Convolutional neural networks on surfaces via seamless toric covers

Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, Yaron Lipman

Problem statement

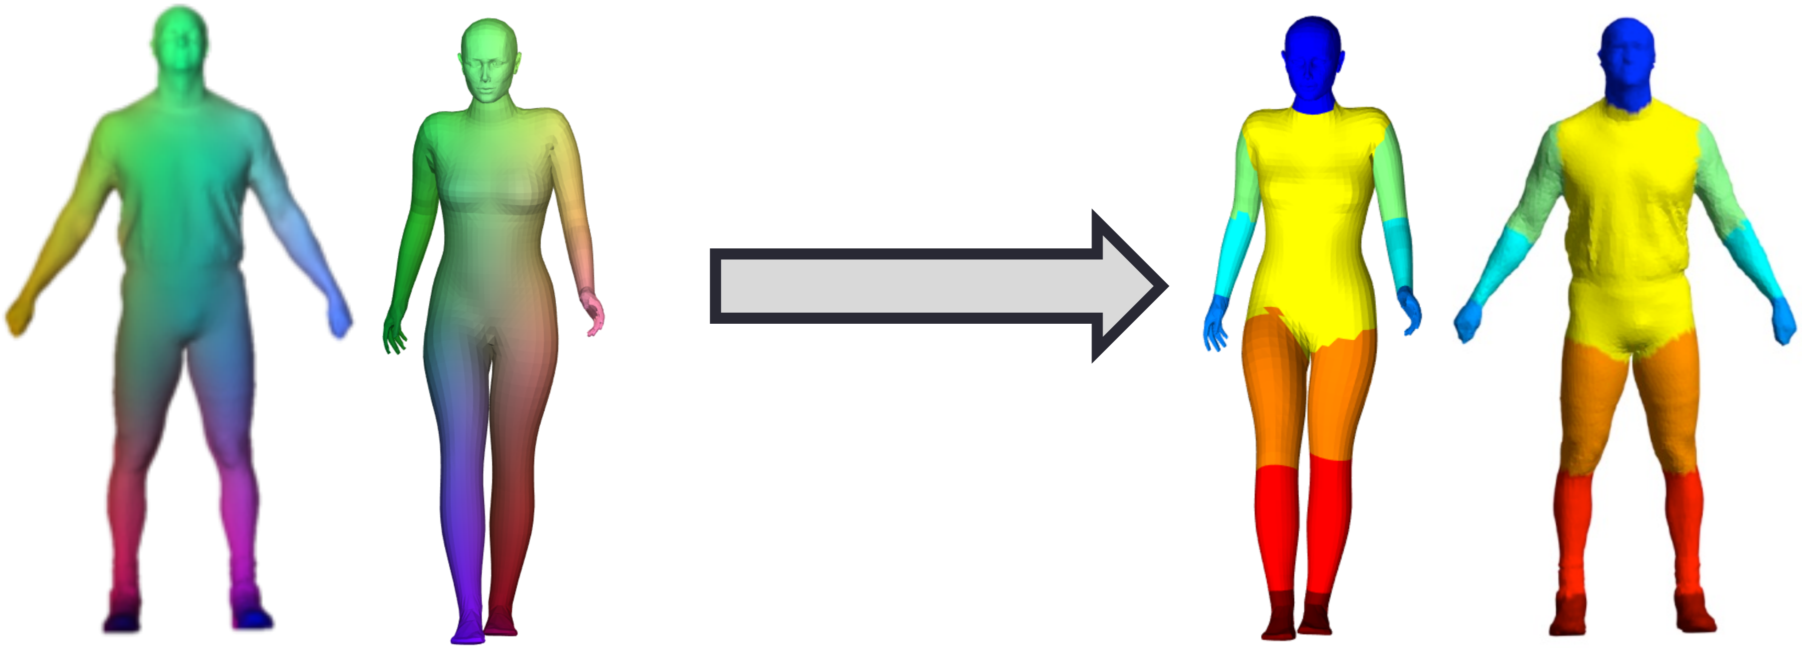

Easy Hard

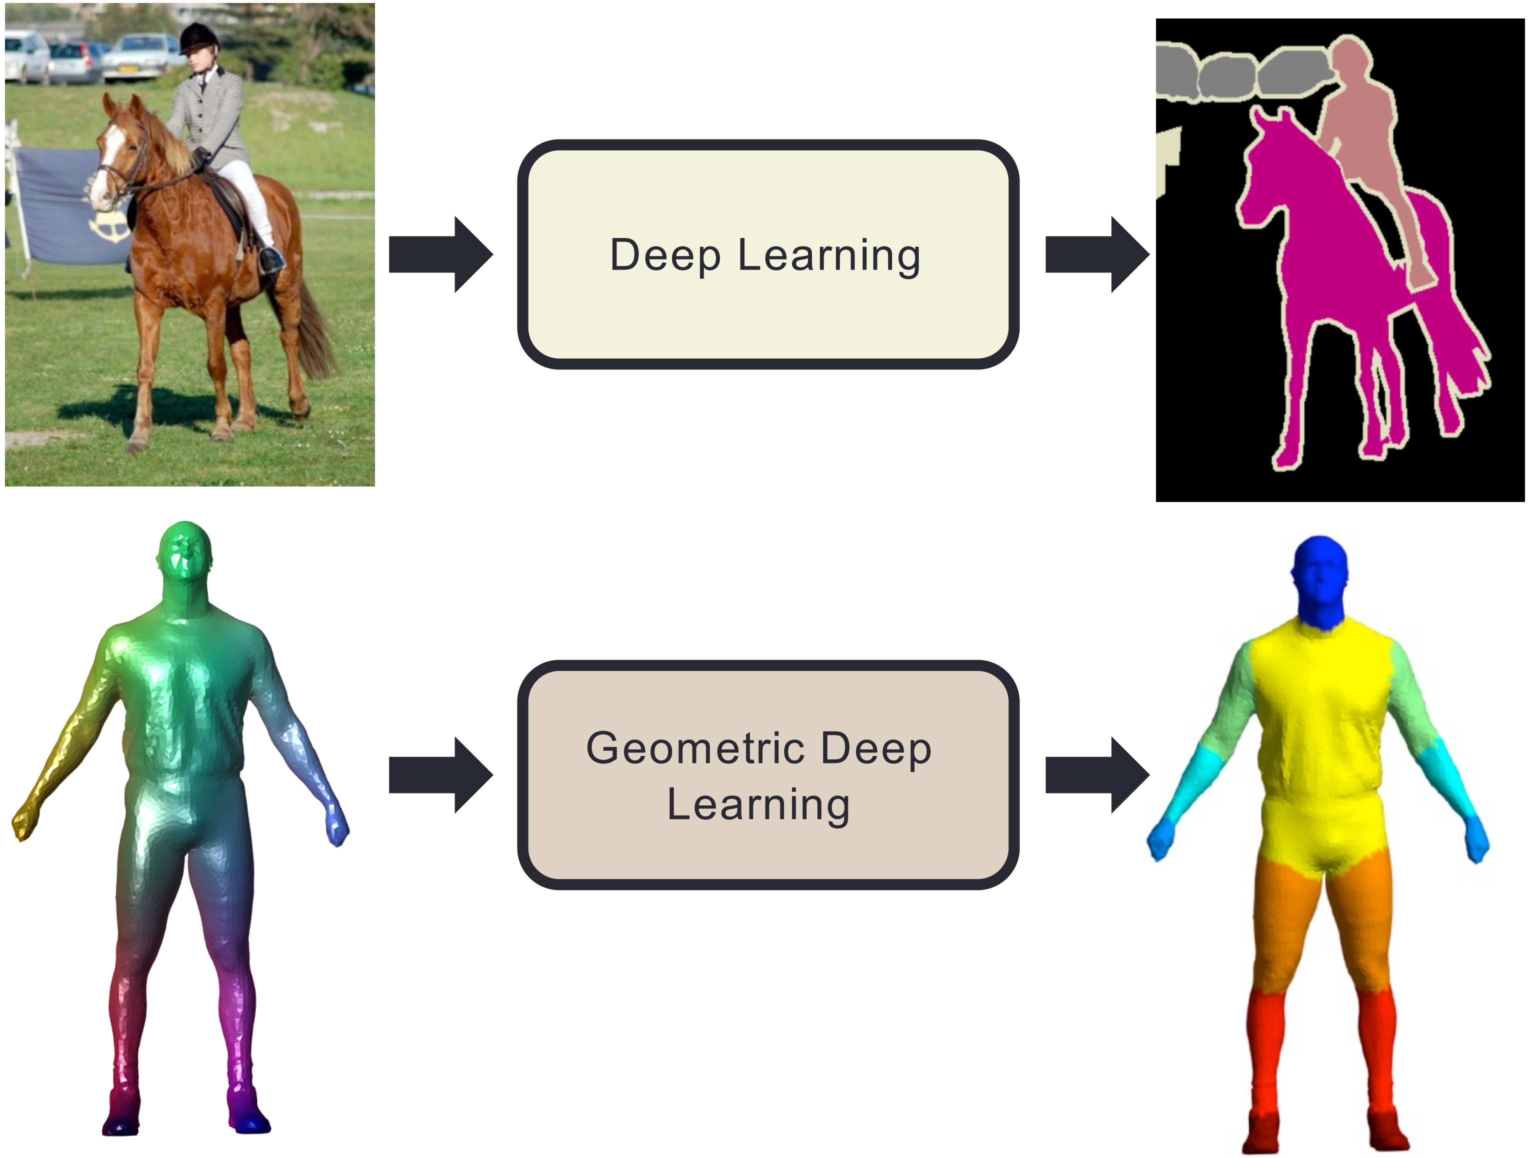
Deep Learning
Geometric Deep Learning

Previous work (1)

- Per-vertex features (Guo et al. 15')

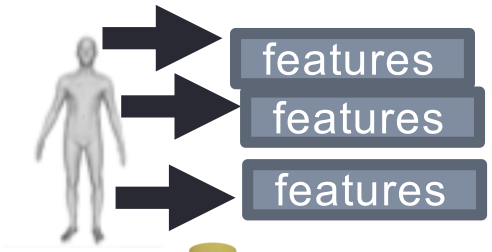

- Volumetric representation (Wu et al. 15')

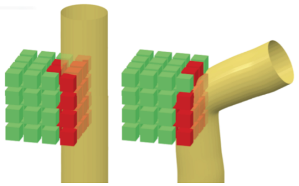

- Rendering based methods (Su et al. 15')

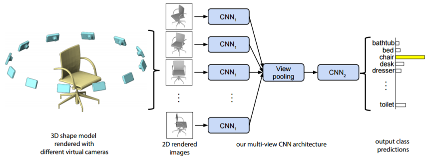

- Patch based methods (Masci et al. 15', Boscaini et al. 16')

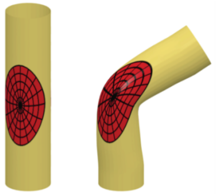

- Learning in the spectral domain (Bruna et al. 13')

Previous Work (2)

- Parameterization based methods (Sinha et al. 2016)
  - High distortion
  - High dimensional parameterization space

