



Example: MNIST digits classification



Regular grid
(fixed graph, different data)



Superpixels
(different graph and data)

Example: MNIST digits classification

Dataset	LeNet-5 ¹	ChebNet ²
* Full grid	99.33%	99.14%
* $\frac{1}{4}$ grid	98.59%	97.51%
300 Superpixels	-	88.05%
150 Superpixels	-	80.94%
75 Superpixels	-	75.62%

Classification accuracy of different methods on MNIST dataset

* All images have the same graph

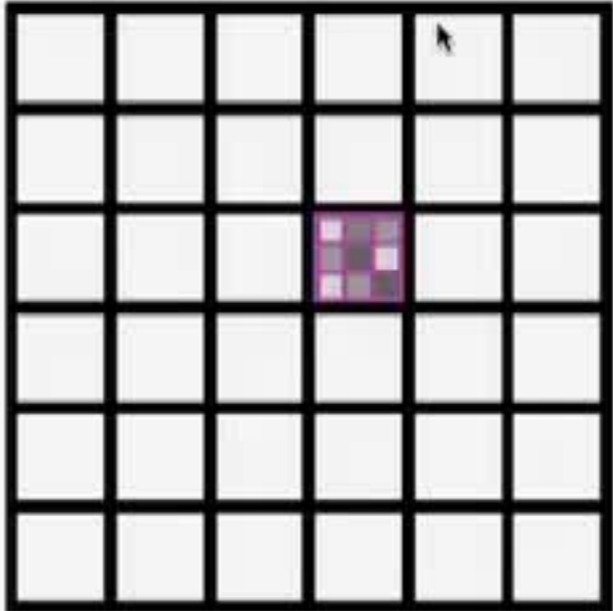
Different formulations of non-Euclidean CNNs



Spectral domain



Spatial domain



Parametric domain

Spatial domain (charting-based)
geometric deep learning methods

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

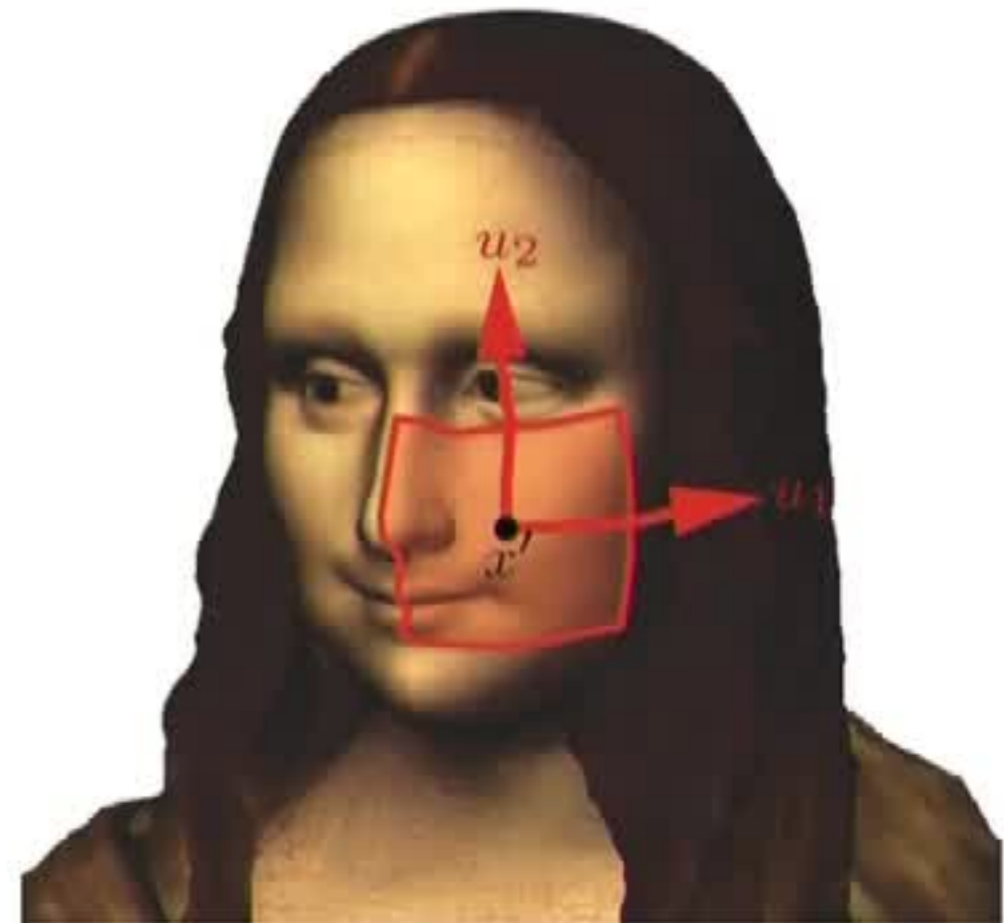
?

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

Patch operators



Image



Manifold

Convolution on manifolds


- Local system of coordinates $\mathbf{u}(x, x')$ around x' (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


$$(f \star g)(x) \propto \sum_{\ell=1}^L g_\ell \underbrace{\int_{\mathcal{X}} w_\ell(\mathbf{u}(x, x')) f(x') dx'}_{\text{patch operator}}$$




Patch operator weight functions



GCNN




ACNN




MoNet


Patch operator weight functions



GCNN

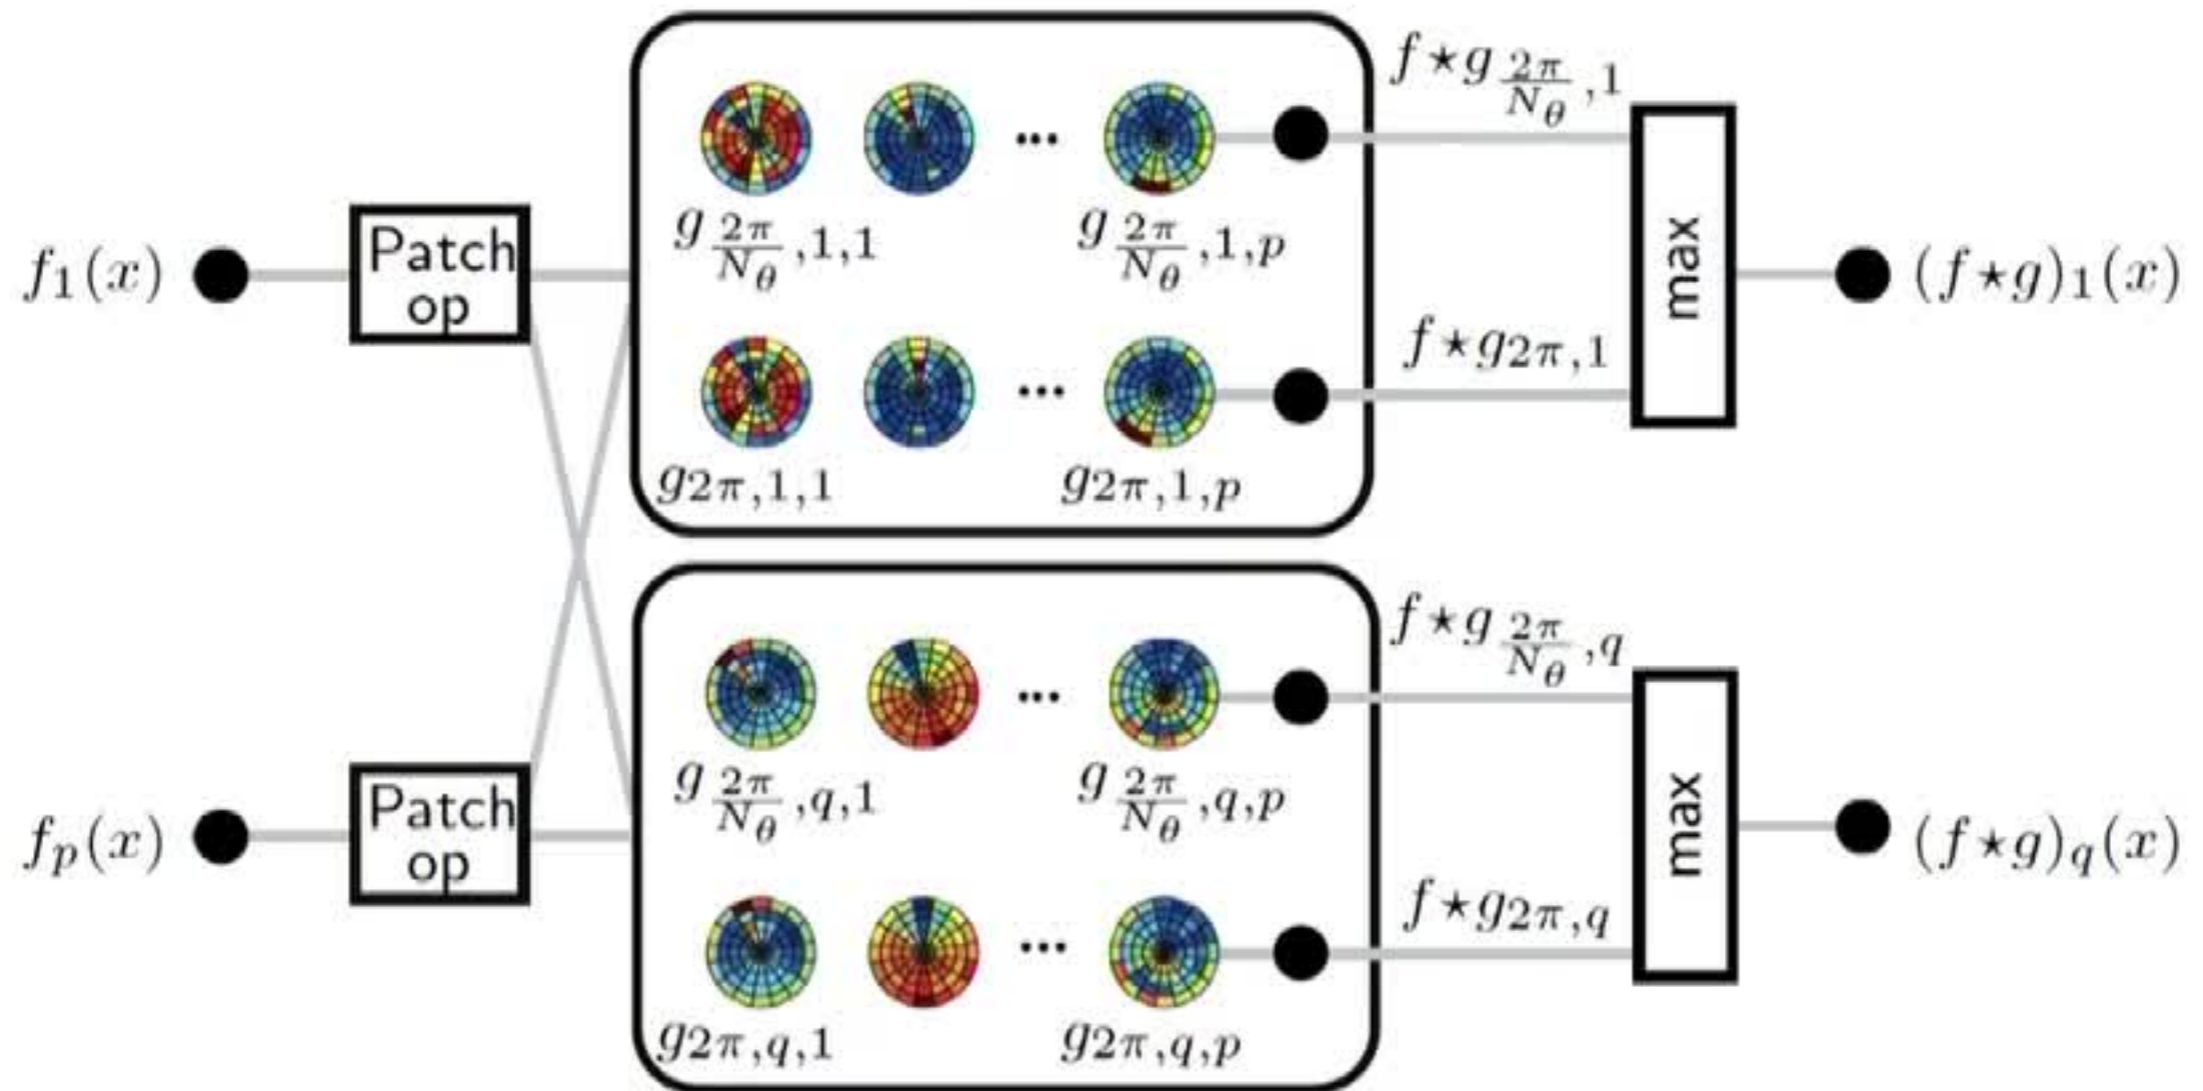


ACNN



MoNet


Geodesic convolution layer




Conv. layer $(f_l \star g)_{\Delta\theta, l}(x) = \xi \left(\sum_{e=1}^p (f_e \star g_{\Delta\theta, l, e})(x) \right)$

Angular max pooling $(f \star g)_l(x) = \max_{\Delta\theta} (f \star g)_{\Delta\theta, l}(x)$


Patch operator weight functions



GCNN



ACNN



MoNet

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




Isotropic



Anisotropic


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'


Patch operator weight functions



GCNN



ACNN



MoNet

Mixture Model Network (MoNet)

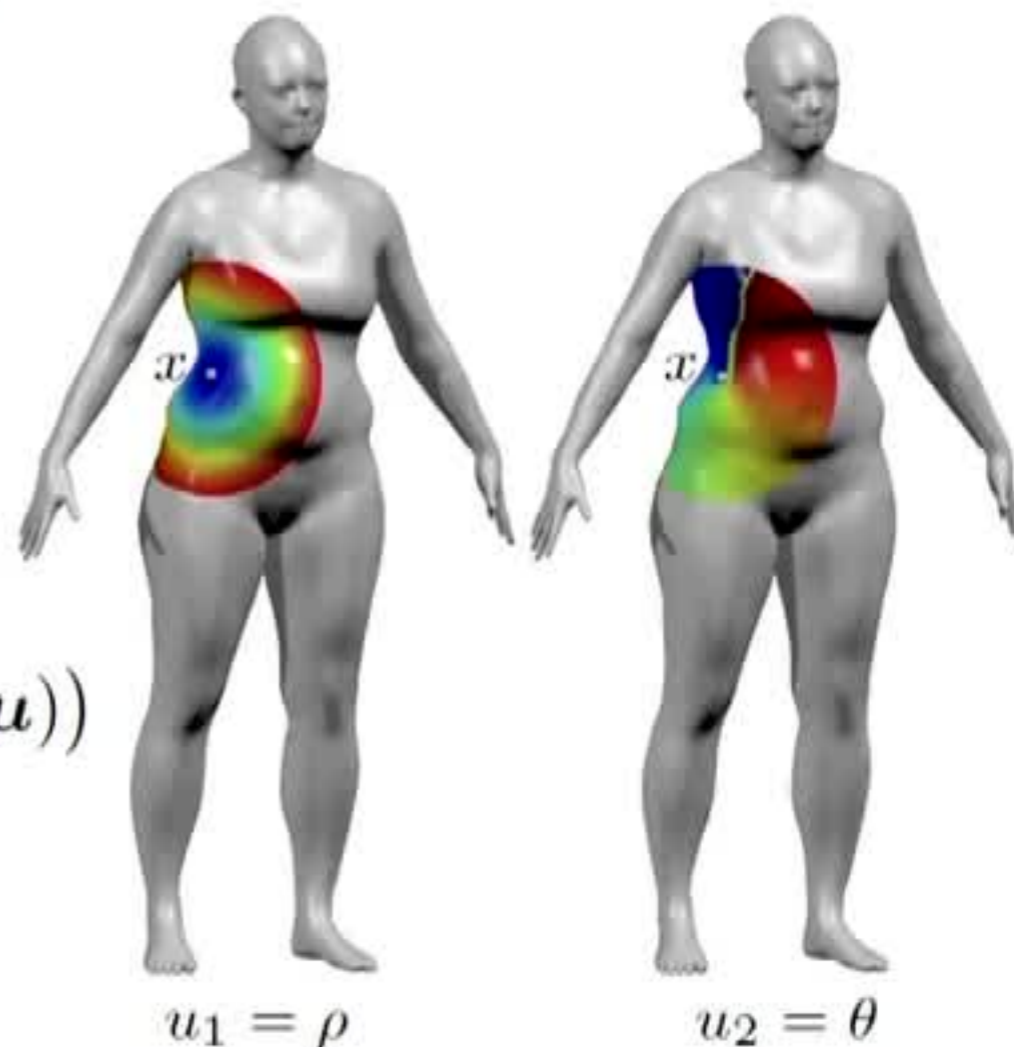
- Geodesic polar coordinates

$$\mathbf{u}(x, y) = (\rho(x, x'), \theta(x, x'))$$

- 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

$$(f \star g)(x) \propto \int_{\mathcal{X}} \underbrace{\sum_{\ell=1}^L g_{\ell} w_{\mu_{\ell}, \Sigma_{\ell}}(\mathbf{u}(x, x'))}_{\text{Gaussian mixture}} f(x') dx'$$


where g_1, \dots, g_L are the spatial filter coefficients and μ_1, \dots, μ_L and $\Sigma_1, \dots, \Sigma_L$ are patch operator parameters

Mixture Model Network on graphs

- Local coordinates \mathbf{u}_{ij} , e.g. vertex degree, geodesic distance,...
- 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



Local coordinates on graph

Spatial convolution

$$(f \star g)_i \propto \sum_{\ell=1}^L g_\ell \sum_{j=1}^n w_{\mu_\ell, \Sigma_\ell}(\mathbf{u}_{i,j}) f_j$$

where g_1, \dots, g_L are the spatial filter coefficients and μ_1, \dots, μ_L and $\Sigma_1, \dots, \Sigma_L$ are patch operator parameters

Spectral vs Spatial methods

ChebNet filter

$$h_i = \sum_{\ell=0}^r \alpha_{\ell} (\Delta^{\ell} \mathbf{f})_i$$

Spatial filter

$$h_i = \sum_{\ell=1}^L g_{\ell} (\mathbf{W}_{\ell} \mathbf{f})_i$$

ChebNet is a particular setting of spatial convolution with local weighting functions given by the powers of the Laplacian $\mathbf{W}_{\ell} = \Delta^{\ell}$

Graph Attention Networks (GAT)

Main idea: neighborhood average

$$\mathbf{f}'_i = \sum_{j:(i,j) \in \mathcal{E}} \alpha_{ij} \mathbf{f}_j$$

weighted by attention score

$$\alpha_{ij} = \frac{e^{\xi([\mathbf{f}_i \mathbf{W}, \mathbf{f}_j \mathbf{W}] \mathbf{a})}}{\sum_{k:(i,k) \in \mathcal{E}} e^{\xi([\mathbf{f}_i \mathbf{W}, \mathbf{f}_k \mathbf{W}] \mathbf{a})}}$$

which is a learnable transformation of the local features with learnable parameters \mathbf{W} , \mathbf{a}

Particular case of MoNet-type architectures!

Dual/Primal Graph CNN (DPGCNN)

Alternate GAT-type convolutions applied on primal and dual graphs

- Dual convolution on $\tilde{\mathcal{G}}$:

$$\tilde{\mathbf{f}}'_{ij} = \xi \left(\sum_{r \in \mathcal{N}_i} \tilde{\alpha}_{ij,ir} \tilde{\mathbf{f}}_{ir} \tilde{\mathbf{W}} + \sum_{t \in \mathcal{N}_j} \tilde{\alpha}_{ij,tj} \tilde{\mathbf{f}}_{tj} \tilde{\mathbf{W}} \right)$$
$$\tilde{\alpha}_{ij,ik} = \frac{e^{\xi([\tilde{\mathbf{f}}_{ij} \tilde{\mathbf{W}}, \tilde{\mathbf{f}}_{ik} \tilde{\mathbf{W}}] \tilde{\mathbf{a}})}}{\sum_{r \in \mathcal{N}_i} e^{\xi([\tilde{\mathbf{f}}_{ij} \tilde{\mathbf{W}}, \tilde{\mathbf{f}}_{ir} \tilde{\mathbf{W}}] \tilde{\mathbf{a}})}} + \sum_{t \in \mathcal{N}_j} e^{\xi([\tilde{\mathbf{f}}_{ij} \tilde{\mathbf{W}}, \tilde{\mathbf{f}}_{tj} \tilde{\mathbf{W}}] \tilde{\mathbf{a}})}}$$

- Primal convolution on \mathcal{G} :

$$\mathbf{f}'_i = \xi \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{f}_j \mathbf{W} \right) \quad \alpha_{ij} = \frac{e^{\xi(\tilde{\mathbf{f}}'_{ij} \mathbf{a})}}{\sum_{k \in \mathcal{N}_i} e^{\xi(\tilde{\mathbf{f}}'_{ik} \mathbf{a})}}$$

Dual/Primal Graph CNN (DPGCNN)

Alternate GAT-type convolutions applied on primal and dual graphs

- Dual convolution on $\tilde{\mathcal{G}}$:

$$\tilde{\mathbf{f}}'_{ij} = \xi \left(\sum_{r \in \mathcal{N}_i} \tilde{\alpha}_{ij,ir} \tilde{\mathbf{f}}_{ir} \tilde{\mathbf{W}} + \sum_{t \in \mathcal{N}_j} \tilde{\alpha}_{ij,tj} \tilde{\mathbf{f}}_{tj} \tilde{\mathbf{W}} \right)$$
$$\tilde{\alpha}_{ij,ik} = \frac{e^{\xi([\tilde{\mathbf{f}}_{ij} \tilde{\mathbf{W}}, \tilde{\mathbf{f}}_{ik} \tilde{\mathbf{W}}] \tilde{\mathbf{a}})}}{\sum_{r \in \mathcal{N}_i} e^{\xi([\tilde{\mathbf{f}}_{ij} \tilde{\mathbf{W}}, \tilde{\mathbf{f}}_{ir} \tilde{\mathbf{W}}] \tilde{\mathbf{a}})}} + \sum_{t \in \mathcal{N}_j} e^{\xi([\tilde{\mathbf{f}}_{ij} \tilde{\mathbf{W}}, \tilde{\mathbf{f}}_{tj} \tilde{\mathbf{W}}] \tilde{\mathbf{a}})}}$$

- Primal convolution on \mathcal{G} :

$$\mathbf{f}'_i = \xi \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{f}_j \mathbf{W} \right) \quad \alpha_{ij} = \frac{e^{\xi(\tilde{\mathbf{f}}'_{ij} \mathbf{a})}}{\sum_{k \in \mathcal{N}_i} e^{\xi(\tilde{\mathbf{f}}'_{ik} \mathbf{a})}}$$

Example: citation networks

Method	Cora¹	CiteSeer²
Manifold Regularization ³	59.5%	60.1%
Semidefinite Embedding ⁴	59.0%	59.6%
Label Propagation ⁵	68.0%	45.3%
DeepWalk ⁶	67.2%	43.2%
Planetoid ⁷	75.7%	64.7%
GCN ⁸	81.6%	70.3%
MoNet ⁹	81.7%	–
GAT ¹⁰	83.0%	72.5%
DPGCN¹¹	83.3%	72.6%

Classification accuracy of different methods on citation network datasets

Data: ^{1,2}Sen et al. 2008; methods: ³Belkin et al. 2006; ⁴Weston et al. 2012; ⁵Zhu et al. 2003; ⁶Perozzi et al. 2014; ⁷Yang et al. 2016; ⁸Kipf, Welling 2016 ; ⁹Monti et al. 2017; ¹⁰Veličković et al. 2018; ¹¹Monti et al. 2018

