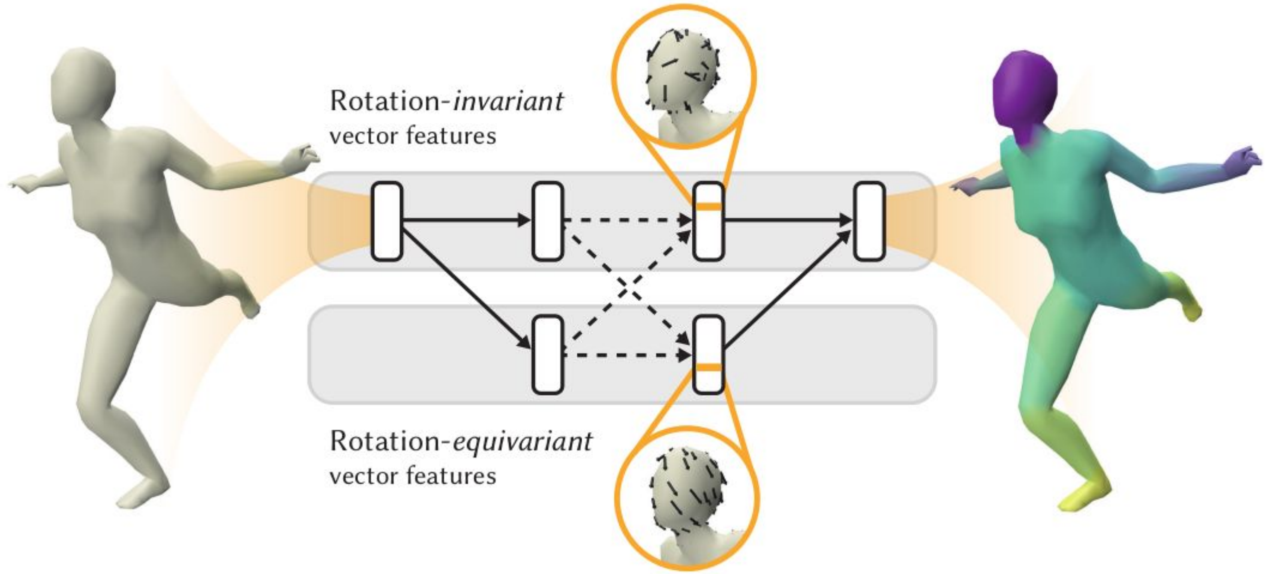

CNNs on Surfaces using Rotation-Equivariant Features

Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt

ACM Transactions on Graphics, 39(4), 2020

Gifs: https://rubenwiersma.nl/hsn

Paper presentation by Thomas Lew, 02/09/2022

Applying CNNs to 3D deep learning

2D data (images)

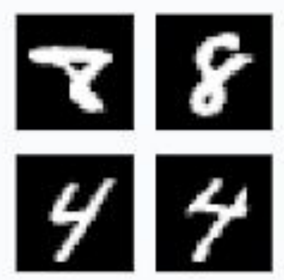

MNIST

https://sthalles.github.io/deep_segmentation_network/

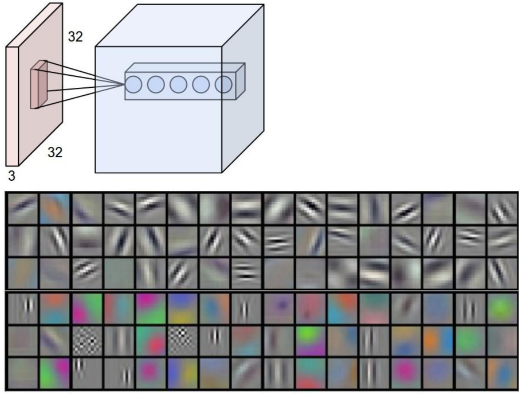

Krizhevsky et al, 2012

3D data (triangle meshes)

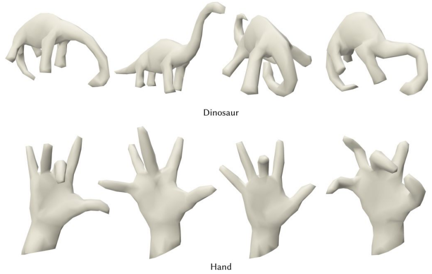

?

