Learning geometry

Ron Kimmel

Geometric Image Processing Lab.

Technion – Israel Institute of Technology

CUHK MATH-IMS Jan 26 2021

Our playground

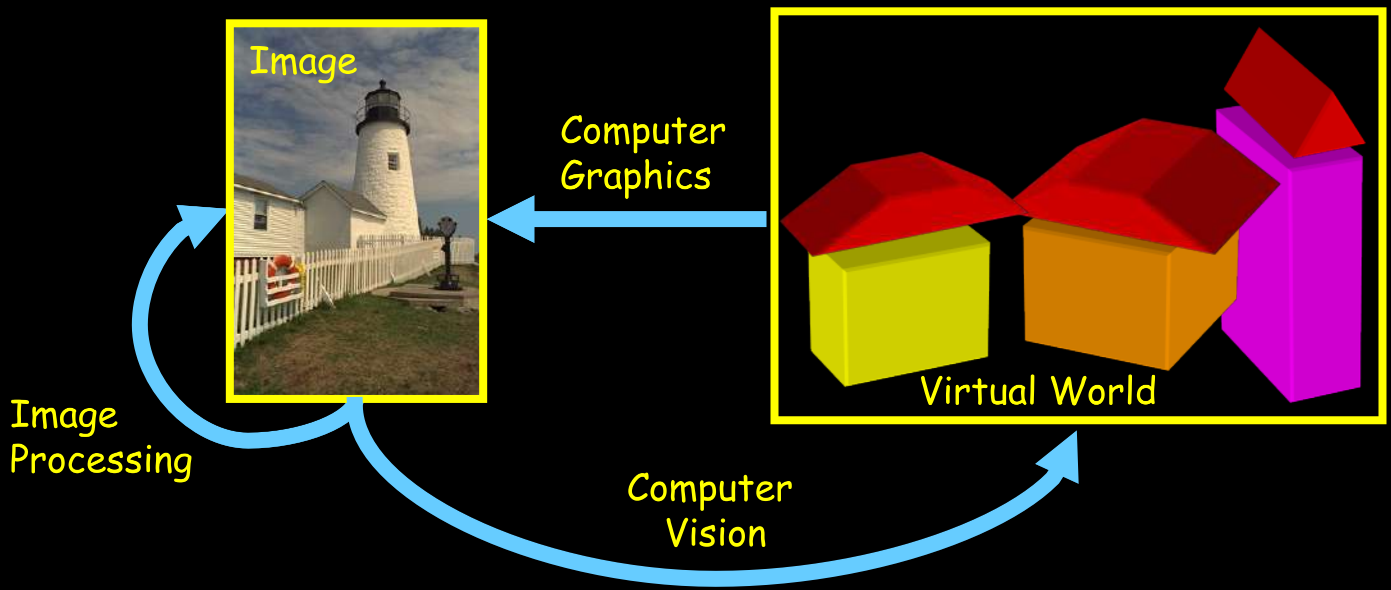

Our playground

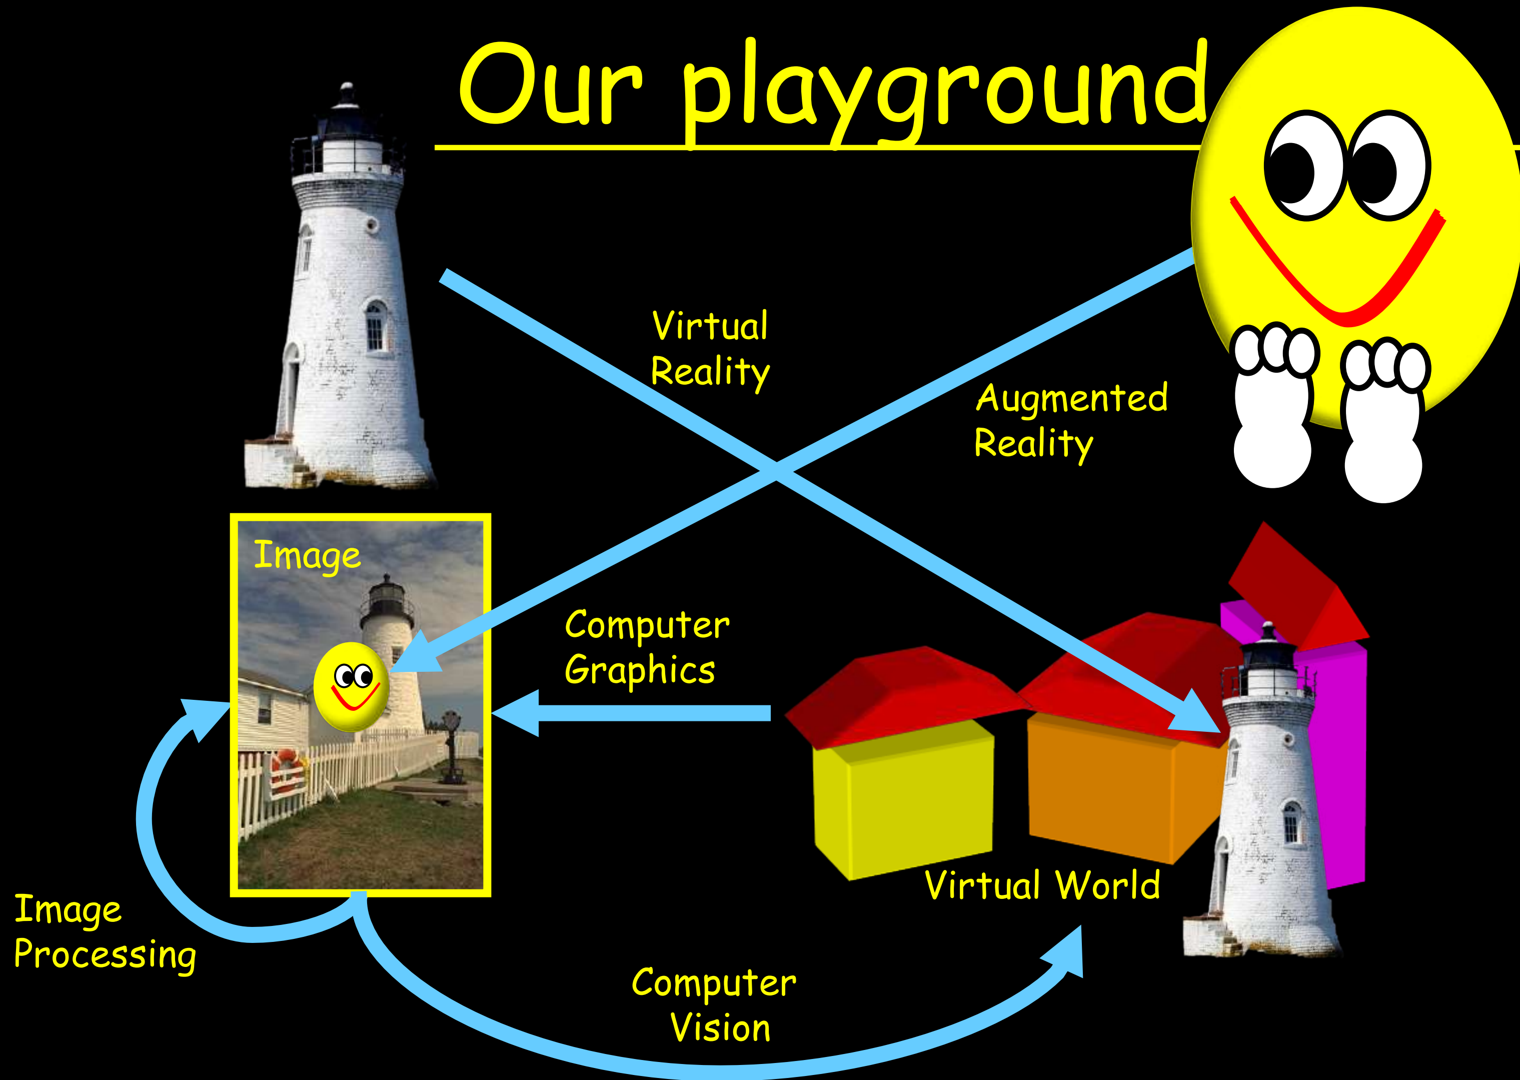

Computer vision

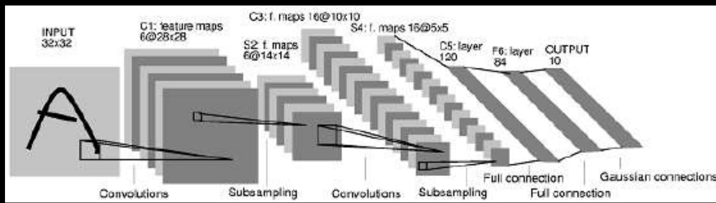

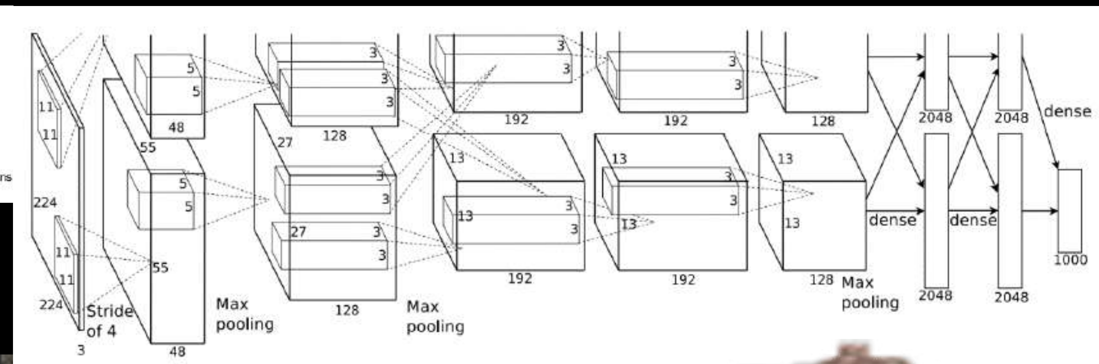

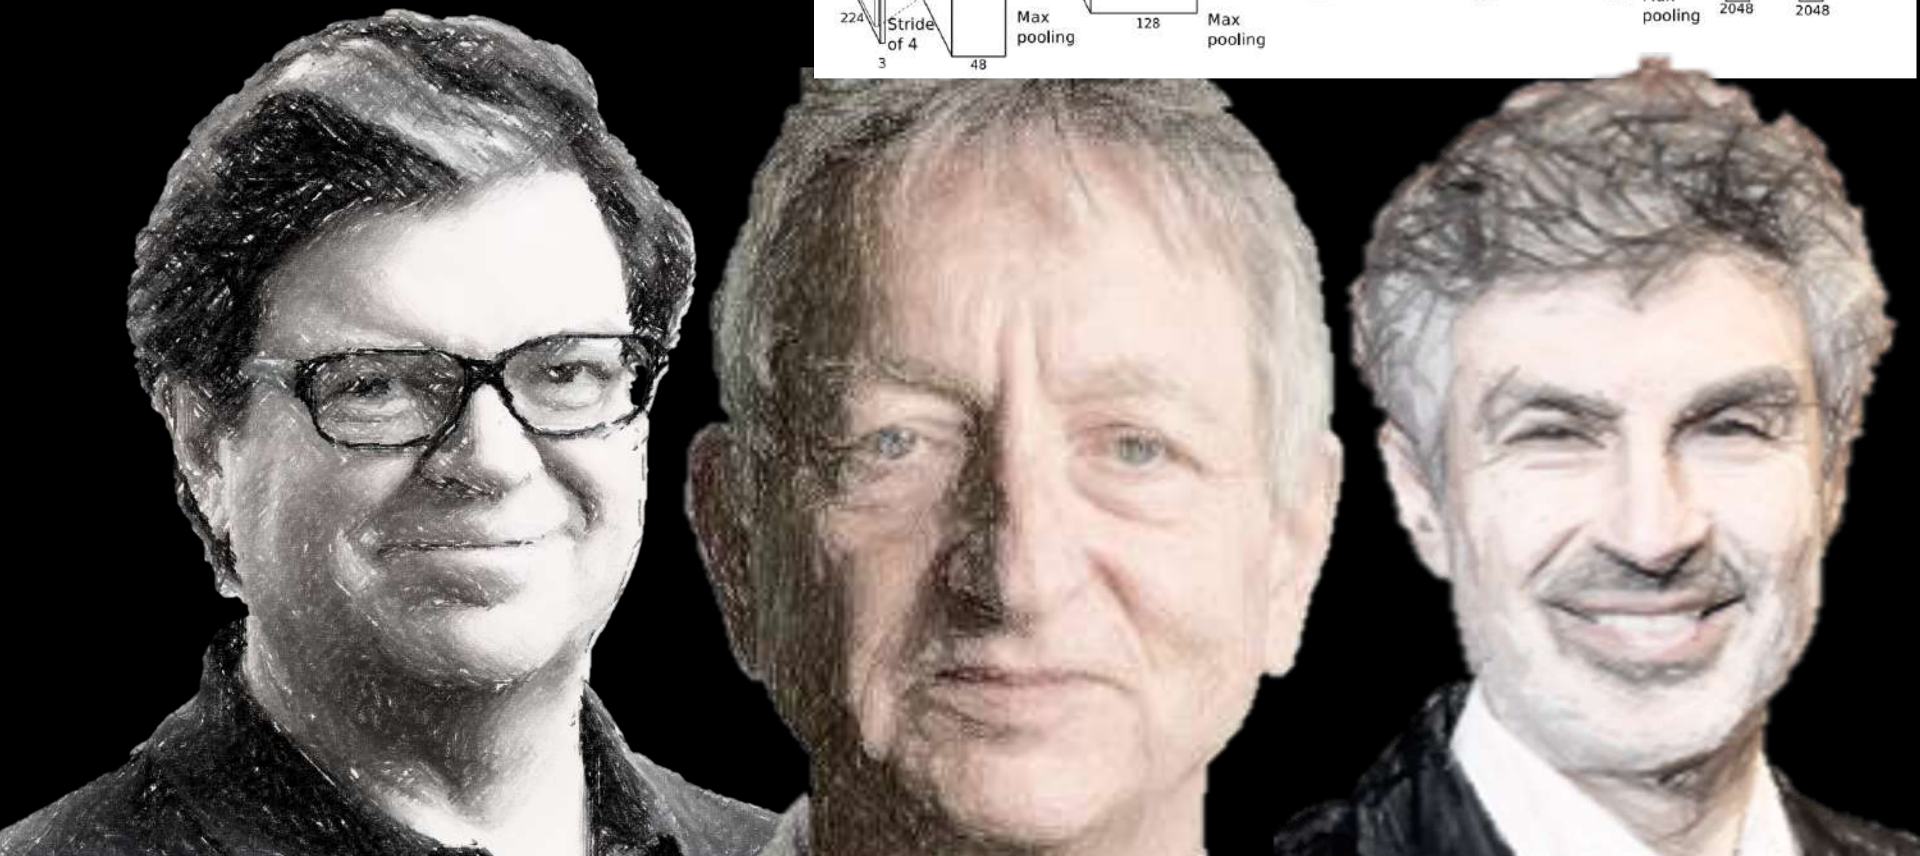

Yann LeCunn Geoff Hinton Yoshua Bengio

ISRAEL SCIENCE FOUNDATION

Ministry of Science, Technology and Space

Israel

SCHMIDT FUTURES

The Lokey Center

Artificial Intelligence Algorithms to Assess Hormonal Status From Tissue Microarrays in Patients With Breast Cancer

Gil Shamai Yoav Binenbaum Ron Slossberg Irit Duek Ziv Gil Ron Kimmel

Technion Israel Institute of Technology

Tel-Aviv Sourasky Medical Center

RAMBAM

Geometric Image Processing lab

JAMA Network Open

Original Investigation | Oncology

Artificial Intelligence Algorithms to Assess Hormonal Status From Tissue Microarrays in Patients With Breast Cancer

Gil Shamai, MSc; Yoav Binenbaum, MD, PhD; Ron Slossberg, MSc; Irit Duek, MD; Ziv Gil, MD, PhD; Ron Kimmel, DSc