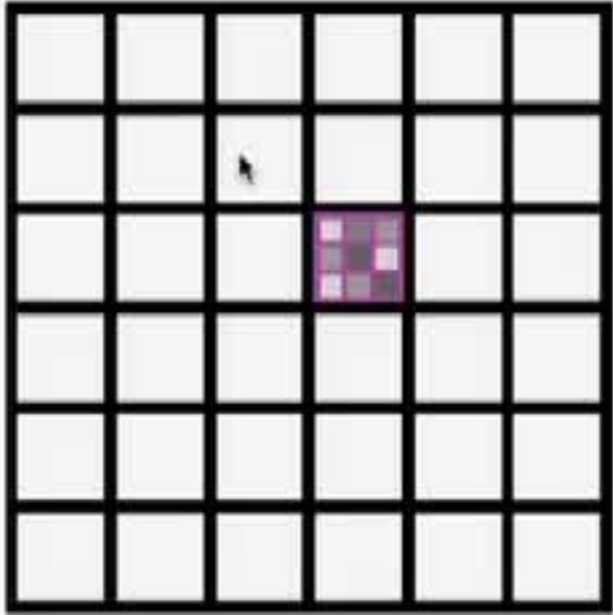
Different formulations of non-Euclidean CNNs



Spectral domain



Spatial domain




Parametric domain

Parametric domain
geometric deep learning methods

Global parametrization

Map the input surface to some parametric domain with shift-invariant structure



- ☺ Allows to use standard CNNs (pull back convolution from the parametric space)
- ☺ Guaranteed invariance to some classes of transformations
- ☹ Parametrization may not be unique
- ☹ Embedding may introduce distortion

Translation invariance on manifolds


Translation on manifold = locally Euclidean translation = flow along a non-vanishing vector field

Poincaré-Hopf Theorem Non-vanishing vector field on a closed orientable compact 2-manifold implies manifold of genus 1 (torus)

Translation invariance on manifolds


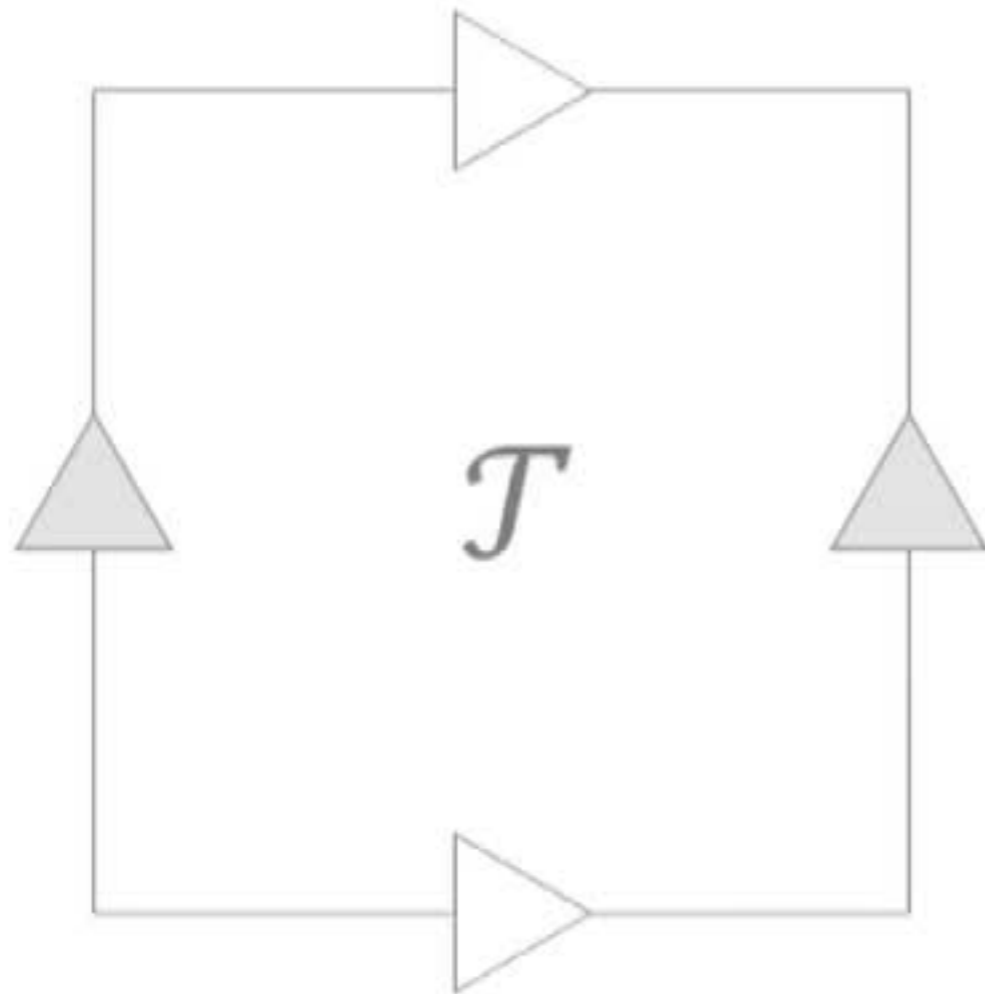
Translation on manifold = locally Euclidean translation = flow along a non-vanishing vector field

Poincaré-Hopf Theorem Non-vanishing vector field on a closed orientable compact 2-manifold implies manifold of genus 1 (torus)



'Hairy Ball Theorem' states that a sphere cannot be combed

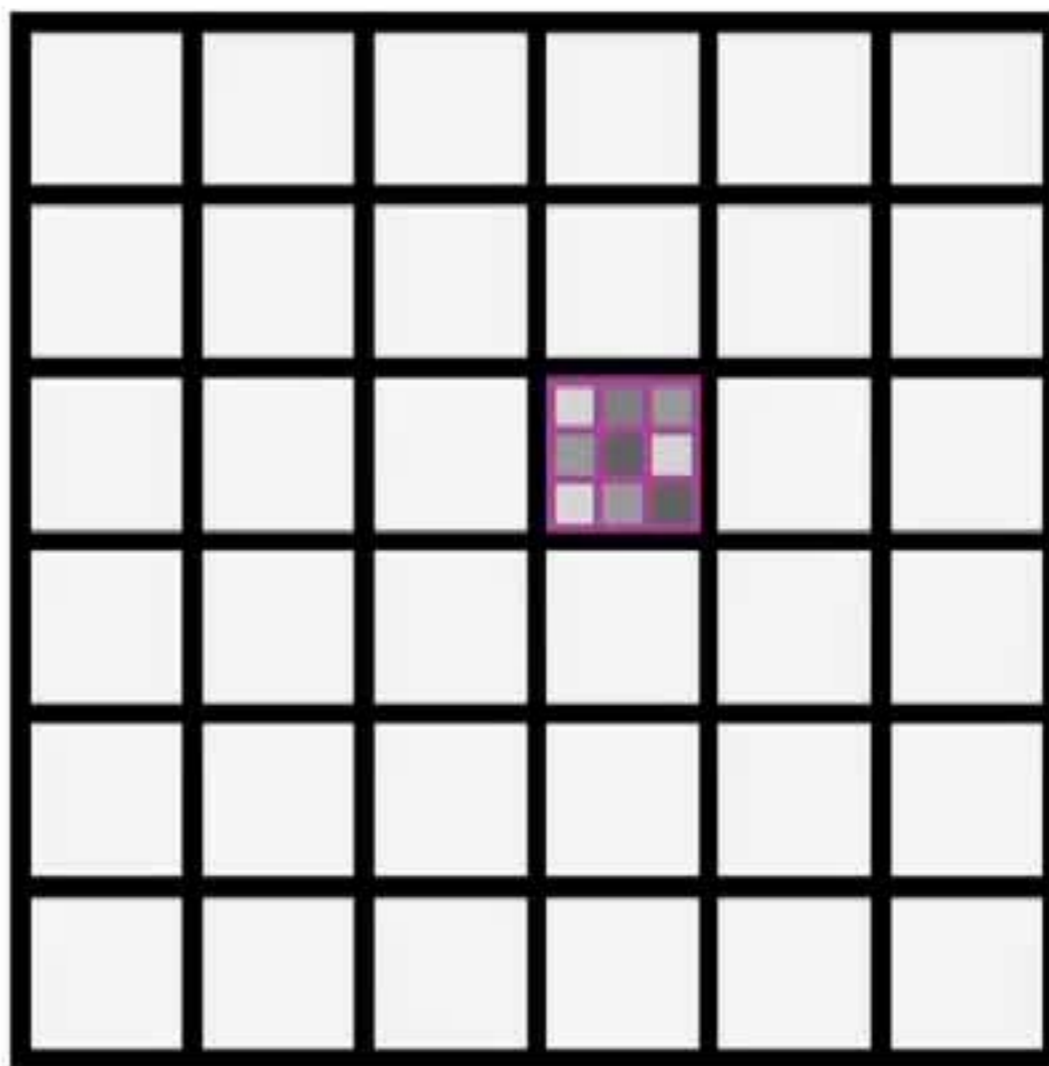
Translation invariance on the torus



Torus is the only closed orientable surface admitting a translation group

Convolution on torus

For any triplet of points on \mathcal{X} , construct [conformal homeomorphism](#) from the 4-cover \mathcal{X}^4 to \mathcal{T} using [orbifold-Tutte](#) method




Conformal zoom


- Embedding depends on the choice of the triplets of points
- 'Conformal zoom' effect

Conformal zoom

- Embedding depends on the choice of the triplets of points
- 'Conformal zoom' effect
- Choose multiple triples and aggregate results in training / test phase




Example: shape segmentation with Toric CNN



Examples of shape segmentation obtained with Toric CNN

Example: shape segmentation with Toric CNN



Examples of shape segmentation obtained with Toric CNN

Application dealing with 3D data



Computer graphics



Virtual/augmented reality



Robotics



Autonomous driving



Medicine



Drug design

Application dealing with 3D data



Computer graphics



Virtual/augmented reality



Robotics



Autonomous driving



Medicine




Drug design

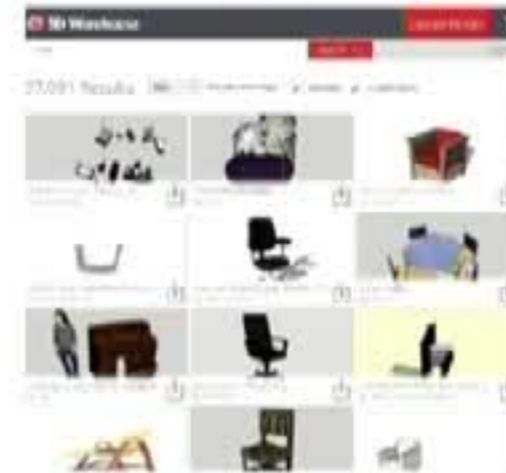
3D data



1998



2008



KINECT



REALSENSE

biobank^{UK}

2018

Clara.io

yobi 3D

Sketchfab

Google TINKER
SketchUp CAD



>\$10K

2005



\$100

2010



\$20

2014



???

2017



>\$10K

2005



\$100

2010



\$20

2014



???

2017



GDC



Faceshift (acquired by Apple in 2015)



Bill Roberts



GDC

Faceshift (acquired by Apple in 2015)



Bill Roberts



GDC

Faceshift (acquired by Apple in 2015)




Bill Roberts



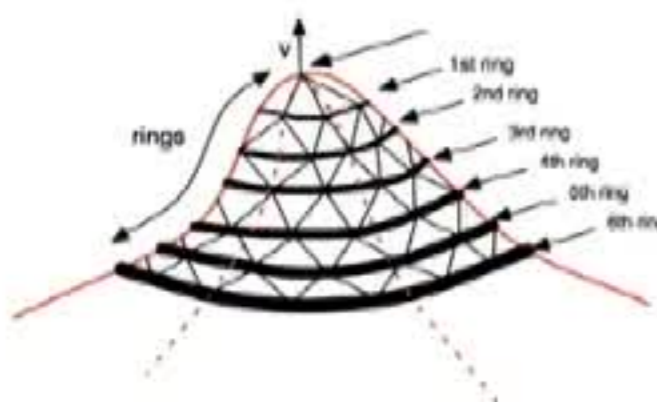
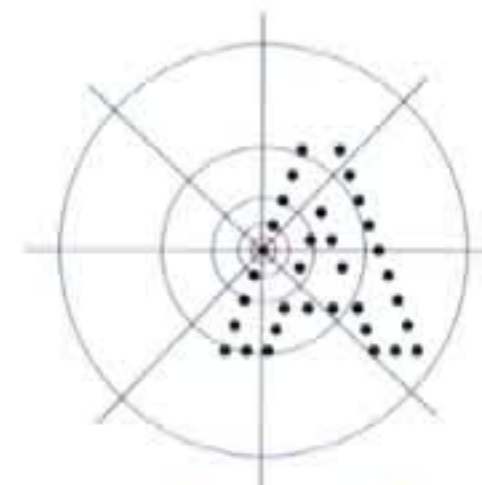


GDC

Faceshift (acquired by Apple in 2015)

Analysis and synthesis



3D feature descriptors



SIFT¹ / MeshHOG²

MSER³ / ShapeMSER⁴

(Intrinsic⁶) Shape context⁵



Spin image⁷


Heat kernel signature⁸

¹Lowe 2004; ²Zaharescu et al. 2009; ³Matas et al. 2002; ⁴Litman et al. 2010;

⁵Belongie et al. 2000; ⁶Kokkinos et al. 2012; ⁷Johnson et al. 1999; ⁸Sun et al. 2009