Applying CNNs to 3D deep learning

2D data (images)

MNIST

https://sthalles.github.io/deep_segmentation_network/

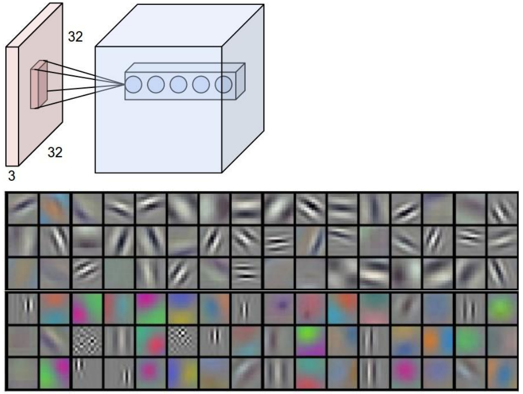

Krizhevsky et al, 2012

3D data (triangle meshes)

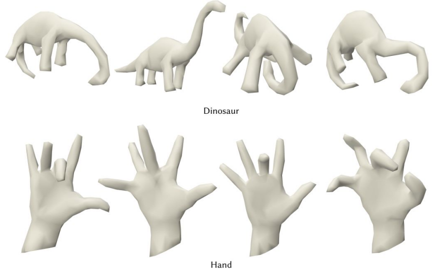

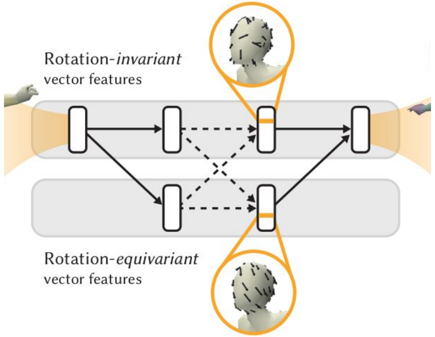

CNNs on meshes / charting approaches

Graph CNNs

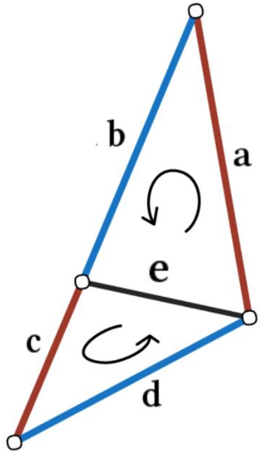

MeshCNN [Hanocka et al, 2019]

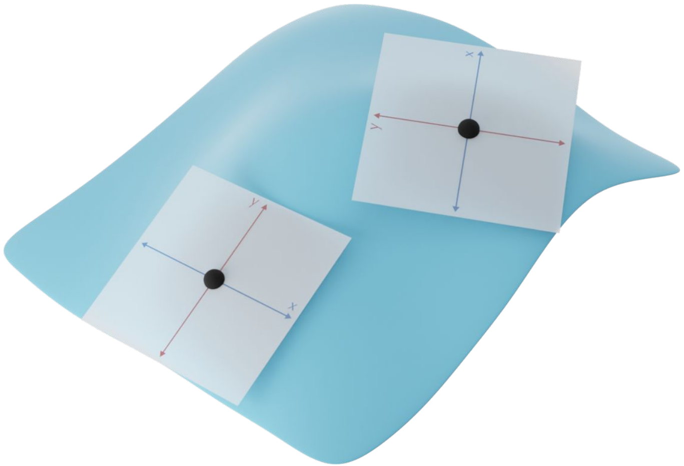

Charting approaches (CNNs: 2D to 3D)