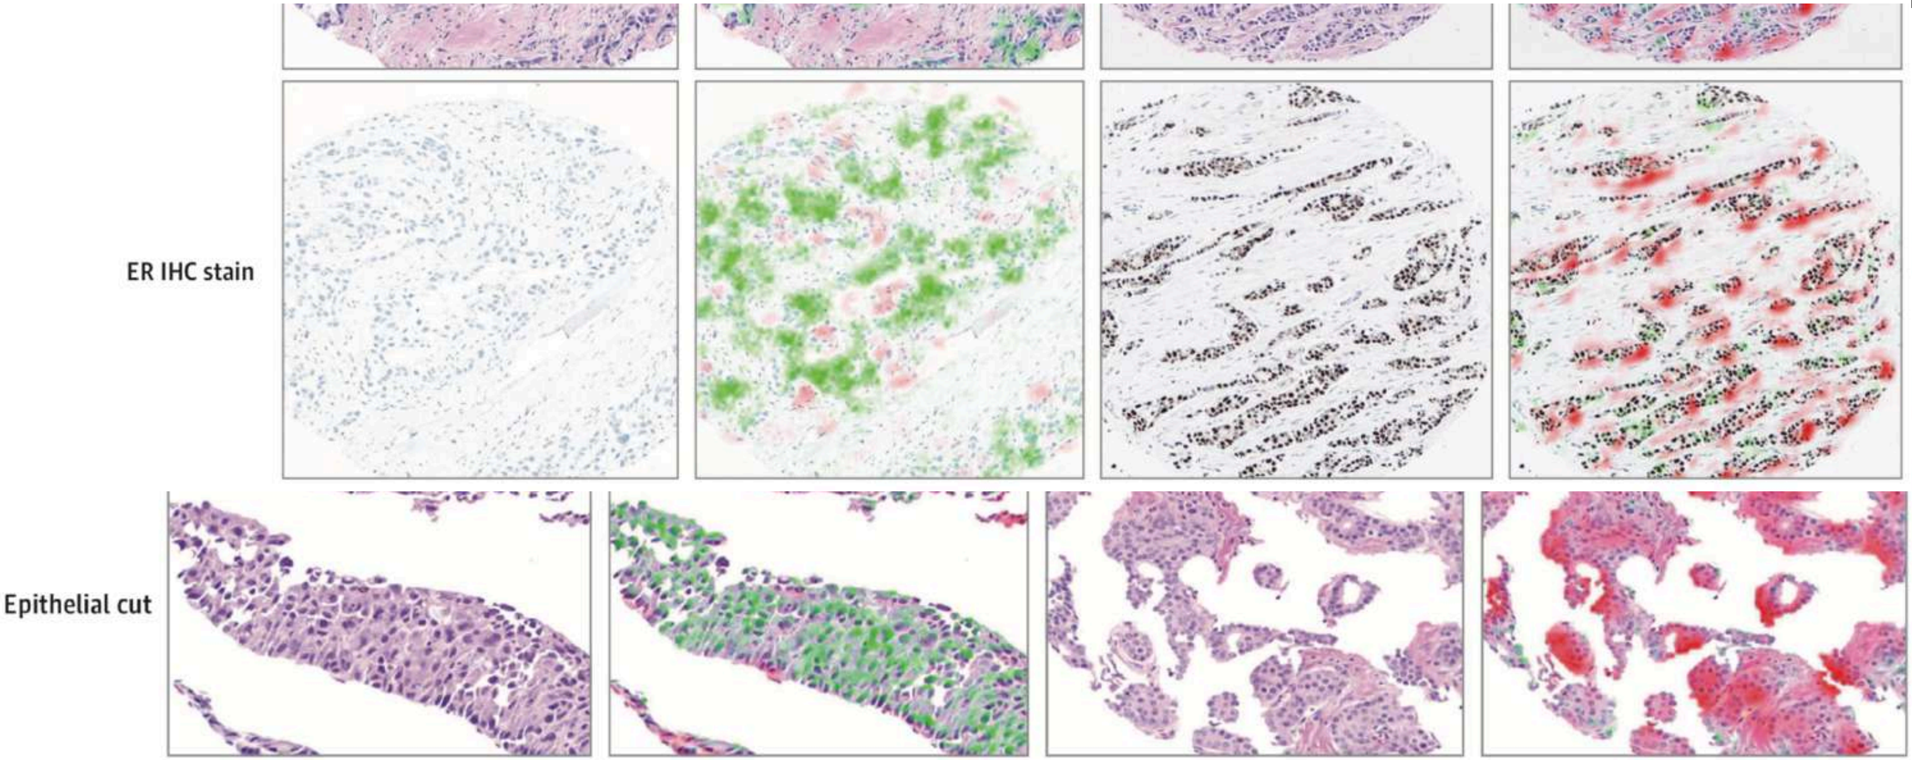

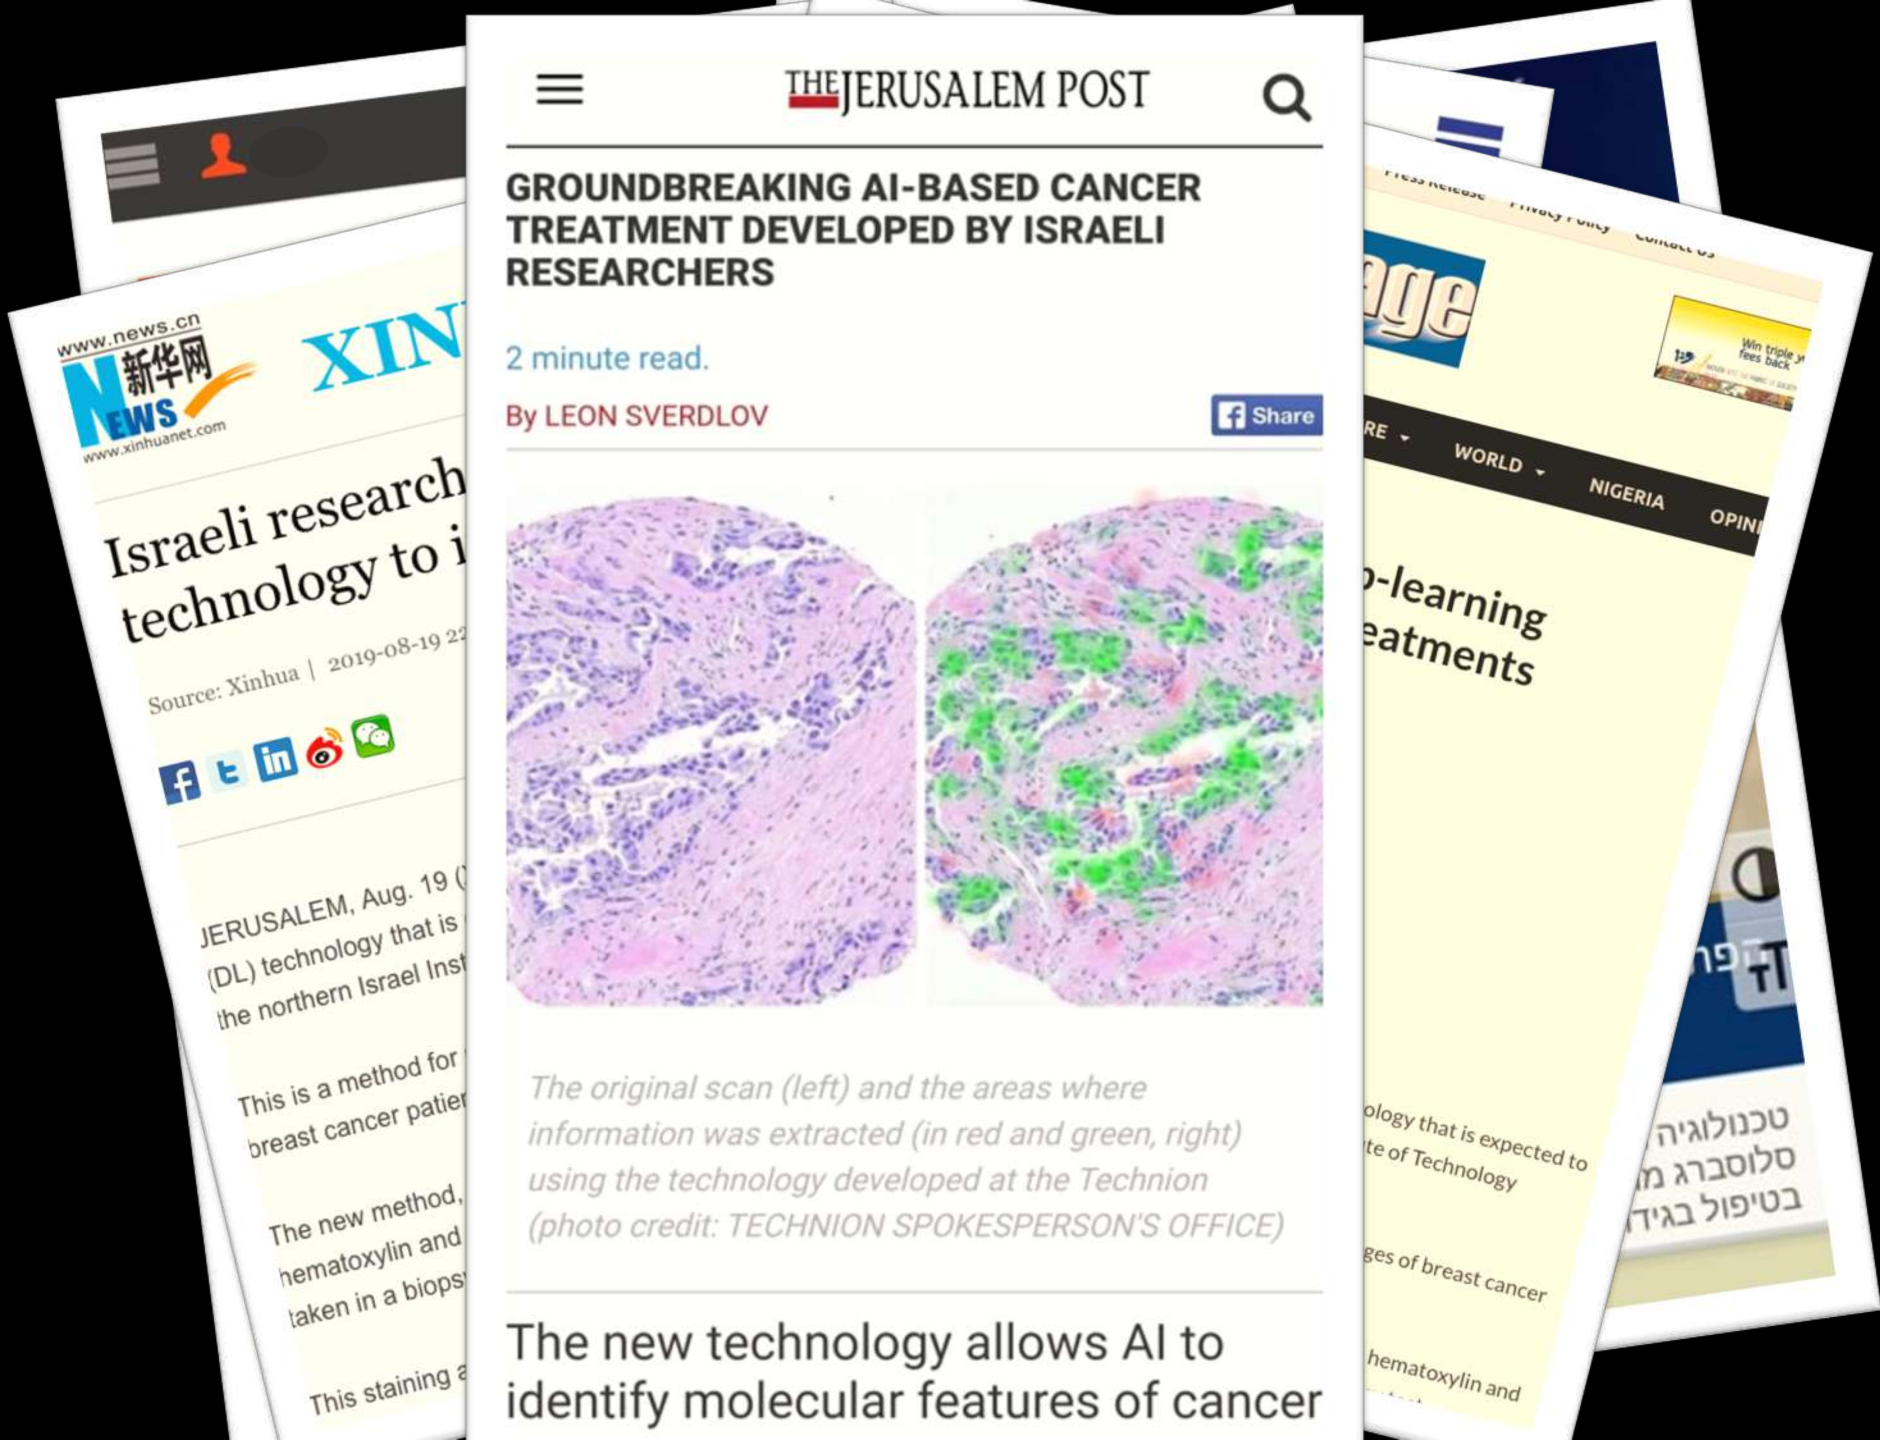

THE JERUSALEM POST

GROUNDBREAKING AI-BASED CANCER TREATMENT DEVELOPED BY ISRAELI RESEARCHERS

2 minute read.

By LEON SVERDLOV

Share

The original scan (left) and the areas where information was extracted (in red and green, right) using the technology developed at the Technion (photo credit: TECHNION SPOKESPERSON'S OFFICE)