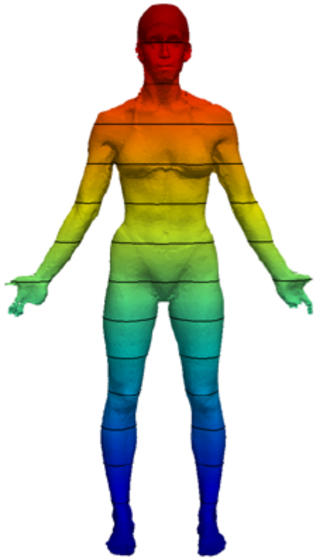

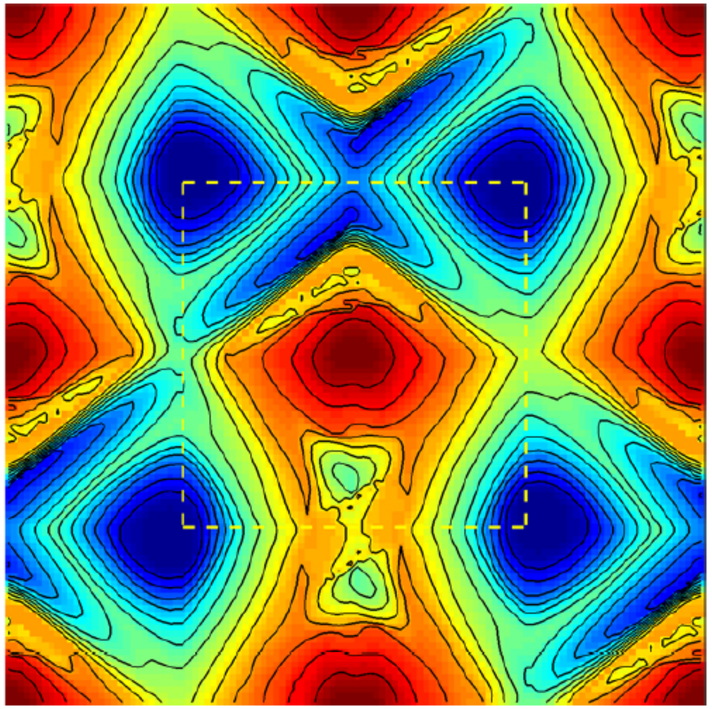

Our approach

- Map the surface to a flat torus

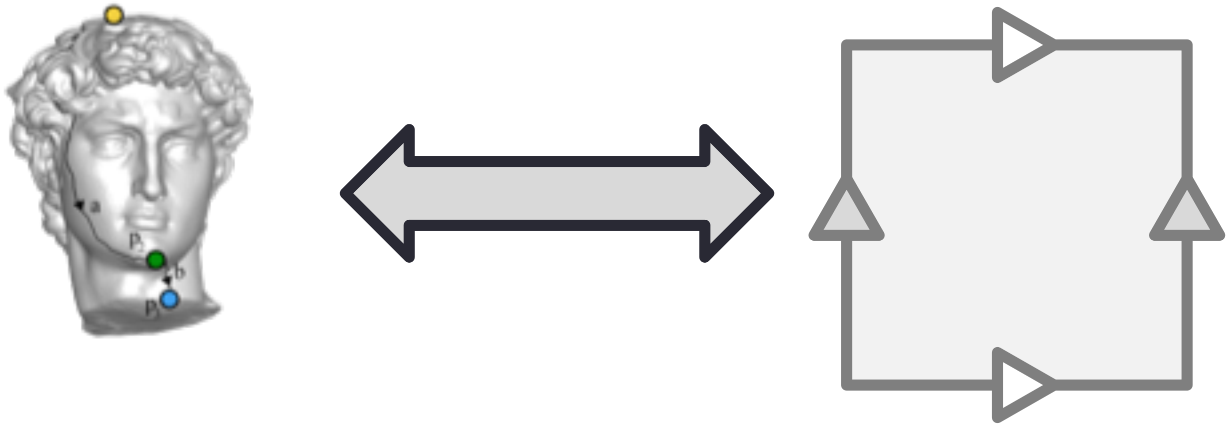

- Use its natural convolution

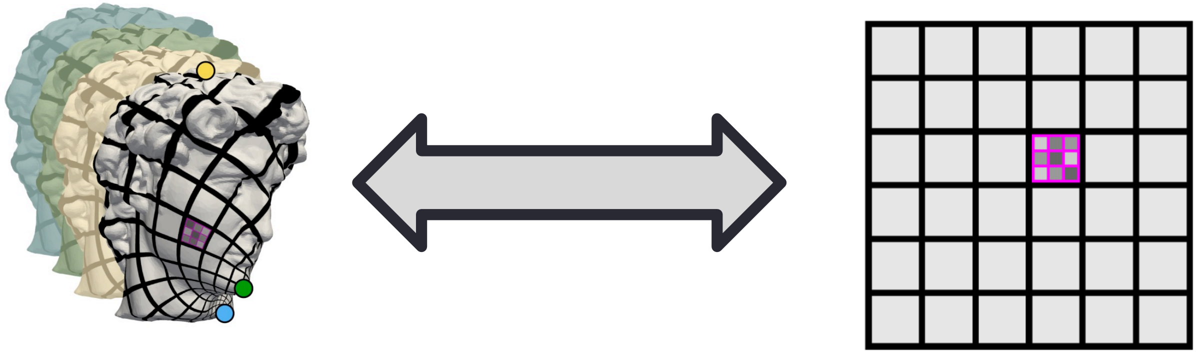

- Use off-the-shelf CNNs for images

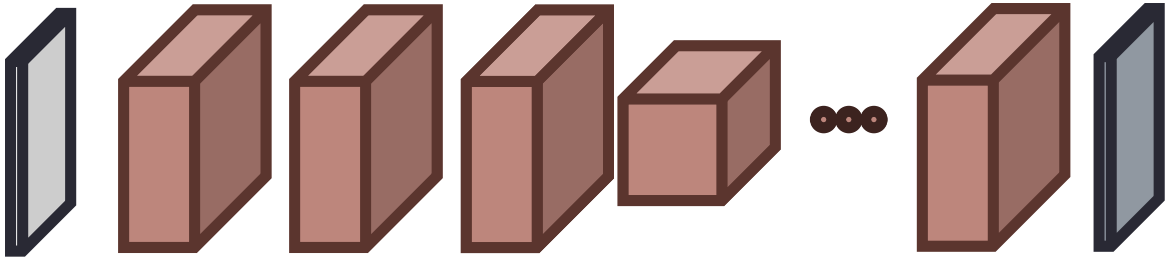

What is required to define a CNN?

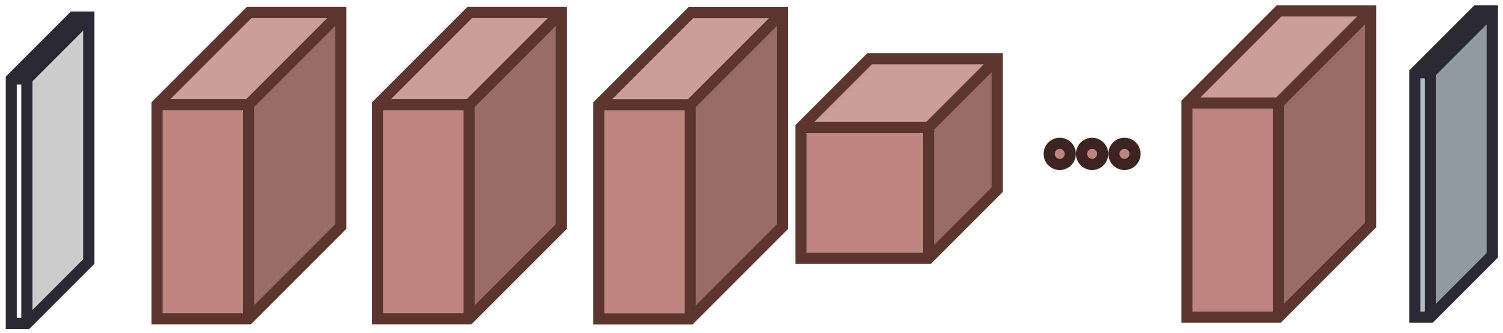

The main idea: how to move on your domain

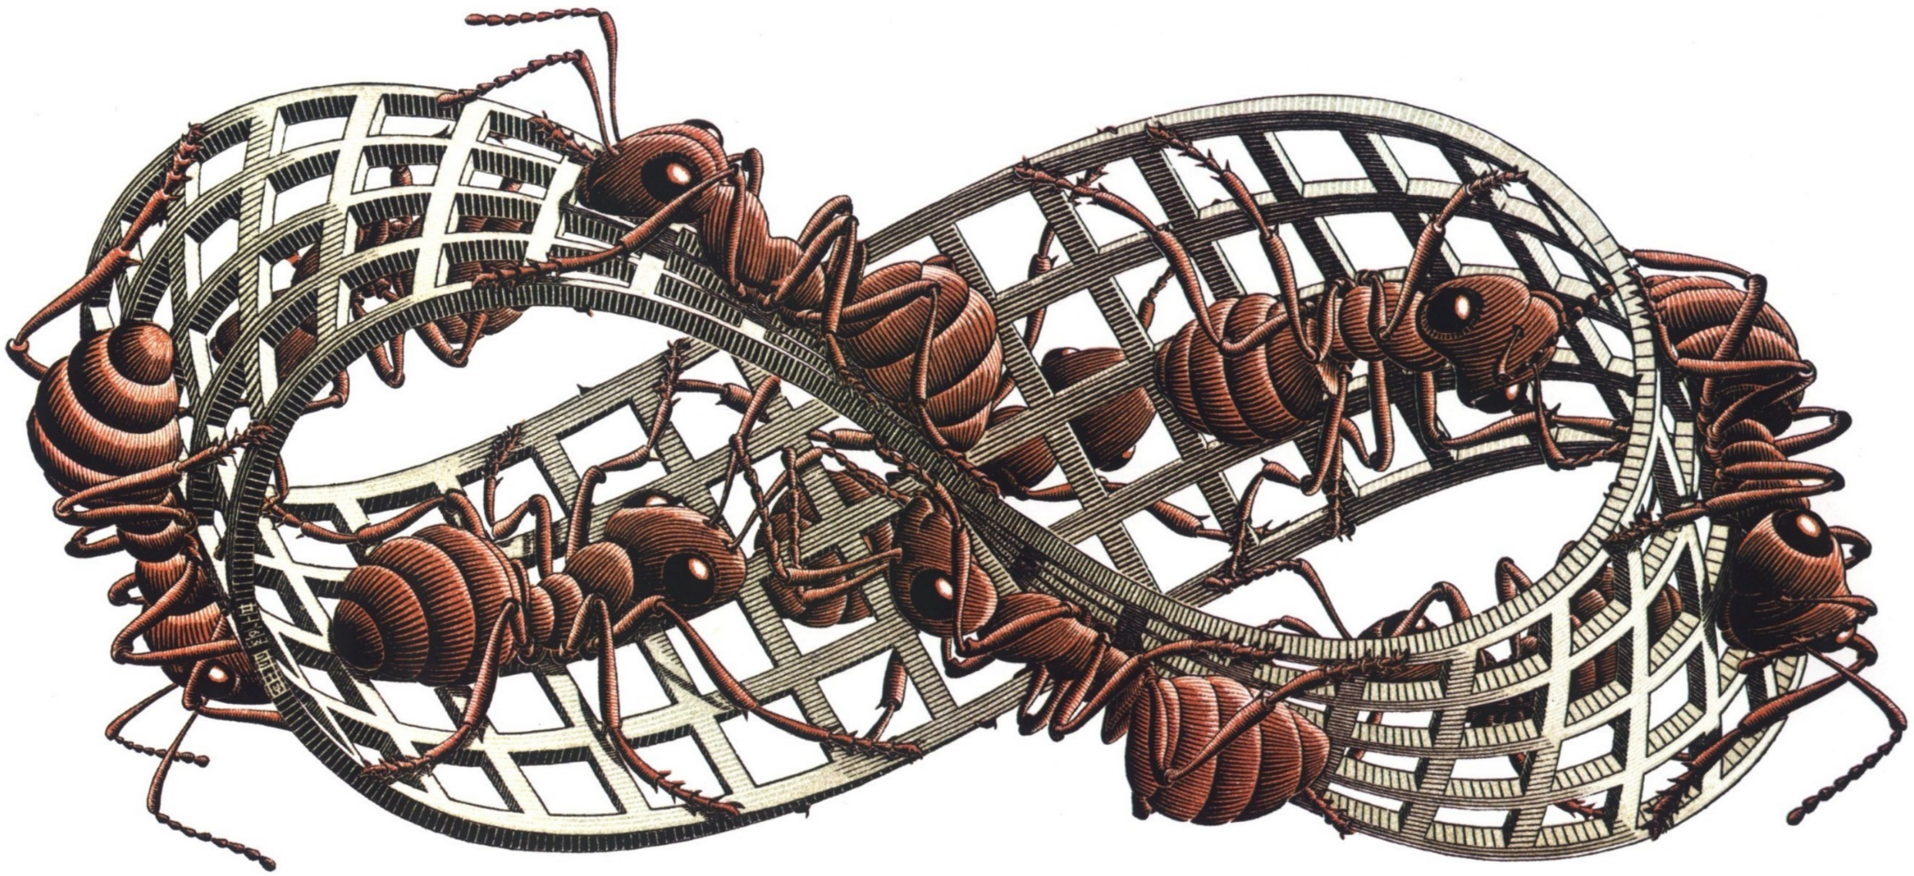

Möbius Strip II (M.C. Escher, 1963, Woodcut)

Translations

Two dimensional, commutative

Isometries of $\mathbb{R}^2$

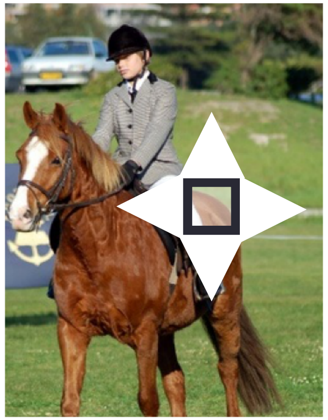

Convolution

Linear

Translation invariant

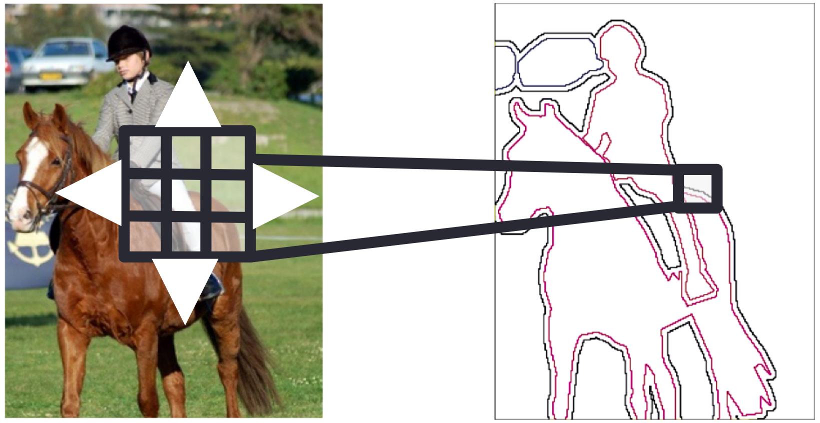

Pooling

Non-linear (max)

Sub-translation invariant

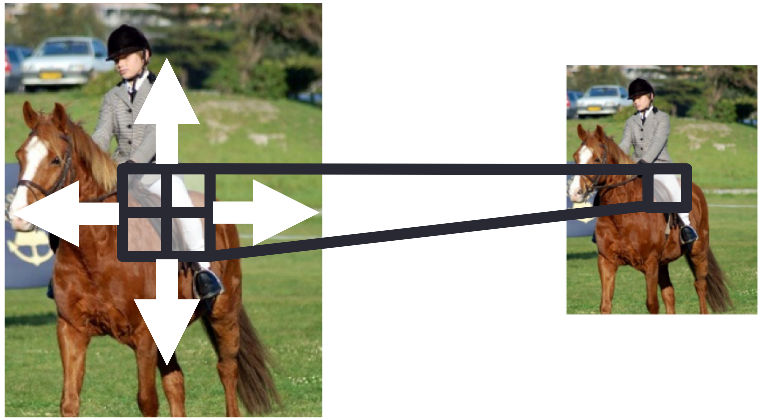

Defining CNNs on surfaces

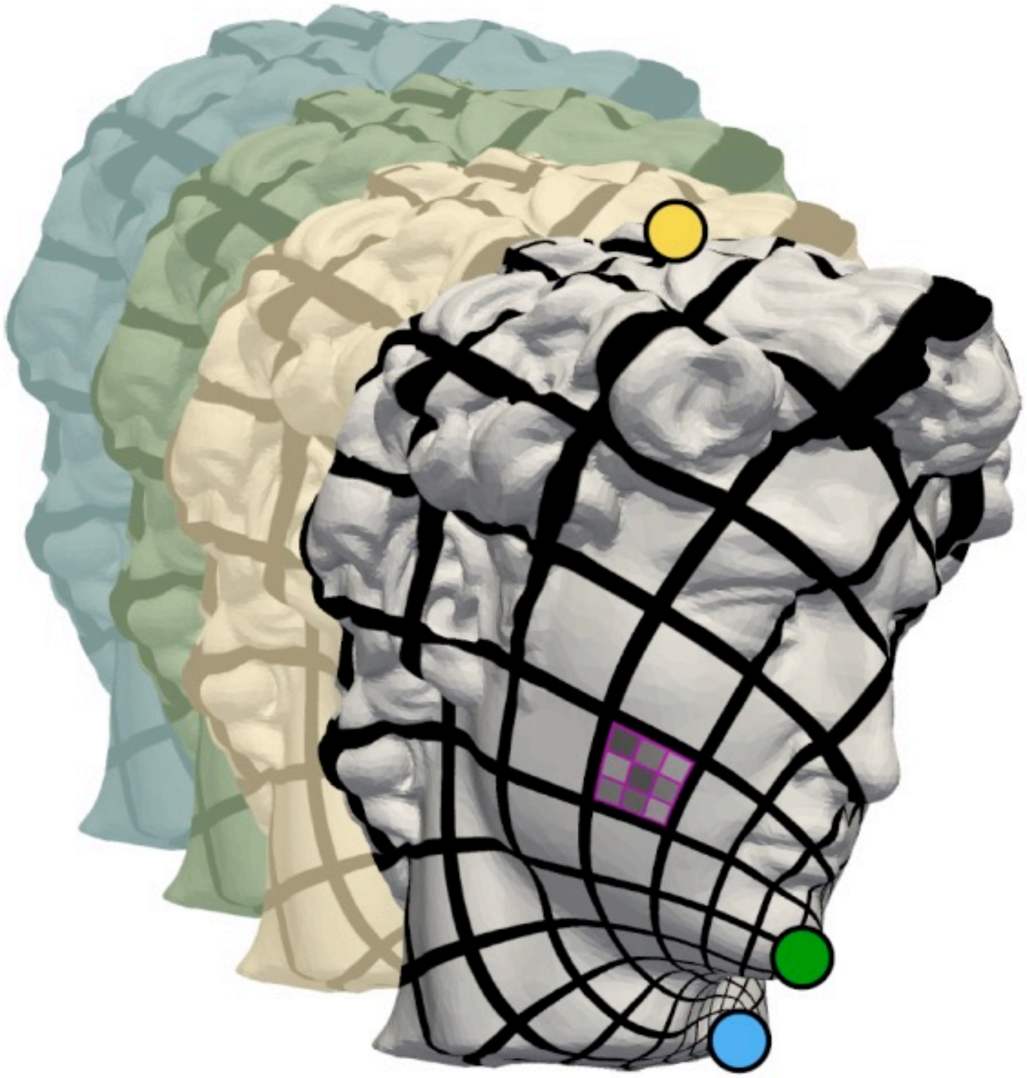

Translations on surfaces?