2D

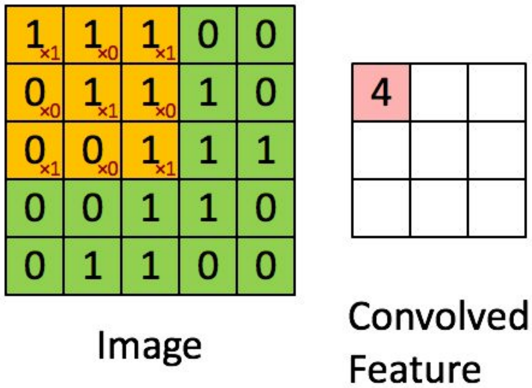

3D

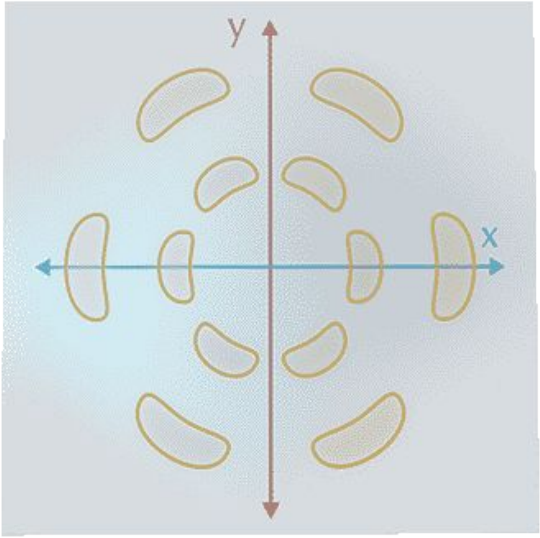

https://danielnouri.org/notes/2014/12/17/using-convolutional-neural-nets-to-detect-facial-keypoints-tutorial/

Charting approaches

a) define a kernel on $\mathbb{R}^2$

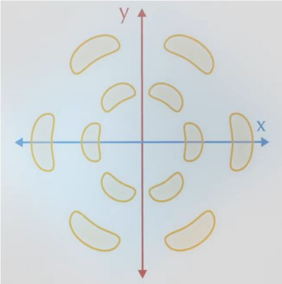

Gaussian kernel on R2
https://www.youtube.com/watch?v=kg1wRBGUYqk

Different kernels are possible, see

- [Poulenard and Ovsjanikov 2018]
- [Boscaini et al. 2016]
- [Monti et al. 2017]

b) apply kernel to tangent plane

$$T_p S \cong \mathbb{R}^2$$

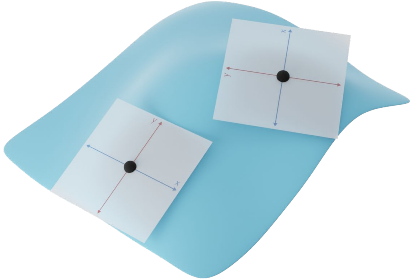

e.g., with exponential map, see GeodesicCNN [Masci et al 2015]

The Vector Heat Method

N. Sharp, Y. Soliman, and K. Crane, ACM Trans. Graph. 38(3), 2019

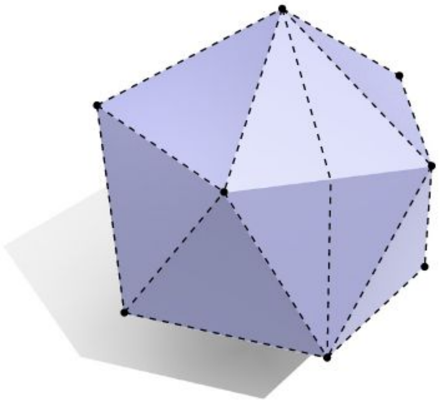

triangle mesh

Vector Heat Method

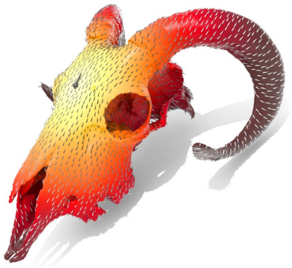

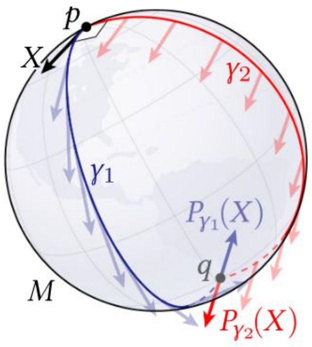

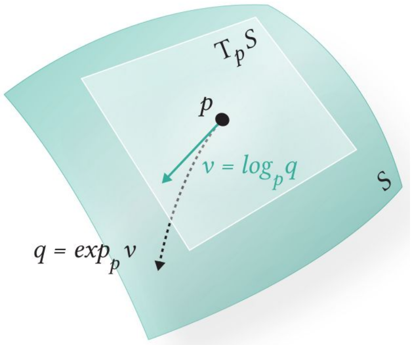

Allows efficiently computing **tangent spaces** on meshes and **parallel transport** maps