Task-specific features




Correspondence



Task-specific features

Correspondence

Similarity



Shape representation




Image-based



Volumetric

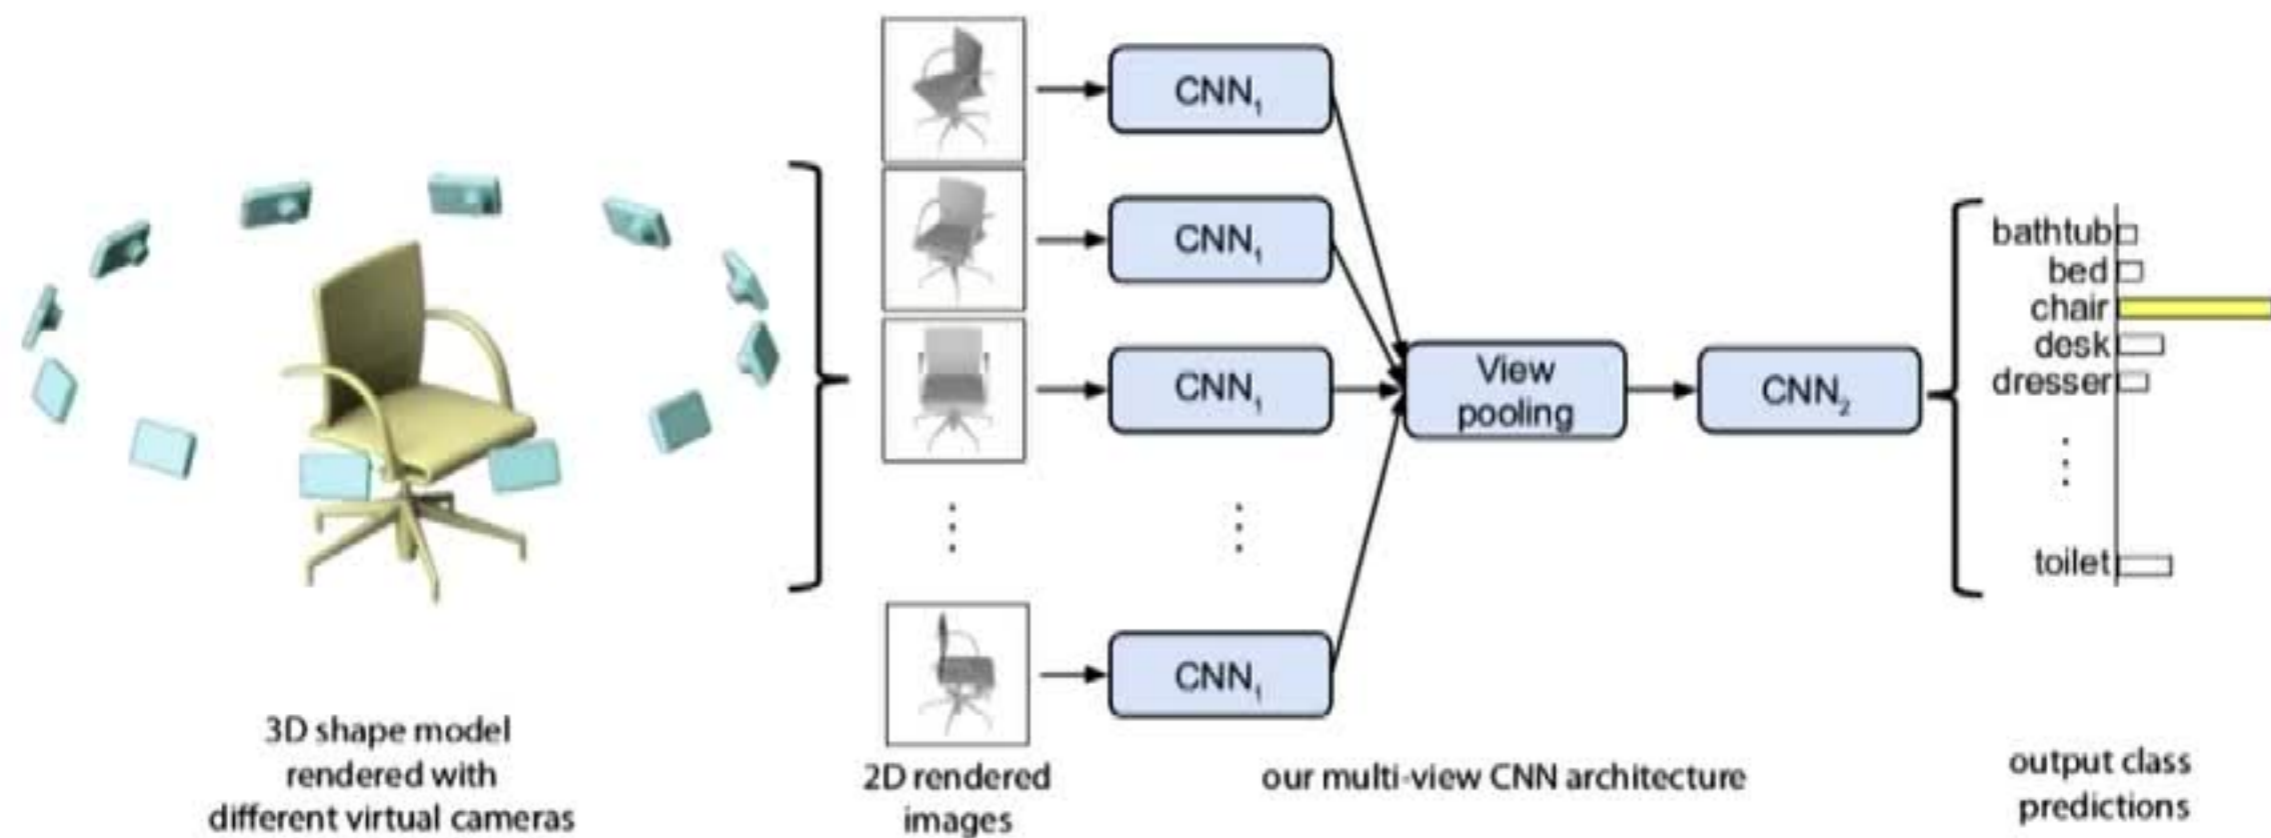


Point-based





Surface-based

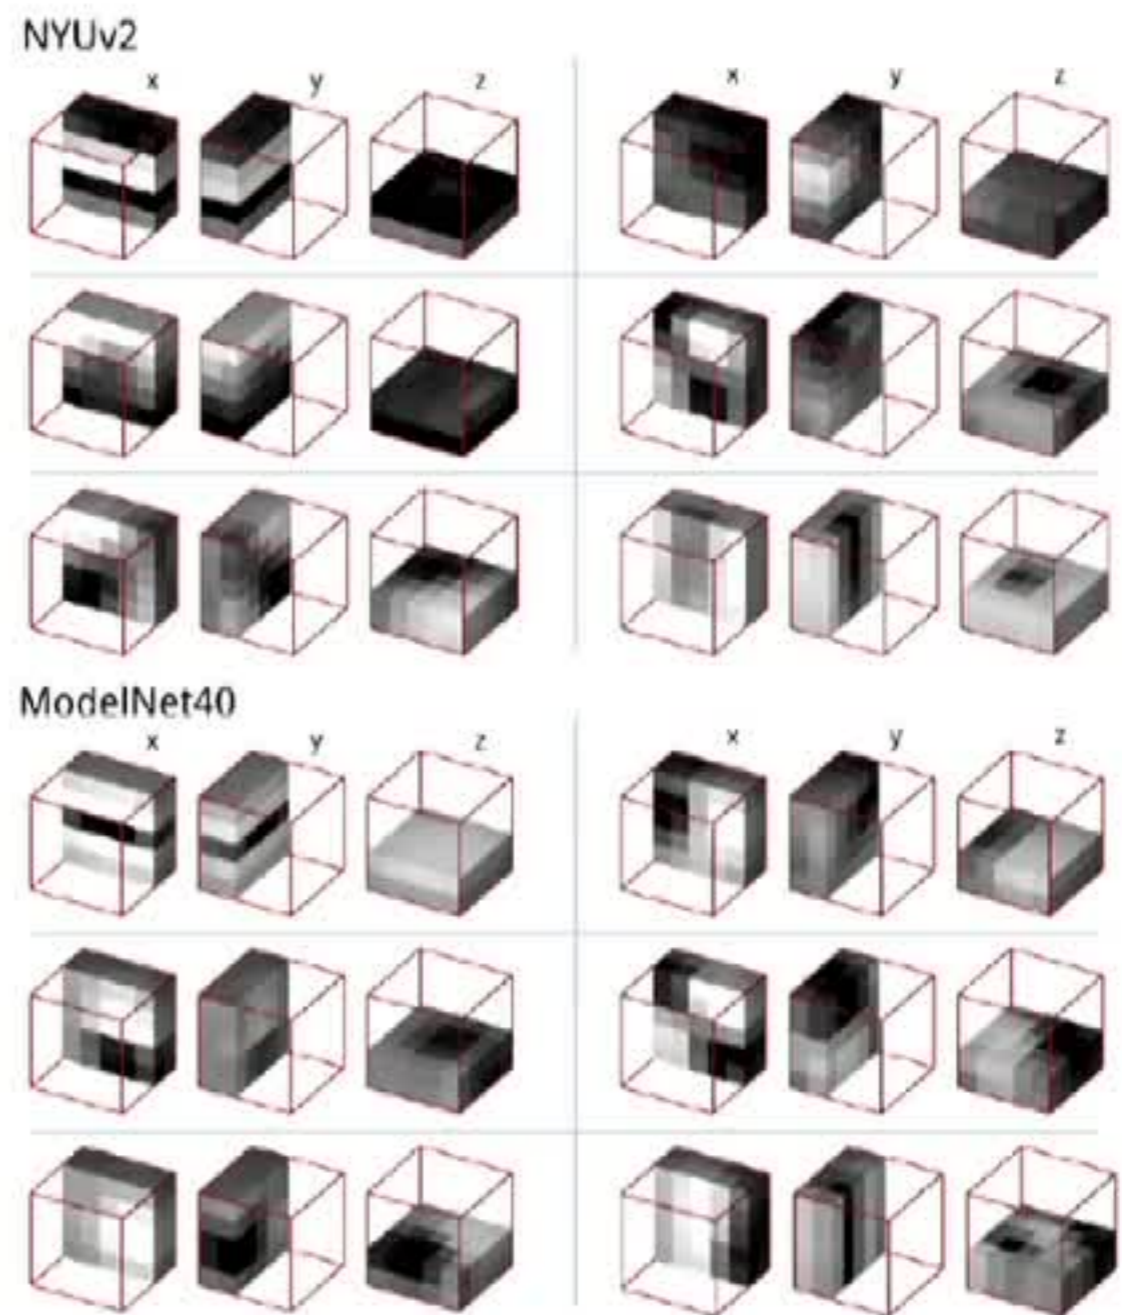
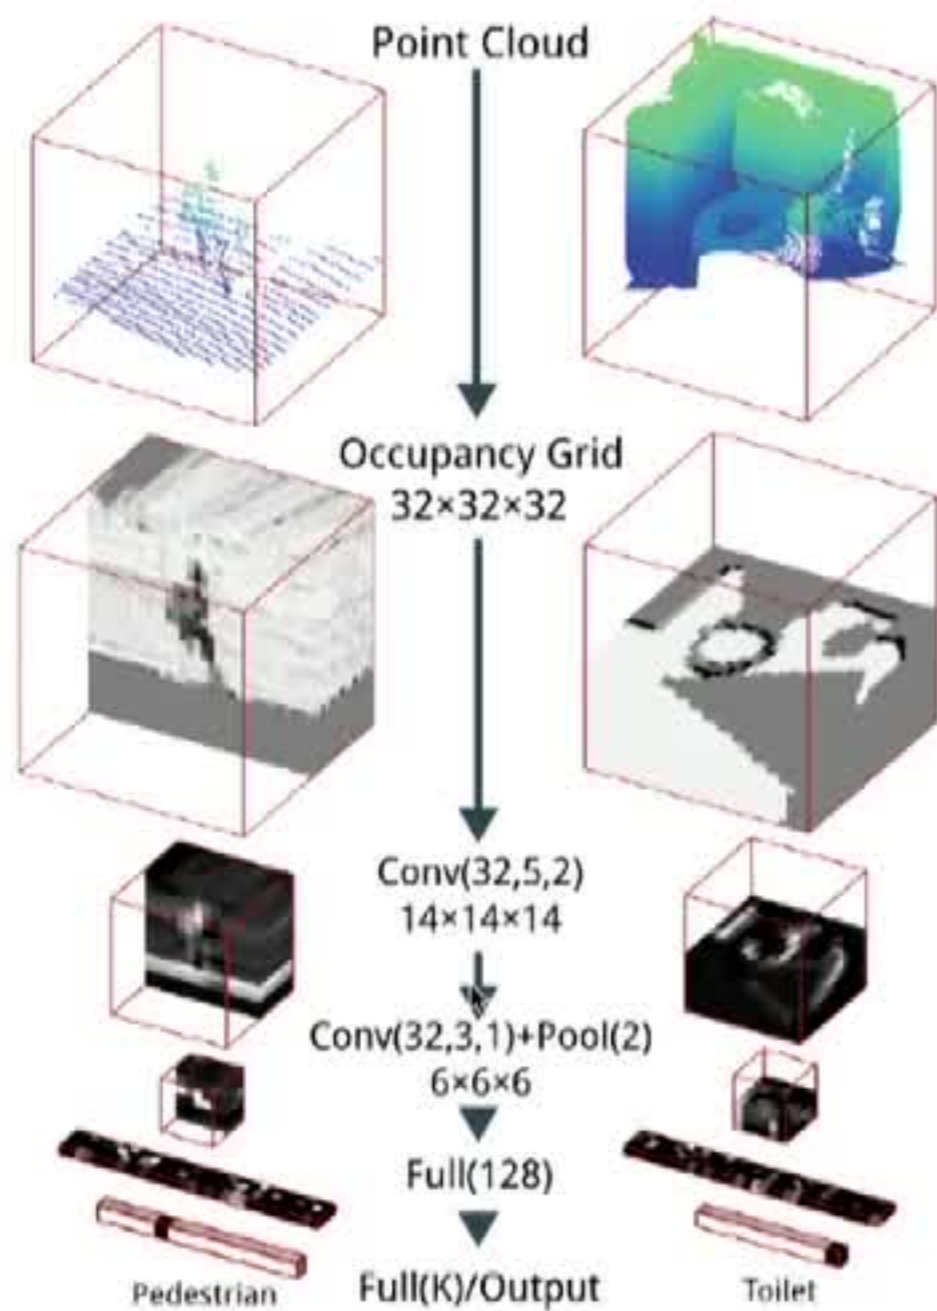
Multi-view CNN





Volumetric methods: 3D ShapeNet



Volumetric methods: VoxNet



Volumetric methods: 3D ShapeNet



Volumetric semantic segmentation




Depth image




Visible surface


- floor
- wall
- bed
- window
- sofa
- objects
- furniture



Semantic segmentation

Efficient volumetric data structures





PointNet: learning on sets

- Permutation-invariant function

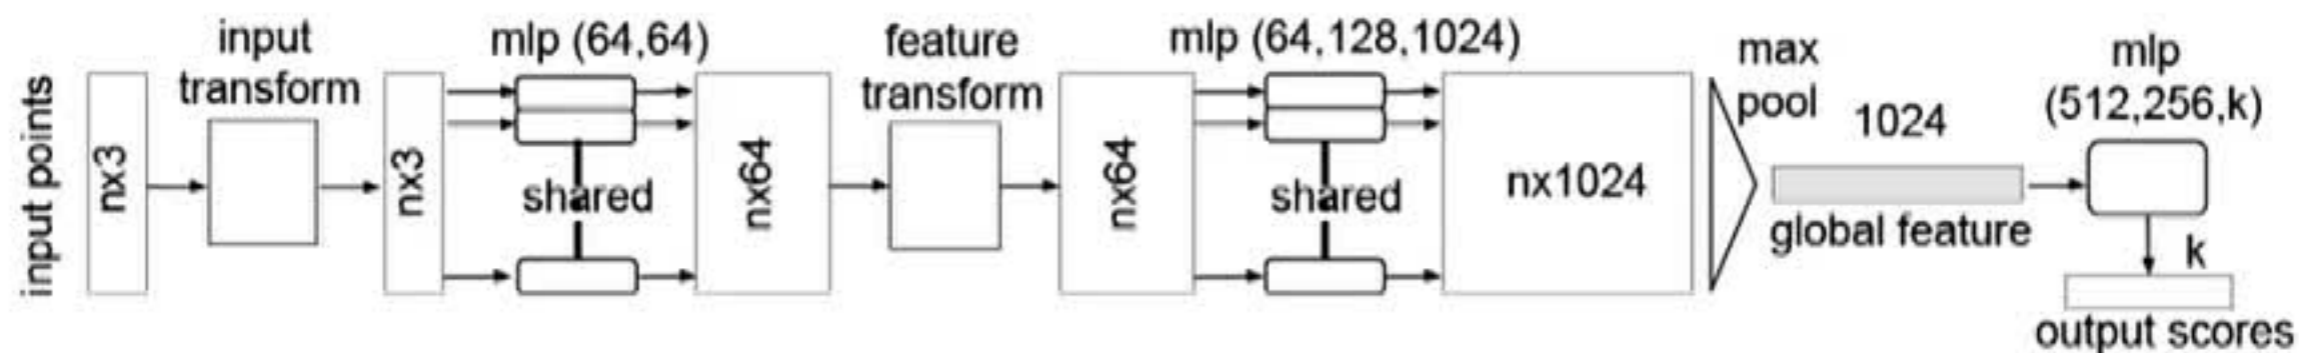
$$f(\mathbf{x}_1, \dots, \mathbf{x}_n) = f(\mathbf{x}_{\pi_1}, \dots, \mathbf{x}_{\pi_n})$$

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

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	$\sum_e g_e w_e(\mathbf{u}_{ij}) \mathbf{x}_j$
PCNN ⁴	\sum	$\sum_{\ell m} c(\mathbf{x}_i \cdot \mathbf{k}_{\ell m}) w_{i q \Theta_\ell}(\mathbf{x}_i, \mathbf{x}_j)$

Dynamic Graph CNN (DynGCNN)

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

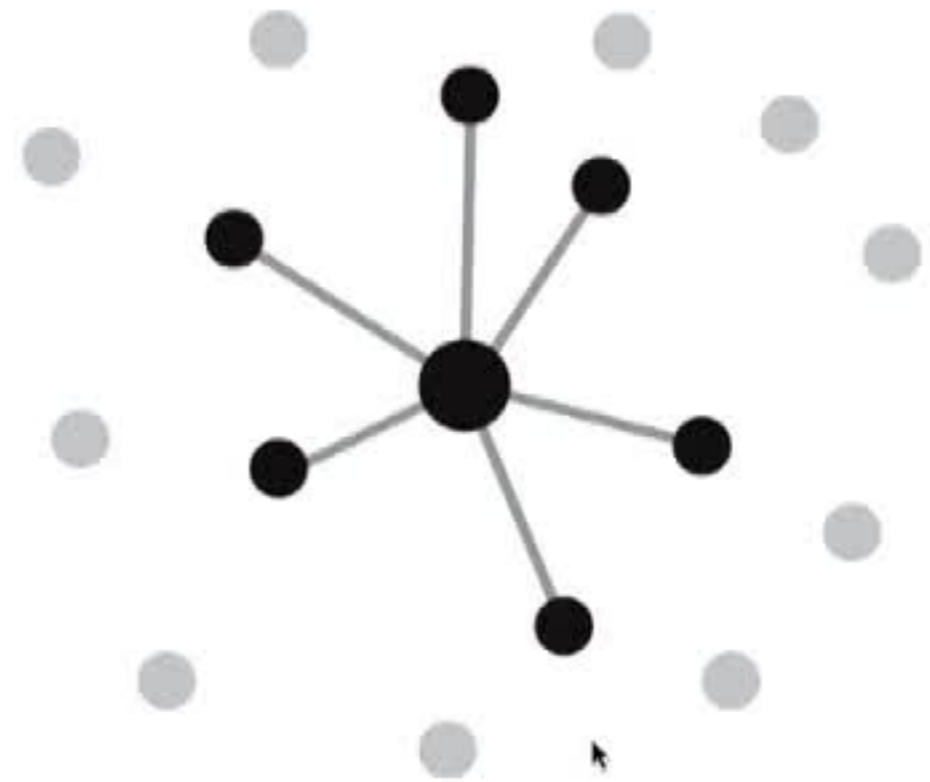


Layer l

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: indoor scans (S3DIS)



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)

Semantic segmentation: indoor scans (S3DIS)



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

Shape representation




Image-based



Volumetric



Point-based



Surface-based

Shape representation




Image-based



Volumetric

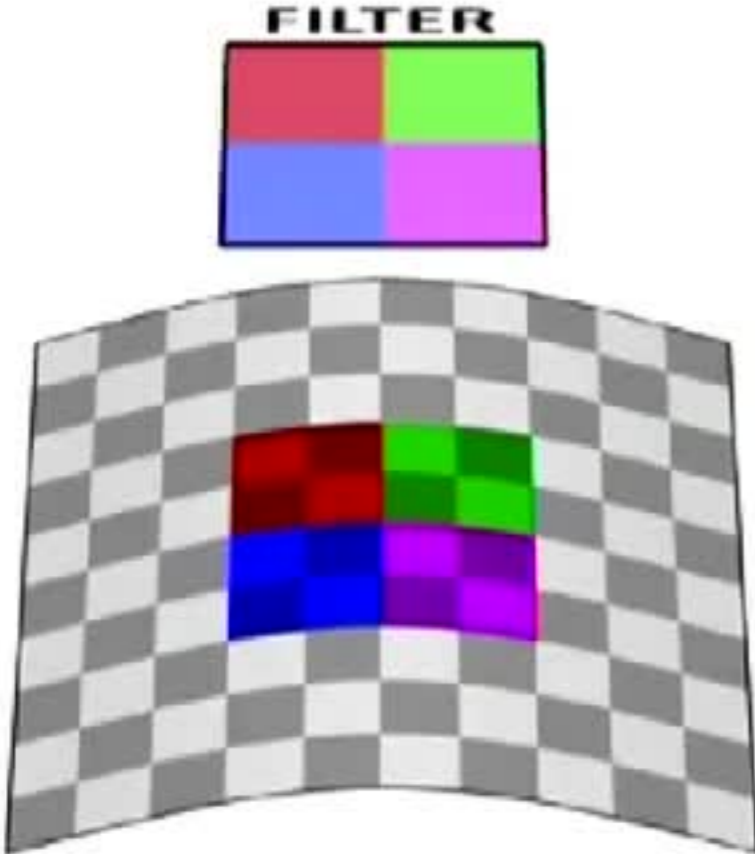


Point-based

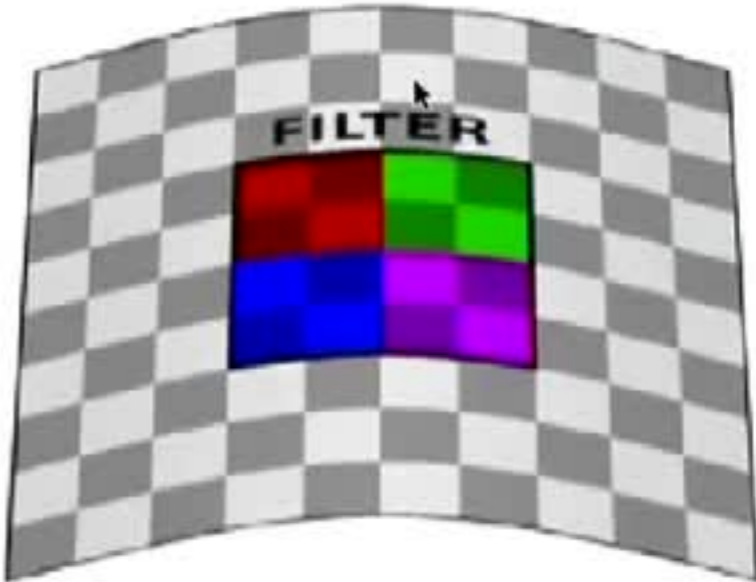


Surface-based

Extrinsic vs Intrinsic CNNs



Extrinsic



Intrinsic


Different formulations of non-Euclidean CNNs



Spectral domain



Spatial domain



Parametric domain

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

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

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



Learning local descriptors with intrinsic CNNs



Training set

Siamese net

Pointwise feature cost

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

two net instances with shared parameters Θ

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


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: Bronstein, Bronstein, Kimmel 2008 (TOSCA); Angelov et al. 2005 (SCAPE); Bogo et al. 2014 (FAUST)