- Translation on surface $\stackrel{\text{def}}{=}$ locally Euclidean translation

- Flow along non-vanishing vector fields

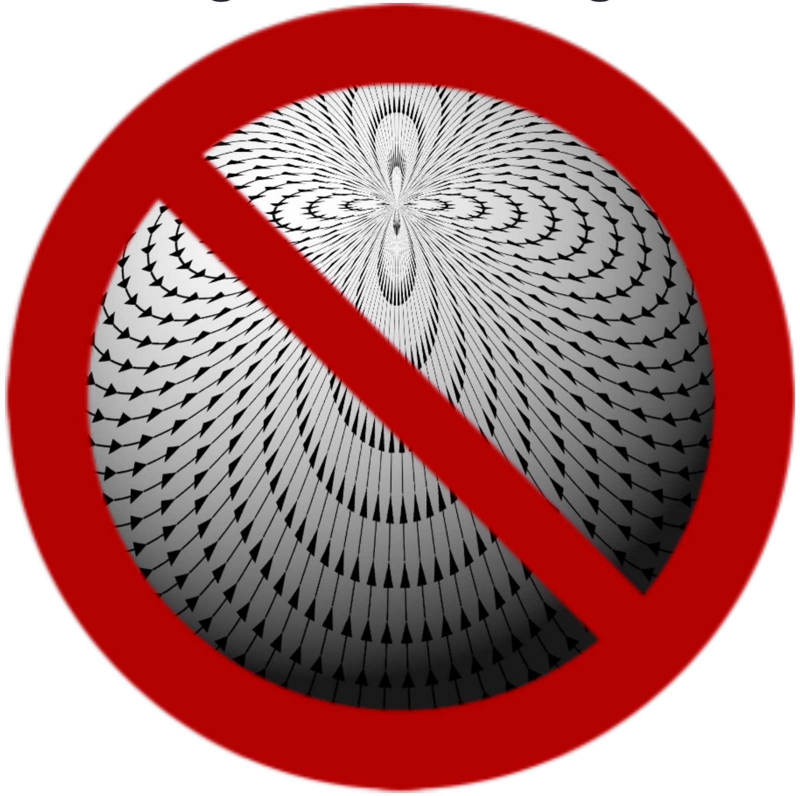

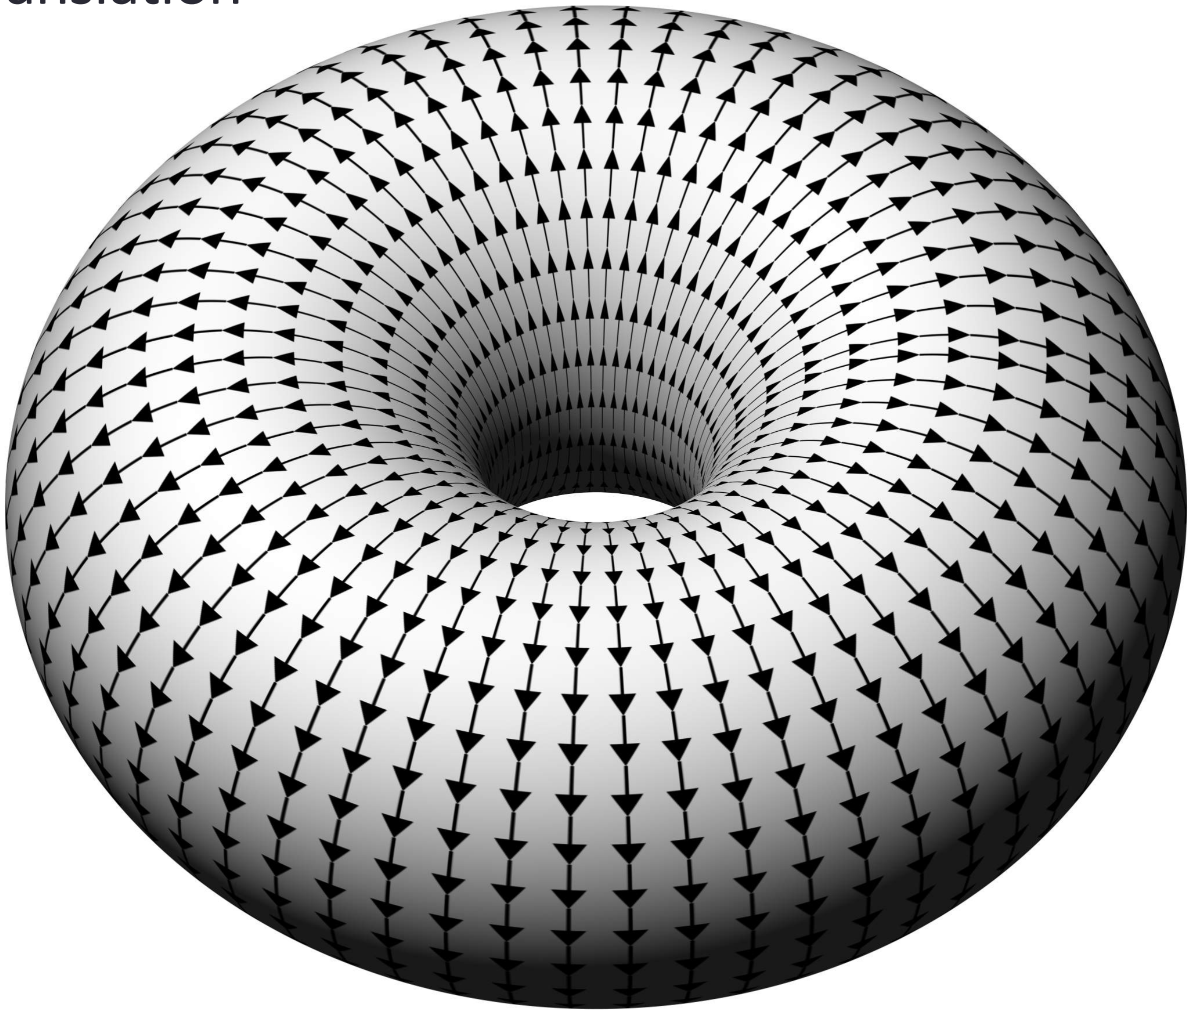

Which compact surfaces admit non-vanishing vector fields?

Flat torus $\mathcal{T}$

- Translations “modulo 1”
- Full translation invariance on the *flat torus*

$\mathcal{T}$

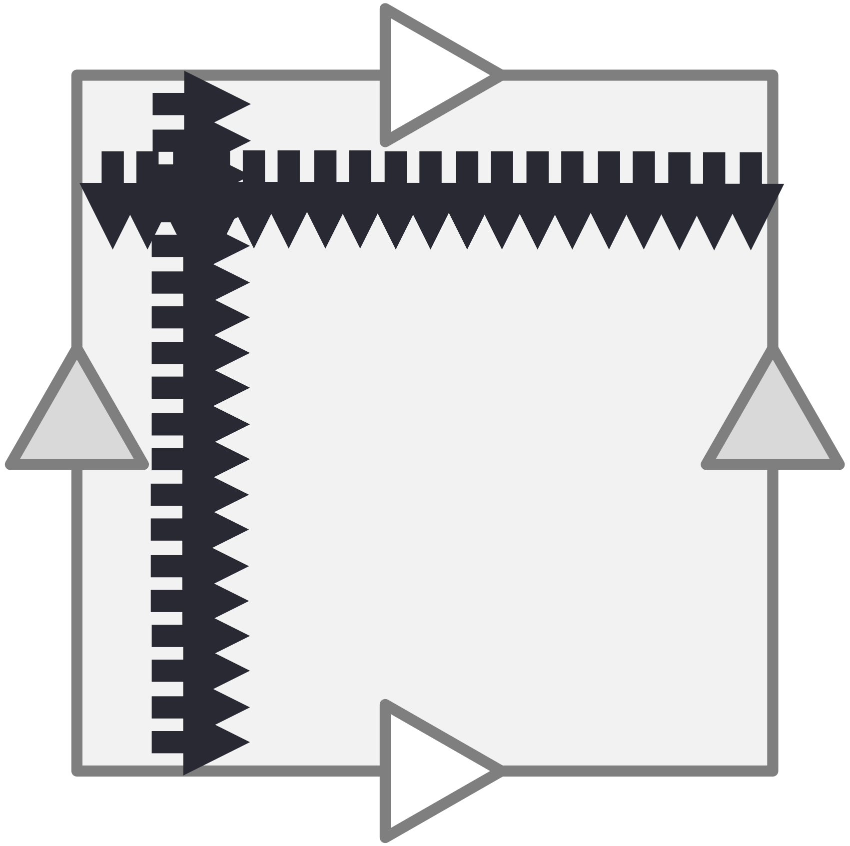

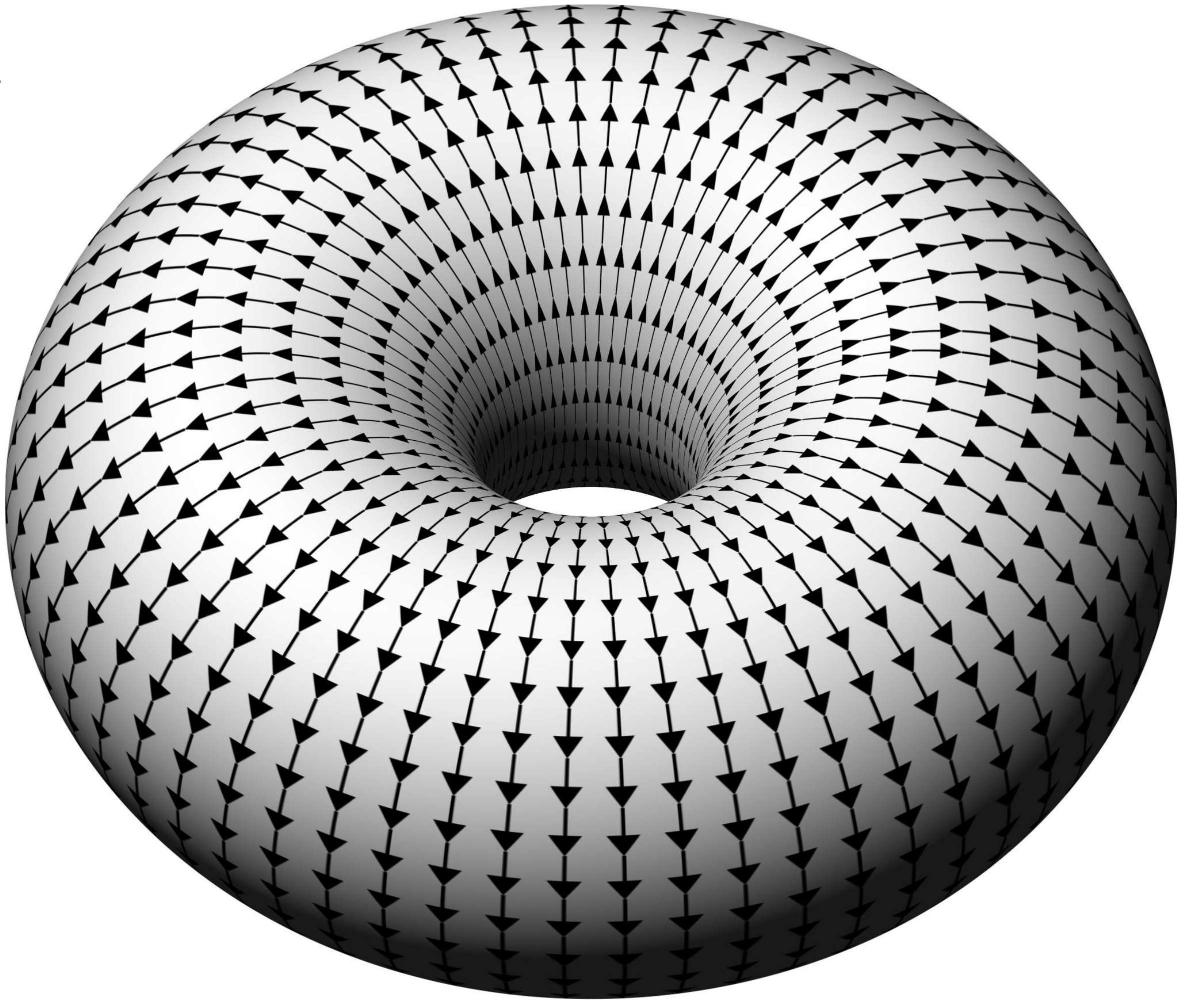

Only the torus!