Charting approaches: limitations

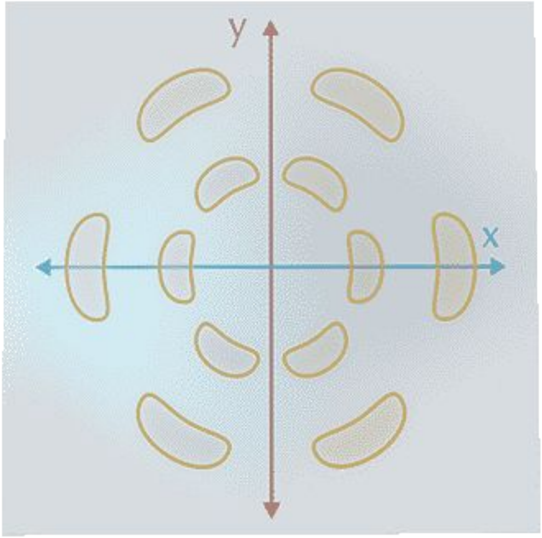

Rotation ambiguity

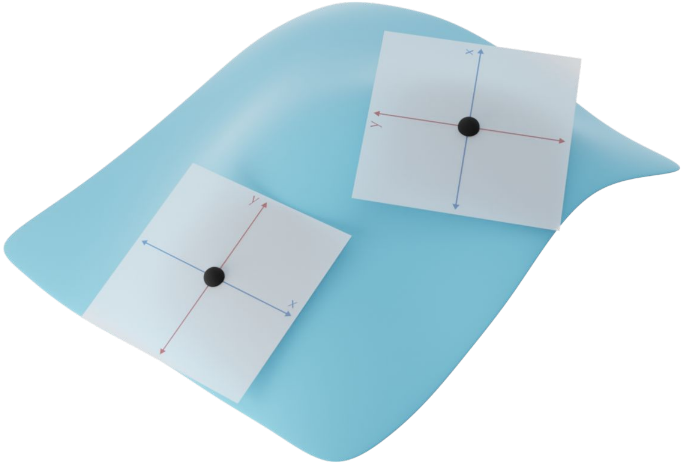

Convolving

how to move the filter along surface manifold without introducing rotations?

Proposed approach

- Features expressed as complex vectors $Xe^{i\theta}$
- Use circular harmonics (harmonic networks: learn radial and angular parameters)
  - rotational-equivariant kernels
- Propose convolutional filters that apply to surfaces
  - Idea: **circular harmonics** + **parallel transport**

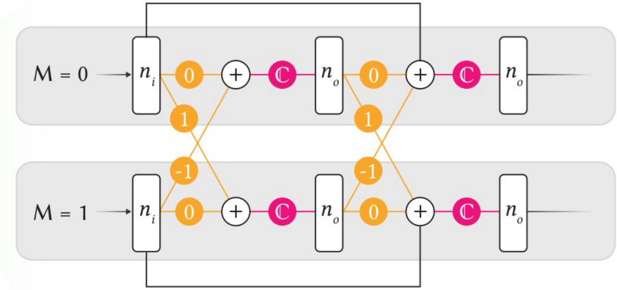

Circular harmonics

circular harmonic filters

$$W_m(r, \theta, R, \beta) = R(r)e^{i(m\theta+\beta)}$$

rotation equivariance

$$[W_m \star x^{\phi}](p) = e^{im\phi}[W_m \star x^0](p)$$

(r, θ)	Polar coordinates
$R : \mathbb{R}_+ \to \mathbb{R}$	Radial profile
β	Phase offset
$m \in \mathbb{Z}$	Rotation order