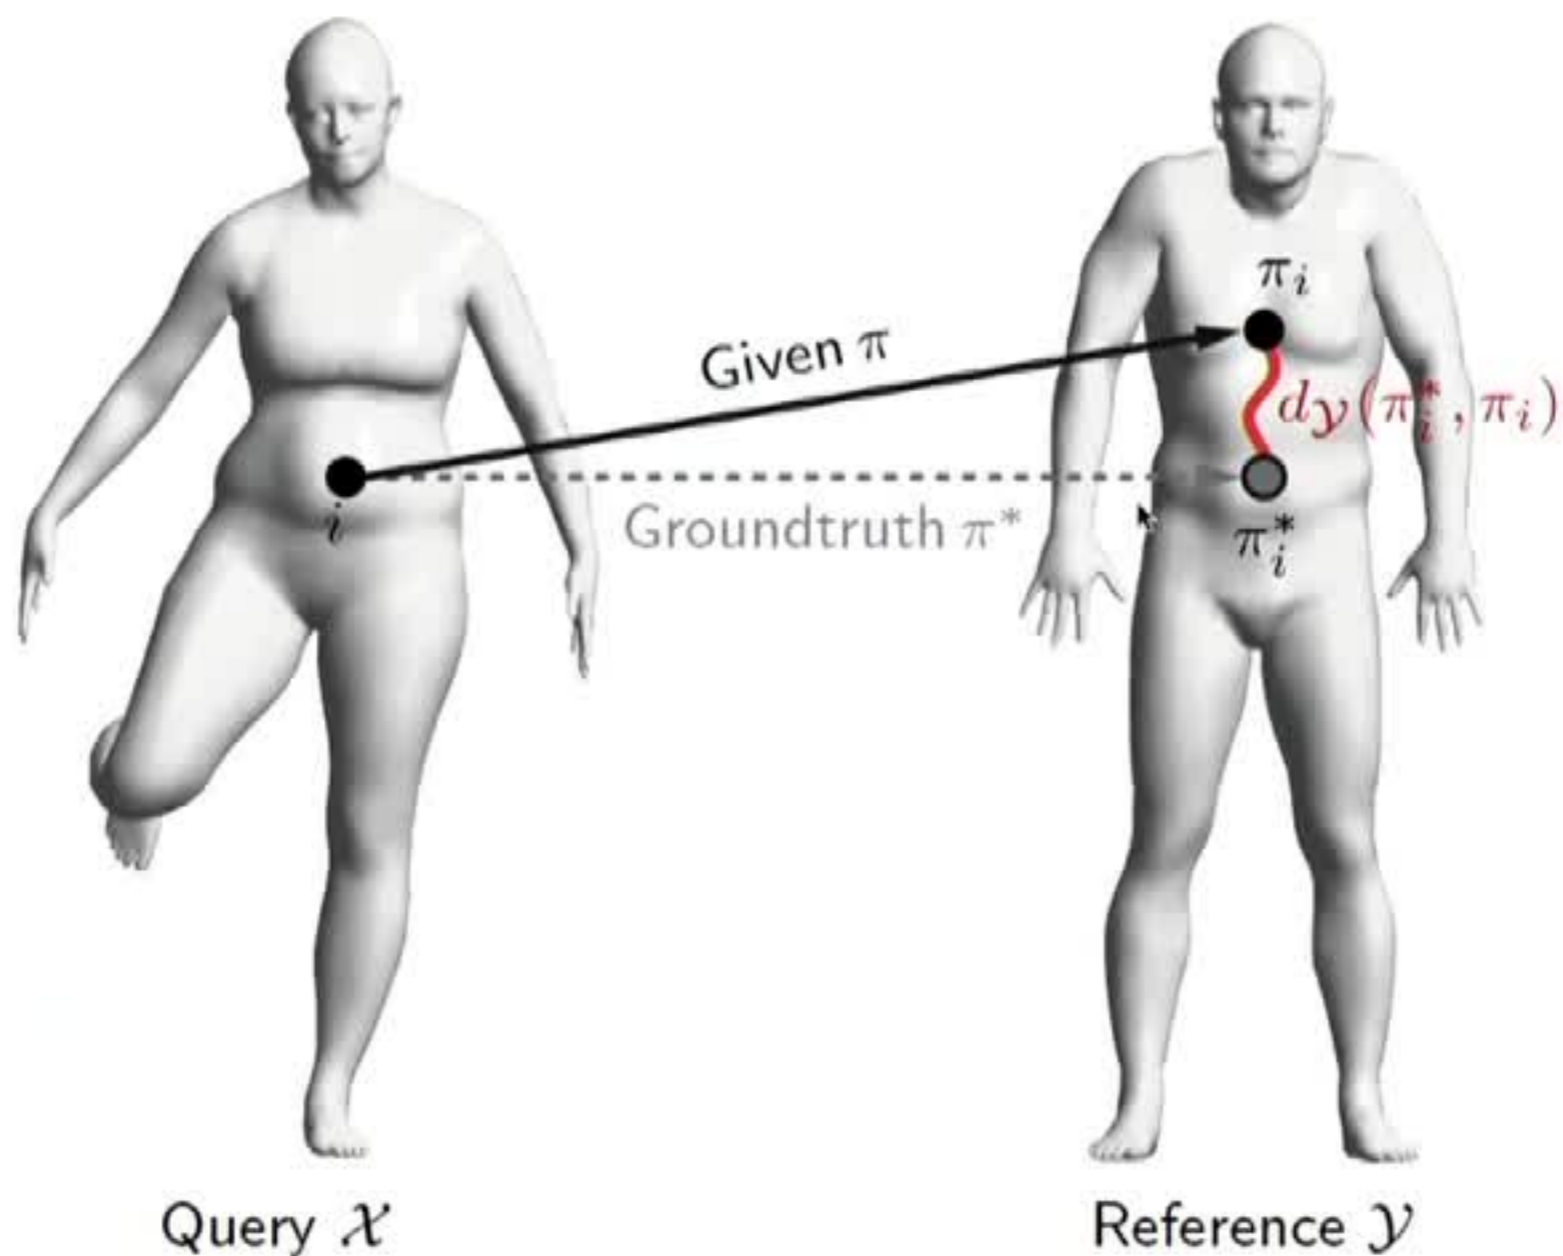
Descriptor quality comparison



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. 2011

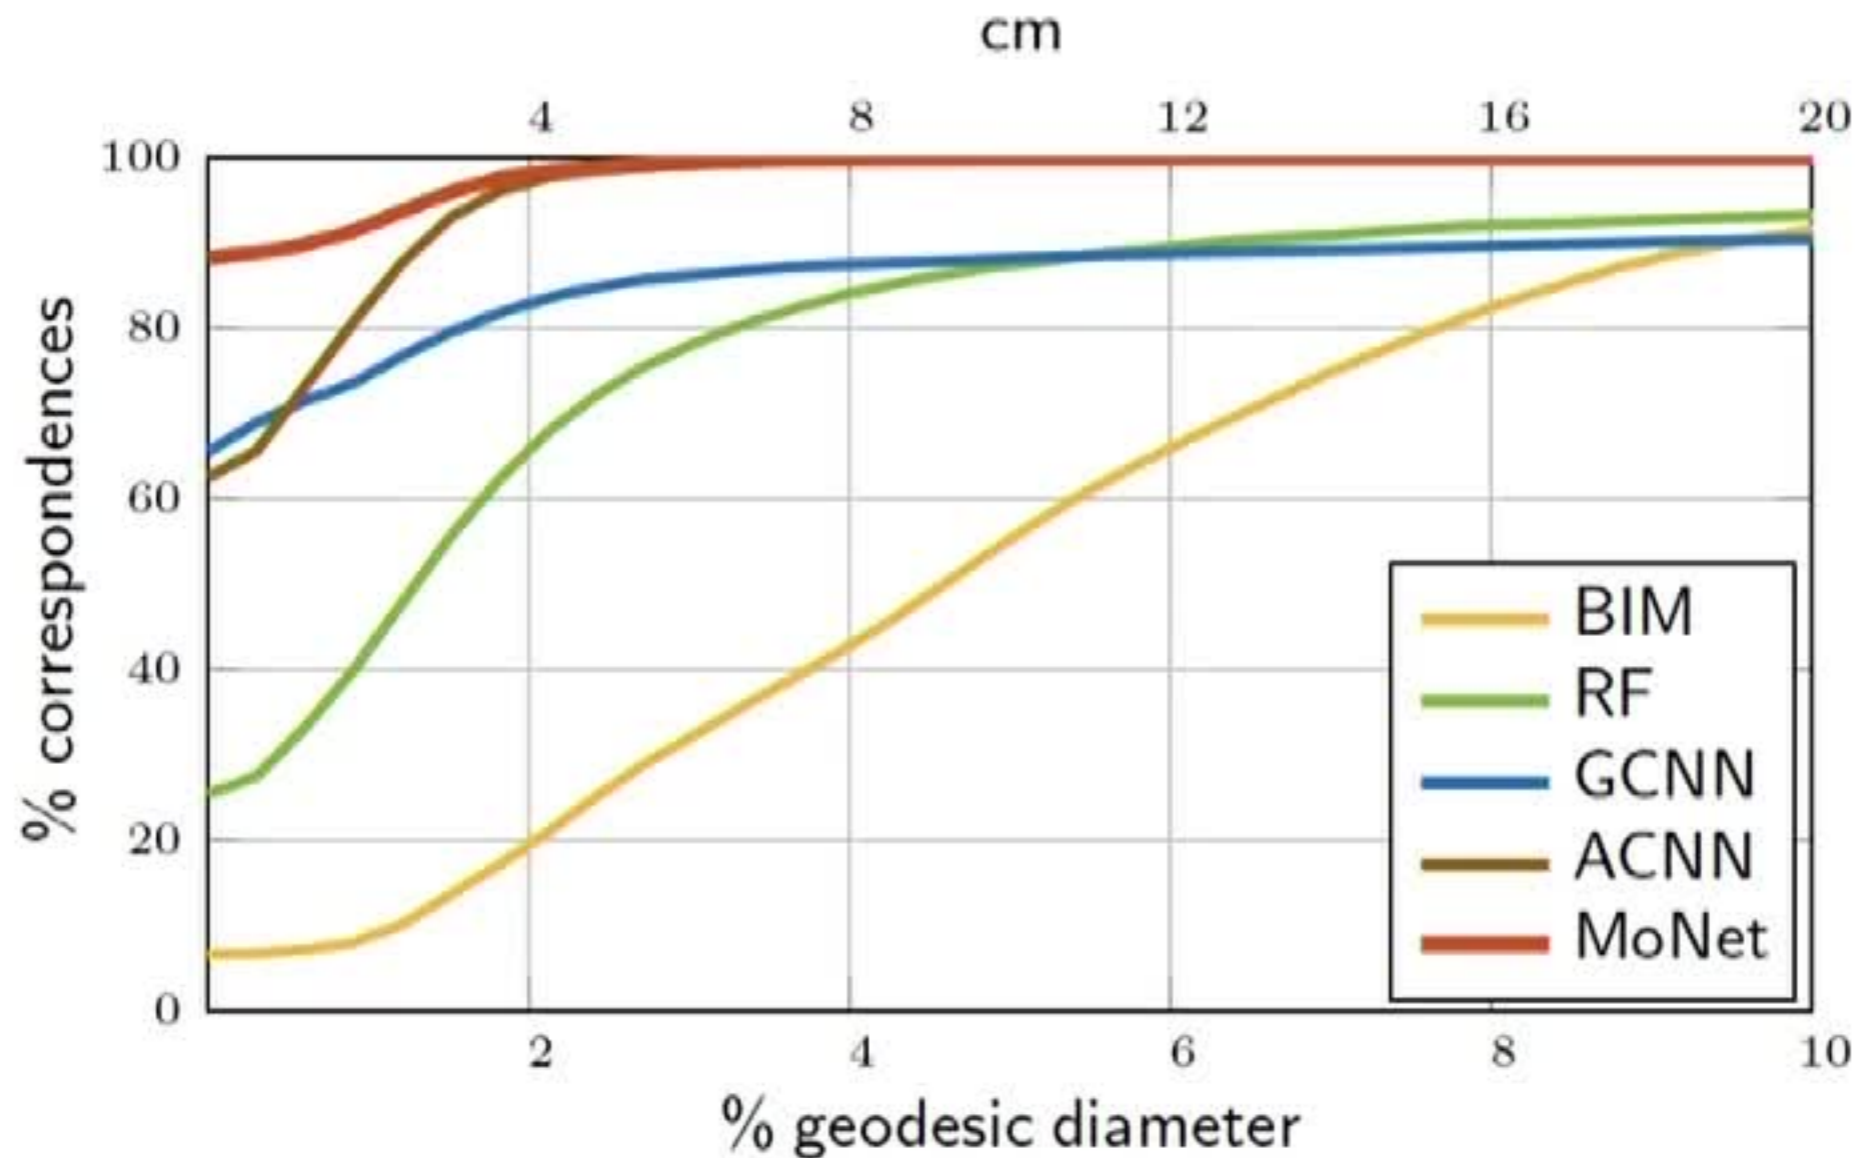
Correspondence evaluation: Princeton benchmark



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: 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

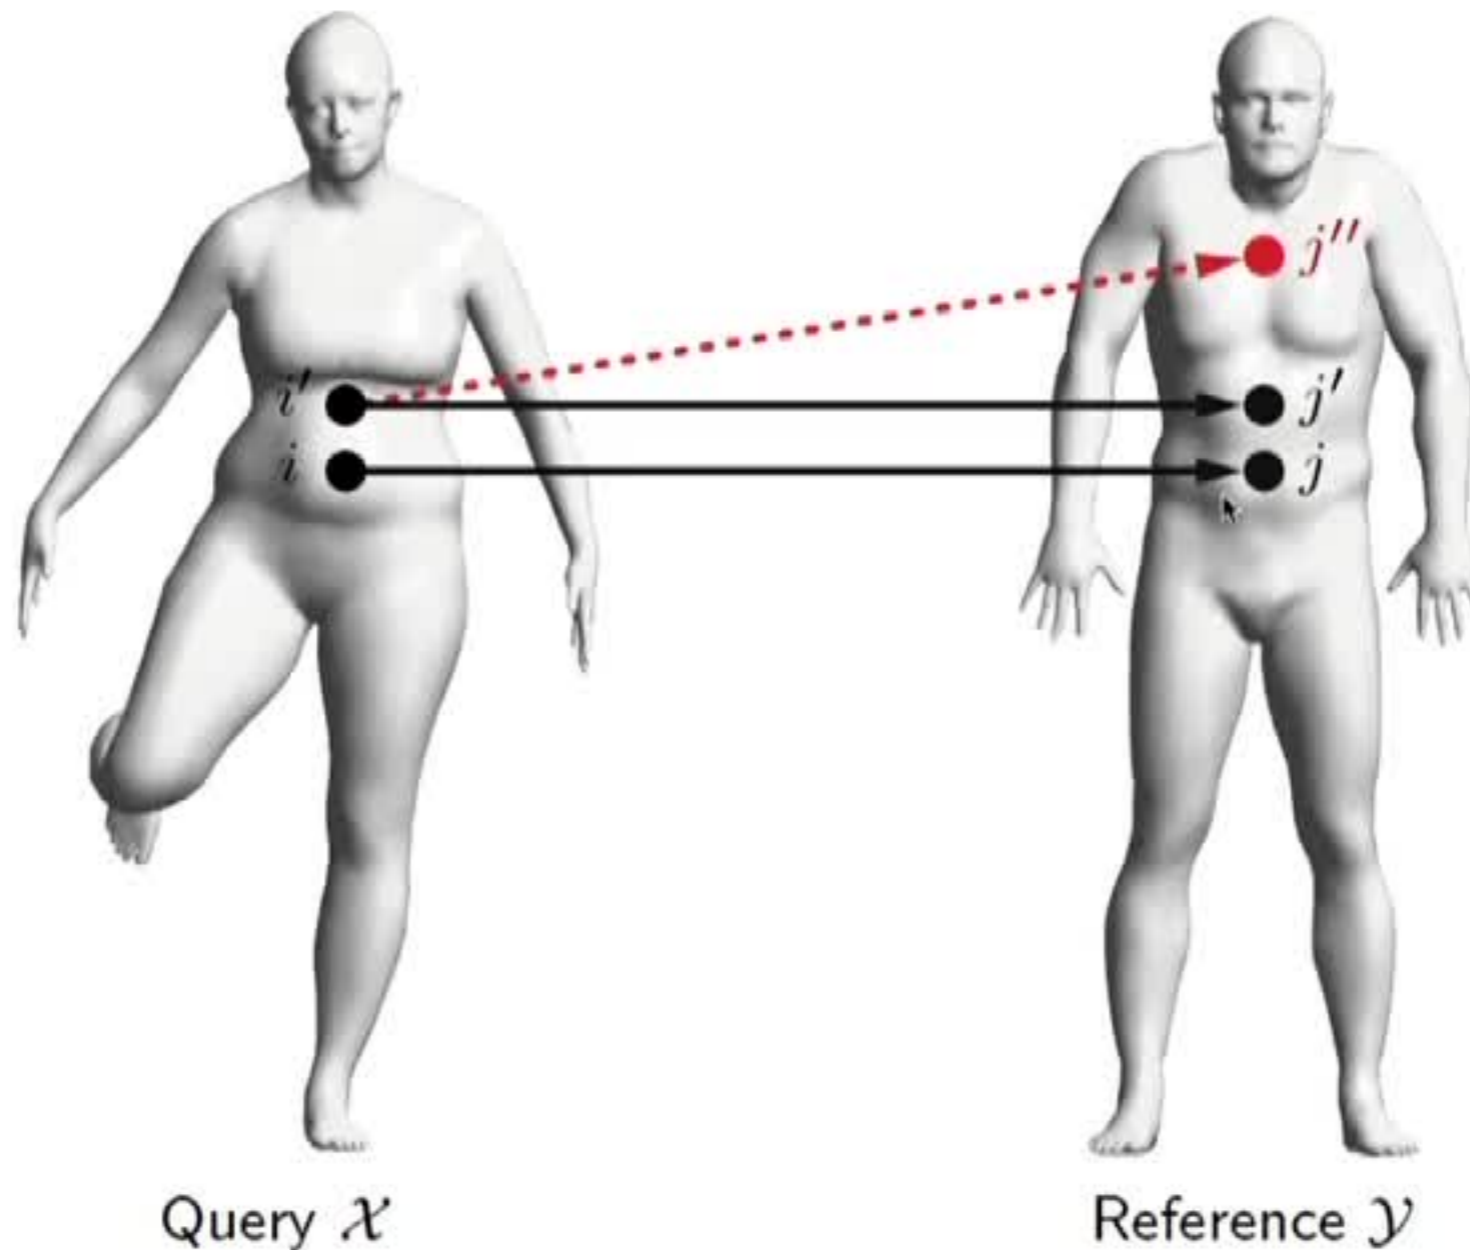


Reference



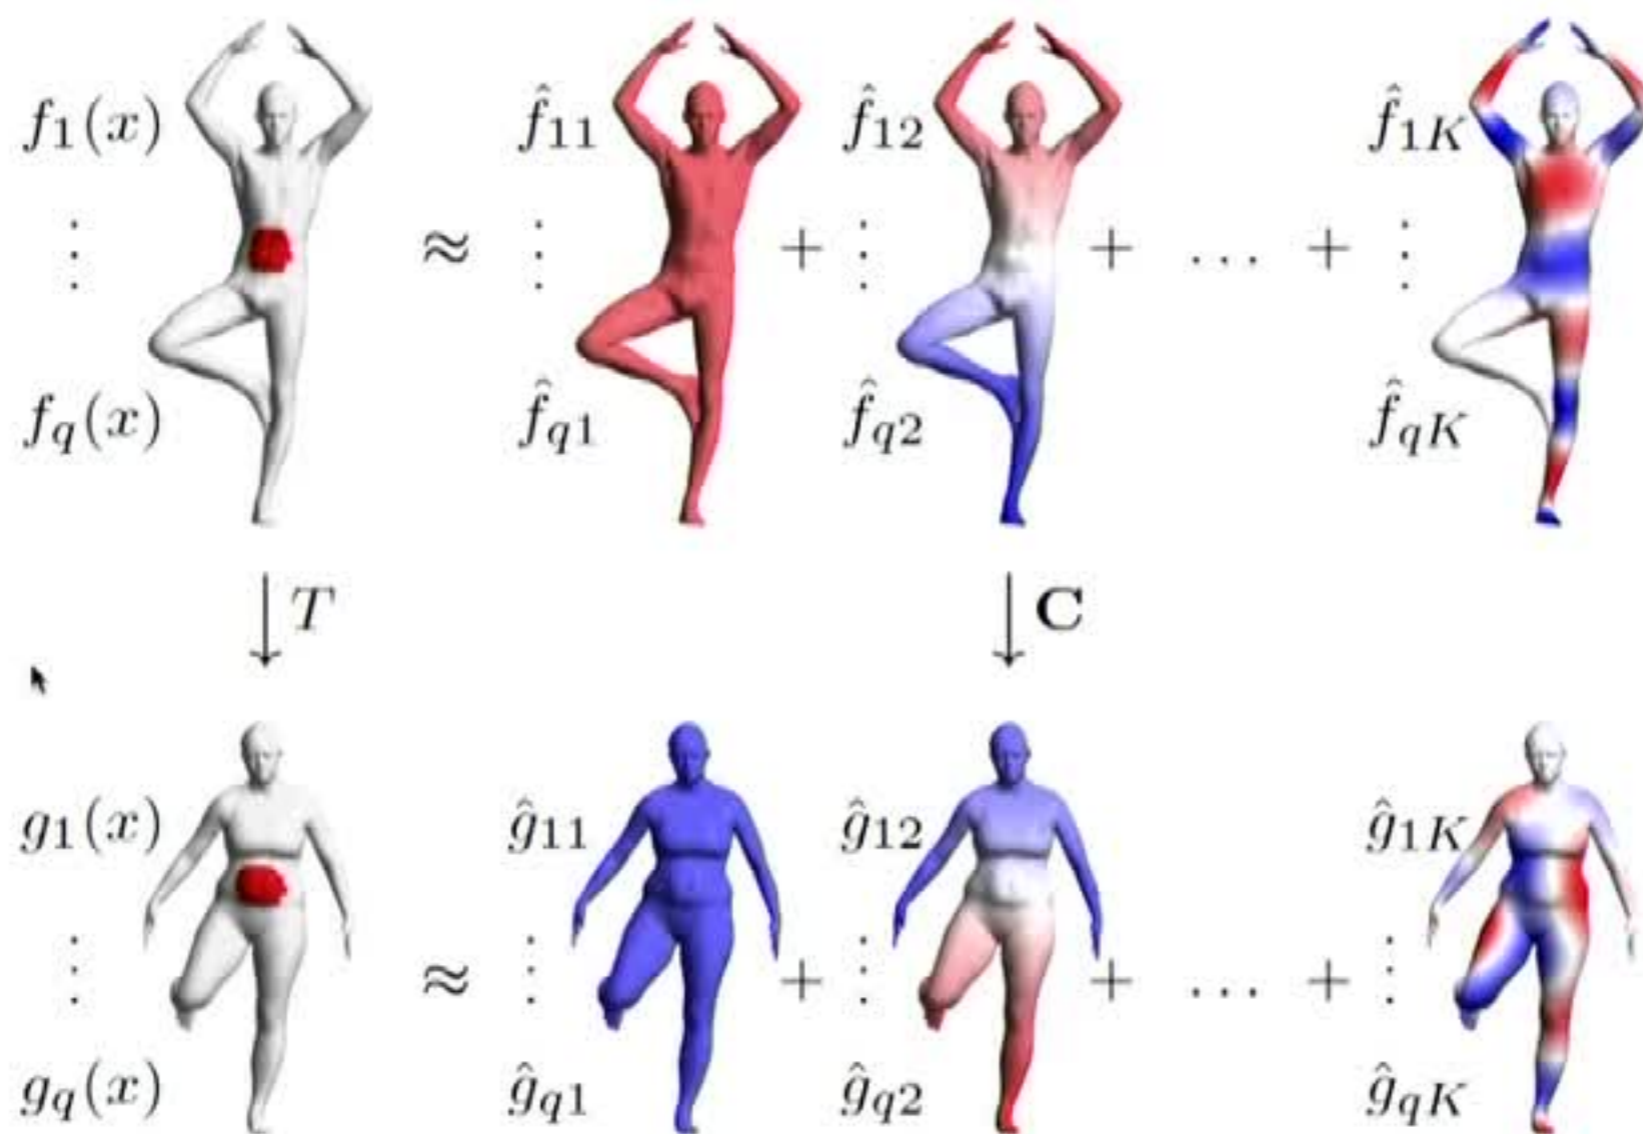
Correspondence visualization (similar colors encode corresponding points)