Index of vector field

Euler characteristic

$$\sum_i \mathrm{index}_{x_i} = \chi$$

- Poincaré-Hopf: For a compact orientable surface

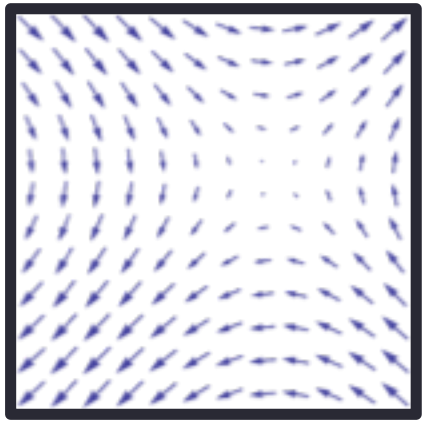

- Index – a measure of the complexity near a vanishing point

- Non-vanishing vector field implies **genus 1 - torus**

CNN on flat torus

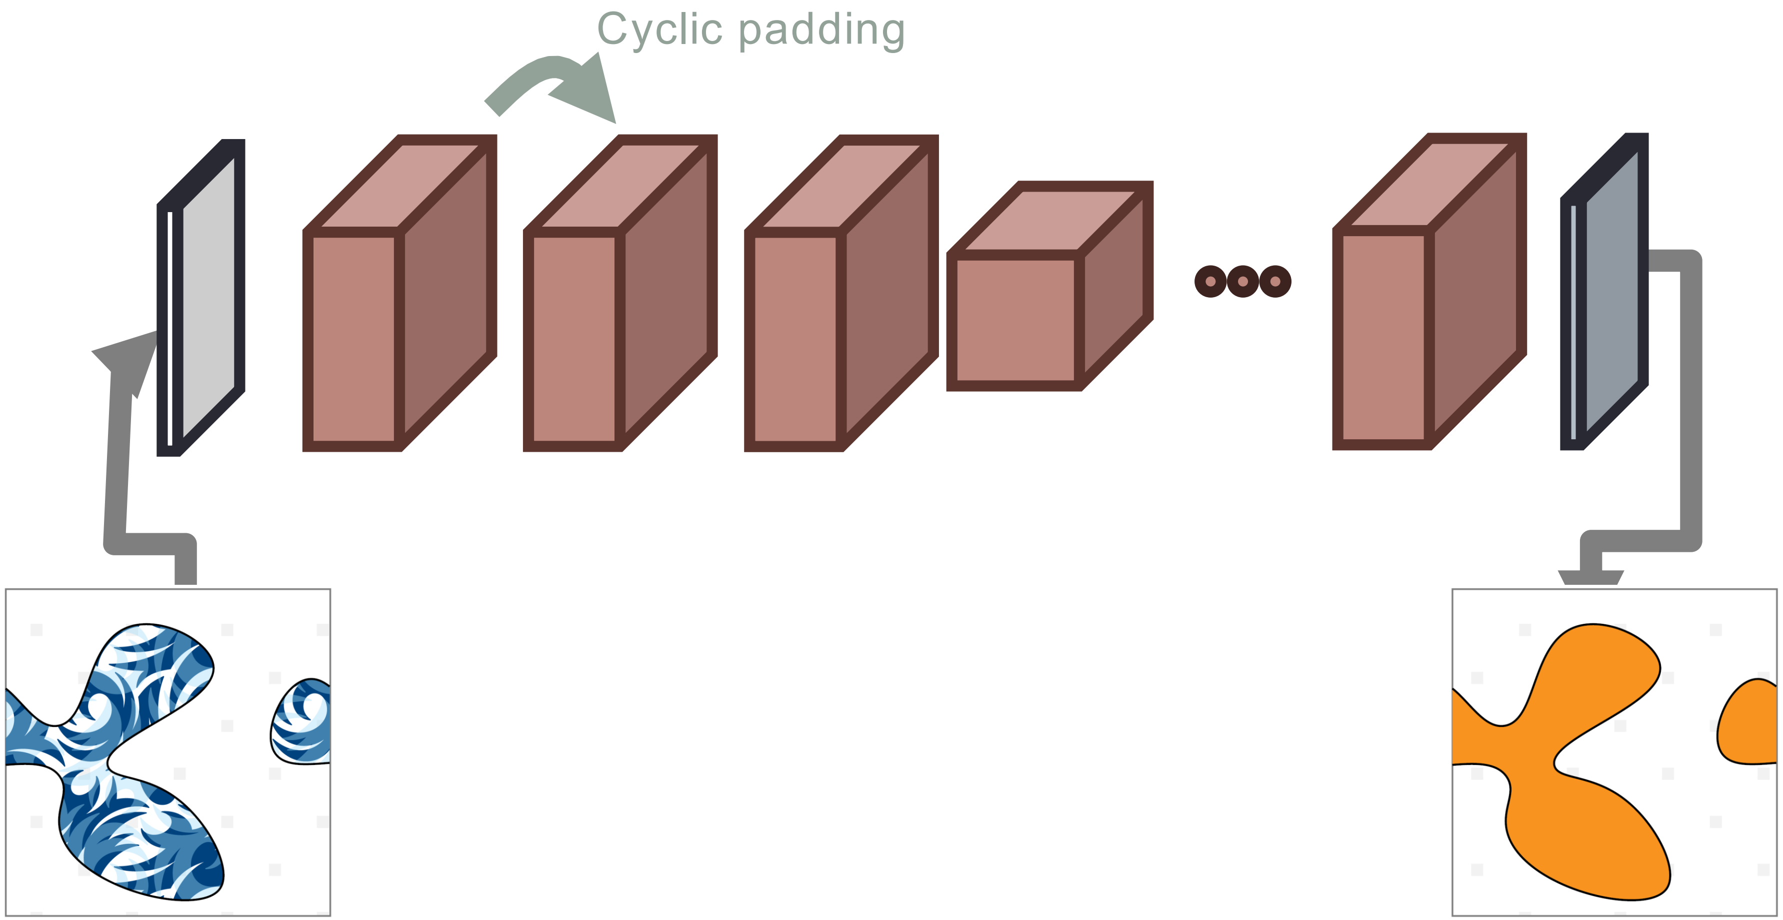

Recap

- CNN is well-defined over flat-torus

- Roadblocks for CNN on sphere-type surfaces
  - ***Topological***: No locally Euclidean translations on spheres
  - ***Geometrical***: The flat torus is flat and our surface is not