The new technology allows AI to identify molecular features of cancer

Coded light 3D scanner

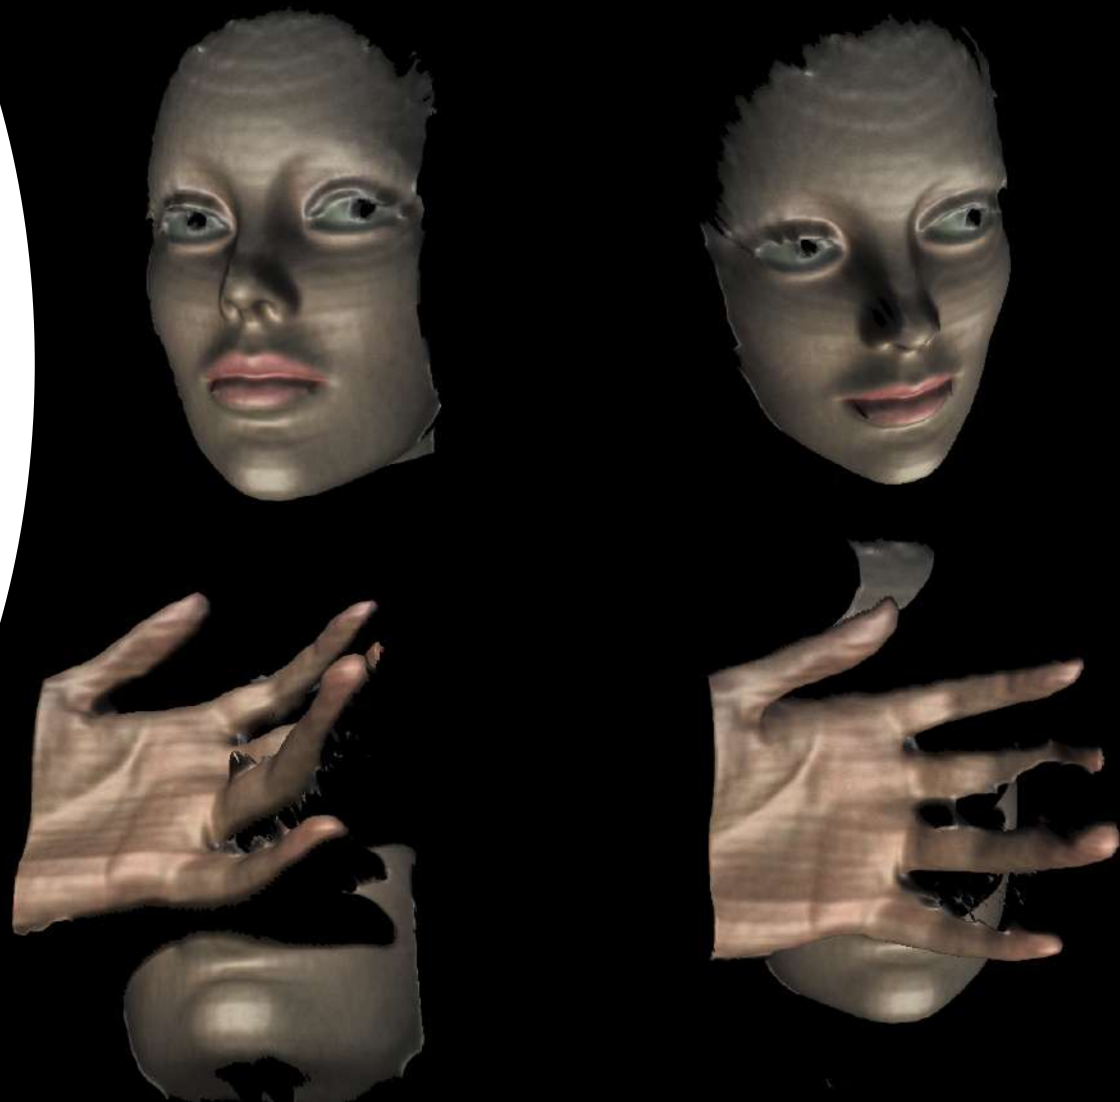

INVISION
Ultimate 3D Experience

intel

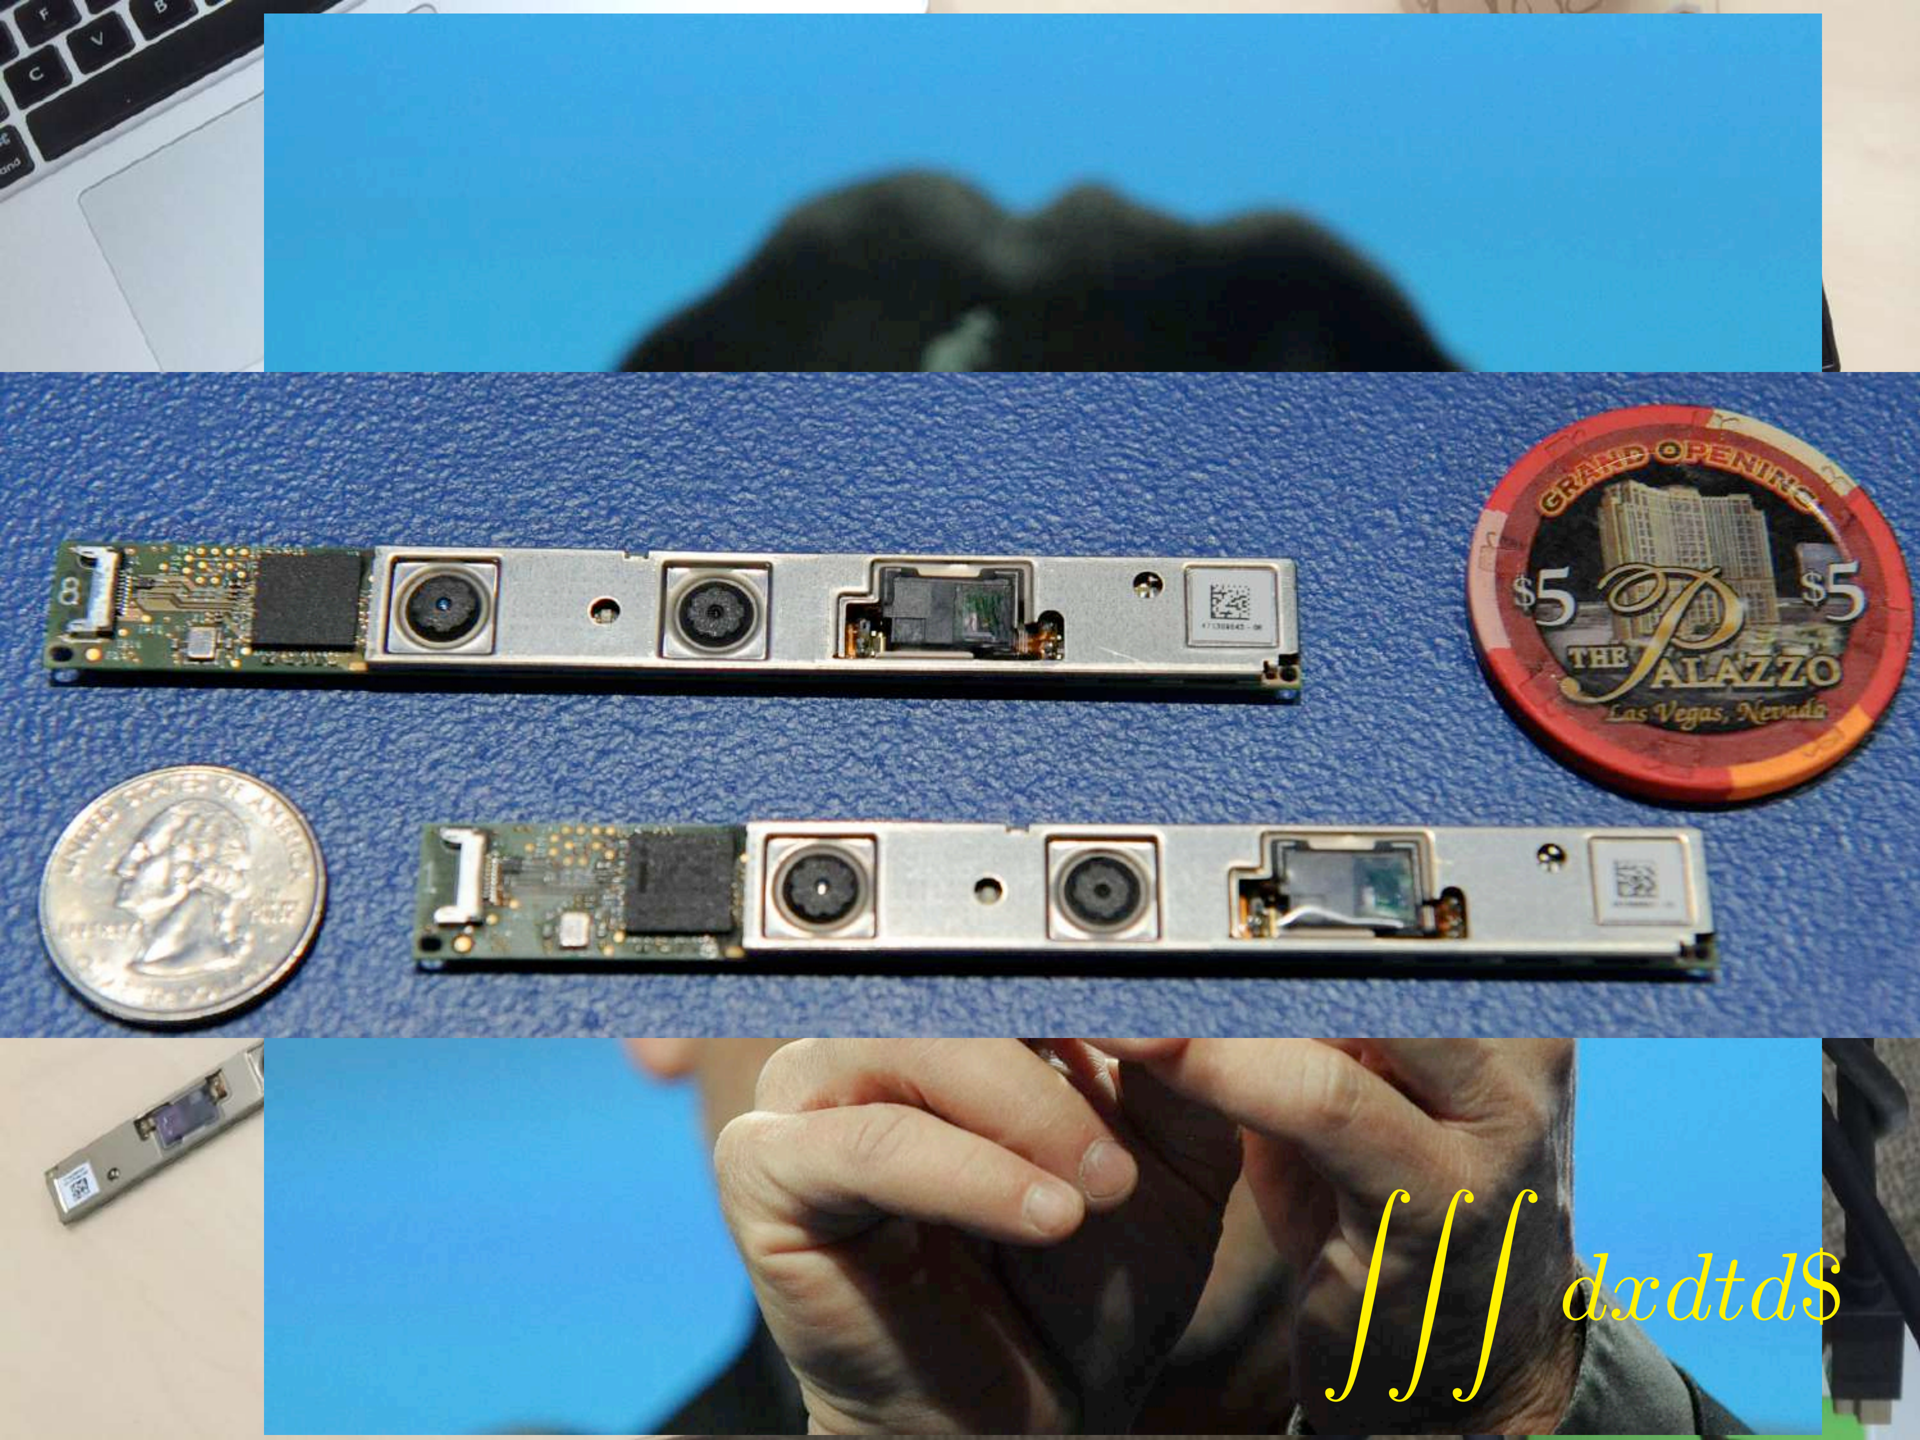
GRAND OPENING
$5
$5
THE PALAZZO
Las Vegas, Nevada
$\iiint dx\,dt\,d\$$

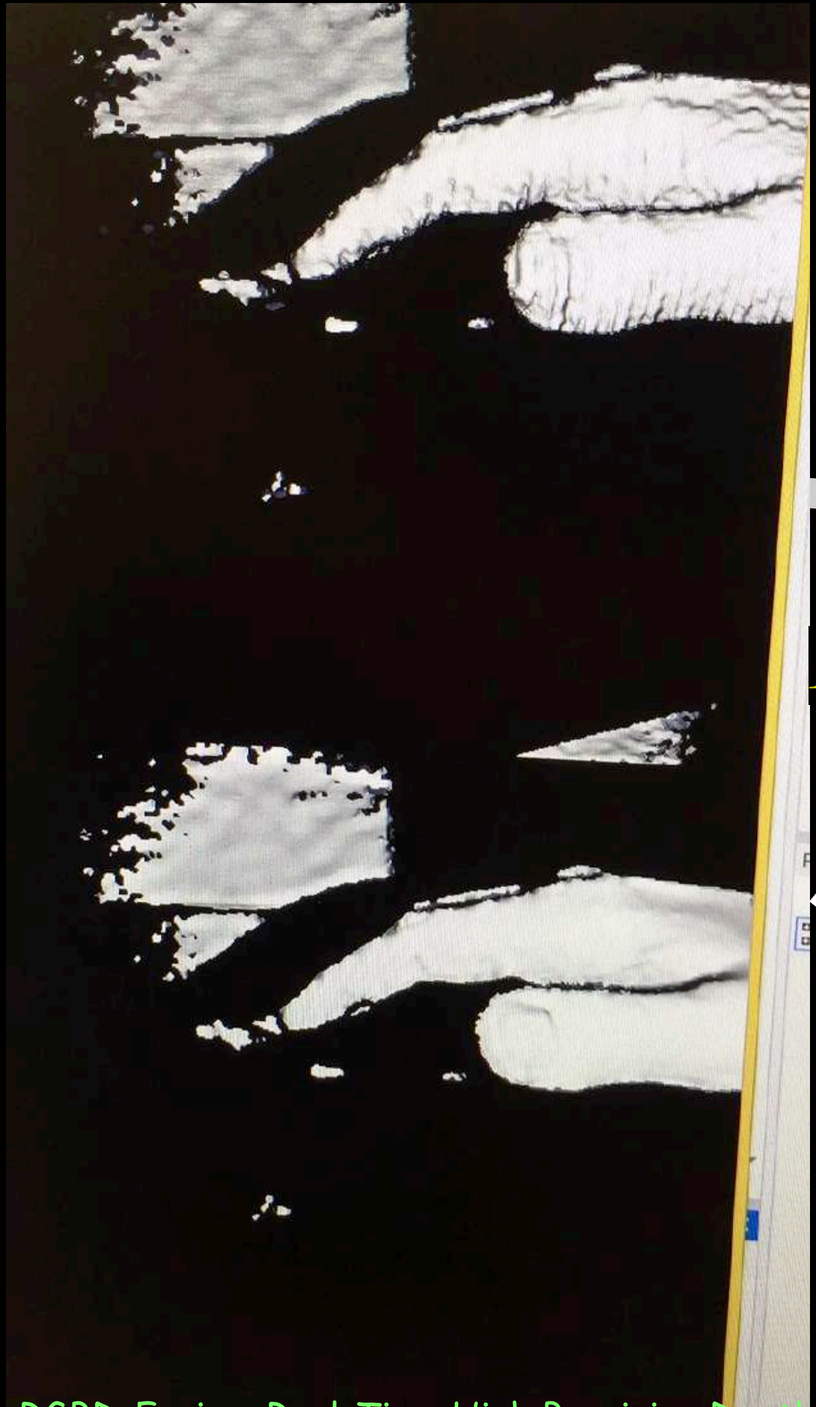

RealSense
refined
via SfS

$$I = \rho\langle\vec{l},\vec{N}\rangle$$

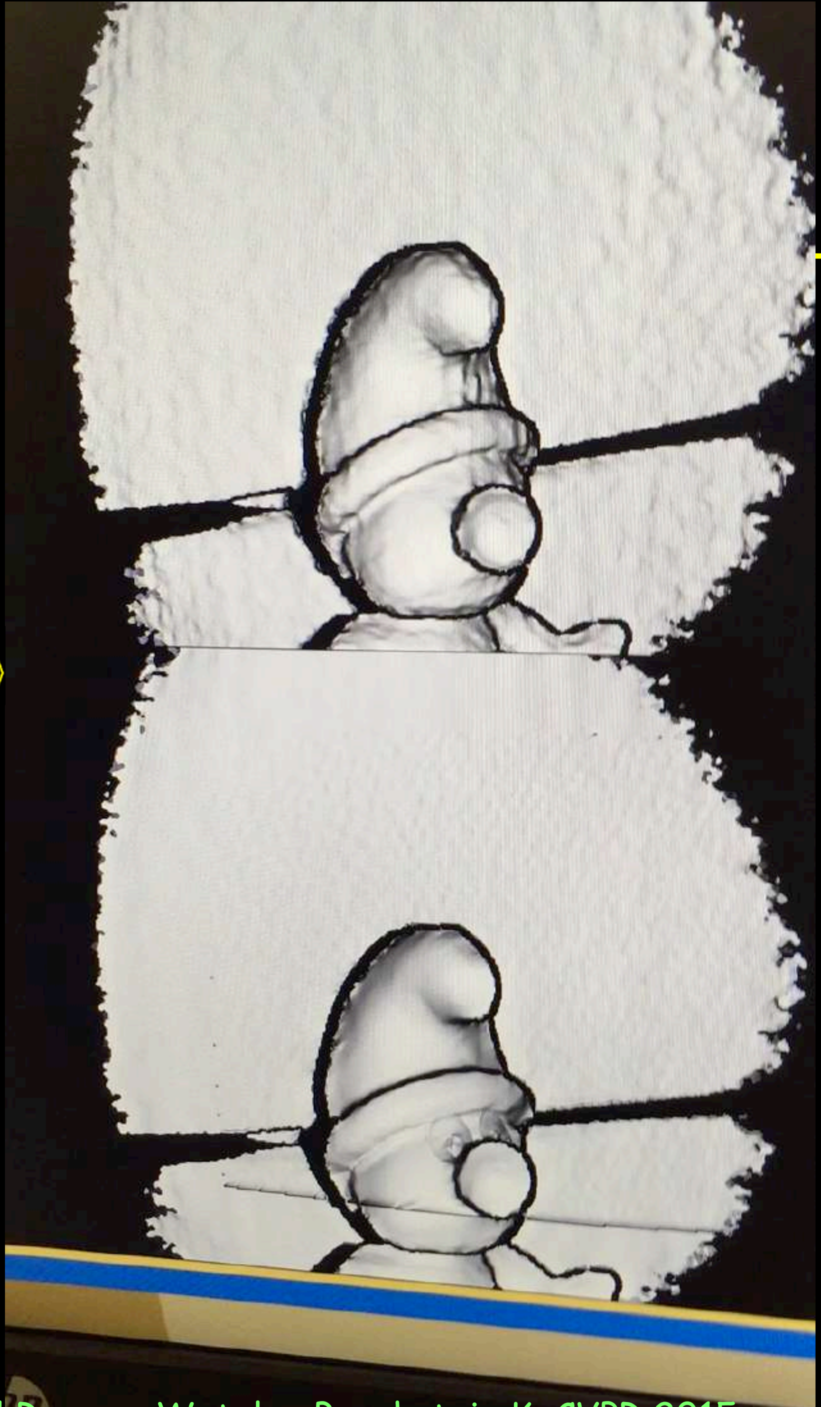

RGBD-Fusion: Real-Time High Precision Depth Recovery Or-El Rosman Wetzler Bruckstein K. CVPR 2015