Input

rotational invariant

rotational equivariant

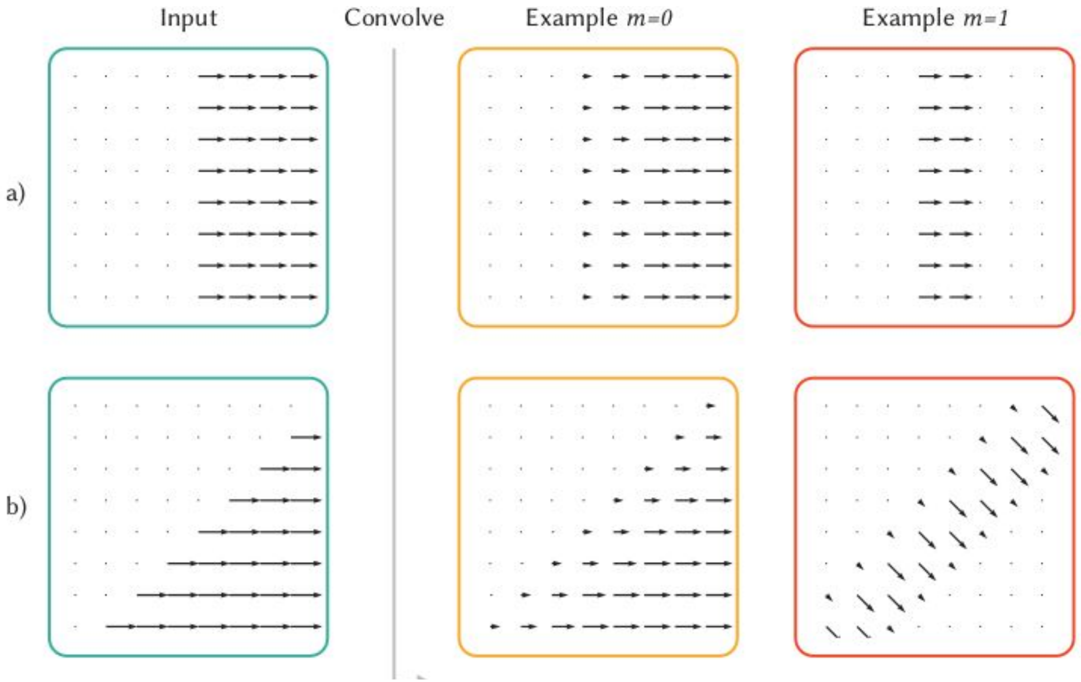

Circular harmonics => Harmonic Nets

$$[W_m \star x^{\phi}](p) = e^{im\phi}[W_m \star x^0](p)$$

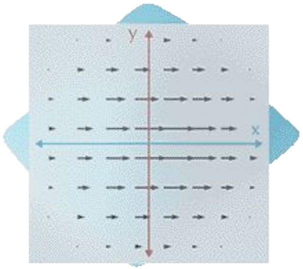

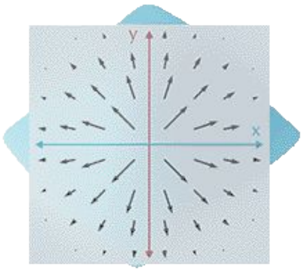

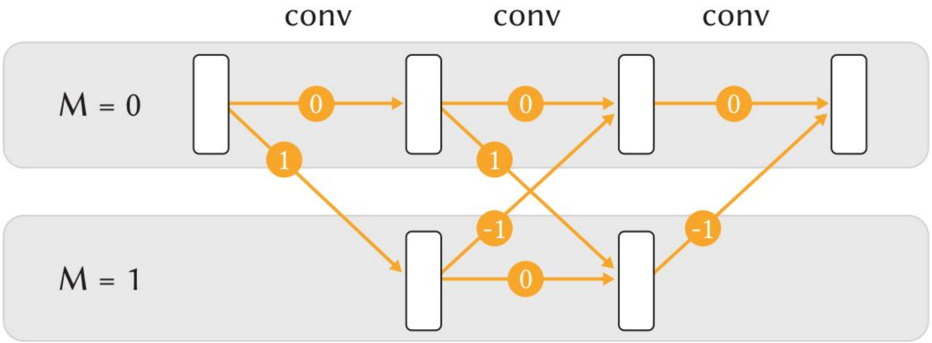

[Poulenard and Guibas 2021] uses “spherical” harmonics instead, since 3D pointcloud

Parallel Transport (exponential map)