Solution: Map the surface to a flat torus

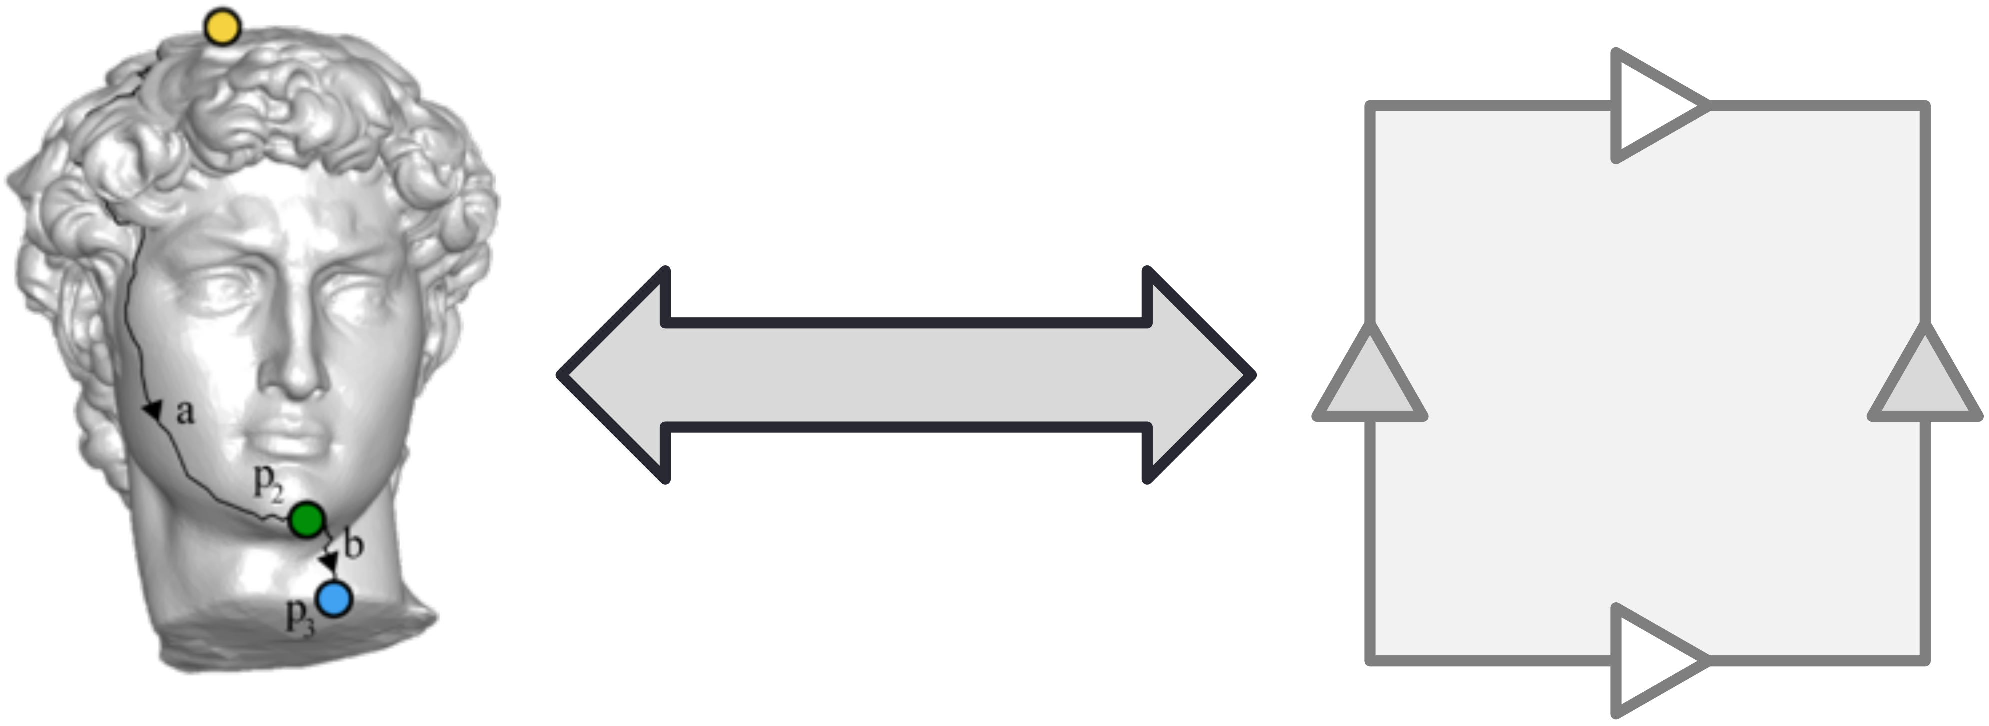

Torus 4-cover

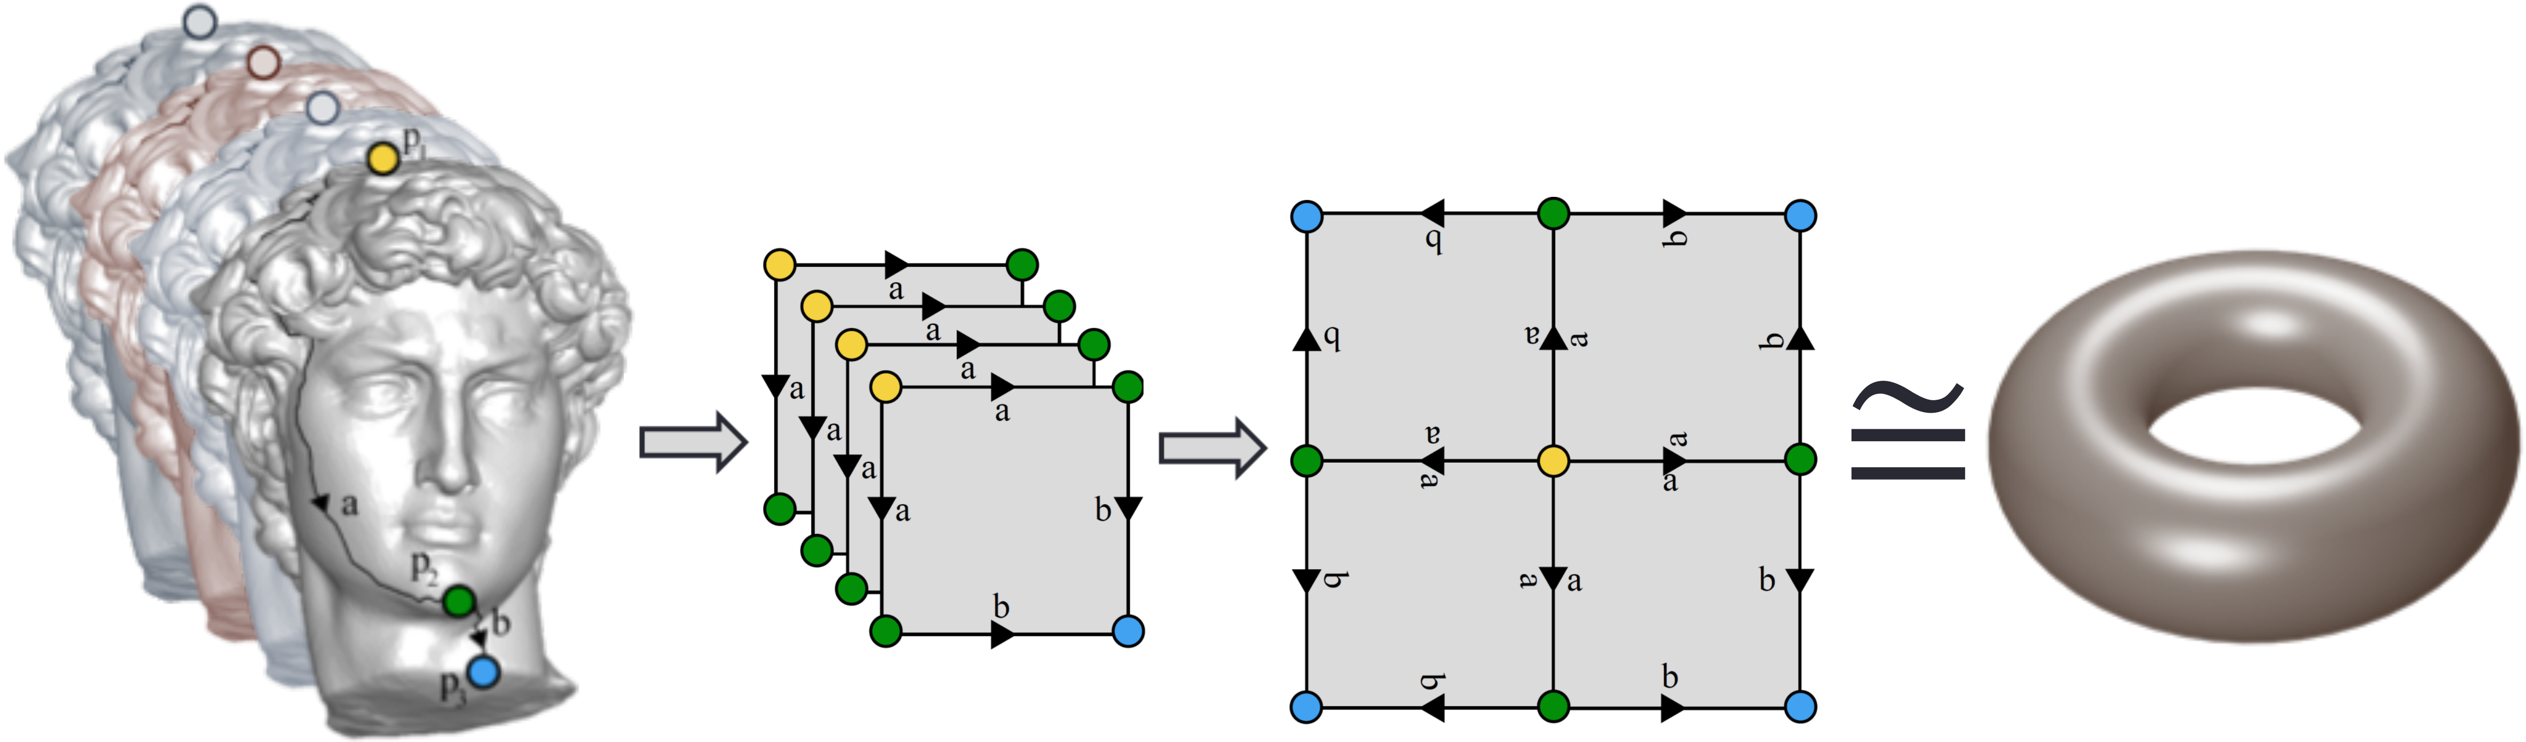

Mapping the Torus to the flat Torus

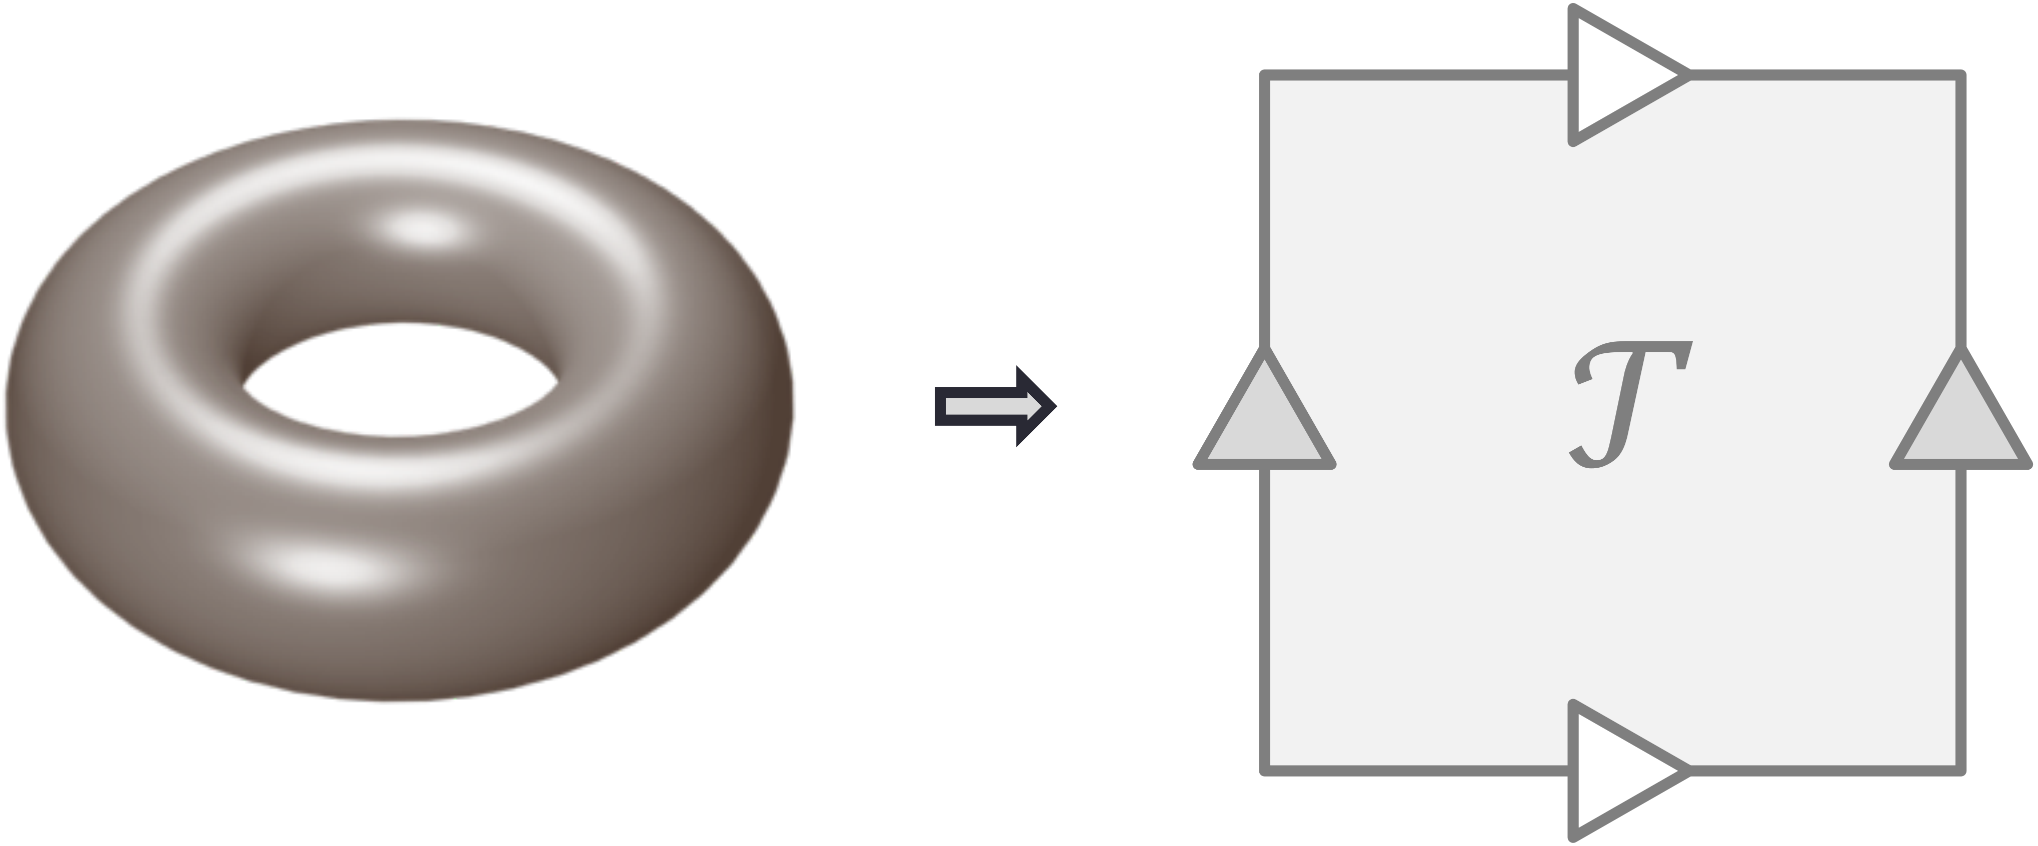

Aigerman and Lipman, 2015

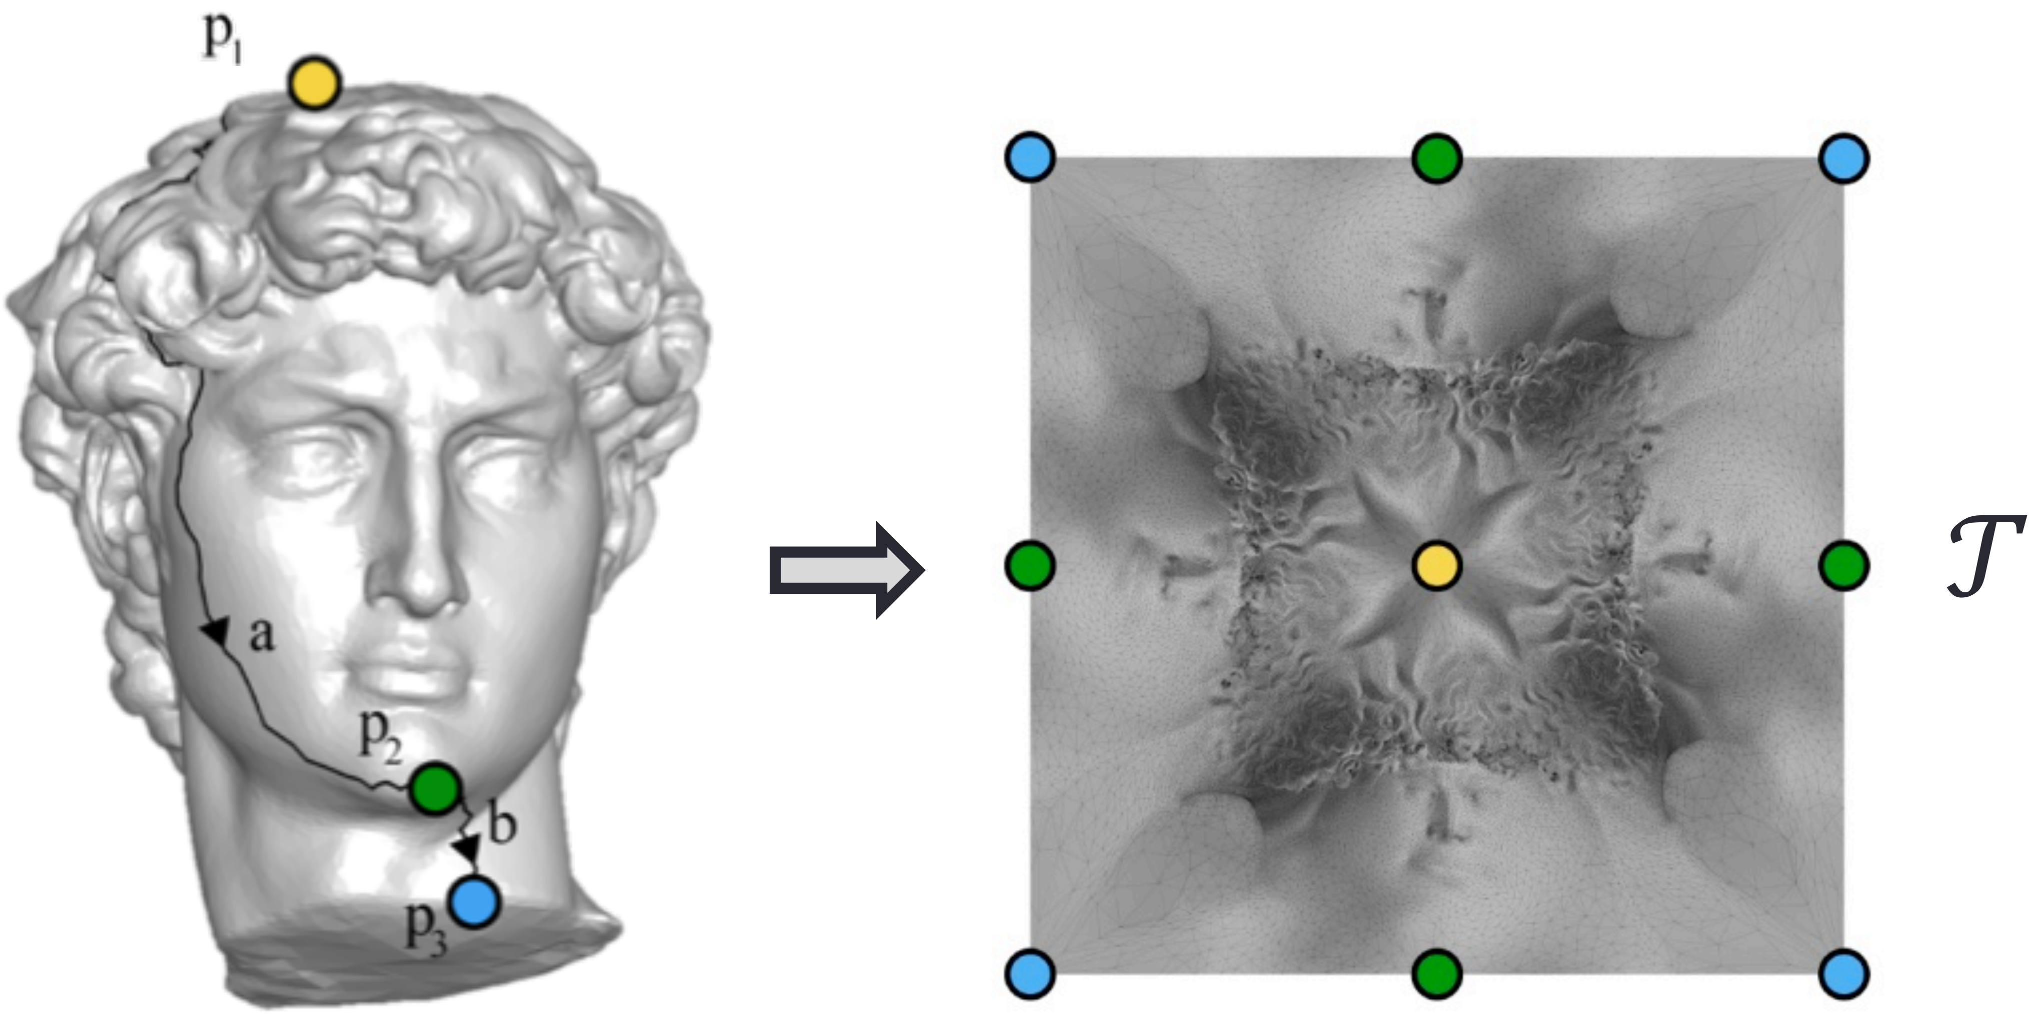
p_1
a
p_2
b
p_3
$\mathcal{T}$

The pullback translation

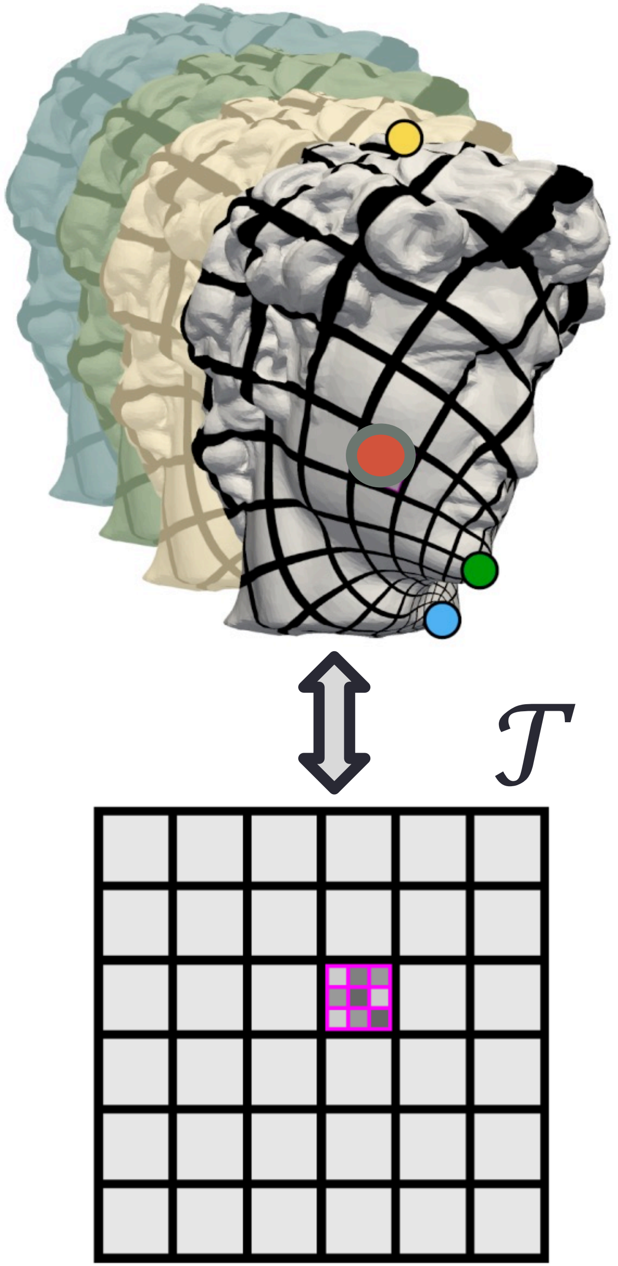

Pull-back

Translations: pull-back Euclidean translations

Two dimensional, commutative

Conformal maps

Pull-back convolution

Linear

Theorem: Translation invariance

Pull-back pooling

Non-linear (max)