Designing facial masks

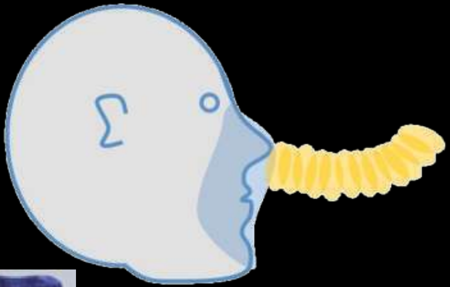

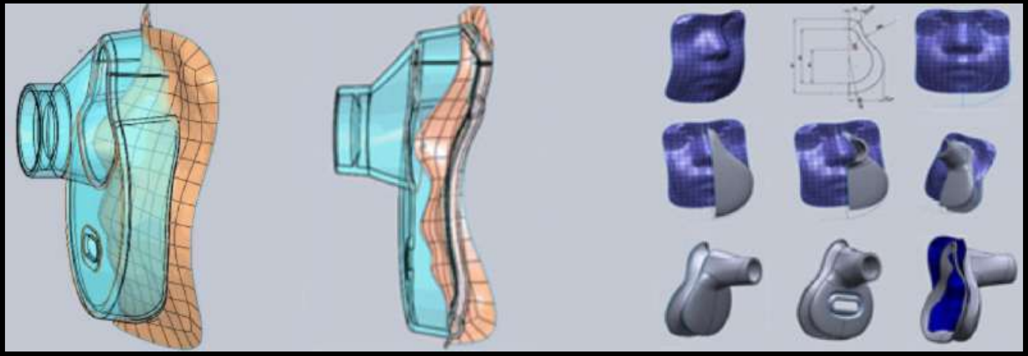

Amirav, Halamish, Raviv, K., Honen, Zvirin, et al.

Amirav, Halamish, Raviv, K., Honen, Zvirin, et al.


René Descartes

How could we use algebra to describe geometry?

