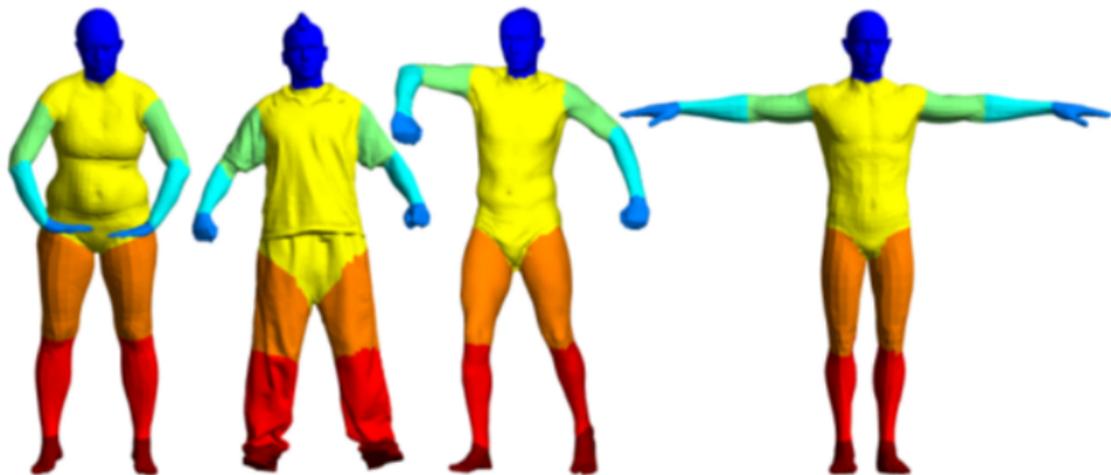
Conformal zoom

Conformality introduces **scale changes**.

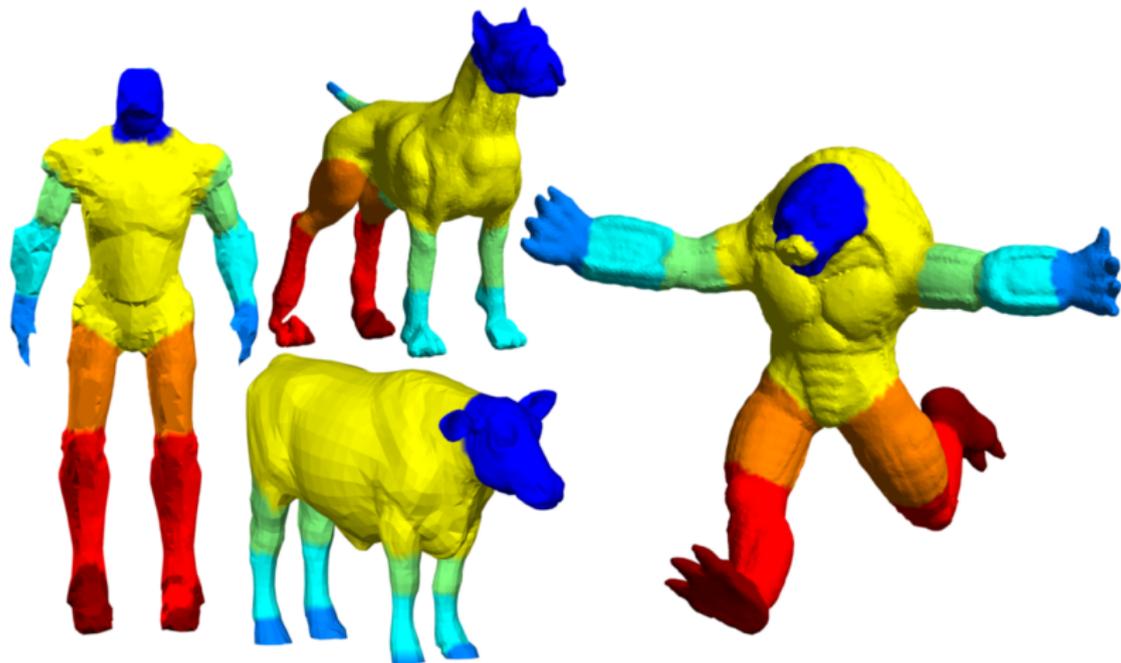
- “Magnifying glass” effect
- Prediction aggregation from different triplets at test time



Segmentation results



Segmentation results



Conclusions

- Application / extension of deep learning to geometric and non-Euclidean structured data
- Deep learning allows designing task-specific features and learning from data **invariance that is impossible to model axiomatically**
- In graphics, an almost virgin field that is very rapidly developing
- Many synergies with other fields
- State-of-the-art results in many 'old' applications
- New applications

Challenges

- Theoretical questions
- Generalization across domains
- Time-varying data / domains
- Directed graphs
- Other shape representations
- Synthesis
- Computation

Shameless advertisement



Shameless advertisement



Geometric and Visual Computing at Sapienza University of Rome

Topics: shape analysis, 3D vision, reconstruction, geometry processing, machine learning, interdisciplinary applications

Thank you!