Sub-translation invariant

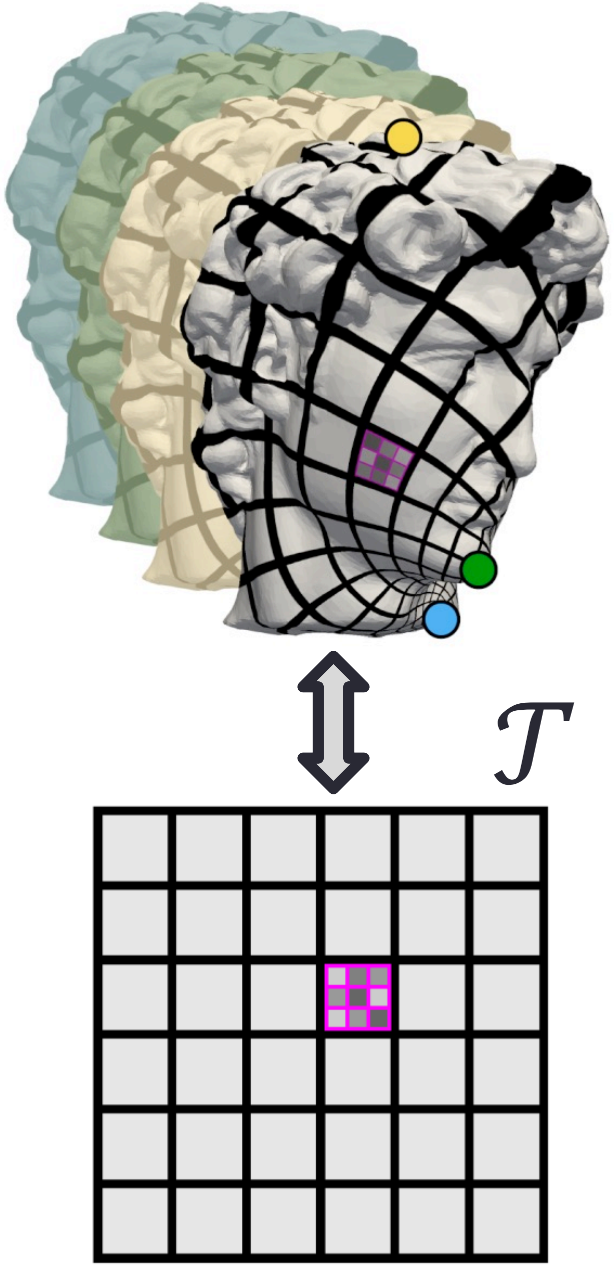

Mapped functions

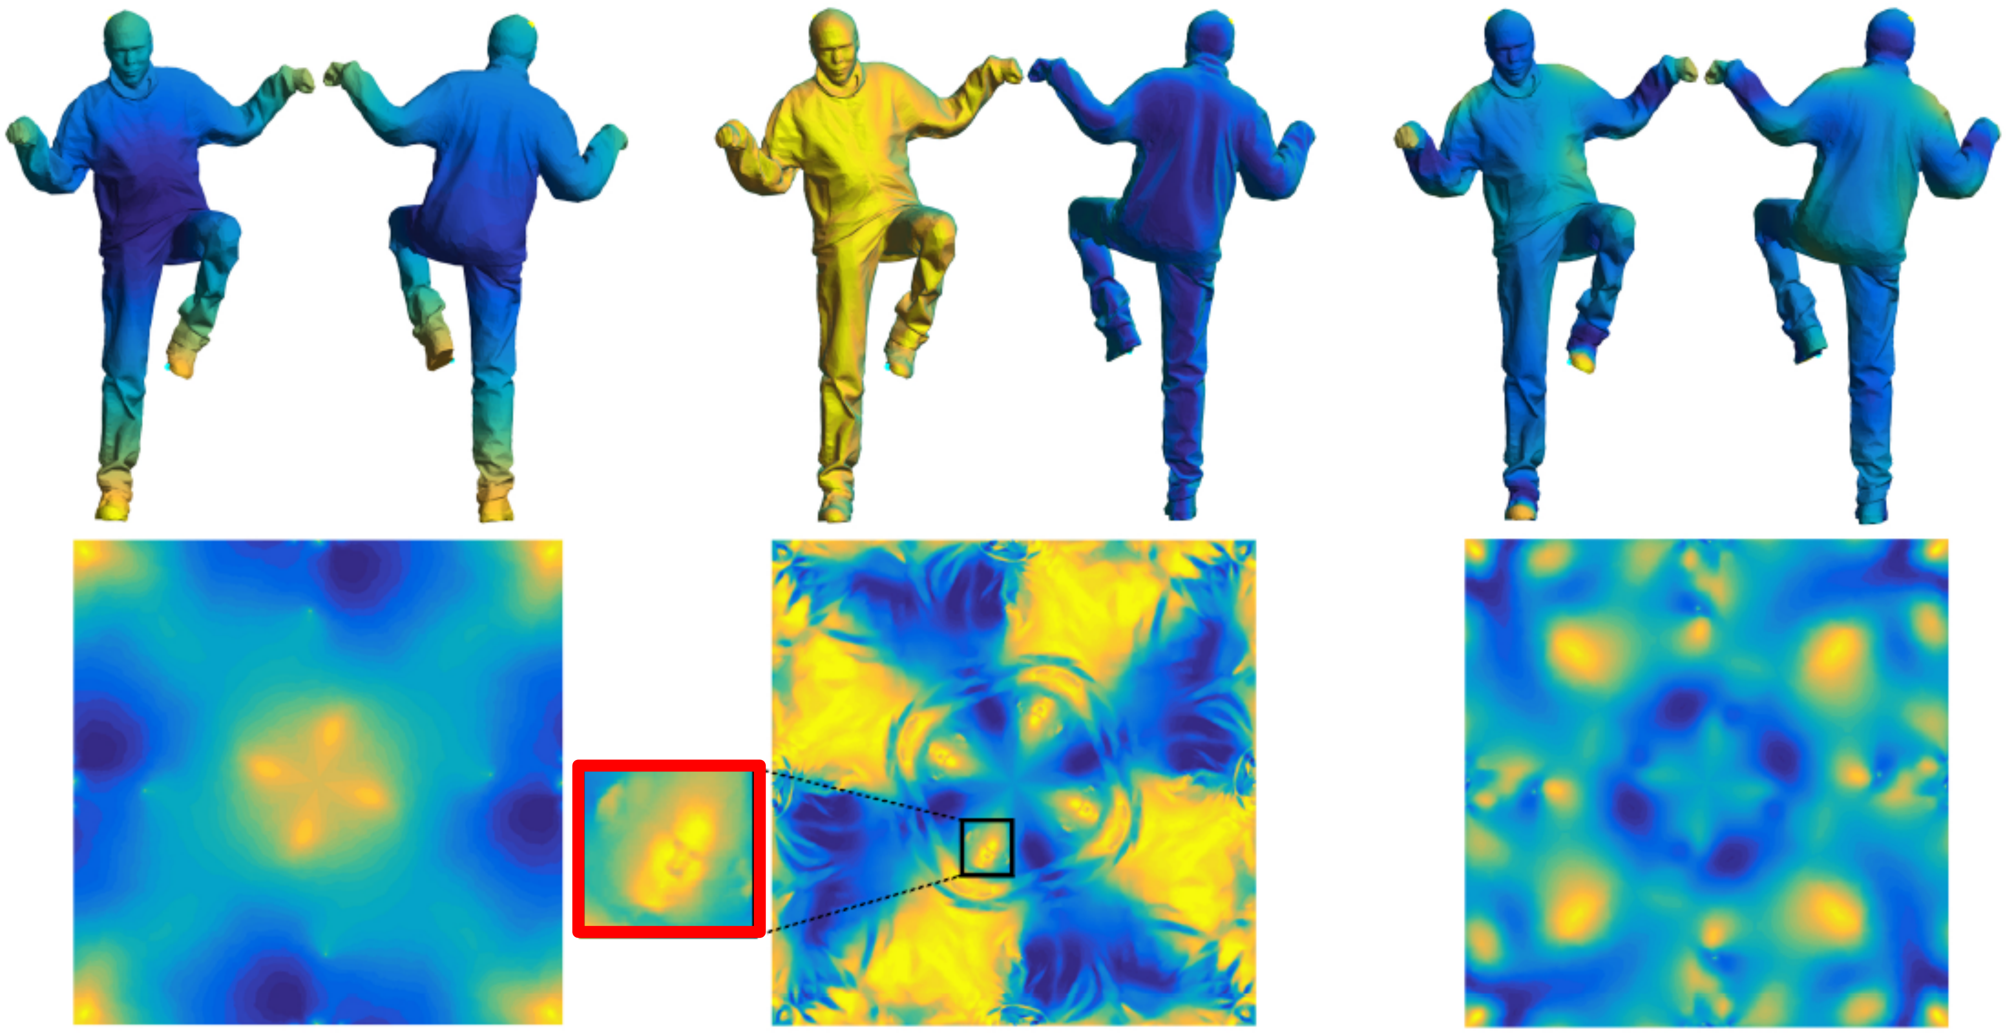

New layers

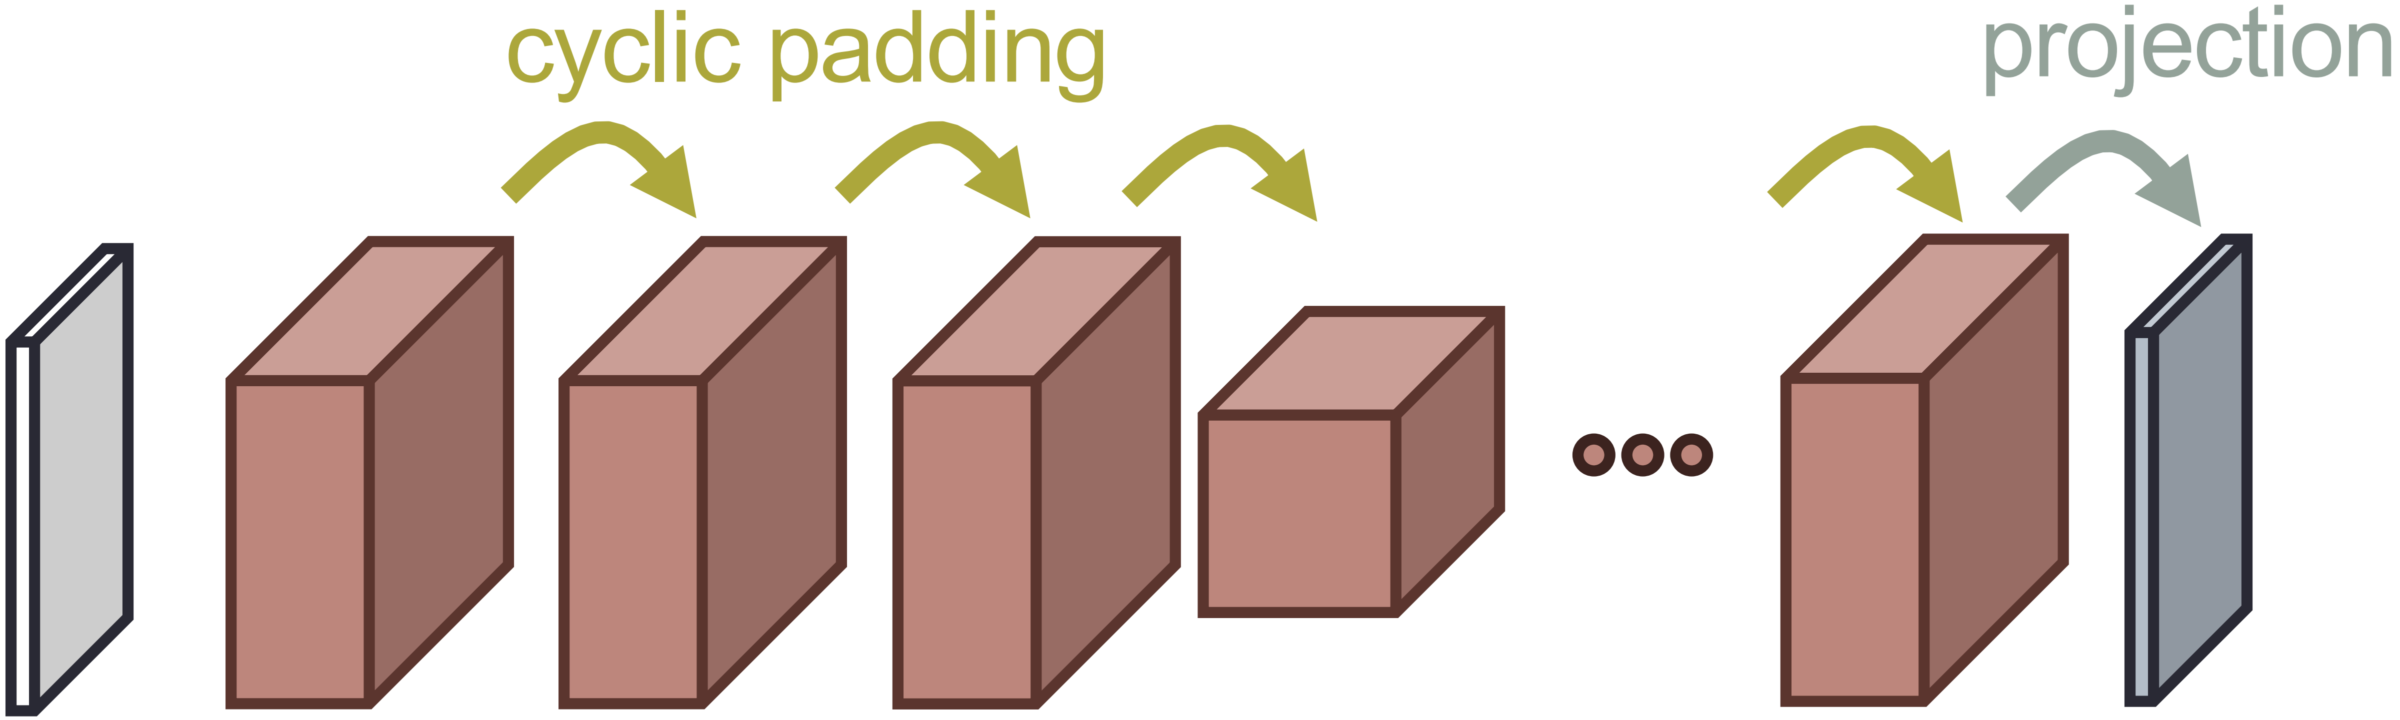

Data generation

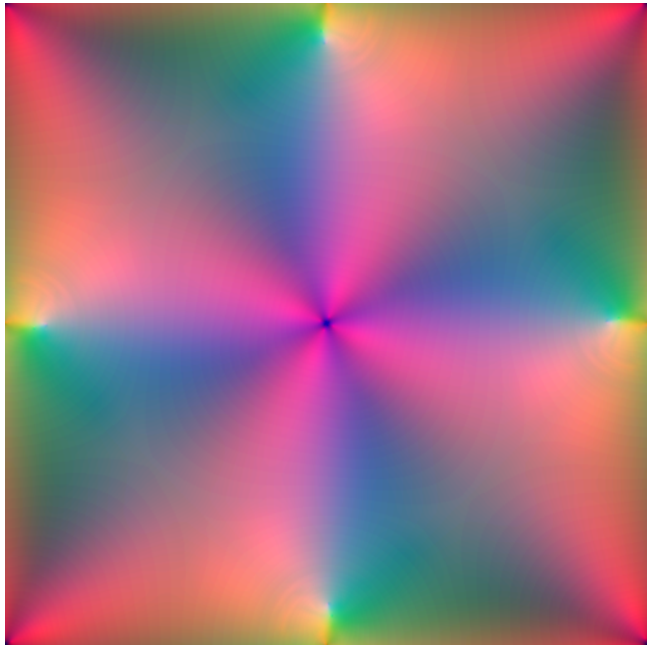

Input image

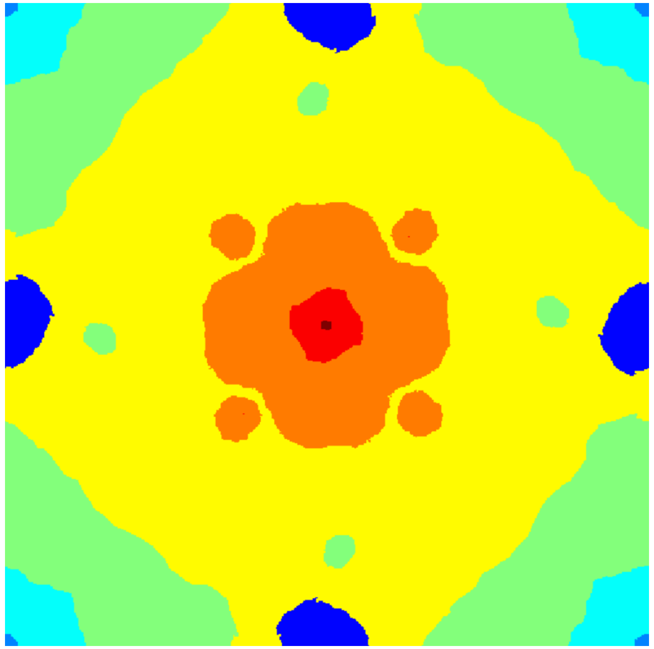

Labels

Test phase

- Aggregation from different triplets
- “Magnifying glass”
- Scale factor as weights