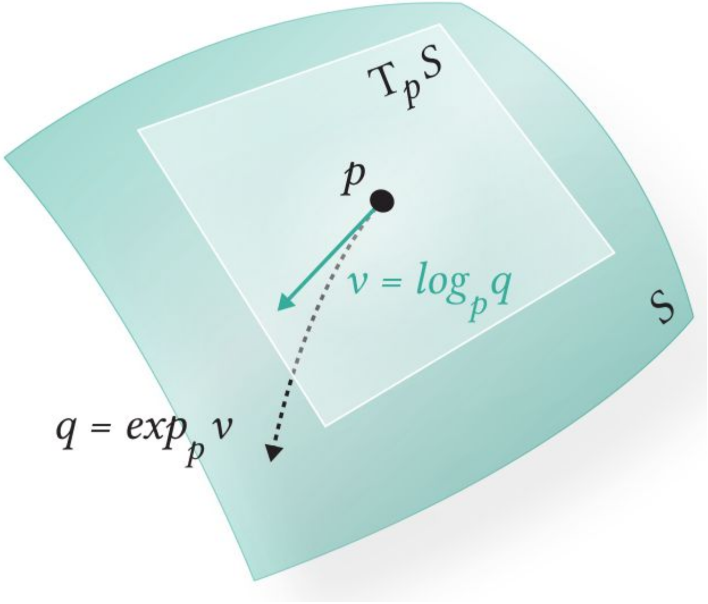

$$P_{j \to i}(x_j) = e^{i\phi_{ij}} x_j$$

Convolution on a surface

$$x_i^{(\ell+1)} = \sum_j w_j \left(R(r_{ij}) e^{i(m\theta_{ij}+\beta)} P_{j\to i}(x_j^{(\ell)}) \right)$$

Convolution on a surface

$$x_i^{(\ell+1)} = \sum_j w_j \left(R(r_{ij}) e^{i(m\theta_{ij}+\beta)} P_{j\to i}(x_j^{(\ell)}) \right)$$

parallel transport

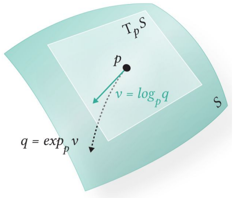

integration weights
(depend on mesh)

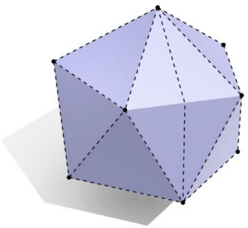

circular harmonics (H-Net)
(rotation equivariant)

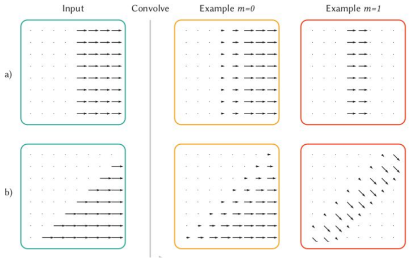

ReLU

$$Xe^{i\theta} \mapsto \mathrm{ReLU}(X + b)e^{i\theta}$$

Only the magnitude (radius) of $Xe^{i\theta}$ is changed

Model architecture

deep U-ResNet architecture from [Poulenard and Ovsjanikov 2018]