Differential Signatures

- Euclidean invariant signature $\{s, \kappa(s)\}$

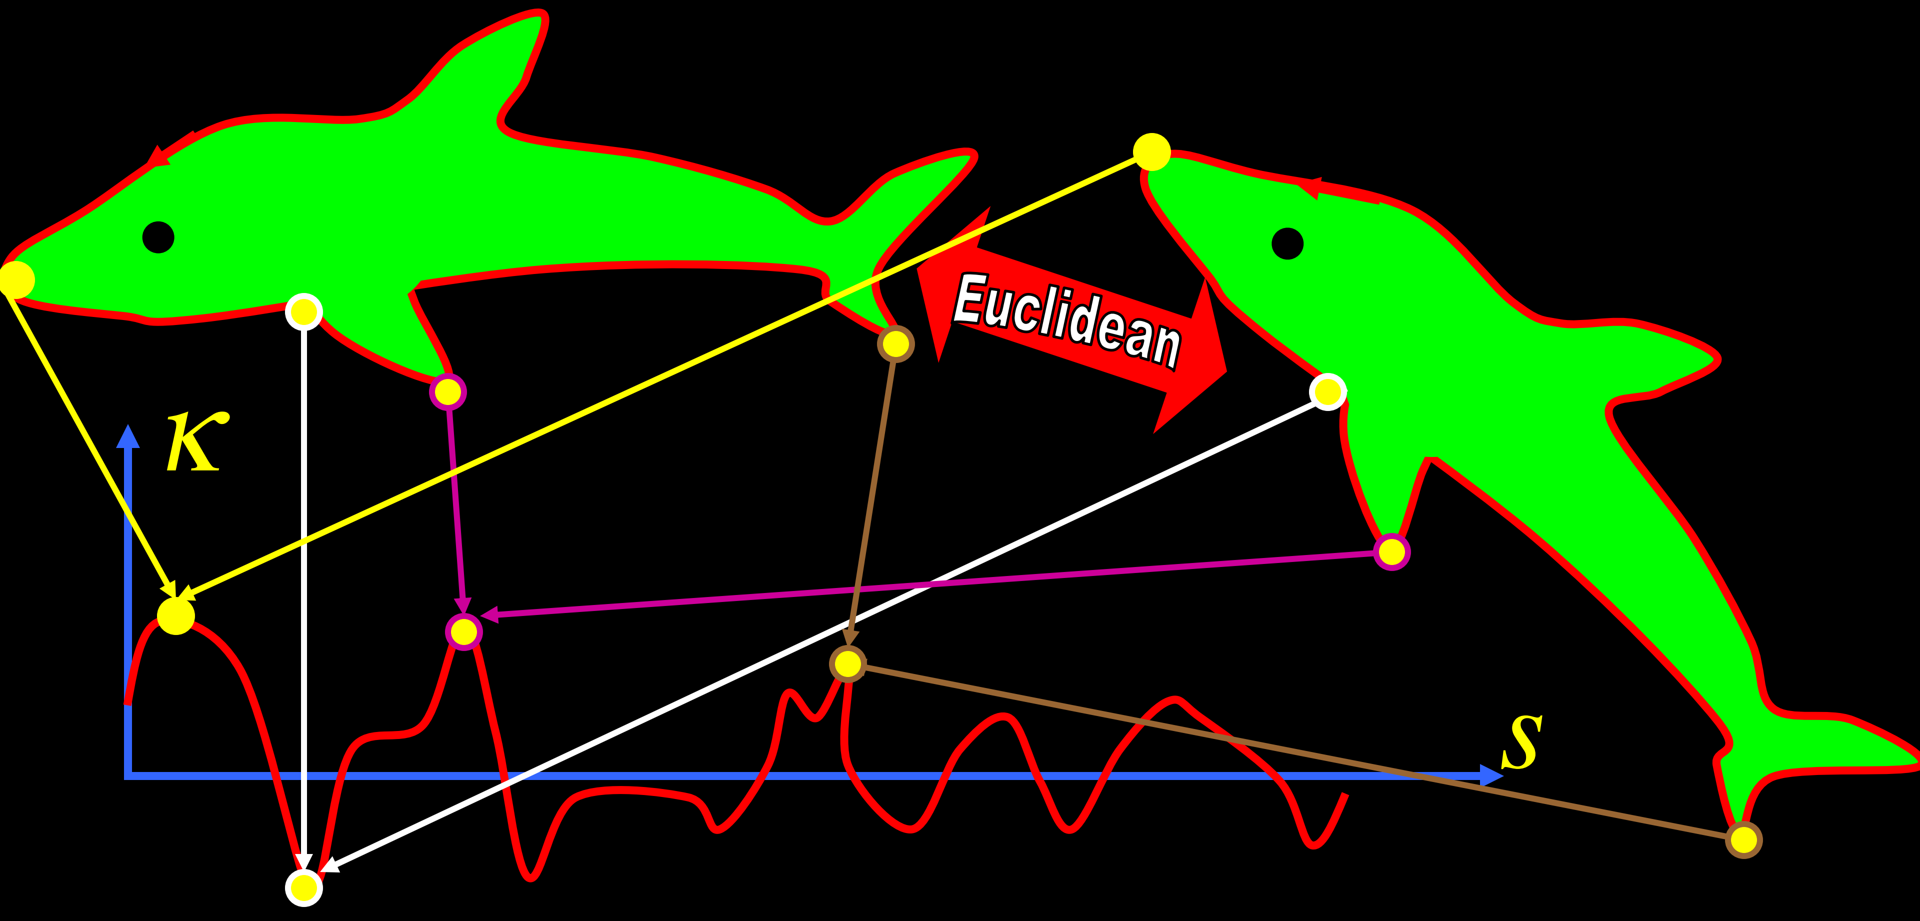

Differential Signatures

- Euclidean invariant signature $\{s, \kappa(s)\}$

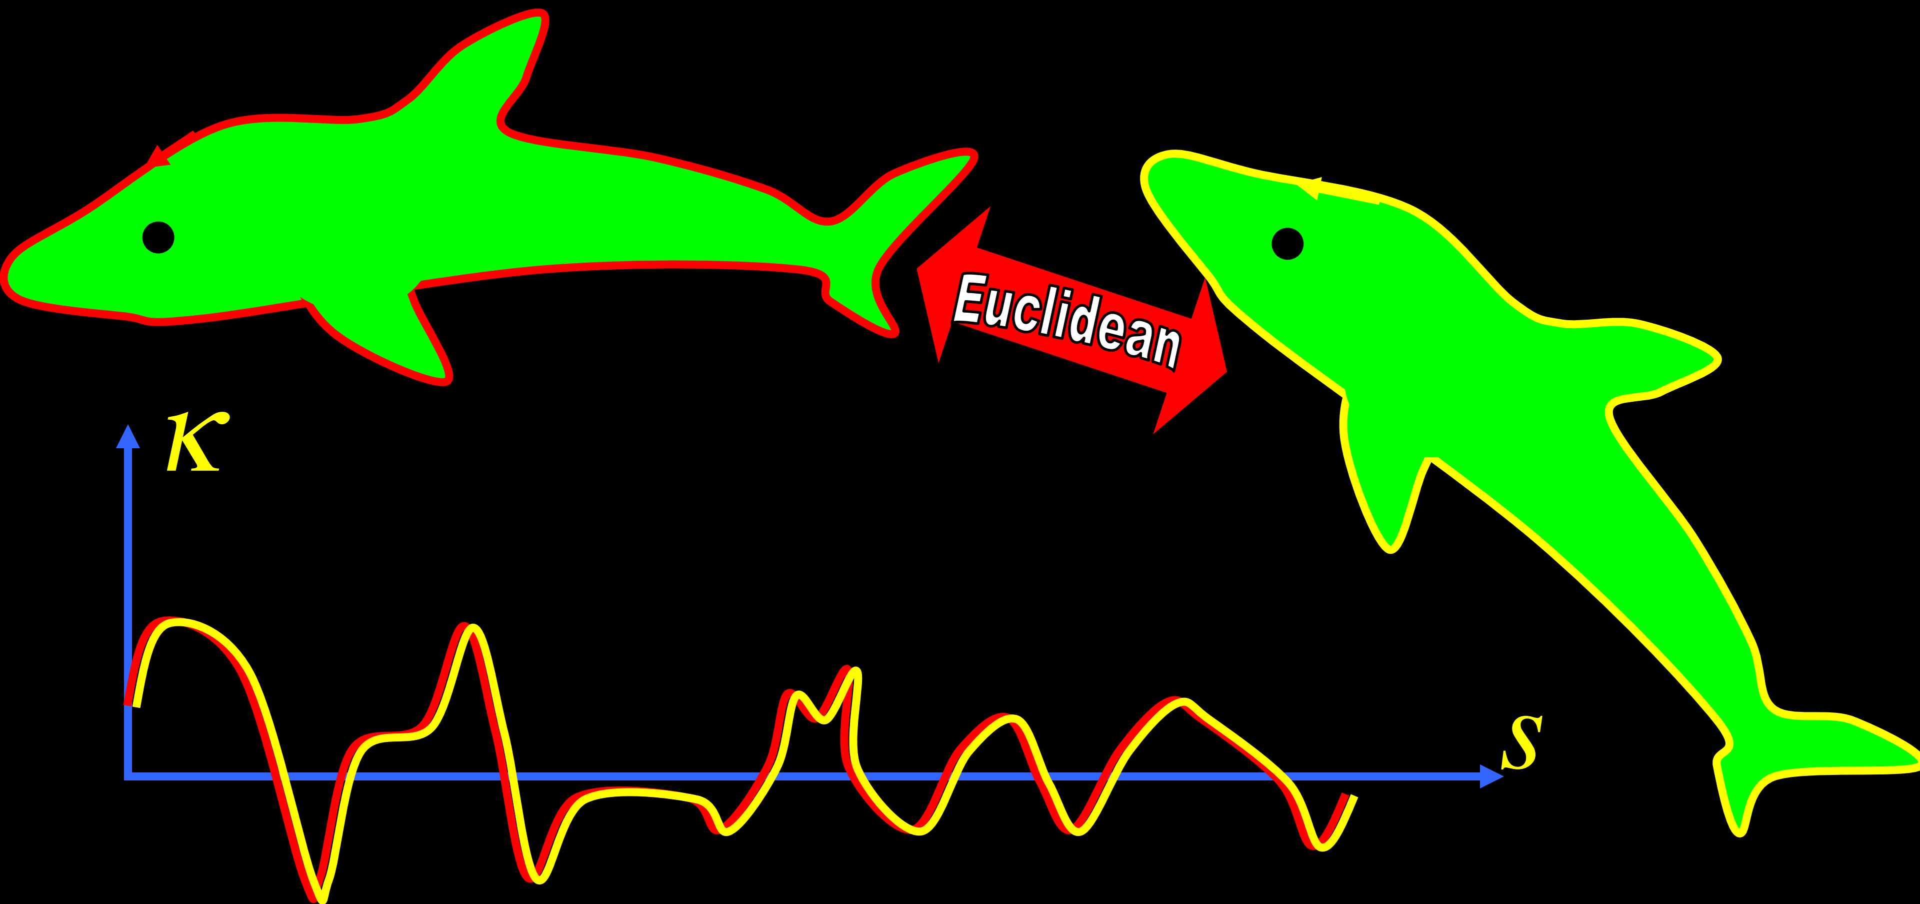

Learning invariants

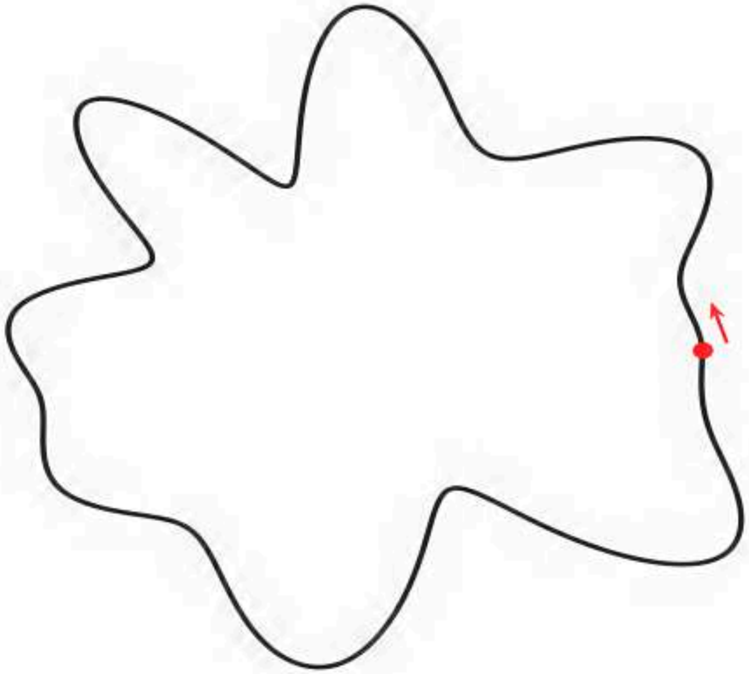

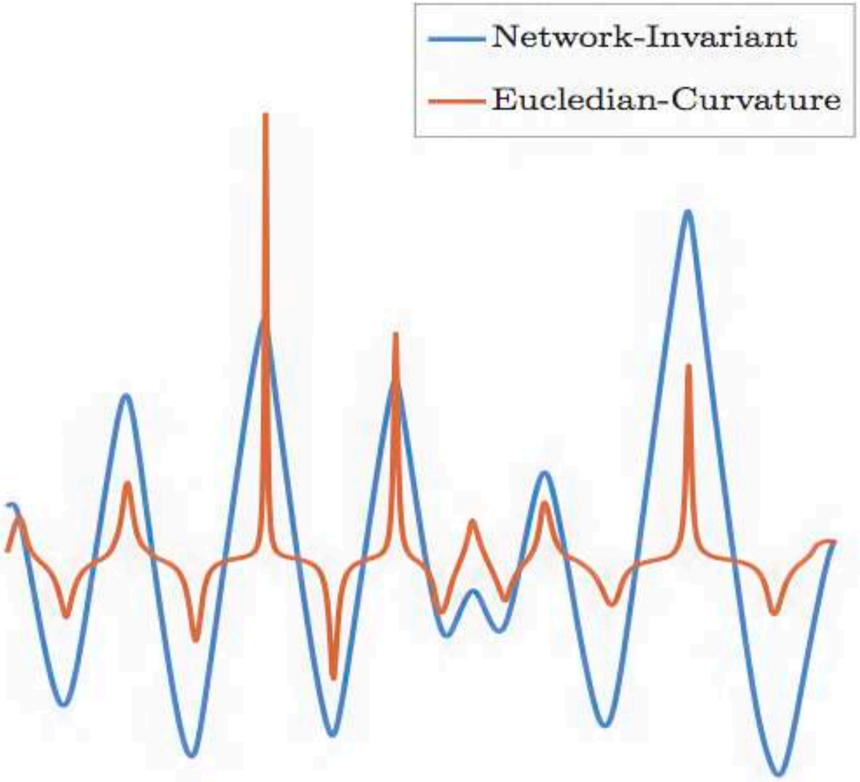

Pai Wetzler K. ICLR 2017

Momenet

Algebraic invariants	3D moment as invariants	High order NN	Deep NN	PointNet	Momenet
1903	1980	1987	2010	2017	2019

Geometric moments

Given a set of points $P \subset \mathbb{R}^3$ and $p_i = (x_i, y_i, z_i)^T \in P$

1st moments $\bar{p} = (\bar{x}, \bar{y}, \bar{z})^T = \frac{1}{n}\sum_{i=1}^{n} p_i$

2nd moments $\frac{1}{n}\sum_{i=1}^{n} p_i p_i^T$