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



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

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

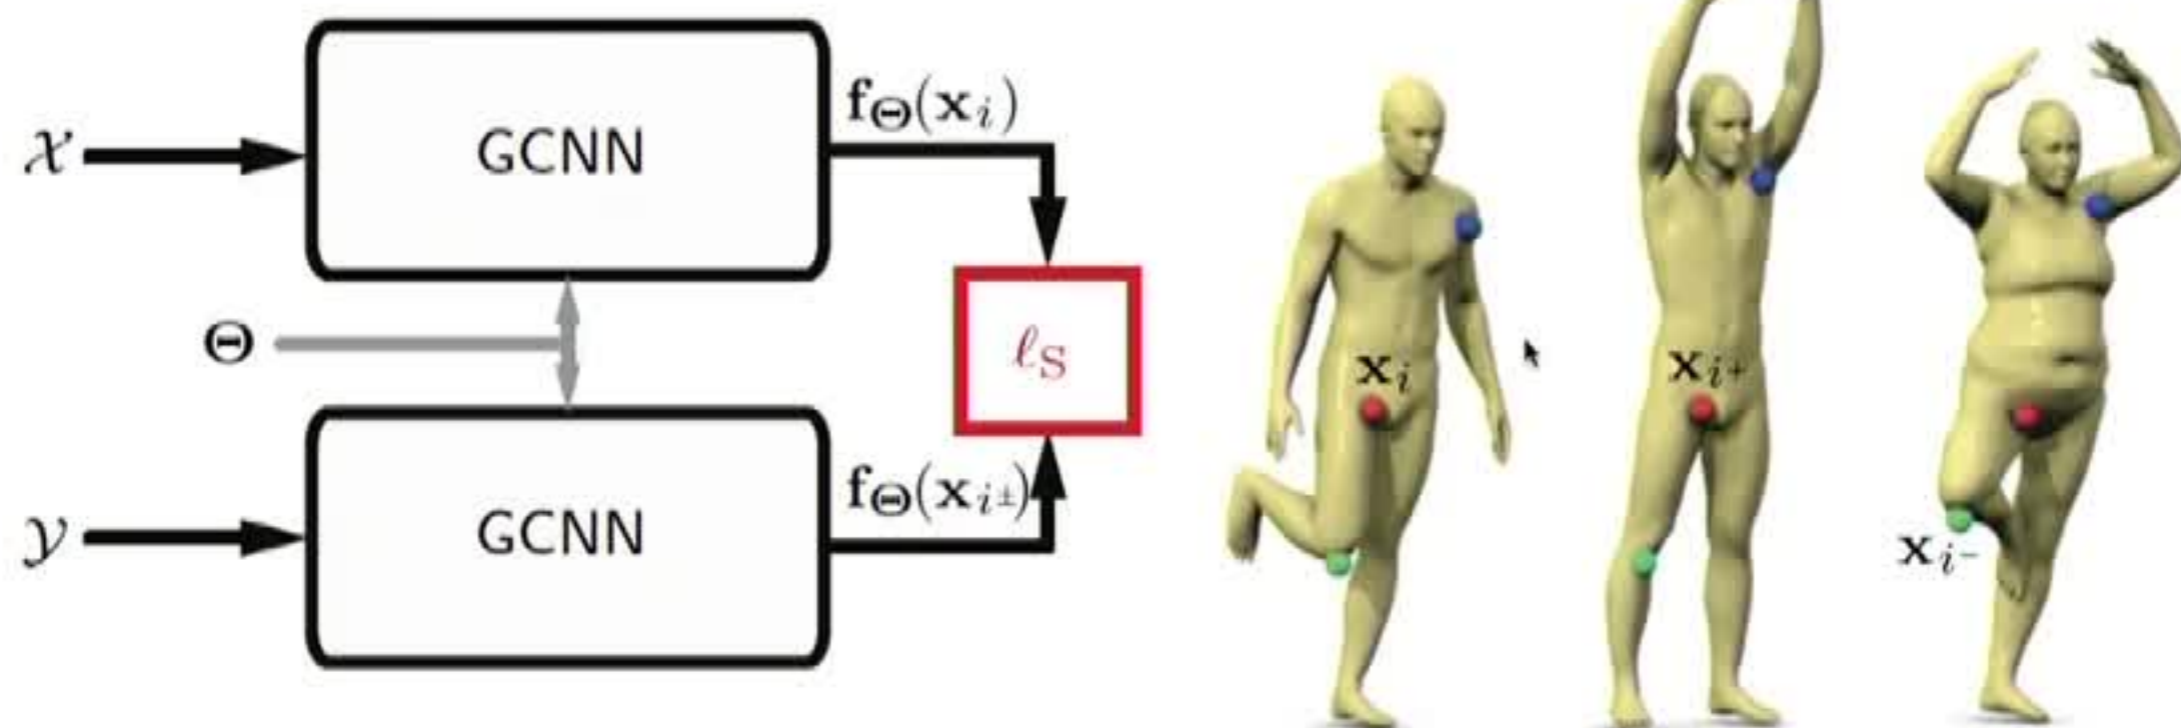
Functional maps: spatial domain



Probability p_{ij} of point j mapping to i

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

Siamese metric learning



Training set

Siamese net

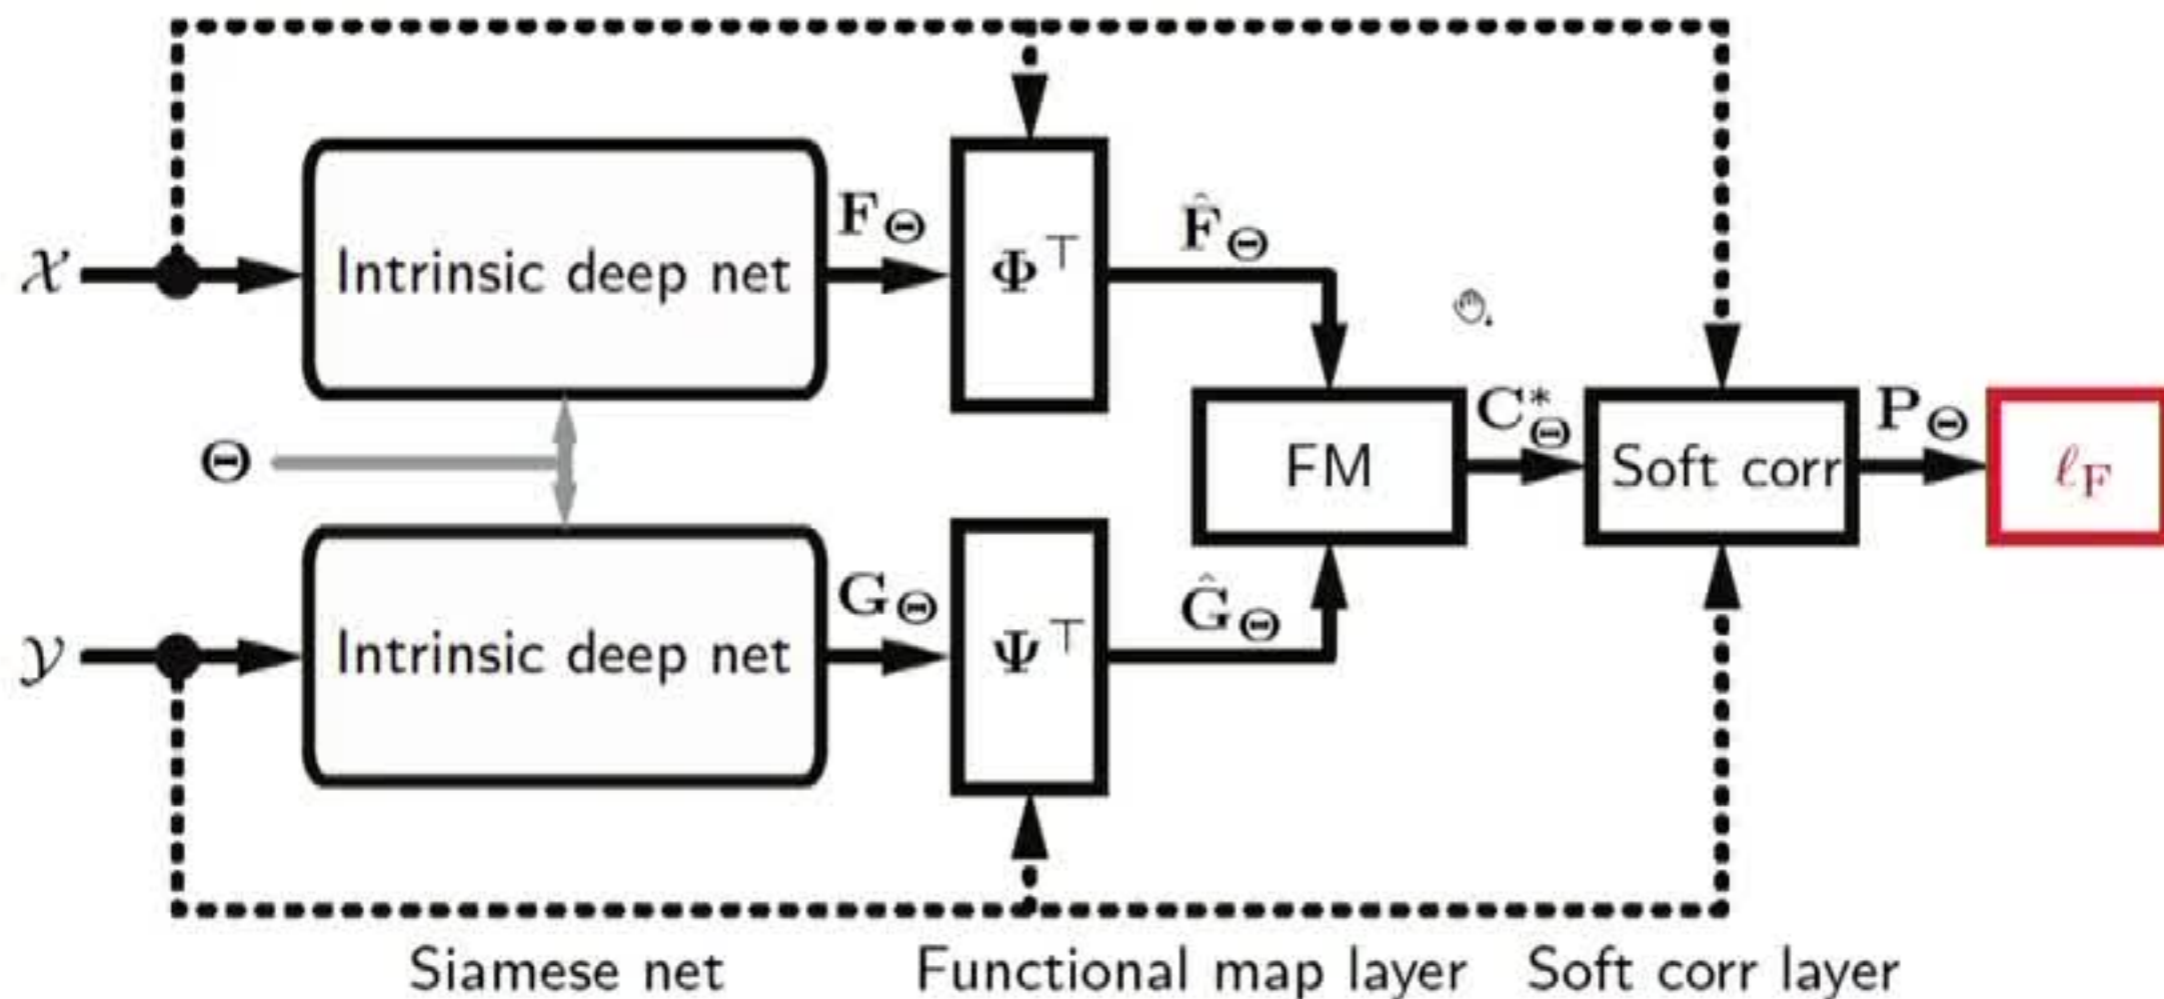
Pointwise feature cost

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

two net instances with shared parameters Θ

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

Structured correspondence with FMNet



Siamese net

two net instances with shared parameters Θ

Functional map layer

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

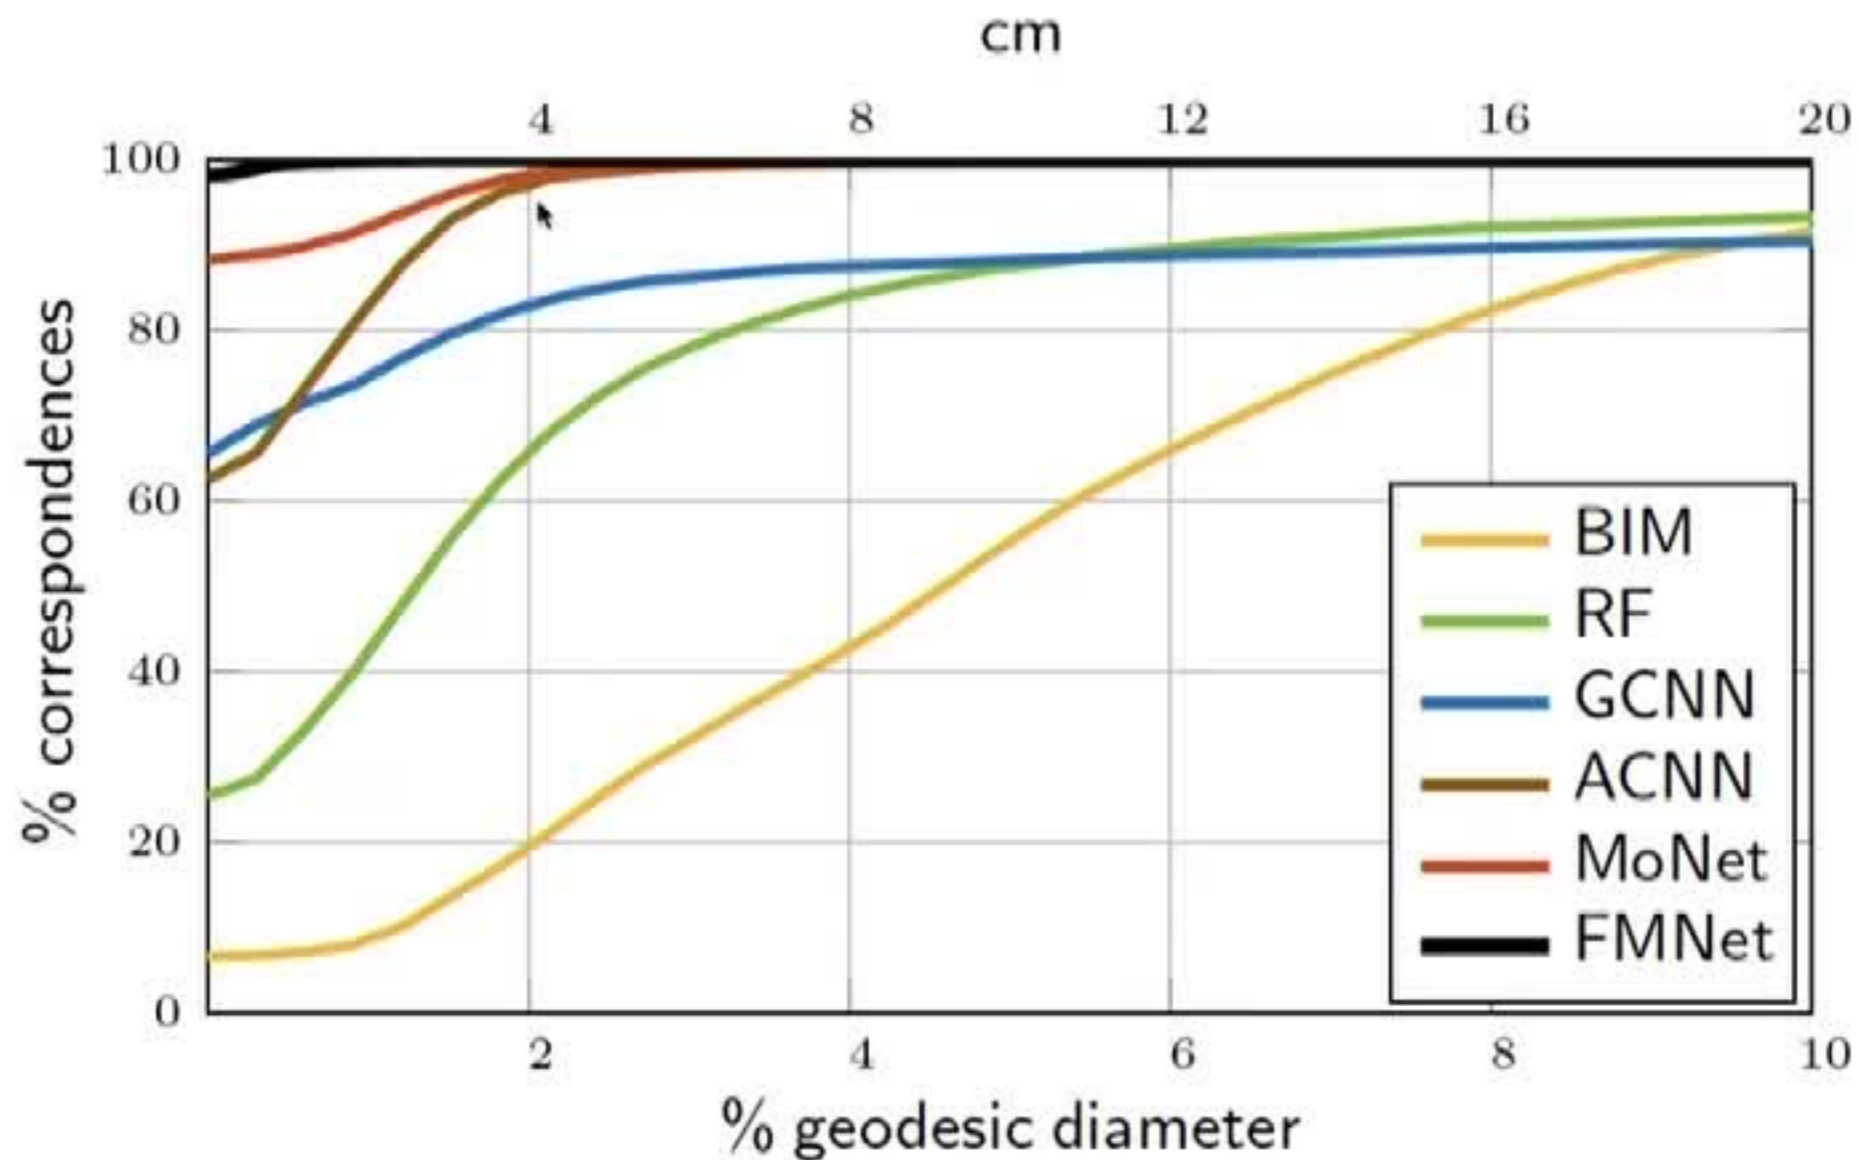
Soft correspondence layer

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

Soft error cost

$$\ell_F(\Theta) = \|P_{\Theta} \circ D_y\|$$

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

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



3D shape analysis and synthesis



Shape completion comparison




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

Shape completion examples



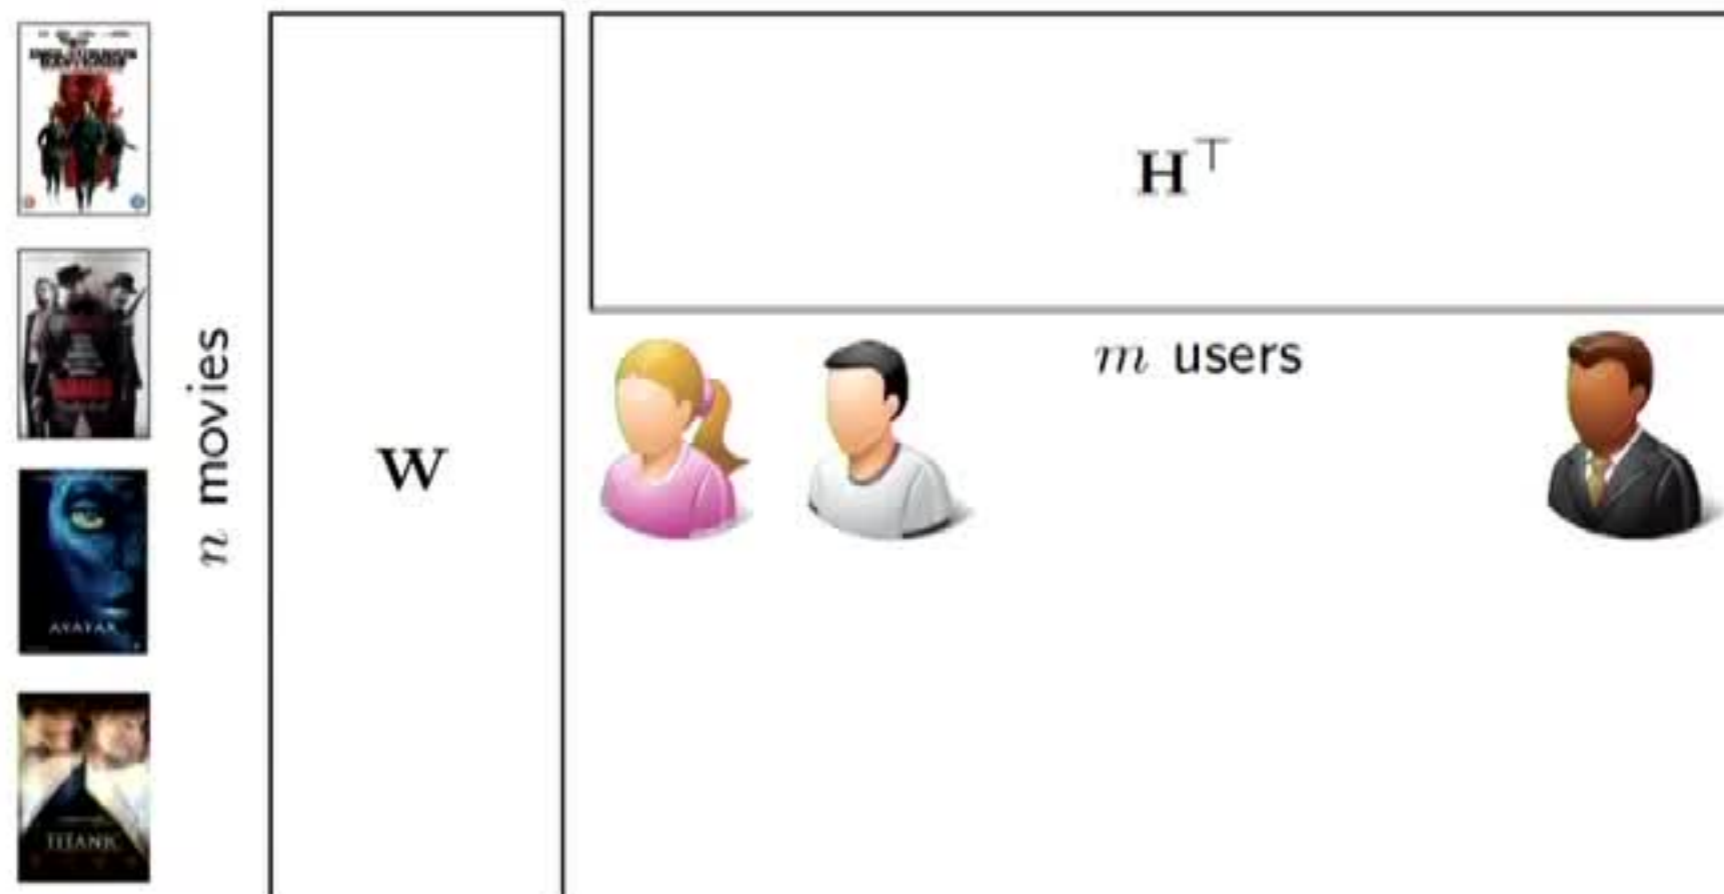
Shape completion examples





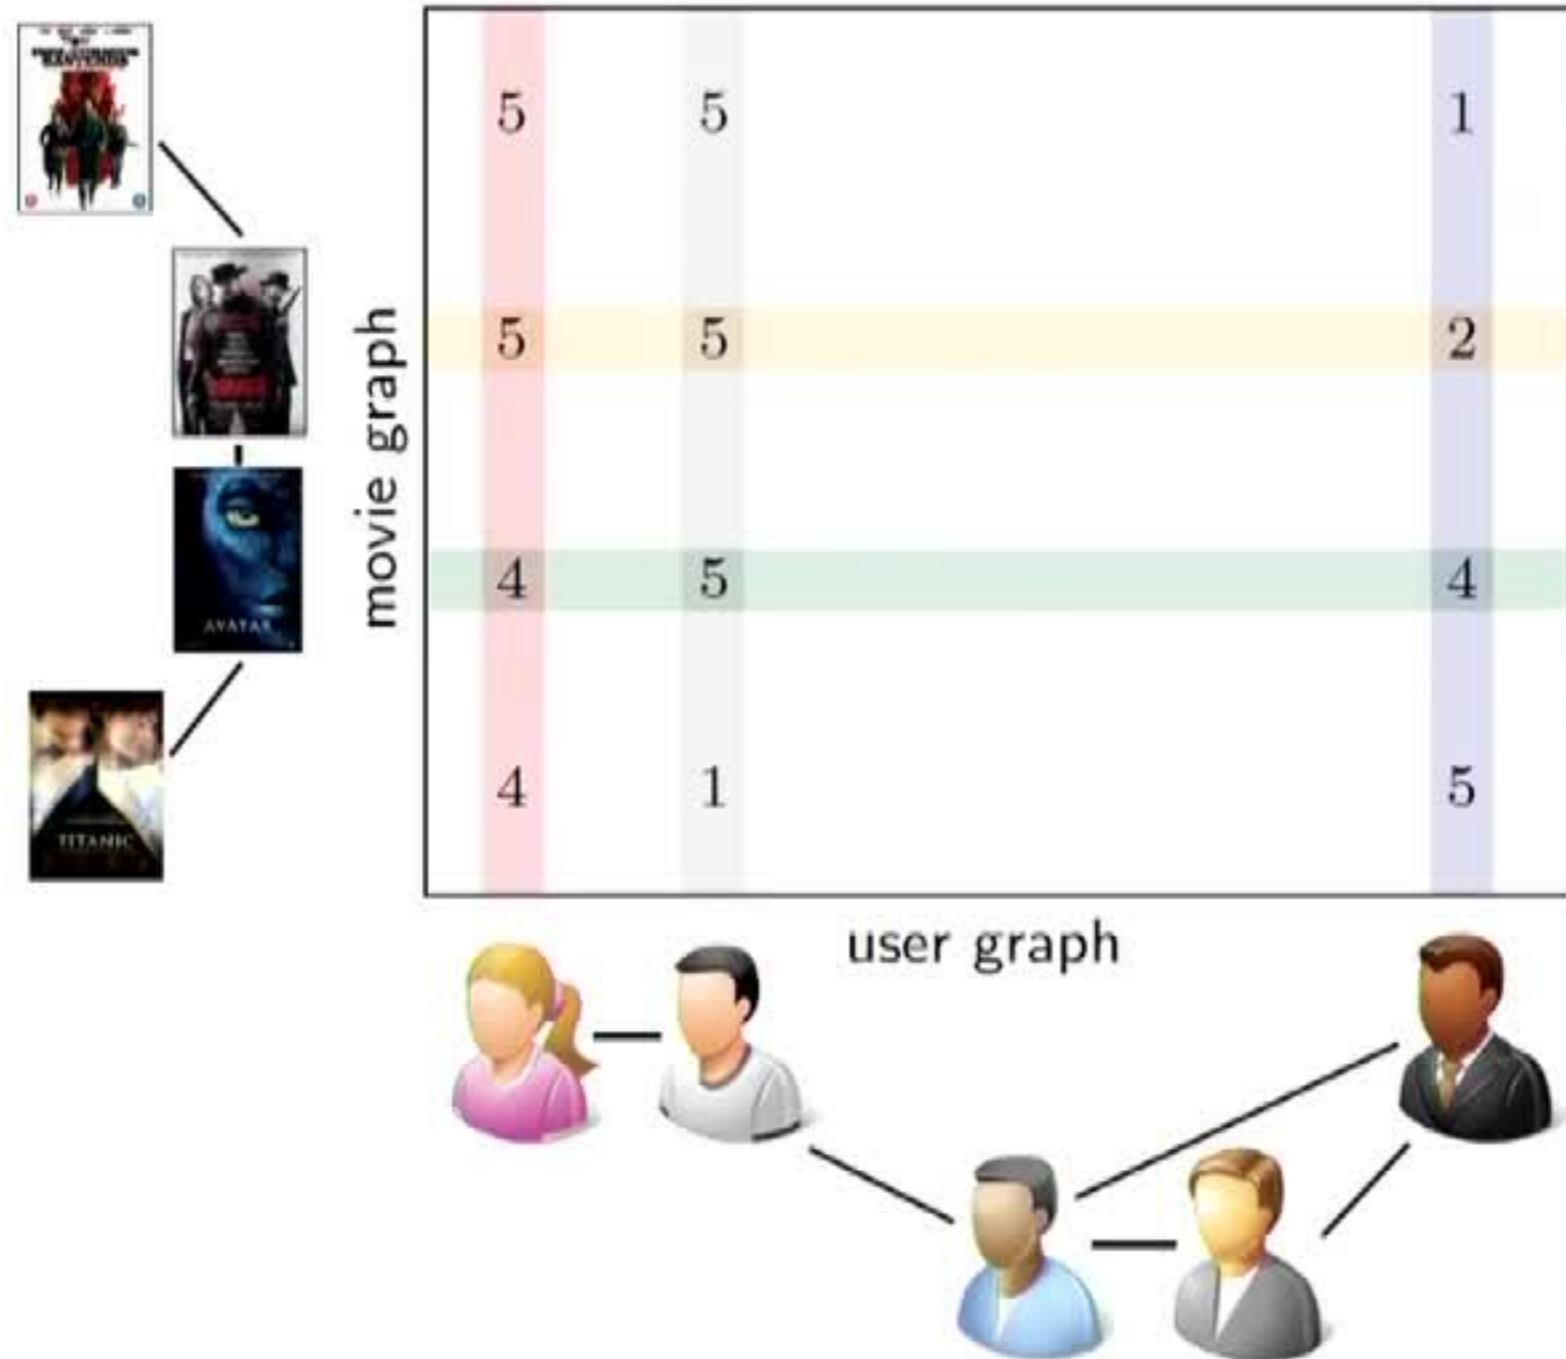


Factorized matrix completion models



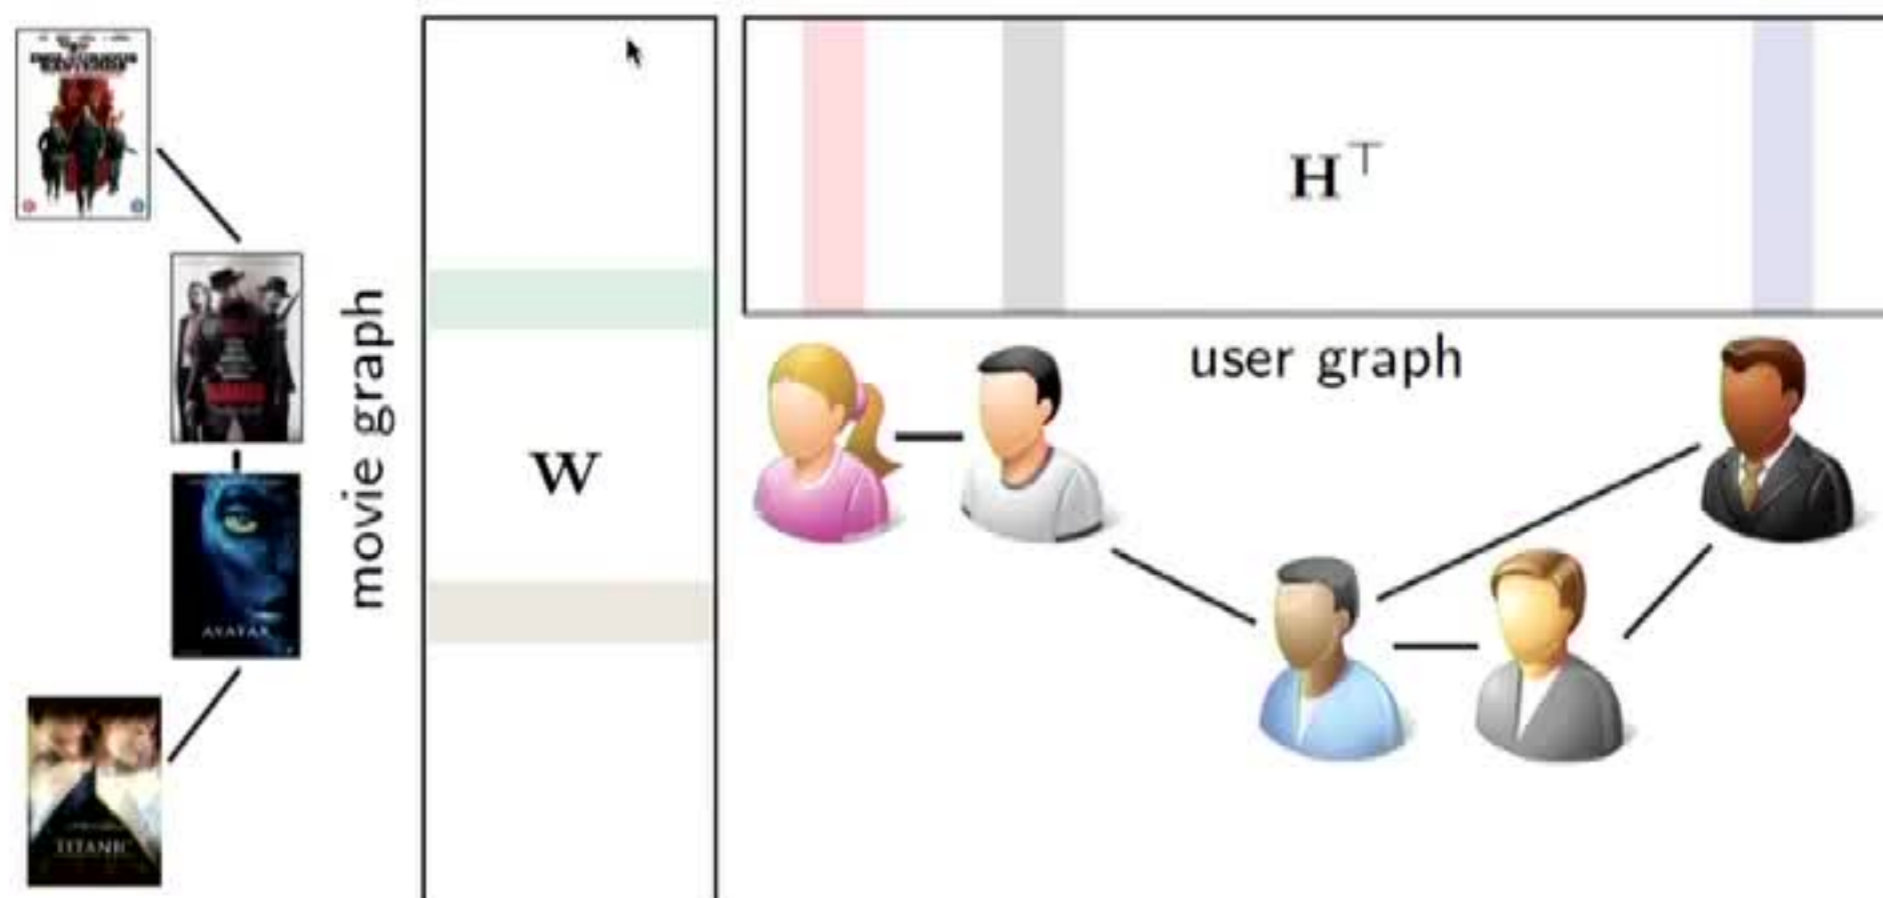
$$\min_{\substack{\mathbf{W} \in \mathbb{R}^{m \times s} \\ \mathbf{H} \in \mathbb{R}^{n \times s}}} \mu \|\Omega \circ (\mathbf{X} - \mathbf{A})\|_F^2 + \mu_c \|\mathbf{W}\|_F^2 + \mu_r \|\mathbf{H}\|_F^2$$

Geometric matrix completion