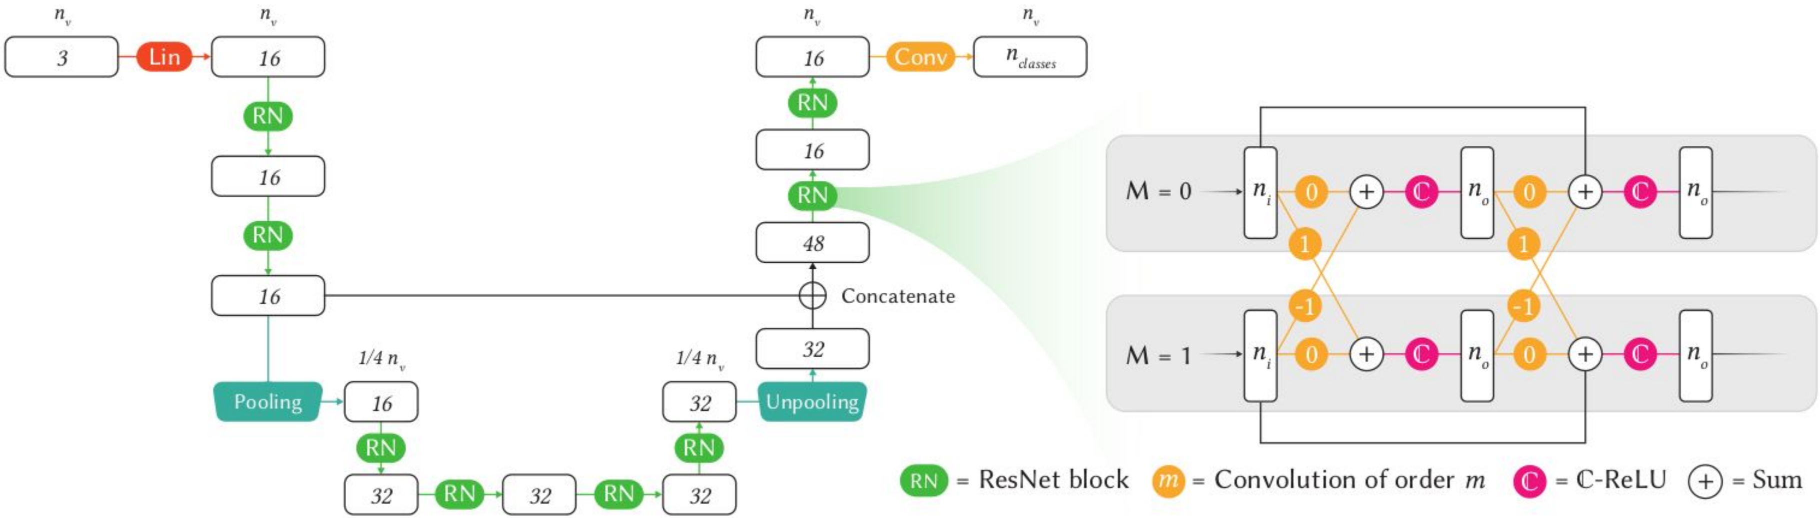

Dataset and metrics

- shape classification: SHREC dataset [Lian et al. 2011],
- correspondence: FAUST dataset [Bogo et al. 2014]
- shape segmentation: human segmentation dataset [Maron et al. 2017]

Results *HSNs perform better than state-of-the-art*

shape classification

Method	Accuracy
HSN (ours)	**96.1%**
MeshCNN	91.0%
GWCNN	90.3%
GI	88.6%
MDGCNN	82.2%
GCNN	73.9%
SG	62.6%
ACNN	60.8%
SN	52.7%

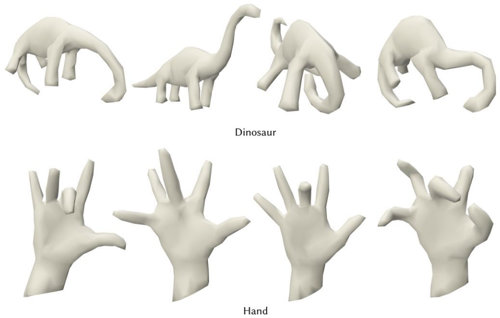

shape segmentation

Method	# Features	Accuracy
HSN (ours)	3	91.14%
MeshCNN	5	**92.30%**
SNGC	3	91.02%
PointNet++	3	90.77%
MDGCNN	64	89.47%
Toric Cover	26	88.00%
DynGraphCNN	64	86.40%
GCNN	64	86.40%
ACNN	3	83.66%

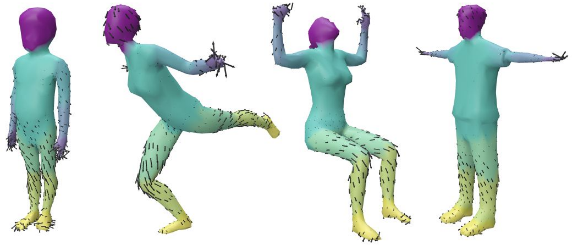

Features visualization + ablation

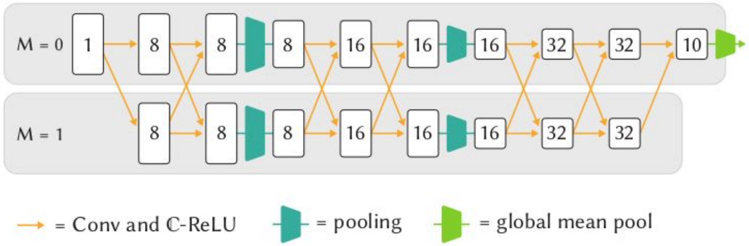

Fig. 13. Architecture for classification of Rotated MNIST.