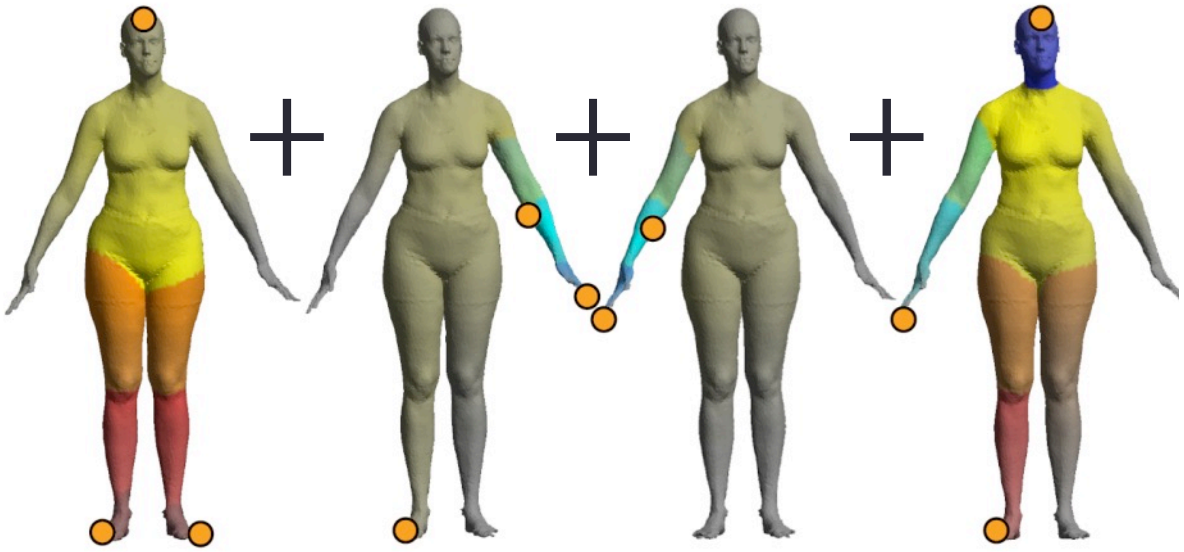

Human body segmentation

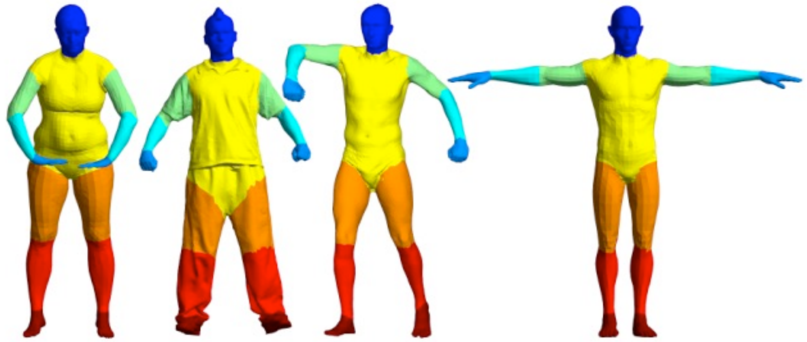

Train: 370 models
FAUST, MIT, SCAPE, ADOBE

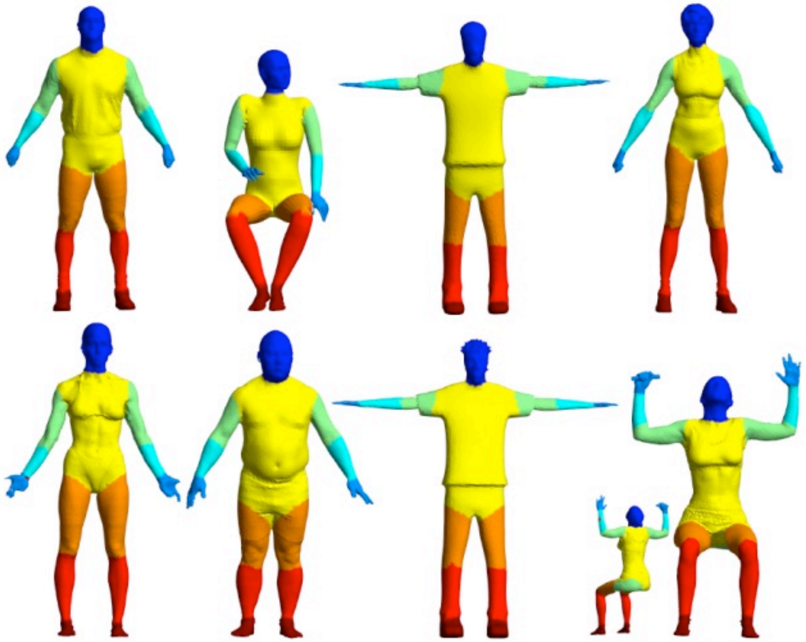

Test: 18 models
SHREC07

Easy functions

- Normals

- Average geodesic distance

- Wave kernel signature

Raw

Complex

Human body segmentation

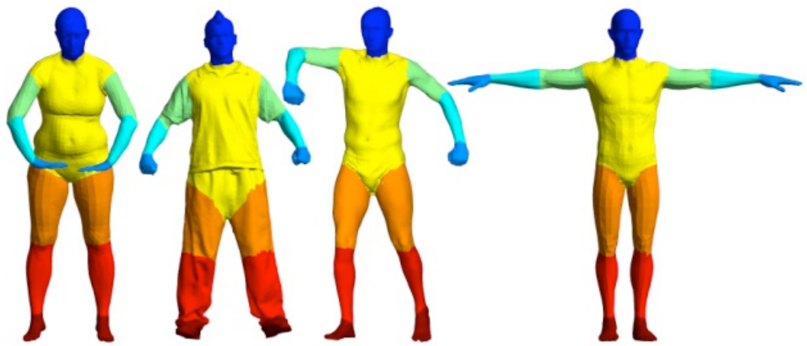

Train: 370 models
FAUST, MIT, SCAPE, ADOBE

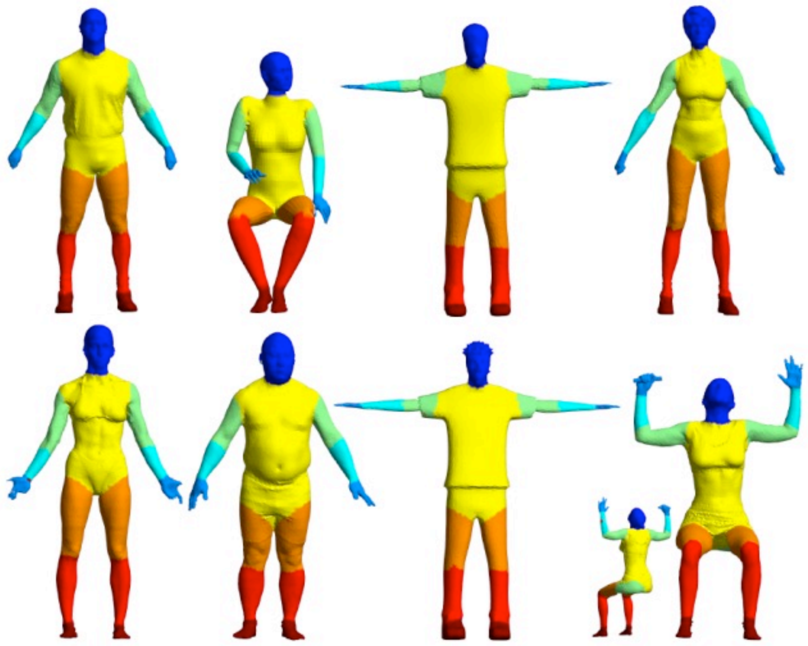

Test: 18 models
SHREC07

CNN applied to other data

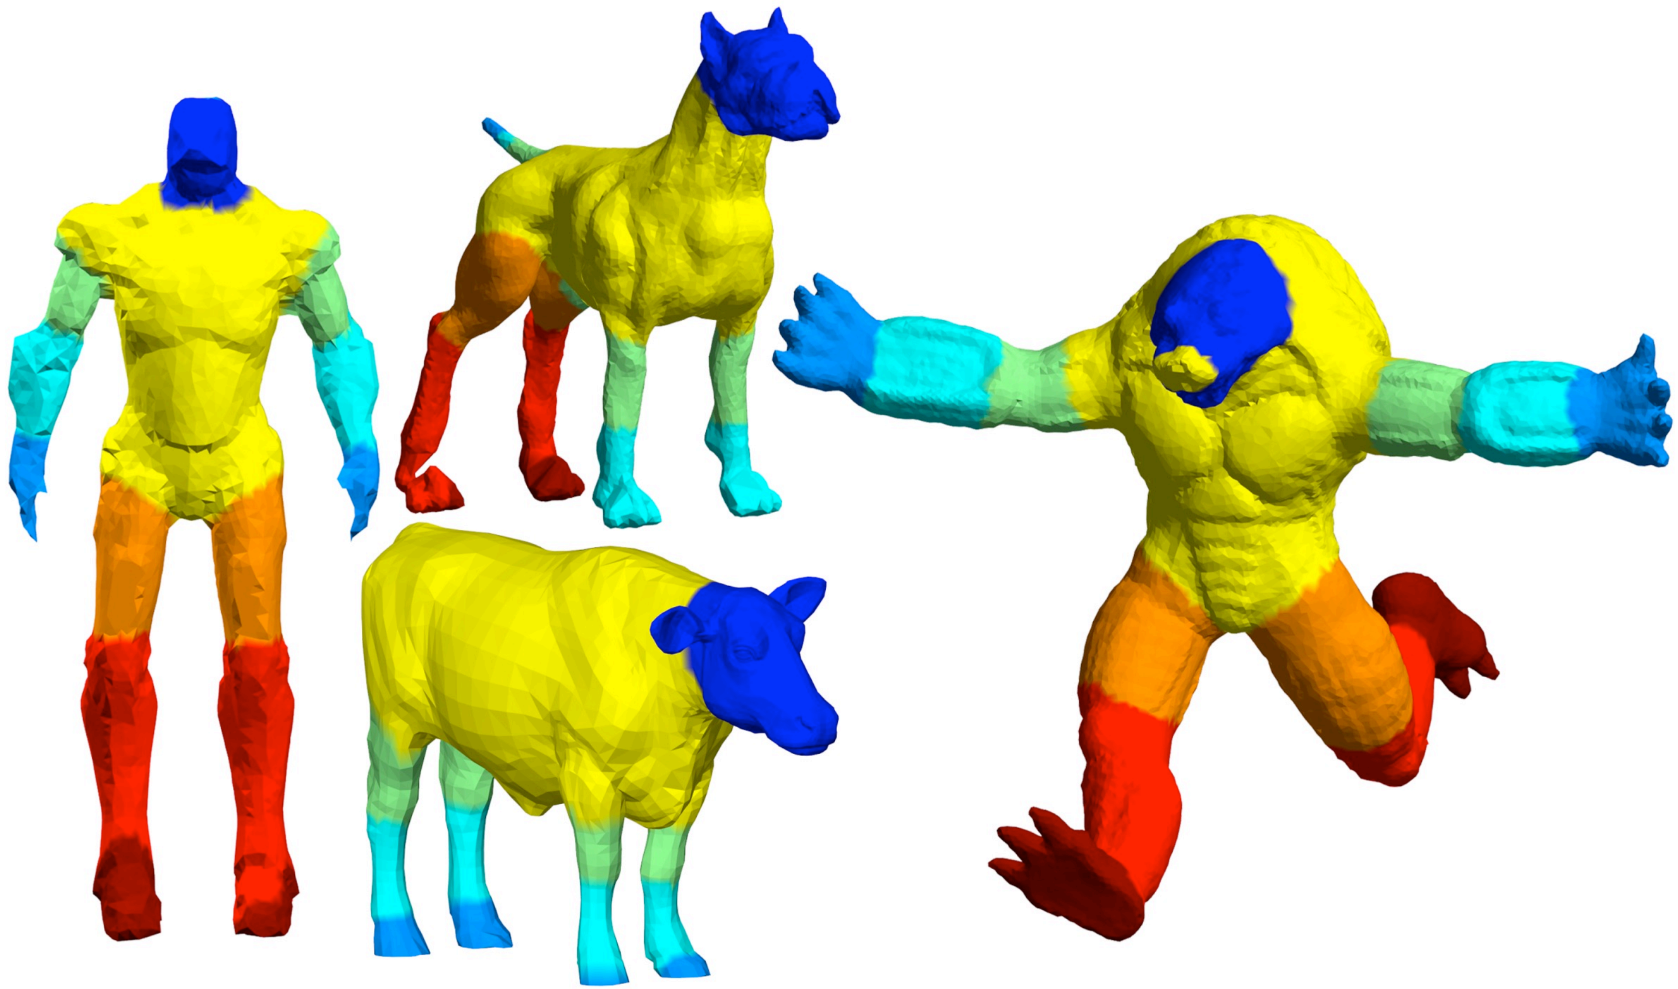

Biological landmarks detection

- Train: 73 teeth from BOYER
- Only curvature and scale factor

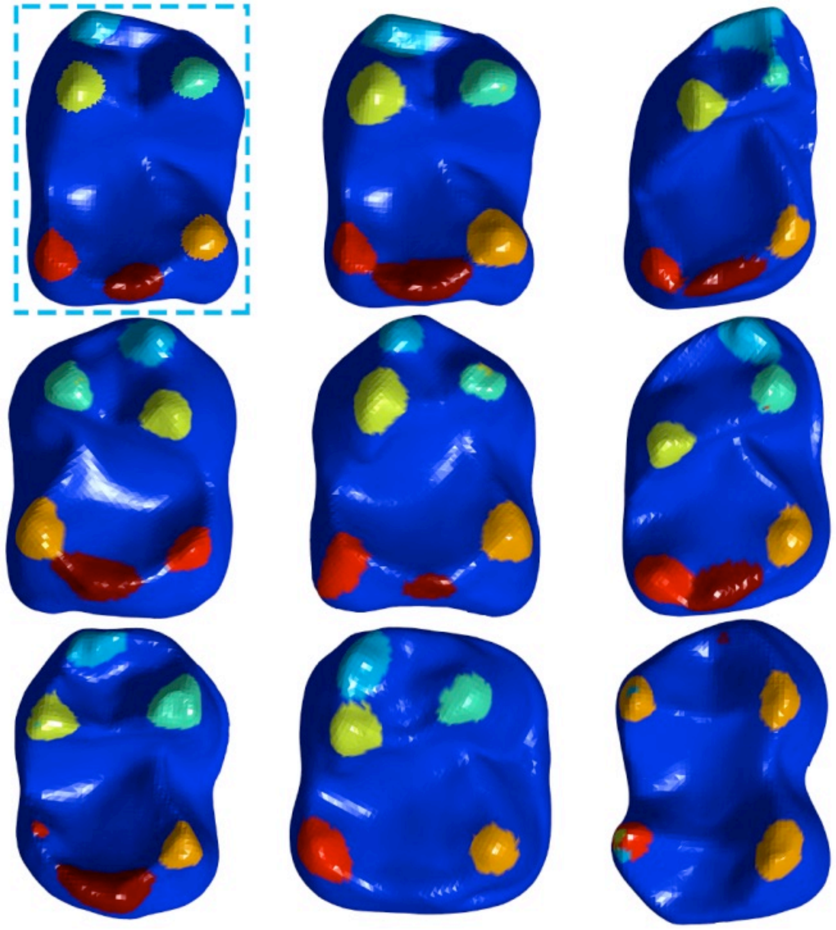

Test: 8 teeth from BOYER

Biological landmarks