$$\min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \mu \|\Omega \circ (\mathbf{X} - \mathbf{A})\|_F^2 + \underbrace{\mu_c \text{tr}(\mathbf{X} \Delta_c \mathbf{X}^\top)}_{\|\mathbf{X}\|_{\mathcal{G}_c}^2} + \underbrace{\mu_r \text{tr}(\mathbf{X}^\top \Delta_r \mathbf{X})}_{\|\mathbf{X}\|_{\mathcal{G}_r}^2}$$

Factorized geometric matrix completion models



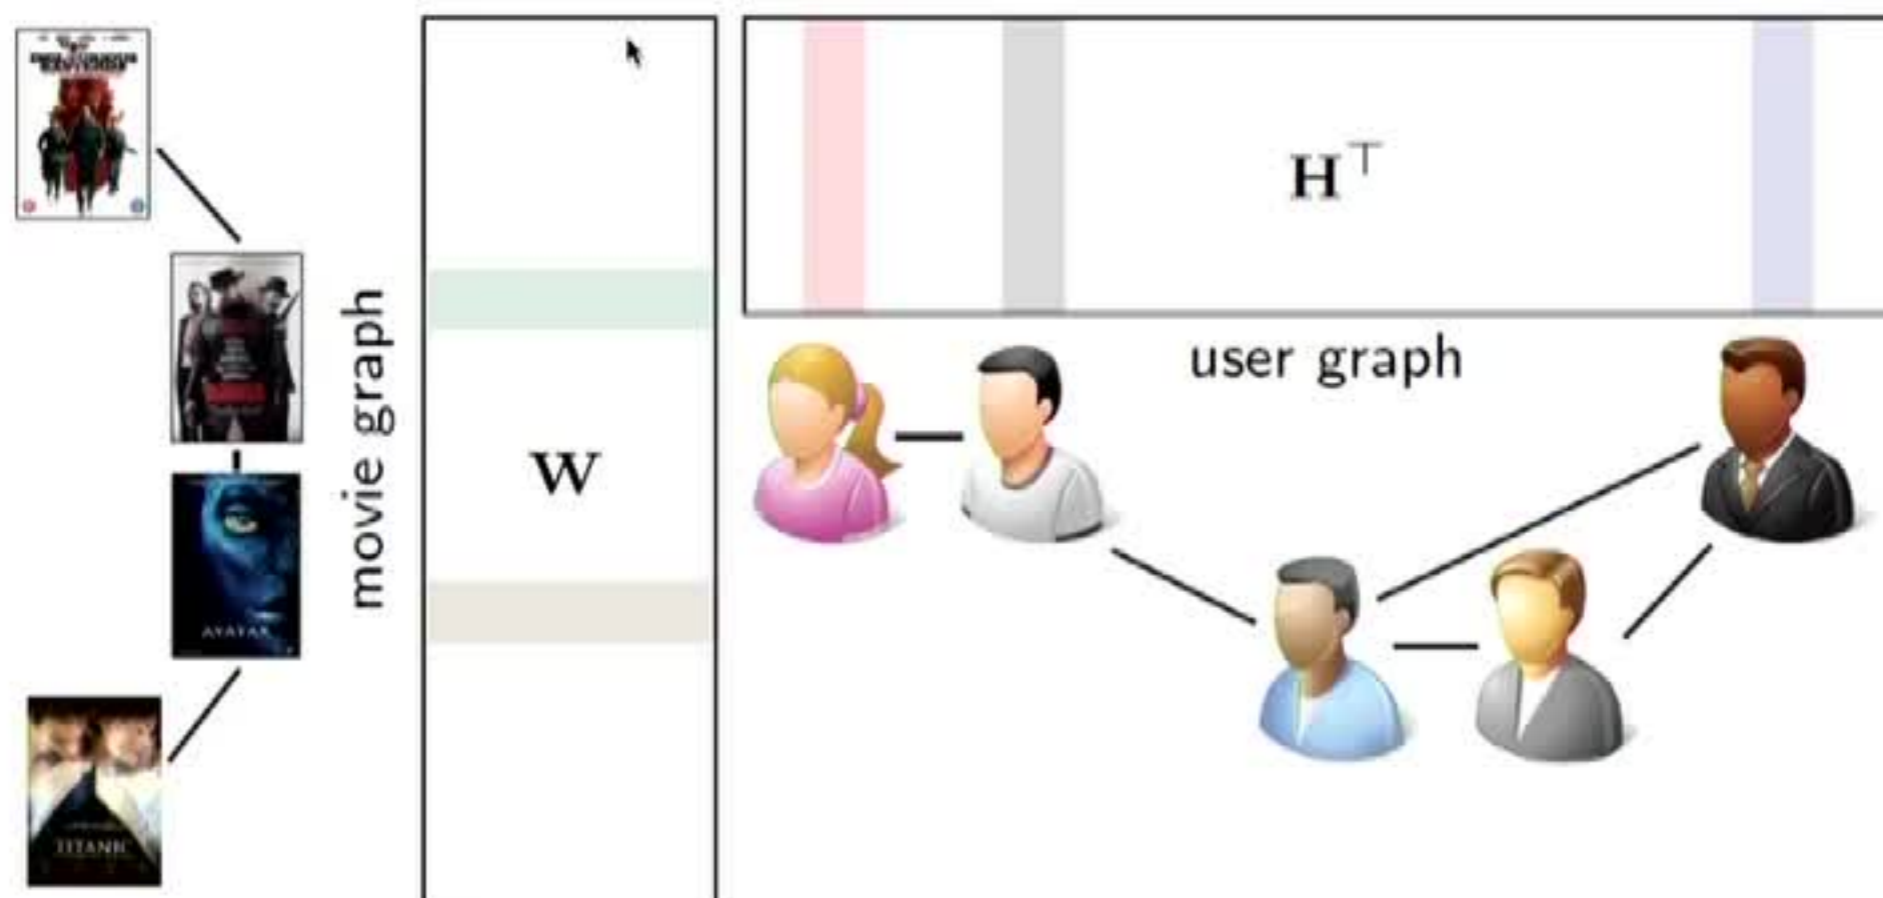
$$\min_{\substack{\mathbf{W} \in \mathbb{R}^{m \times s} \\ \mathbf{H} \in \mathbb{R}^{n \times s}}} \mu \|\Omega \circ (\mathbf{X} - \mathbf{A})\|_F^2 + \mu_c \text{tr}(\mathbf{H}^T \Delta_c \mathbf{H}) + \mu_r \text{tr}(\mathbf{W}^T \Delta_r \mathbf{W})$$

☺ $\mathcal{O}(n + m)$ variables instead of $\mathcal{O}(nm)$

☺ $\text{rank}(\mathbf{X}) = \text{rank}(\mathbf{W}\mathbf{H}^T) \leq s$ by construction

How to learn filters on
multiple graphs?

Factorized geometric matrix completion models



$$\min_{\substack{\mathbf{W} \in \mathbb{R}^{m \times s} \\ \mathbf{H} \in \mathbb{R}^{n \times s}}} \mu \|\Omega \circ (\mathbf{X} - \mathbf{A})\|_F^2 + \mu_c \text{tr}(\mathbf{H}^T \Delta_c \mathbf{H}) + \mu_r \text{tr}(\mathbf{W}^T \Delta_r \mathbf{W})$$

☺ $\mathcal{O}(n + m)$ variables instead of $\mathcal{O}(nm)$


☺ $\text{rank}(\mathbf{X}) = \text{rank}(\mathbf{W}\mathbf{H}^T) \leq s$ by construction

2D Fourier transform



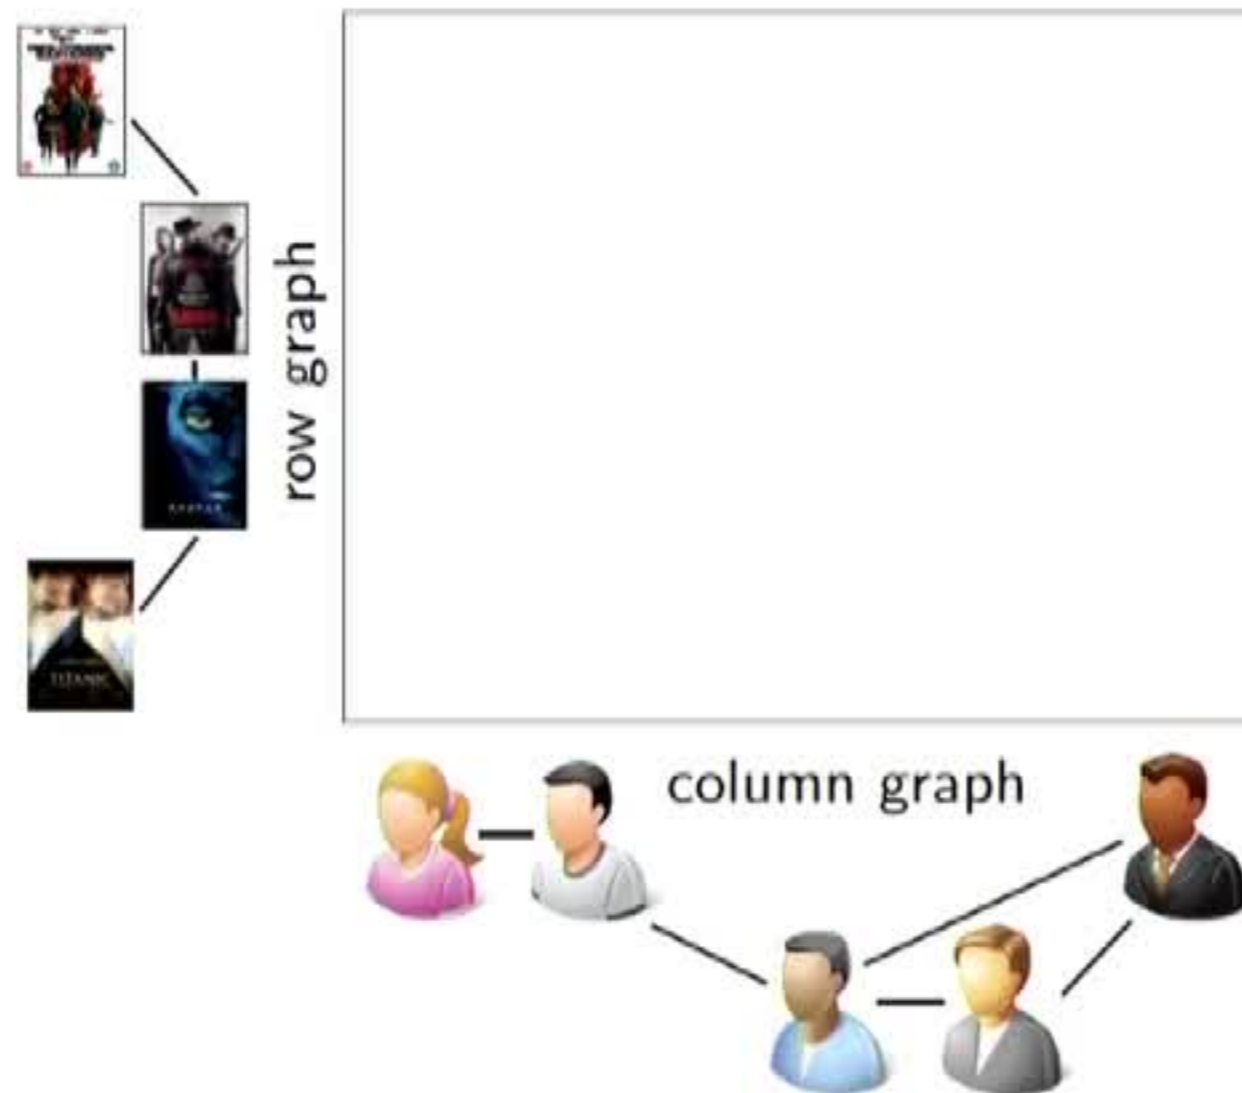
X

2D Fourier transform



Column-wise transform + Row-wise transform = 2D transform

Multi-graph spectral filters




Multi-graph bi-variate polynomial filter

$$\mathbf{Y} = \tau_{\Theta}(\mathbf{X}) = \sum_{j,j'=1}^r \theta_{jj'} \Delta_r^j \mathbf{X} \Delta_c^{j'}$$

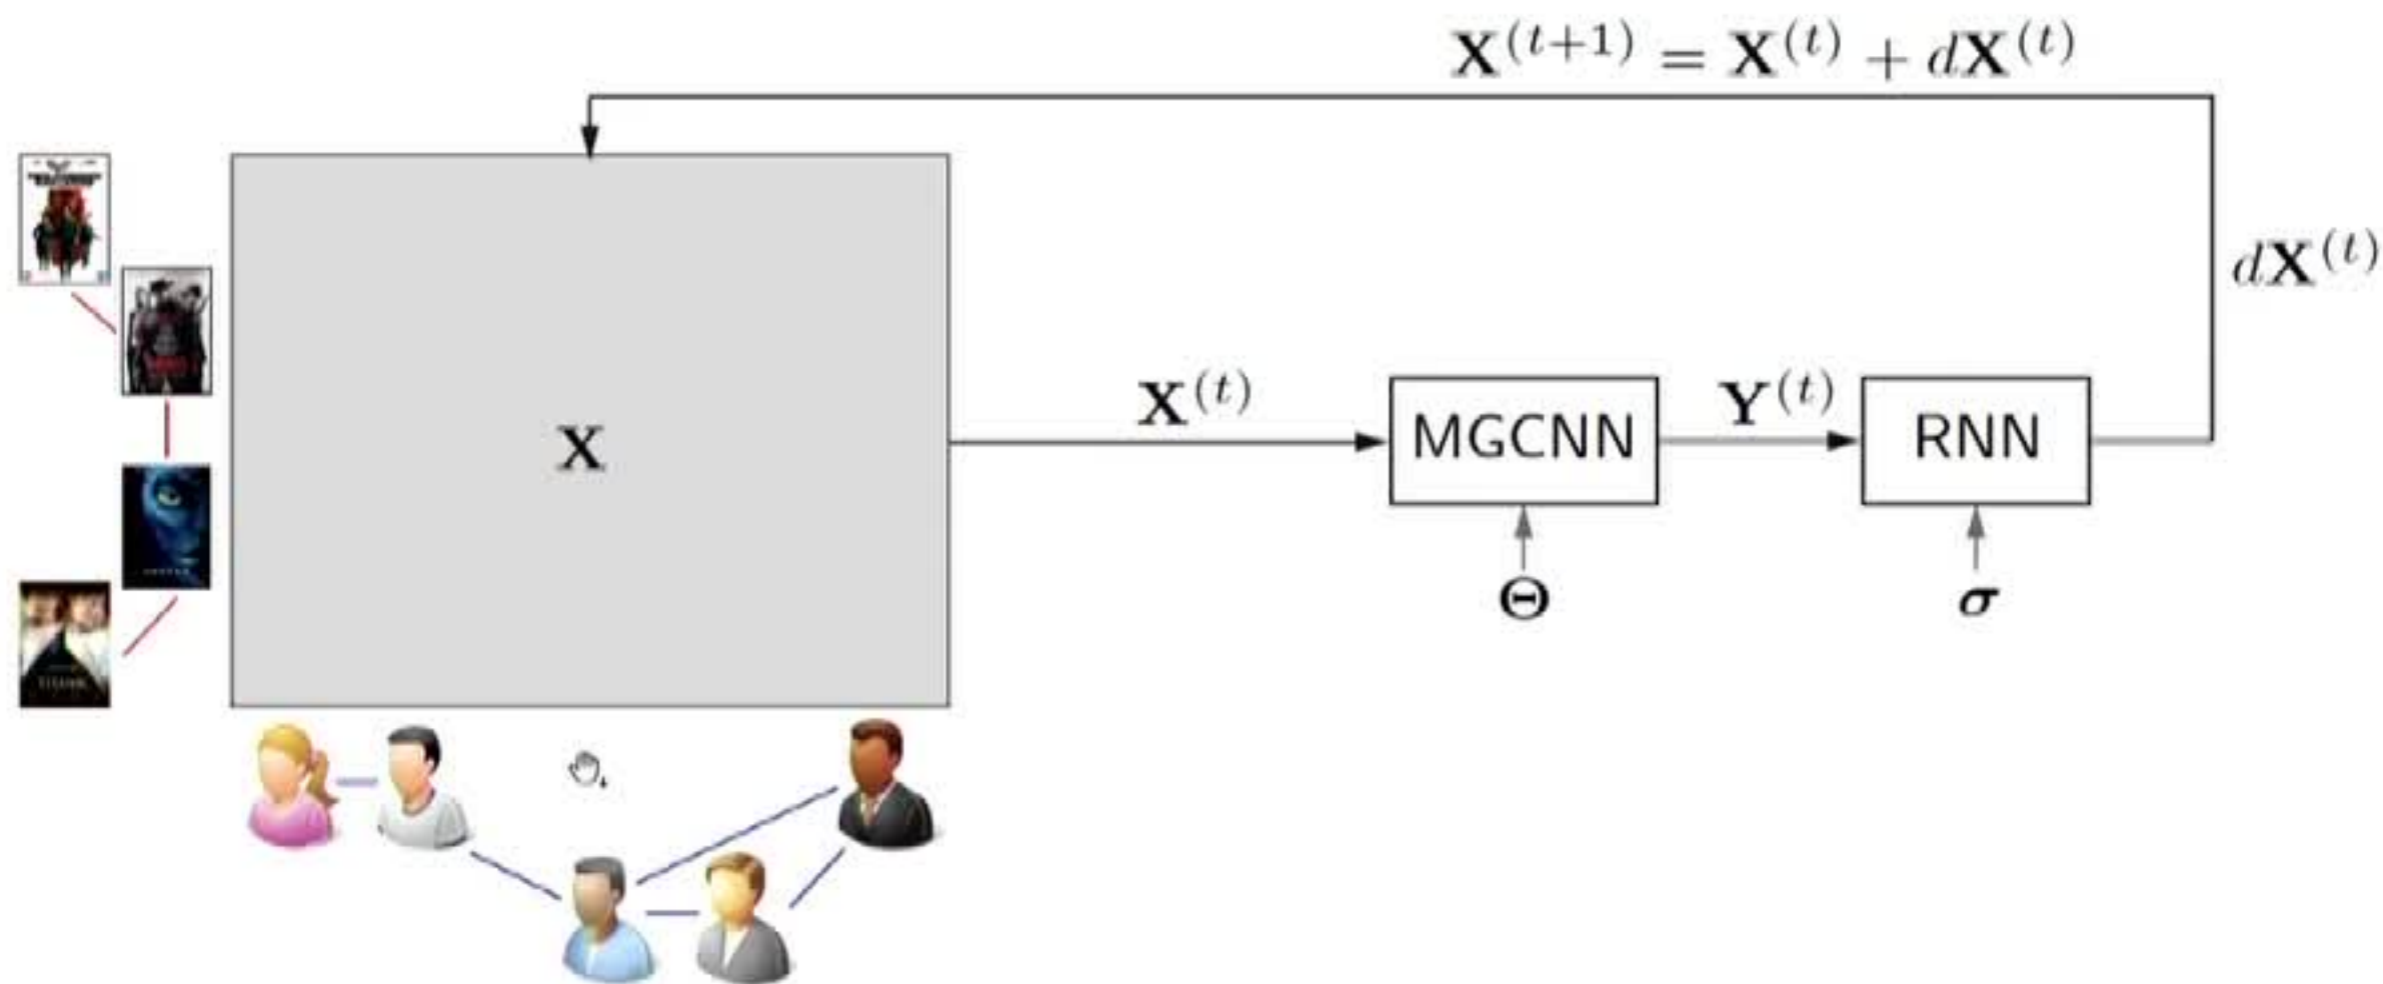
where $\Theta = (\theta_{jj'})$ is the $r \times r$ matrix of filter parameters