Table 4. Results of HSN tested on shape segmentation for multiple configurations.

Method	Streams ($M = \ldots$)	Accuracy
HSN	0, 1	**91.14%**
HSN	0	88.74%
HSN (parameters ×4)	0	87.25%
HSN (pc aligned)	0, 1	86.22%

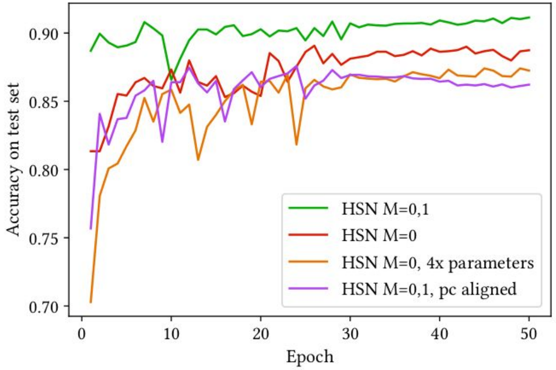

Fig. 14. Validation accuracy per training epoch several configurations of HSN on shape segmentation.

Conclusion

- Proposed convolutional filters that apply to surfaces
  - Idea: circular harmonic kernels + parallel transport

  Rotational invariant/equivariant depending on filter order M
- Better performance and requires less parameters than other approaches
- Next:
  - using the learned features / representations for other tasks
  - extensions to pointclouds

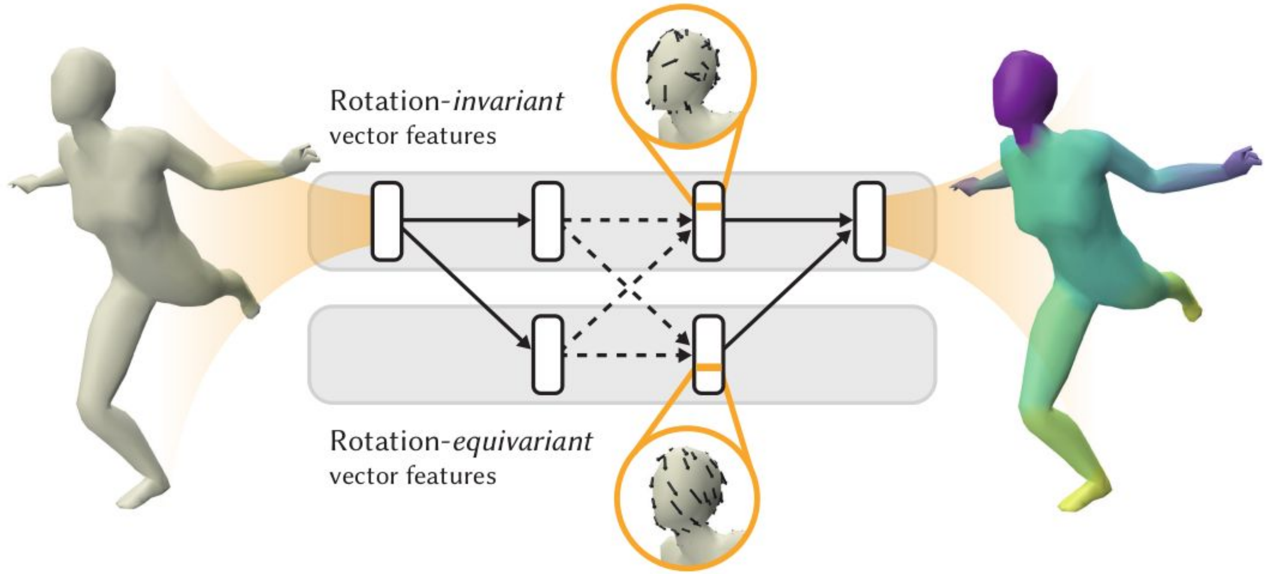

CNNs on Surfaces using Rotation-Equivariant Features

Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt

ACM Transactions on Graphics, 39(4), 2020

Paper presentation by Thomas Lew, 02/09/2022