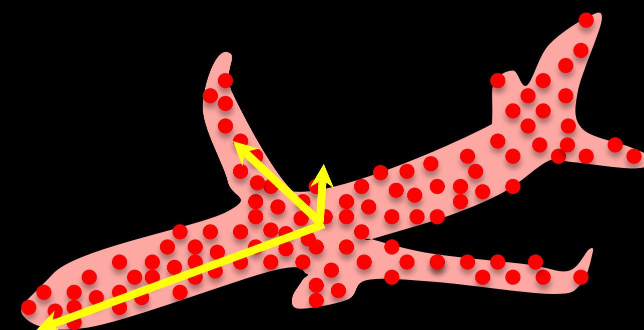

Young. The algebra of invariants, 1903

Hall. Three-dimensional moment invariants. PAMI, 1980

Maxwell, Learning, invariance, and generalization in high-order neural networks. Applied optics, 1987

Su, Mo, Guibas. Pointnet: Deep learning on point sets. CVPR'17

Joseph-Rivlin, Zvirin, K., GMDL workshop, ICCV'19

Moments

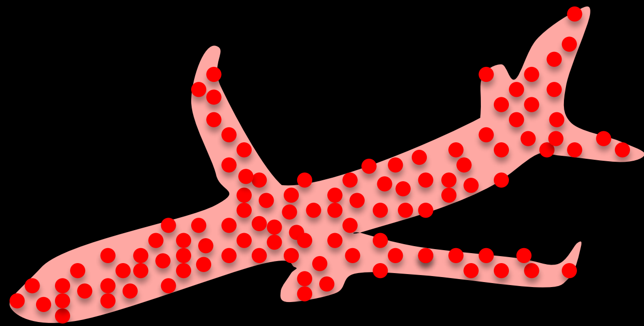

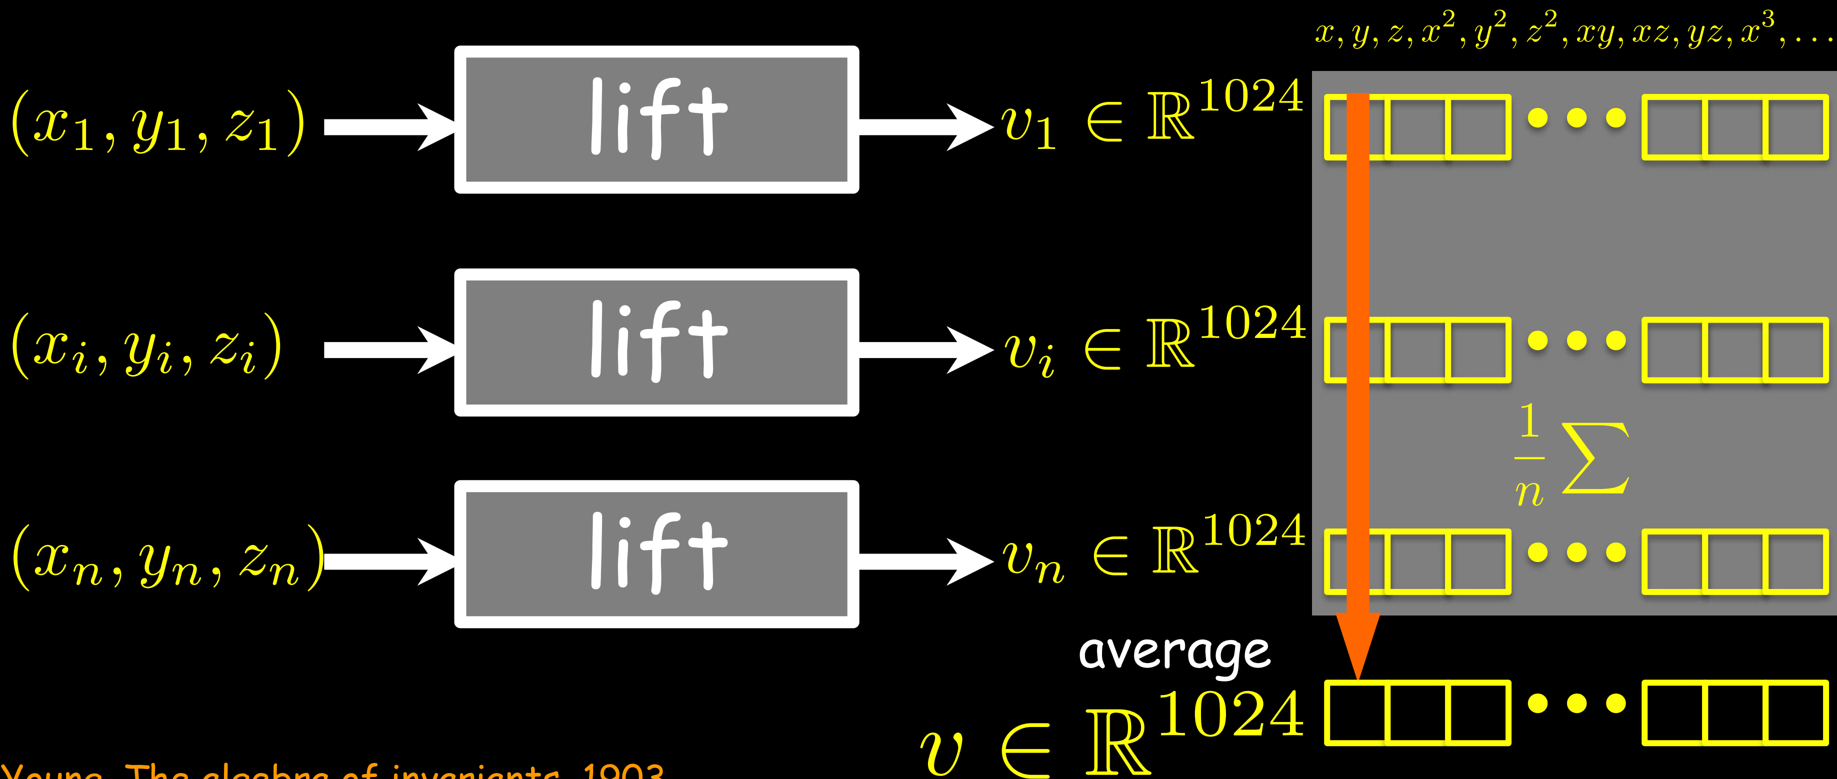

Young. The algebra of invariants, 1903
Hall. Three-dimensional moment invariants. PAMI, 1980
Maxwell, Learning, invariance, and generalization in high-order neural networks. Applied optics, 1987
Su, Mo, Guibas. Pointnet: Deep learning on point sets. CVPR'17
Joseph-Rivlin, Zvirin, K., GMDL workshop, ICCV'19