Multi-graph spectral filters



Examples of multi-graph spectral filters (shown is $|\tau(\lambda_c, \lambda_r)|$)


Matrix completion with Recurrent Multi-Graph CNN



Recurrent multigraph CNN (RMCNN) architecture for matrix completion

$$\min_{\Theta, \sigma} \|\mathbf{X}_{\Theta, \sigma}^{(T)}\|_{\mathcal{G}_r}^2 + \|\mathbf{X}_{\Theta, \sigma}^{(T)}\|_{\mathcal{G}_c}^2 + \frac{\mu}{2} \|\Omega \circ (\mathbf{X}_{\Theta, \sigma}^{(T)} - \mathbf{A})\|_{\mathbb{F}}^2$$


Incremental updates with RNN



RMSE=2.26

Matrix completion results on a synthetic dataset (initialization)

Incremental updates with RNN



RMSE=0.01

Matrix completion results on a synthetic dataset after $t = 10$ instances


Matrix completion methods comparison

Method	MovieLens ¹	Flixster ²	Douban ³	Yahoo ⁴
IMC ⁵	1.653	–	–	–
GMC ⁶	0.996	–	–	–
MC ⁷	0.973	–	–	–
GRALS ⁸	0.945	1.245	0.833	38.042
sRGCNN (Cheb)⁹	0.929	0.926	0.801	22.415
sRGCNN (Cayley) ¹⁰	0.922	–	–	–
sRGCNN (Primal/Dual)¹¹	0.915	0.902	0.789	21.970


Accuracy (RMS error) of matrix completion methods on real data


Data: ¹Miller et al. 2003; ²Jamali, Ester 2010; ³Ma et al. 2011; ⁴Dror et al. 2012
Methods: ⁵Jain, Dhillon 2013; ⁶Kalofolias et al. 2014; ⁷Candès, Recht 2012; ⁸Rao et al. 2015; ⁹Monti, Bresson, Bronstein 2017; ¹⁰Levie et al. 2017; ¹¹Monti et al. 2018

Application in High-Energy Physics





ATLAS detector in the Large Hadron Collider (CERN)







ATLAS detector in the Large Hadron Collider (CERN)






IceCube neutrino observatory



IceCube neutrino observatory

Neutrino detection in IceCube

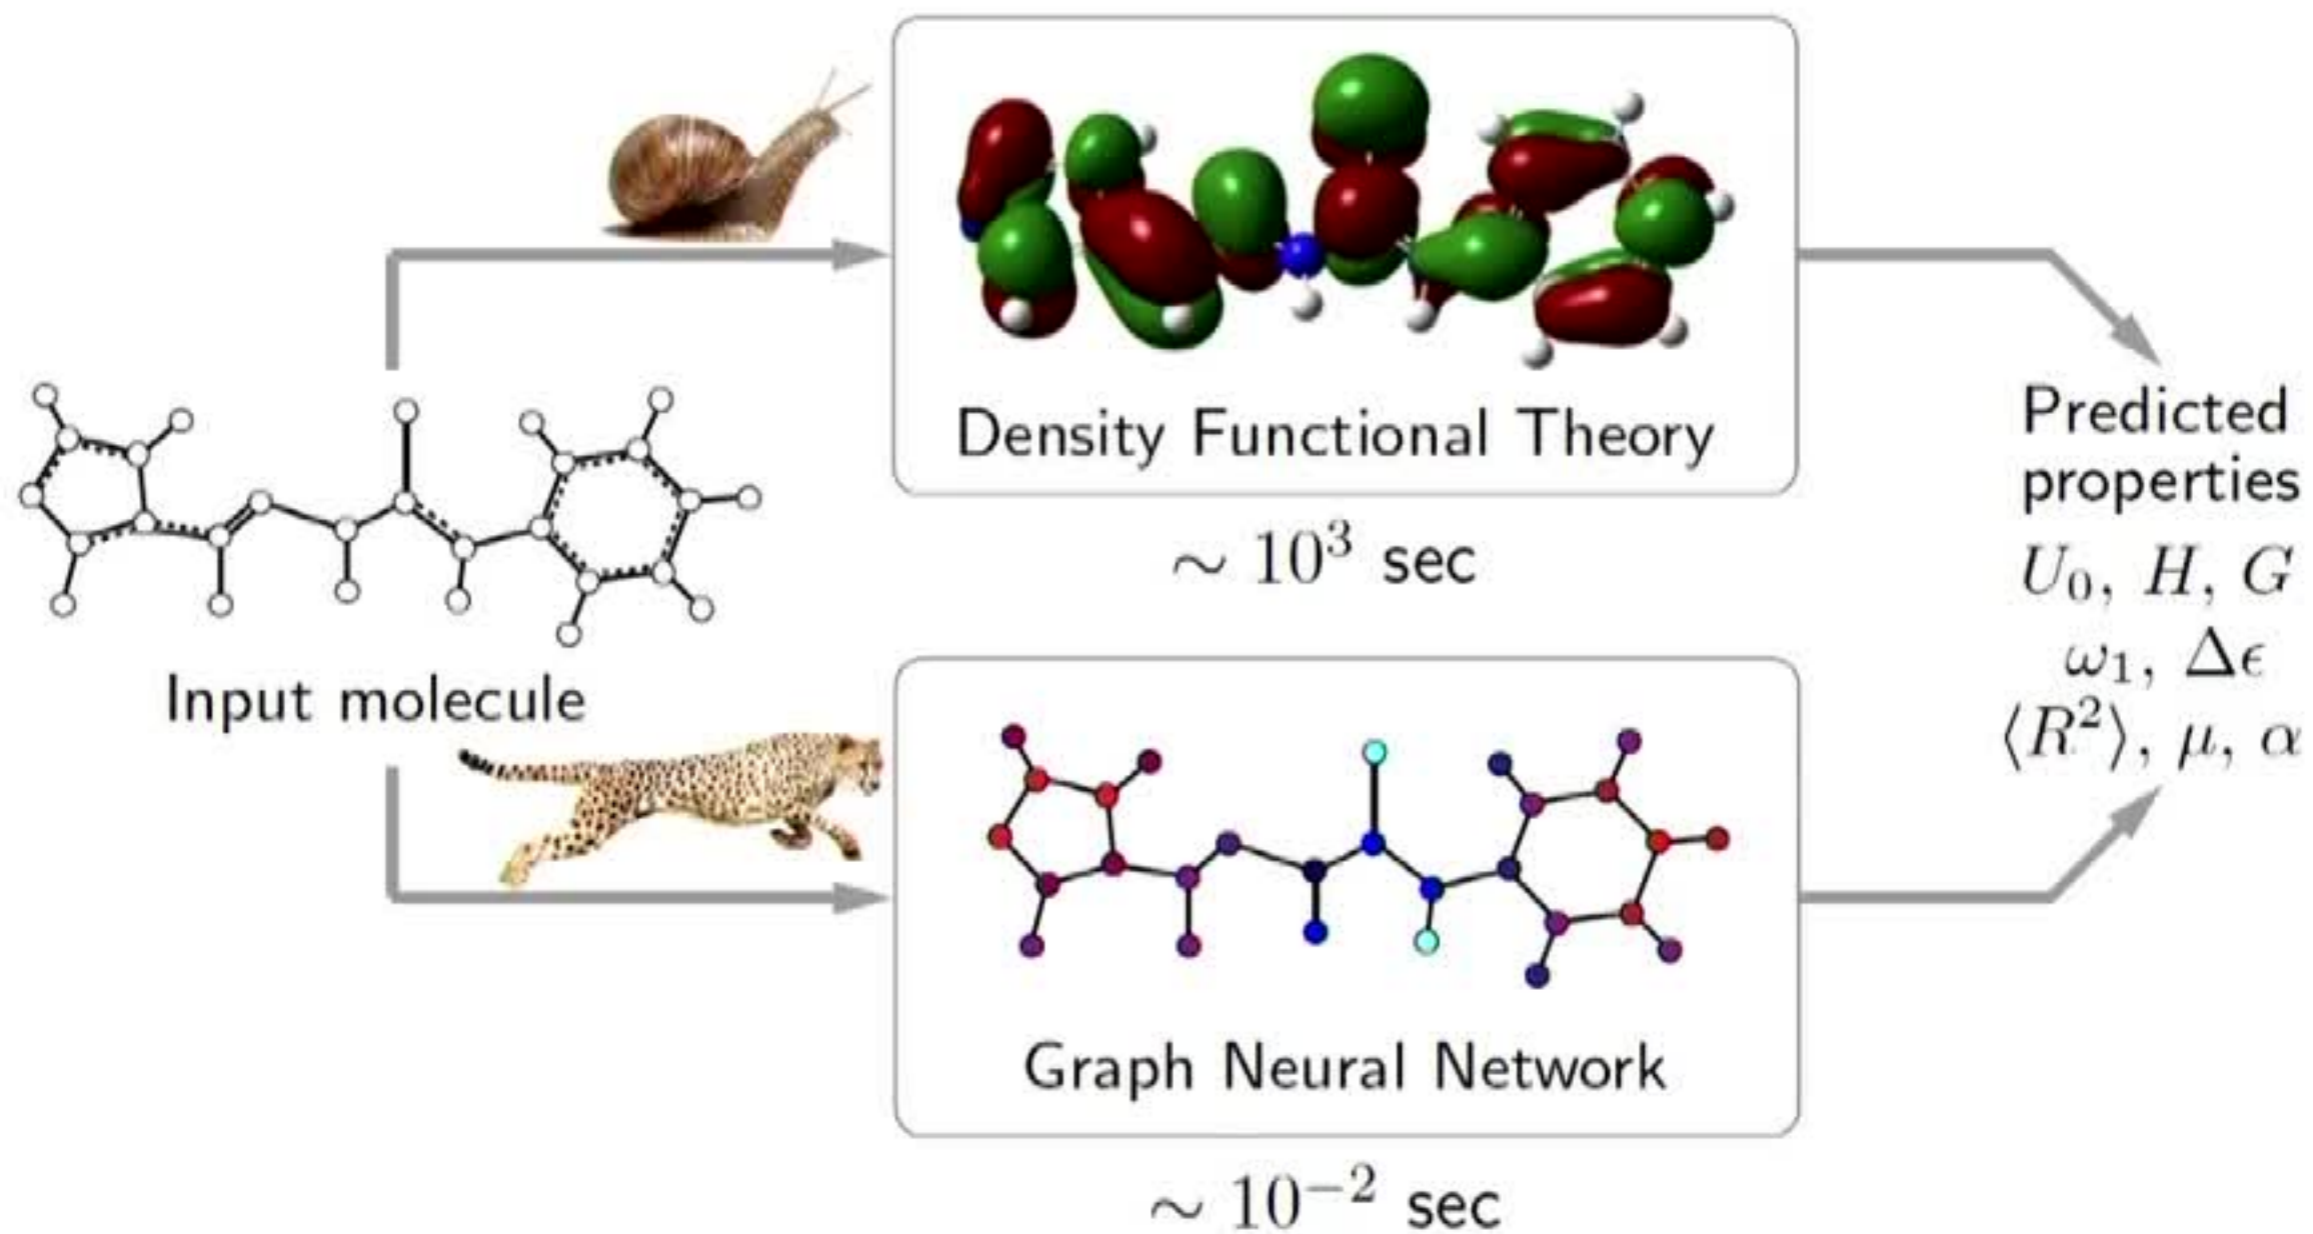


ROC curve comparing different methods for neutrino detection


Choma, Monti et al. 2018 for the IceCube collaboration
(joint work with NERSC+LBL, Berkeley)

Applications in Chemistry and Drug Design

Molecule property prediction



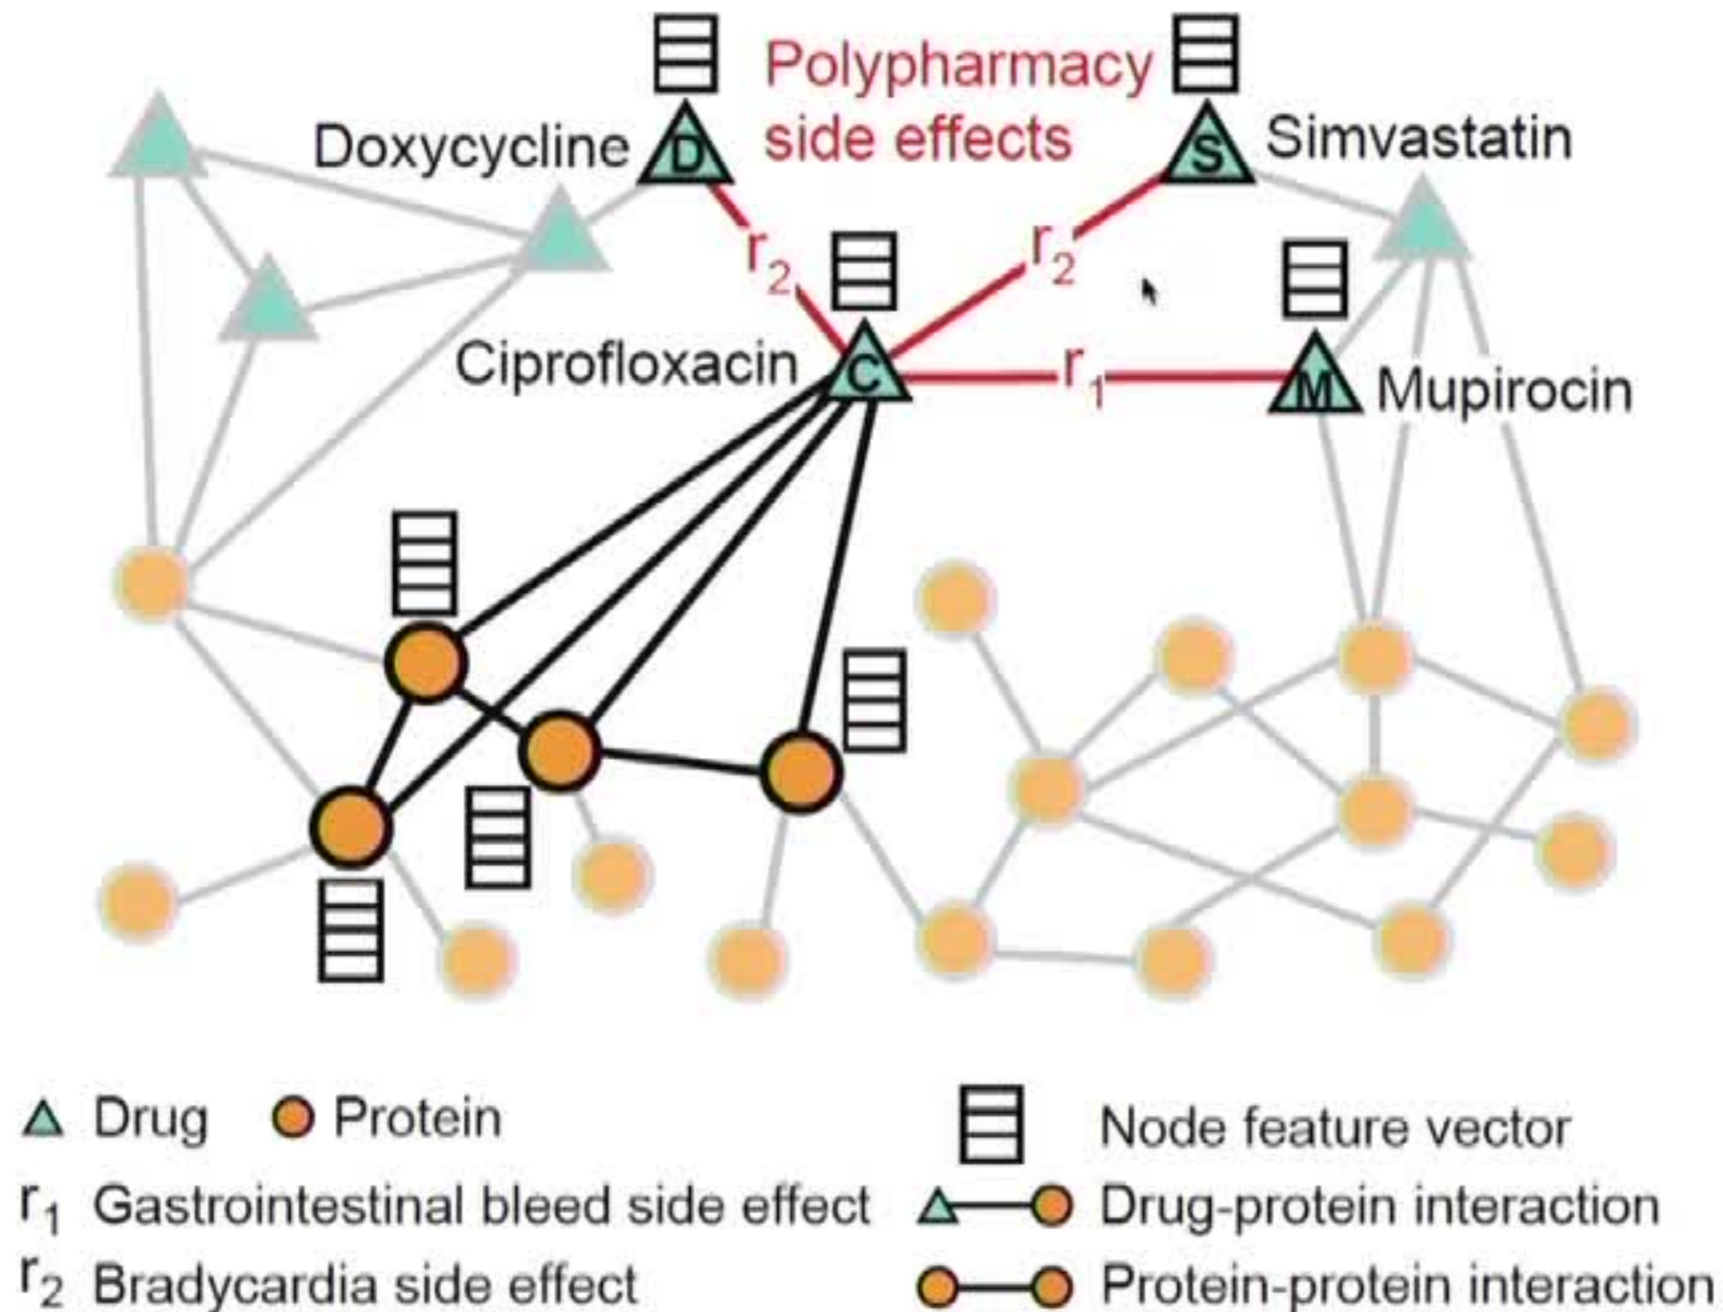
Molecule generation



Molecules generated with a graph VAE

Simonovsky, Komodakis 2017; You et al. 2018; De Cao, Kipf 2018 (MolGAN)