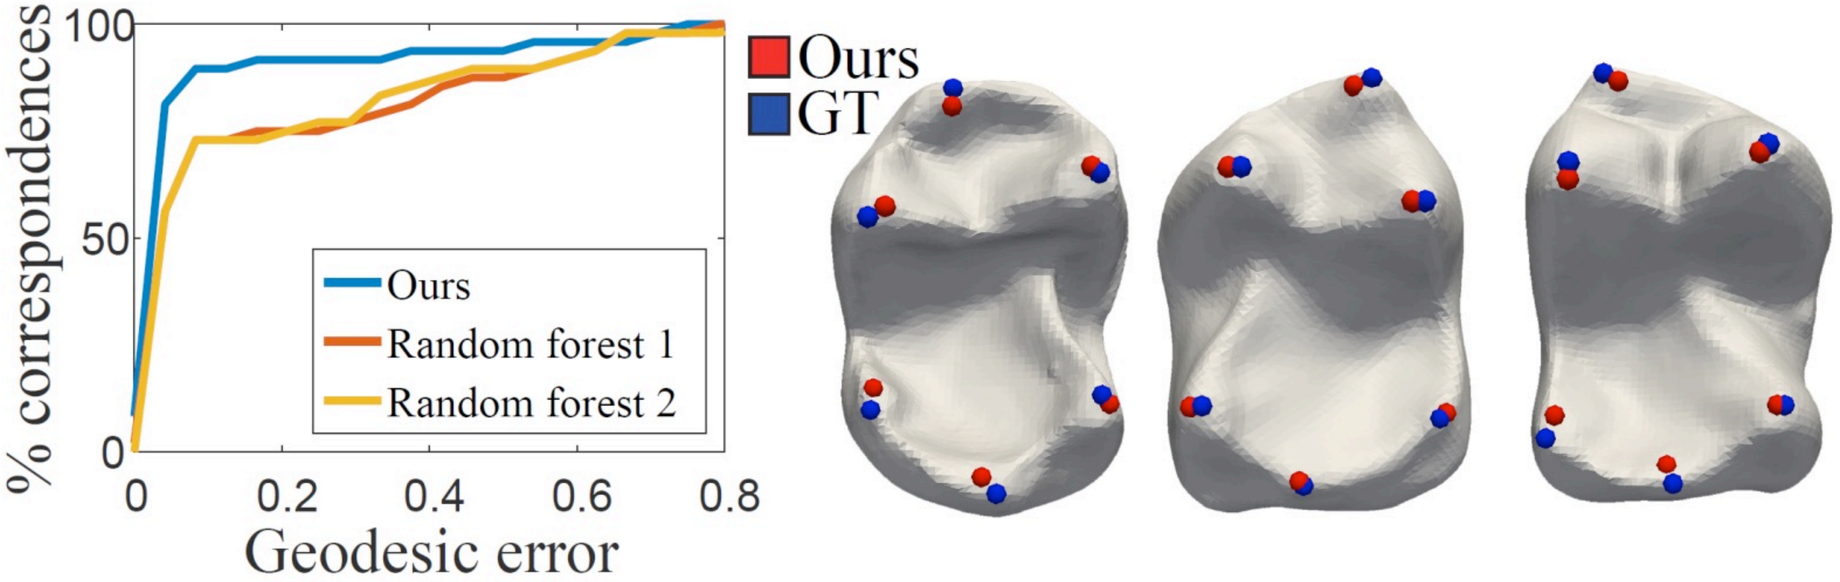

Biological landmarks

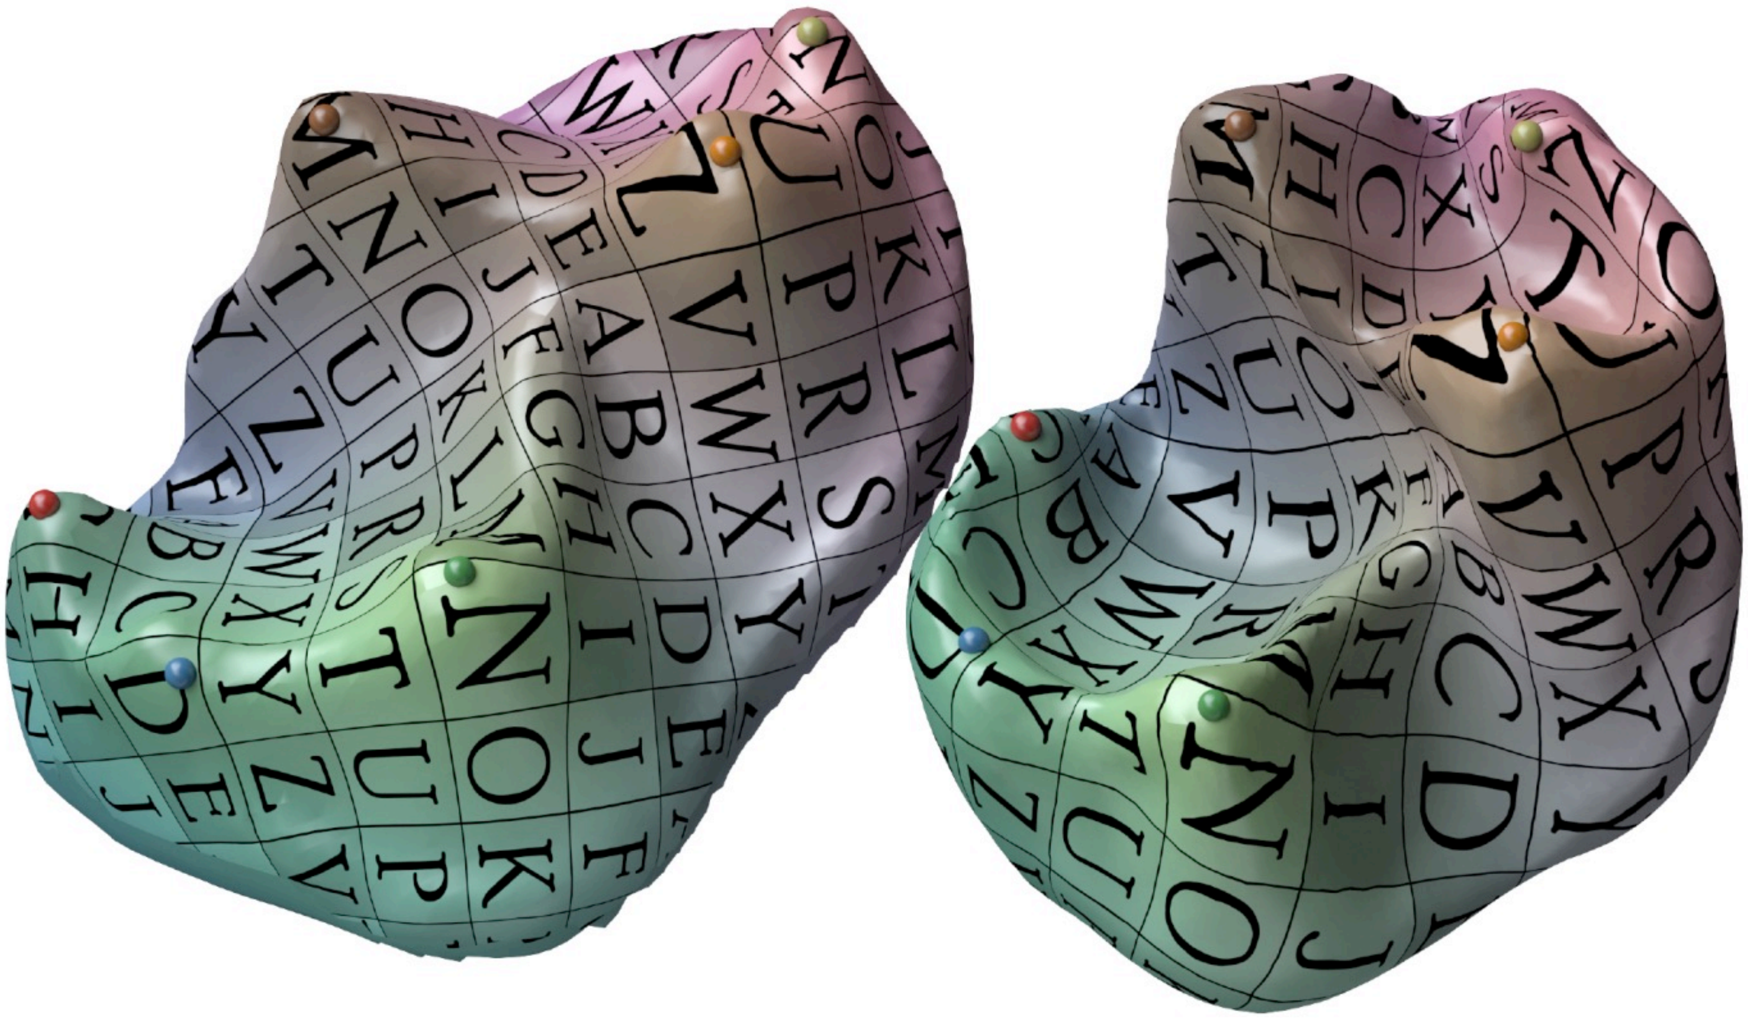

[Aigerman and Lipman 16']

Future applications

- Texture prediction on surfaces

- BRDF prediction

- UV coordinates learning

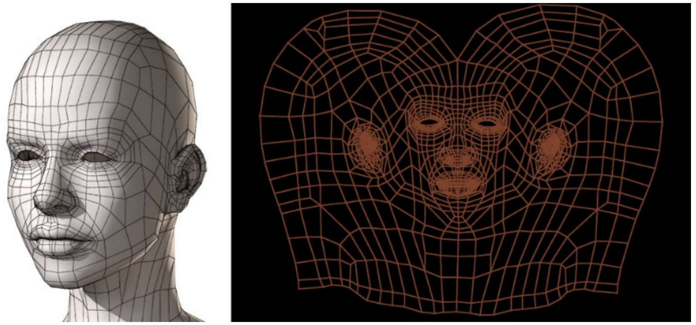

Conclusion

- CNN of sphere-type surfaces
  - We defined a meaningful convolution on surfaces
  - Learns from raw features
  - Reusing CNN software for images

Limitations and future work

- Scope: Only sphere type surfaces
- No canonical choice for triplets (and convolutions)
- Learn aggregation operator

The End

- Code is available online: http://www.wisdom.weizmann.ac.il/~haggaim/

- Support
  - ERC Starting Grant (SurfComp)
  - Israel Science Foundation
  - I-CORE

- Thanks for listening!