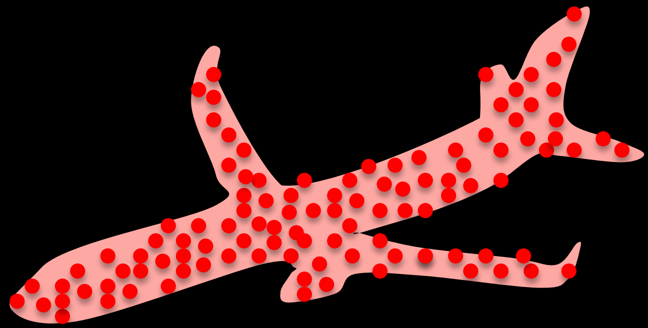

PointNet

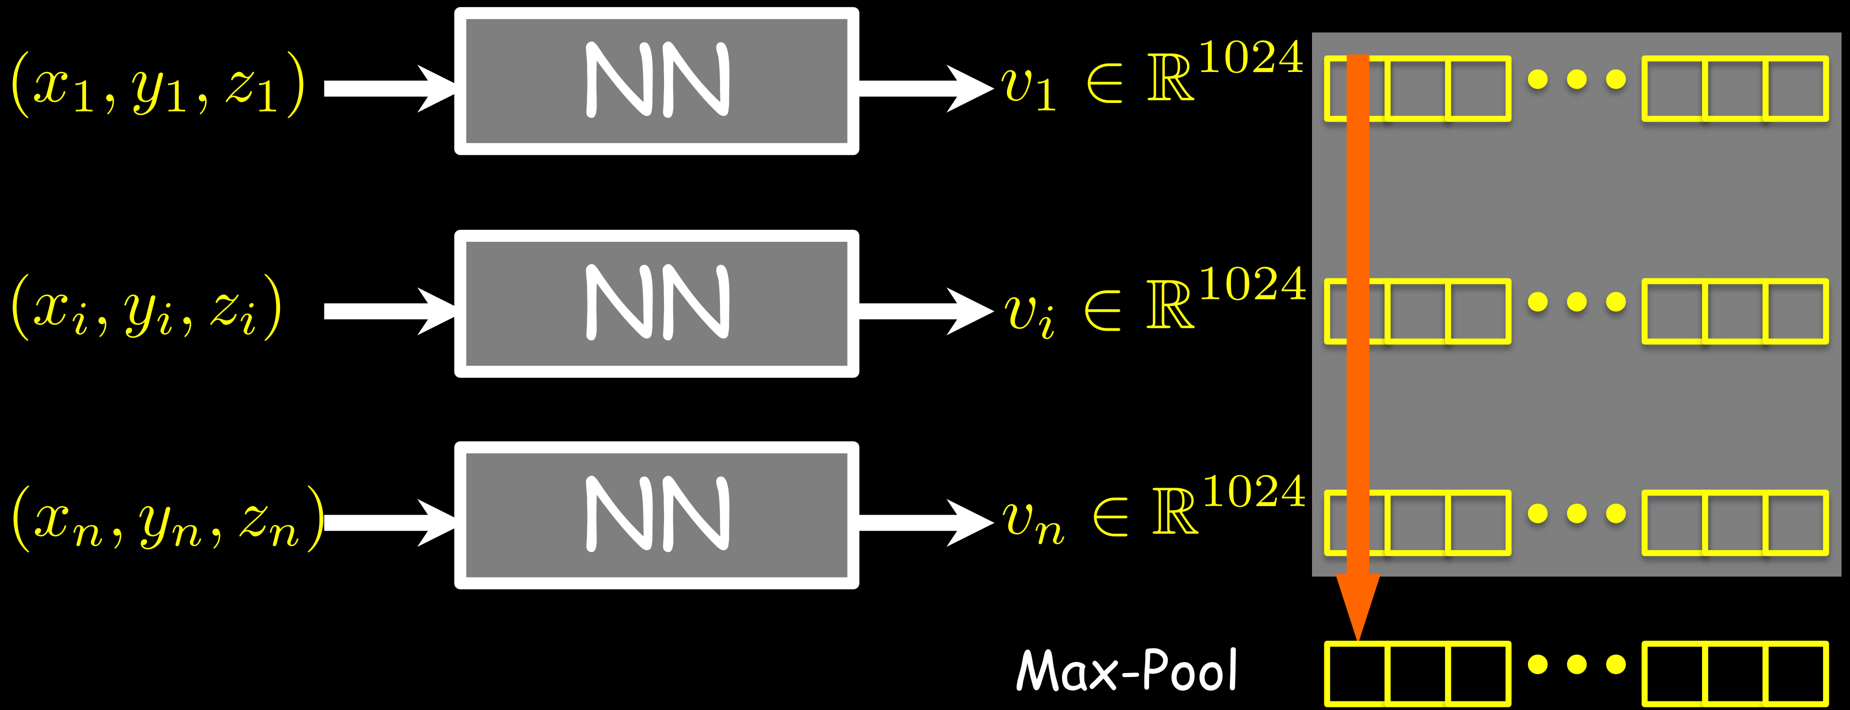

Su, Mo, Guibas. CVPR'17

Momenet

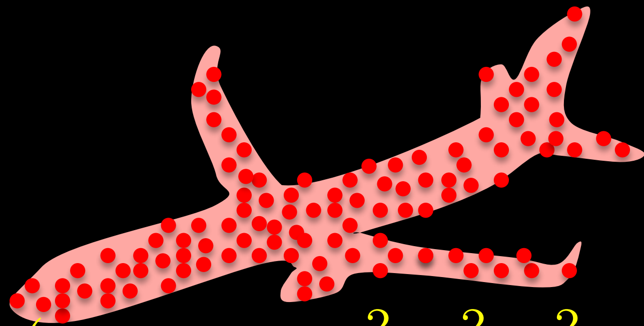

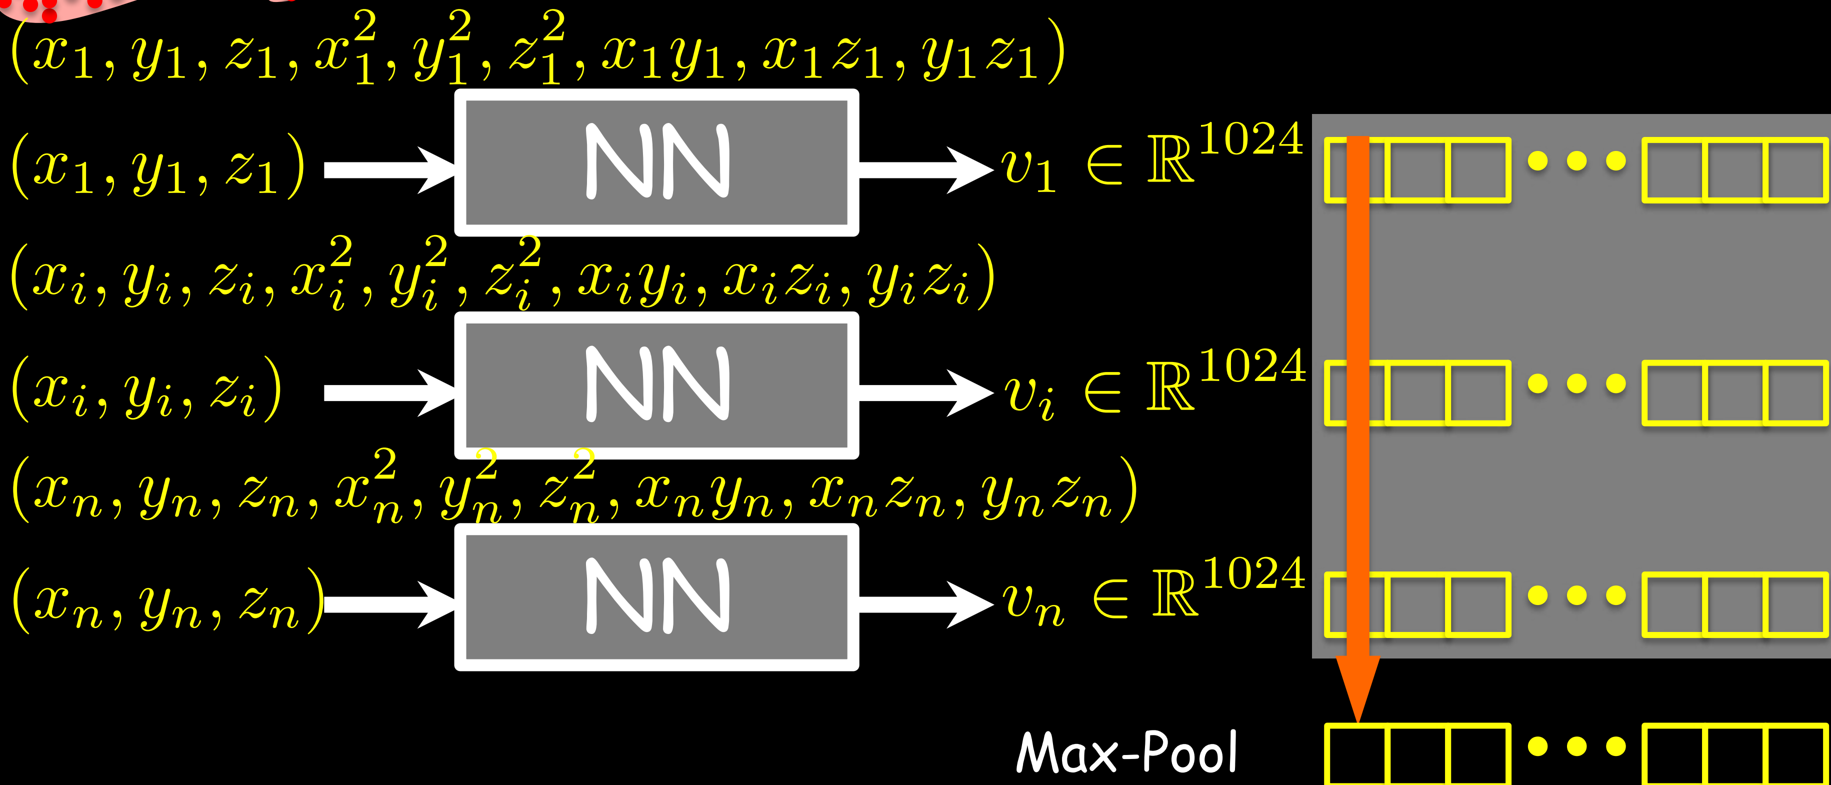

Joseph-Rivlin, Zvirin, K., Flavor the moments, GMDL workshop, ICCV'19

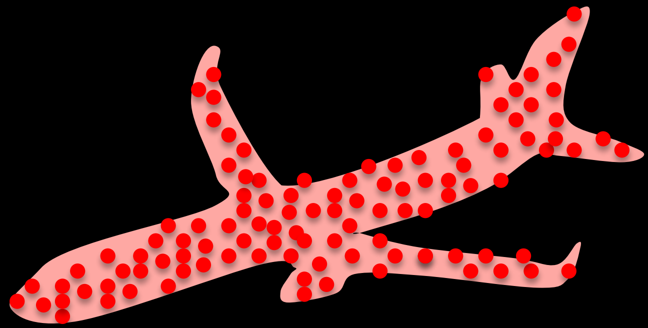

Momenet

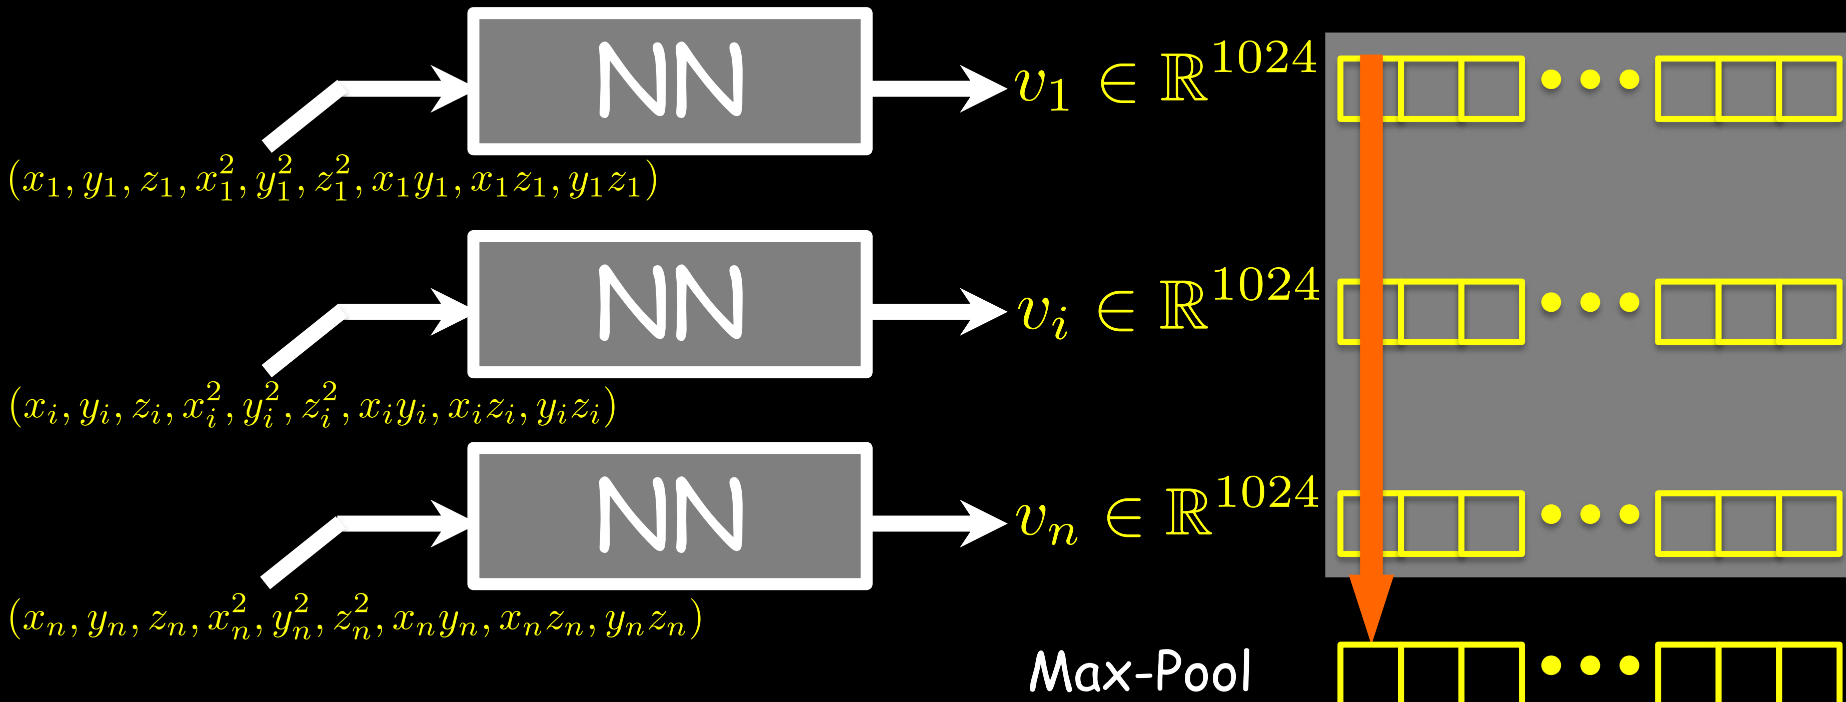

Joseph-Rivlin, Zvirin, K., Flavor the moments, GMDL workshop, ICCV'19

Results on ModelNet40

	Memory (MB)	Inference Time (msec)	Overall Accuracy (%)