Polypharmacy and Drug repurposing

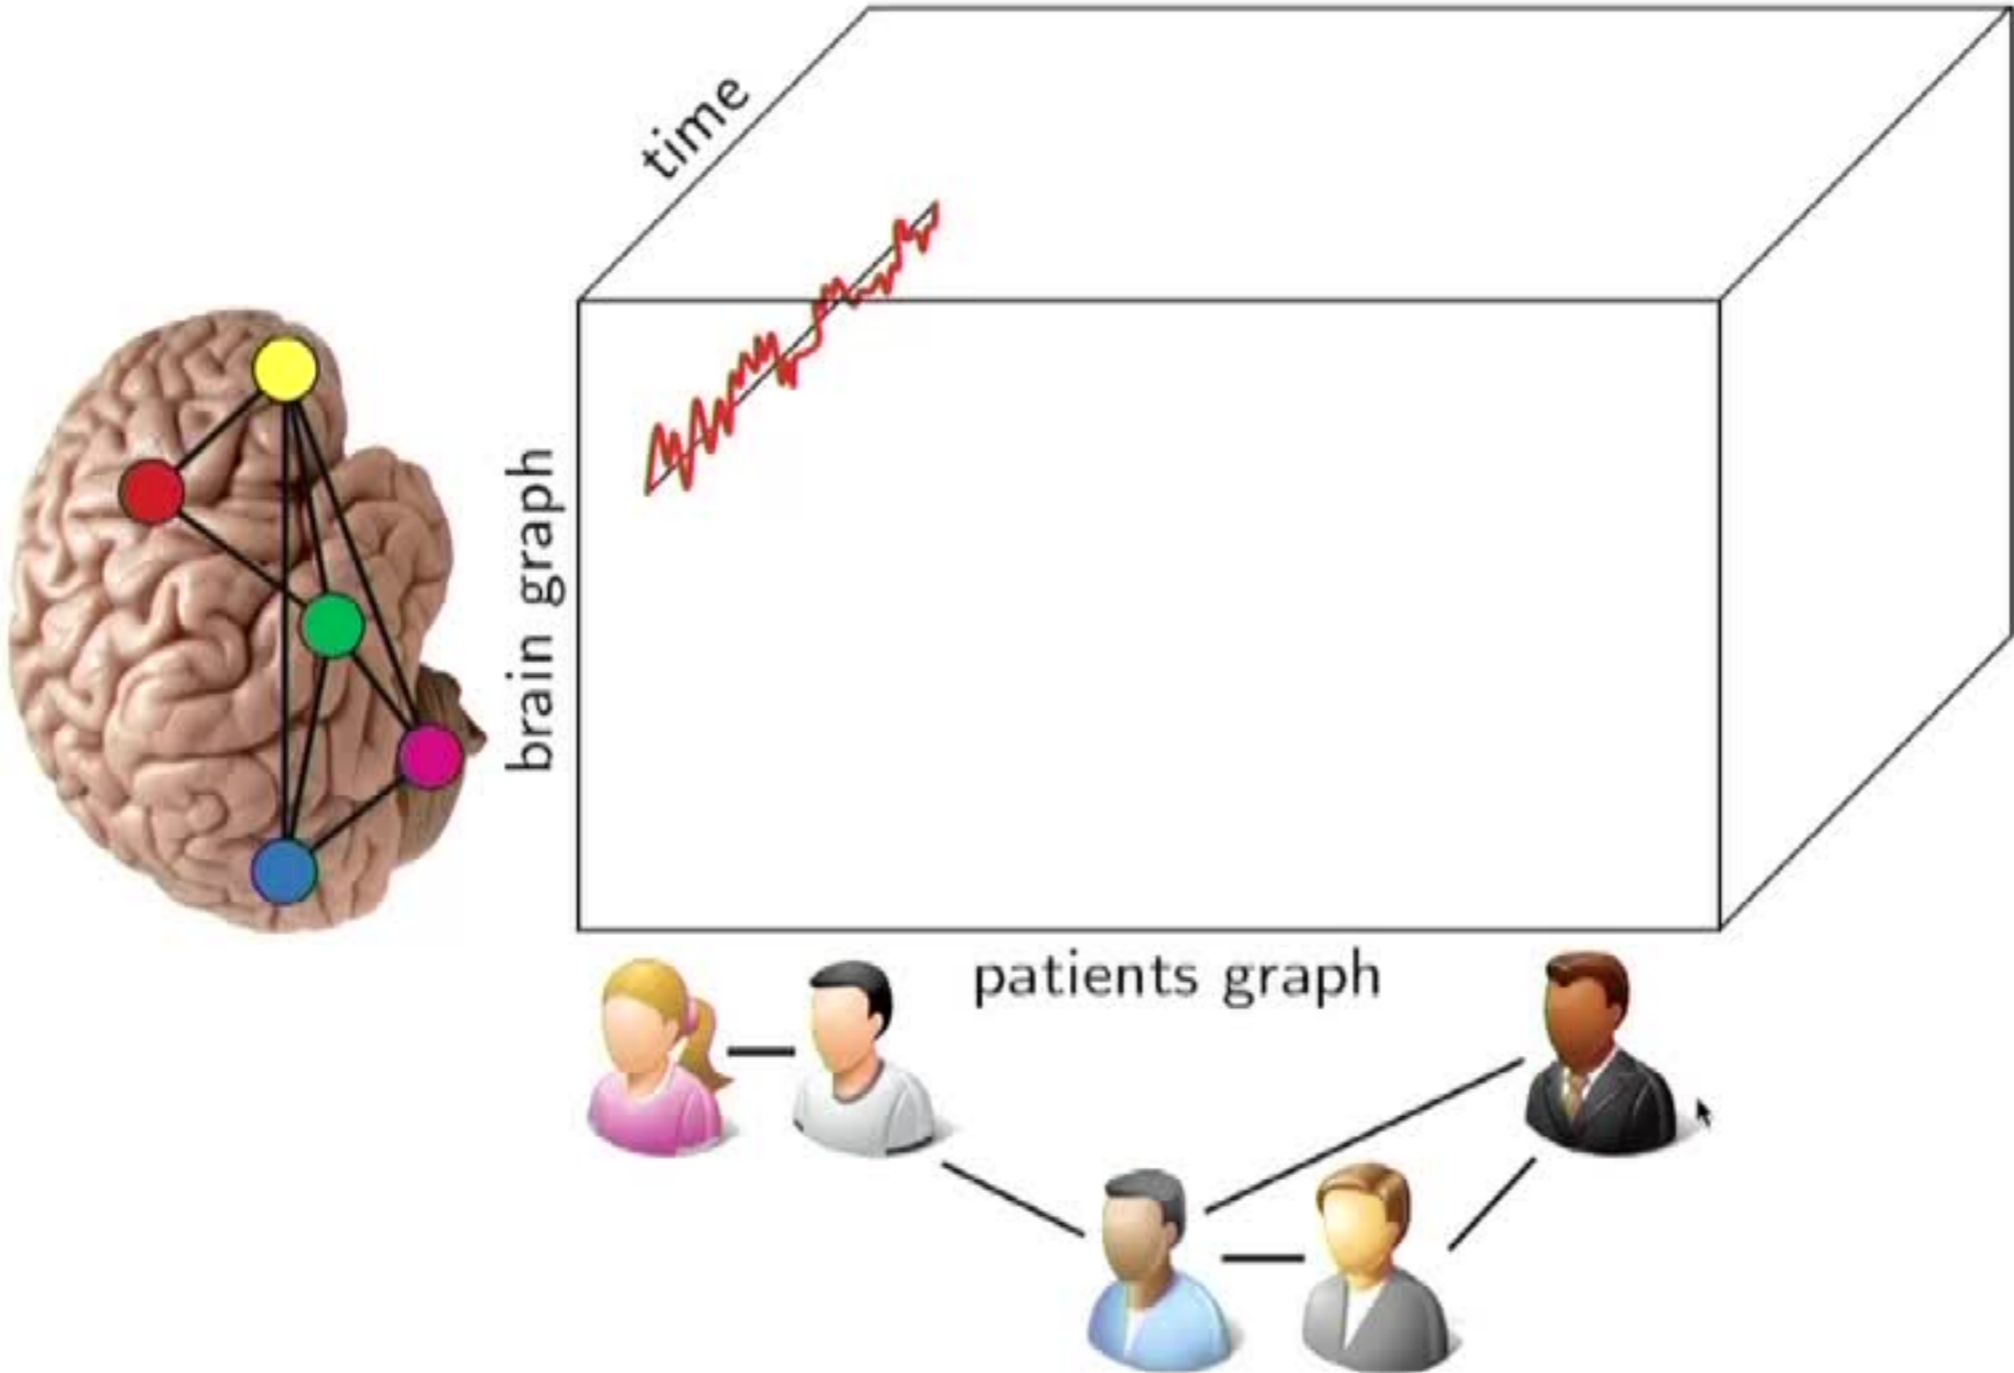


Prediction of side effects of drug combinations using deep learning on multimodal graph

Other Applications



Brain imaging




Adversarial noise

Targeted



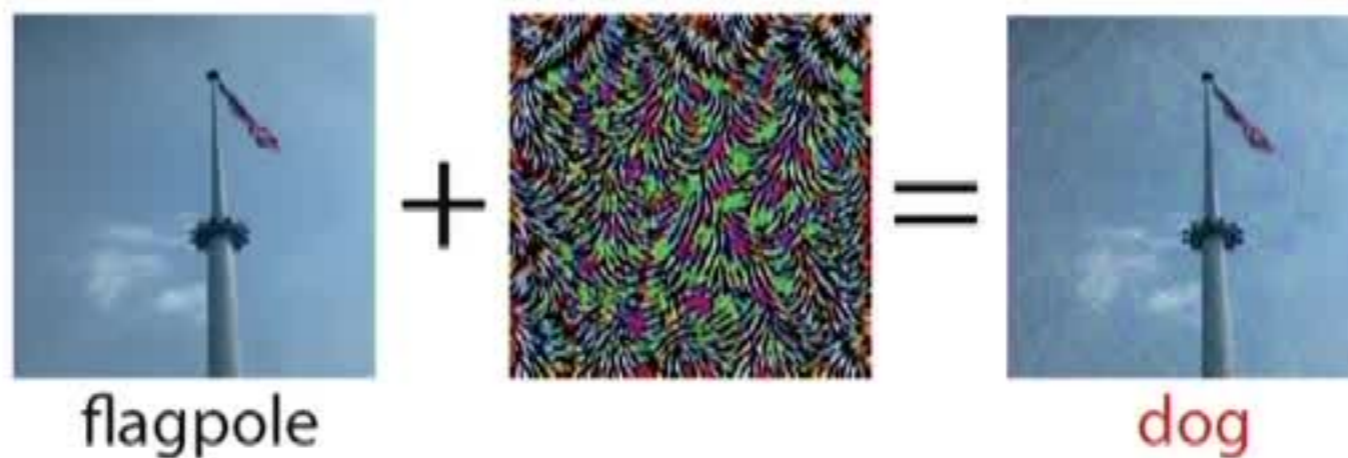
Eykholt et al. 2017

One pixel



Su, Vargas, Sakurai 2018

Universal



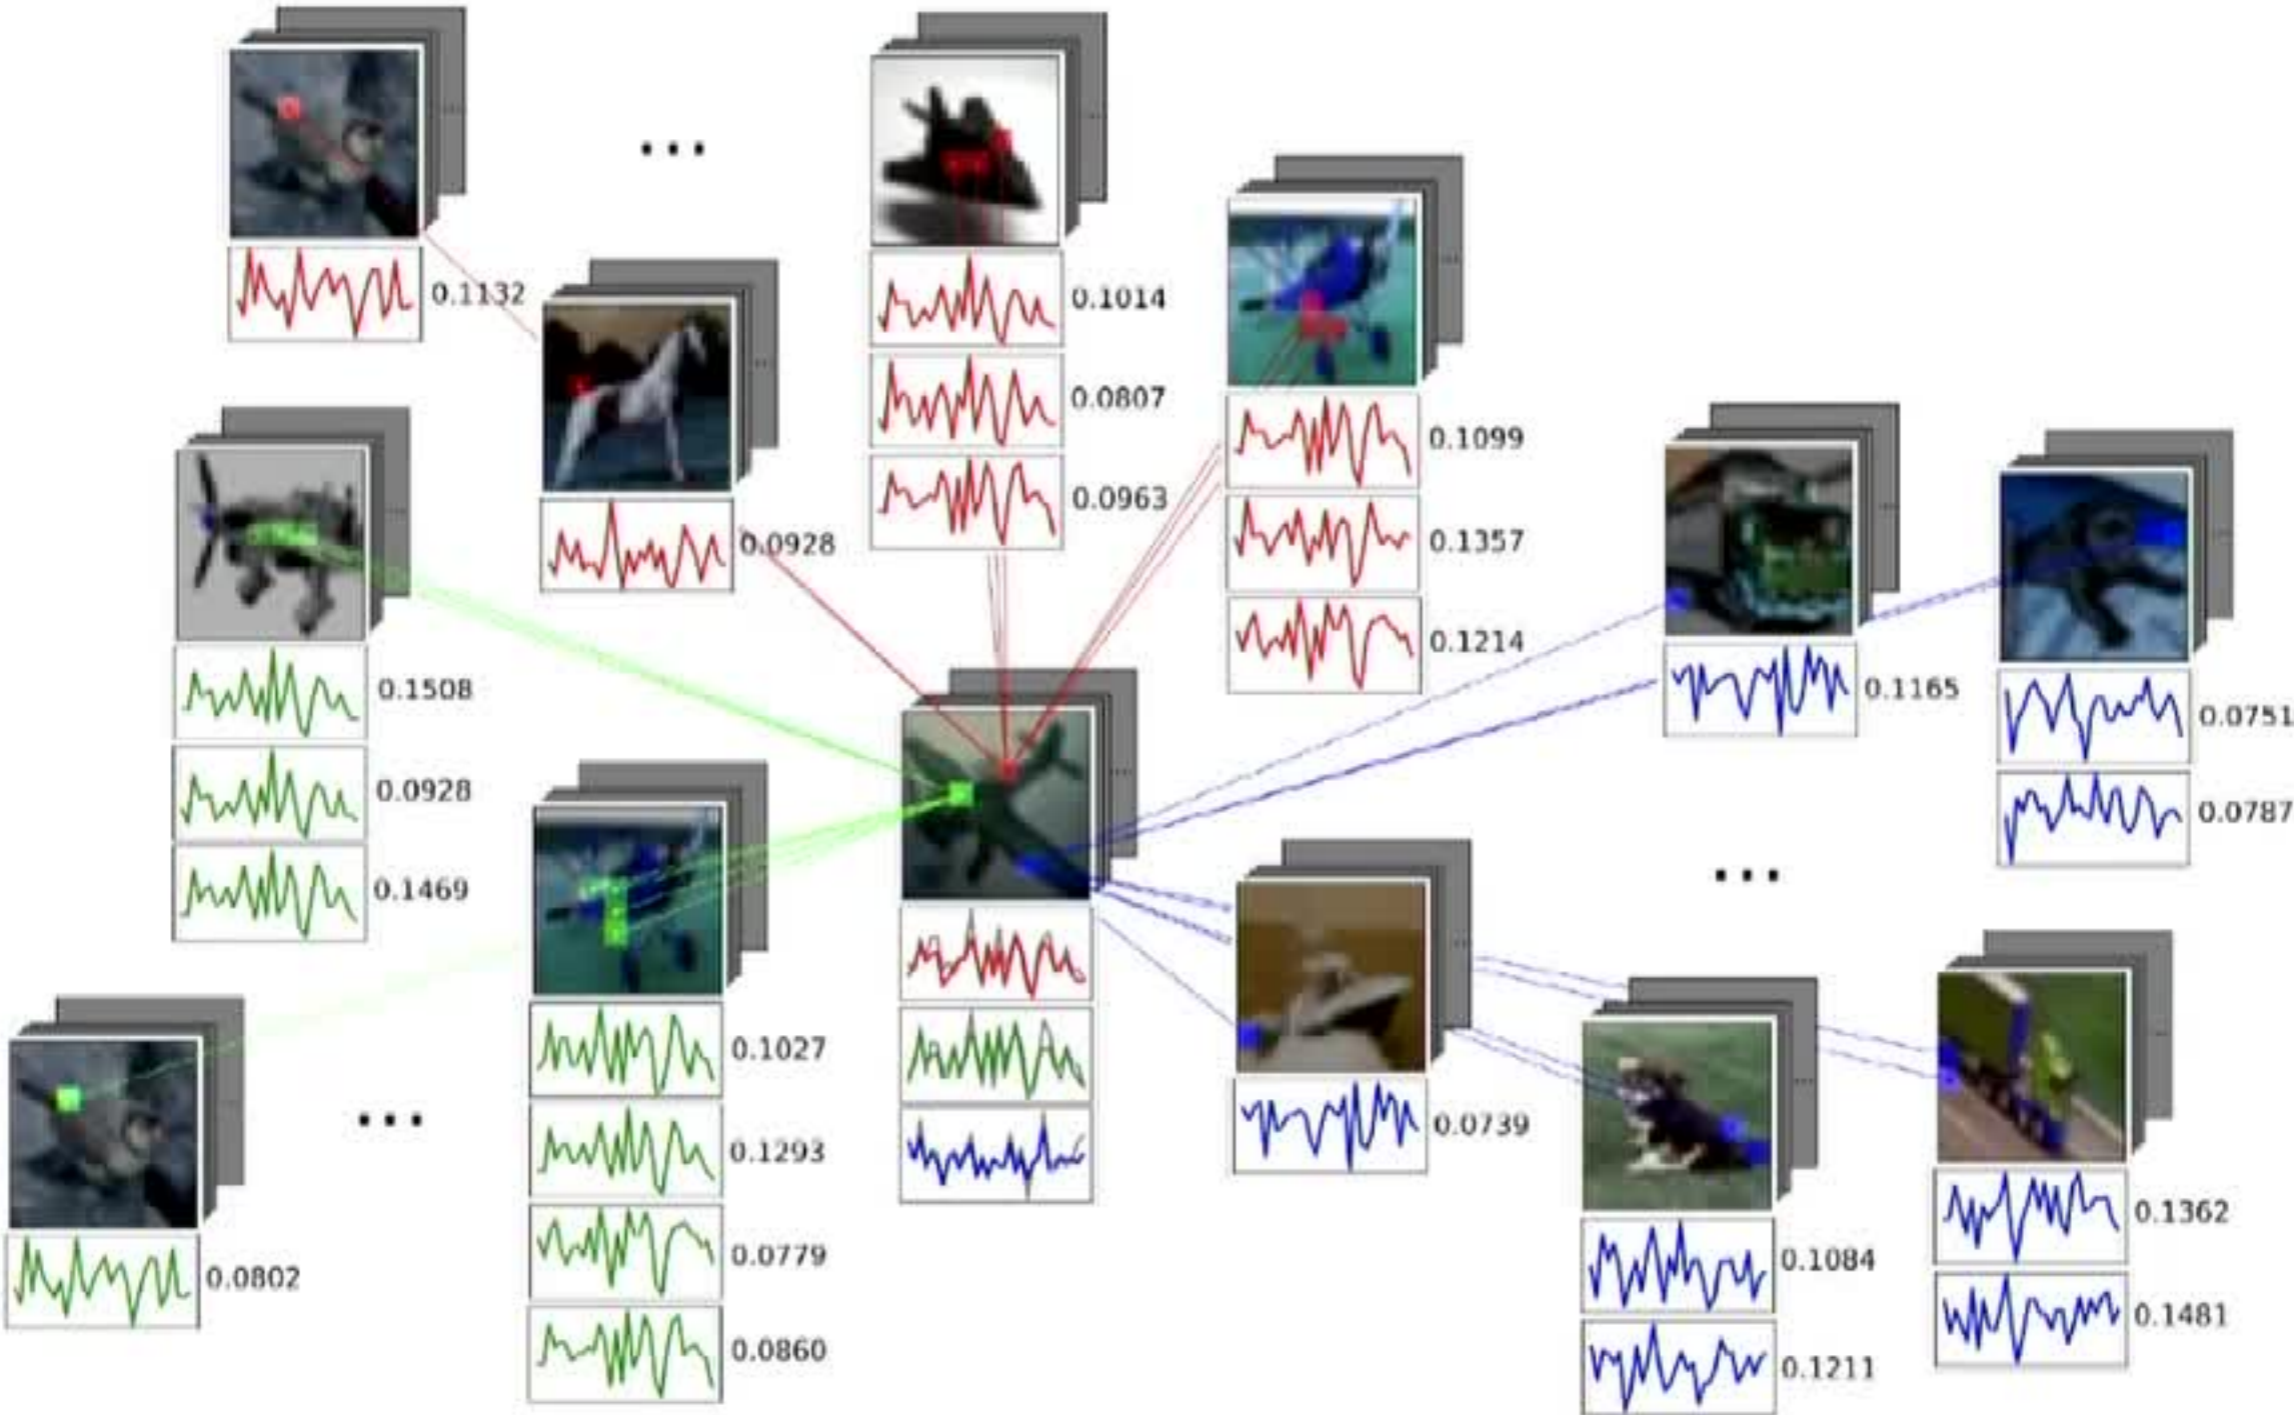
Moosavi-Dezfooli et al. 2017

Adversarial noise




Adversarial attack on a face recognition system makes it mistake faces with specially crafted glasses (top) for other people (bottom)


Peer regularization



Peer regularization

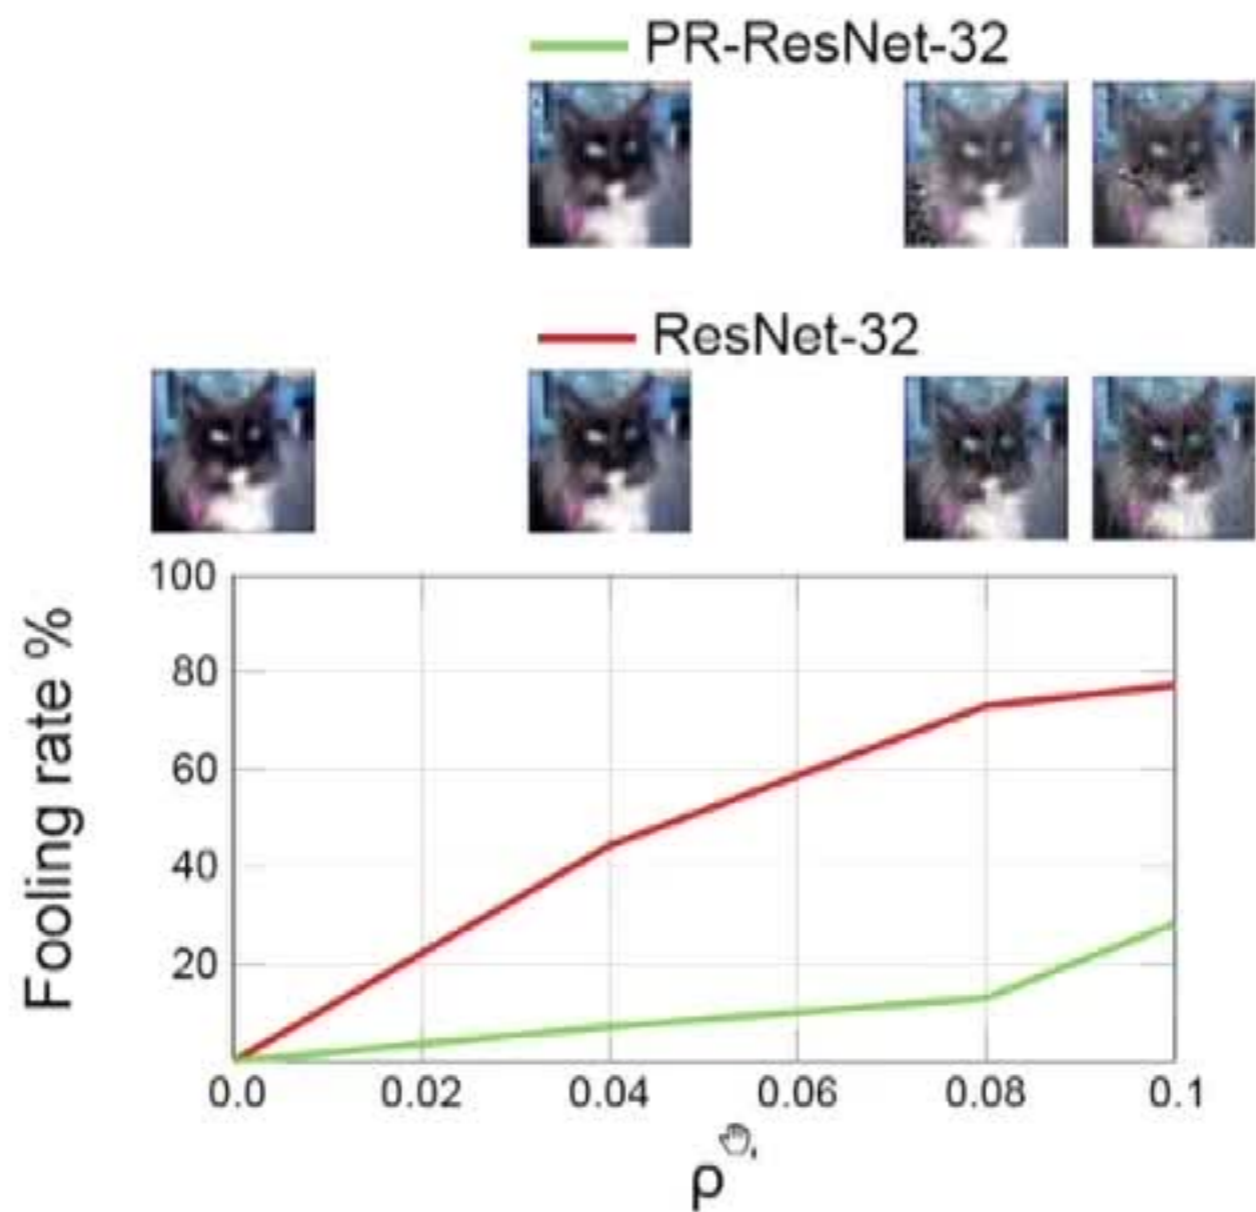


Peer-Regularized LeNet-5



Peer-Regularized ResNet-32

Universal perturbation



Fooling rate on CIFAR-10 for universal adversarial perturbation of different strength ρ

Universal perturbation

Original



$\rho = 0.04$ truck



$\rho = 0.08$ bird



$\rho = 0.1$ frog



Standard ResNet-32

$\rho = 0.04$ truck



$\rho = 0.08$ truck




$\rho = 0.1$ truck




Peer-Regularized ResNet-32

Non-targeted perturbation (gradient descent)



Standard ResNet-32



Peer-Regularized ResNet-32

Much stronger noise is needed to fool PeerNet!

Safe spaces for South
Asia's cultures *p. 1086*

Synthetic Notch receptors
enhance T cell therapies *p. 1112*

Dietary fiber fights
diabetes *p. 1151*

Science

\$15
9 MARCH 2018
sciencemag.org

AAAS


HOW LIES SPREAD

On social media,
false news beats the truth

pp. 1094 & 1146

Vosoghi et al. 2018

Non-targeted perturbation (gradient descent)



Standard ResNet-32



Peer-Regularized ResNet-32

Much stronger noise is needed to fool PeerNet!

Safe spaces for South
Asia's cultures *p. 1086*

Synthetic Notch receptors
enhance T cell therapies *p. 1112*

Dietary fiber fights
diabetes *p. 1151*

Science

\$15
9 MARCH 2018
sciencemag.org

AAAS


HOW LIES SPREAD

On social media,
false news beats the truth

pp. 1094 & 1146

Vosoghi et al. 2018

Fake news detection



Fake news



True news



We are hiring in Switzerland and UK!

Conclusions

- Geometric Priors exist beyond Euclidean domains
- They can be exploited by modeling locality and weight sharing
- Spectral and Spatial models obtain locality and sharing through different routes
- Resulting neural models related to attention mechanisms
- Challenge: Scaling up models to millions of nodes
- Large areas of application: physical sciences, machine translation, graphics, relational learning