PointNet	40	5.6	89.2
PointNet(baseline)	20	5.1	87.9
Momen^{e}t	**20**	**5.1**	**89.6**
PointNet++	12	10.4	90.7
DGCNN	21	17.3	92.2
PCNN	17	54.1	92.3
Momen^{e}t (+kNN)	**21**	**9.6**	**92.4**

Wu, et al.. 3D shapenets. CVPR2015.
Wang et al. . Dynamic graph CNN. Arxiv 2018.
Atzmon, Maron, & Lipman. Point convolutional NN. ACM T Graph. 2018
Joseph-Rivlin, Zvirin, K., Flavor the moments, GMDL workshop, ICCV'19

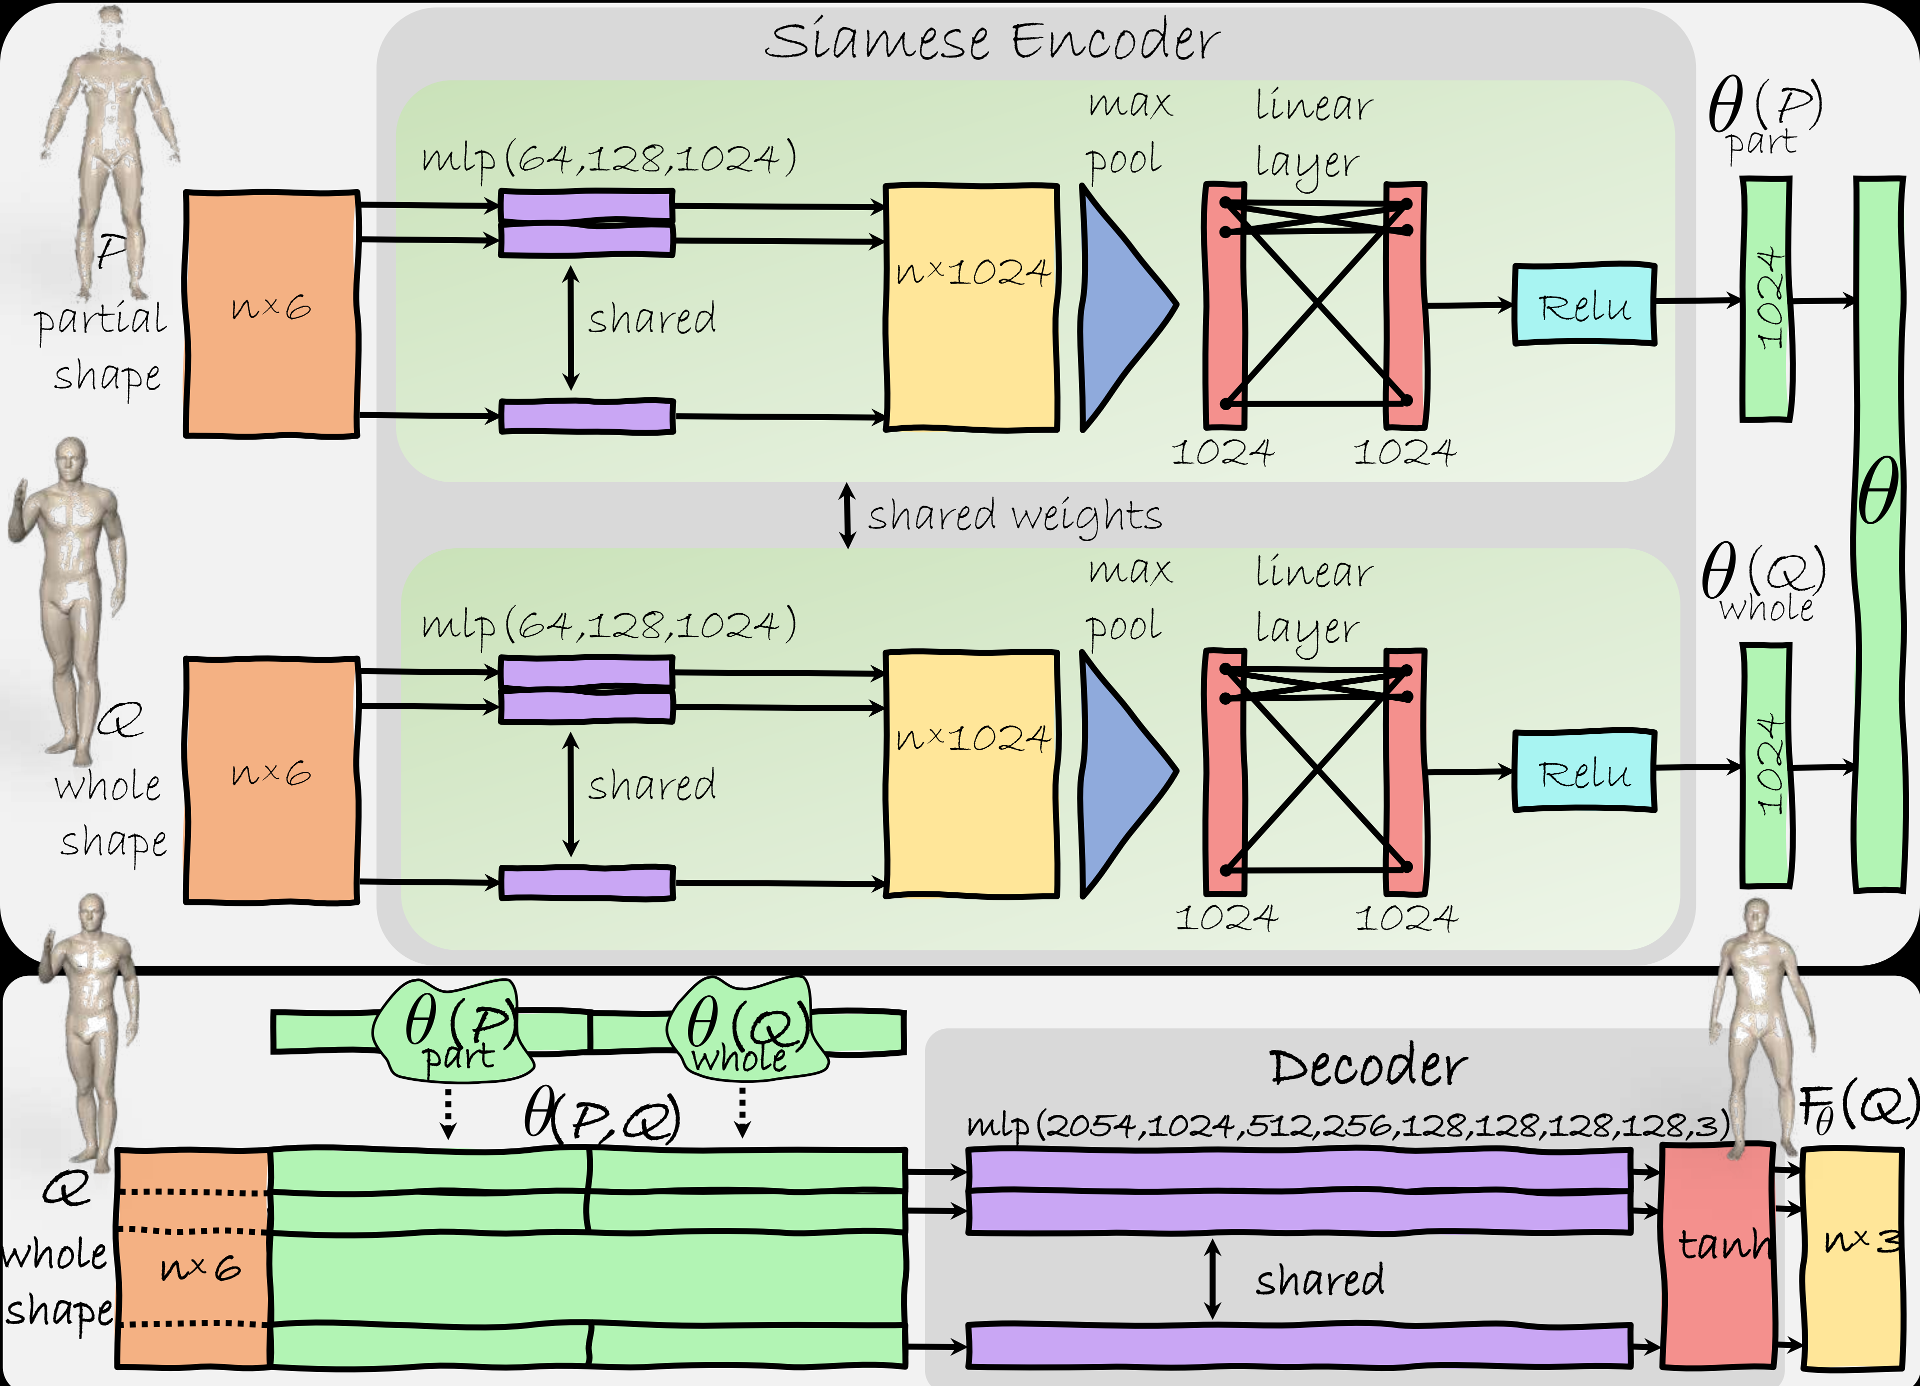

Halimi et al. ECCV 2020

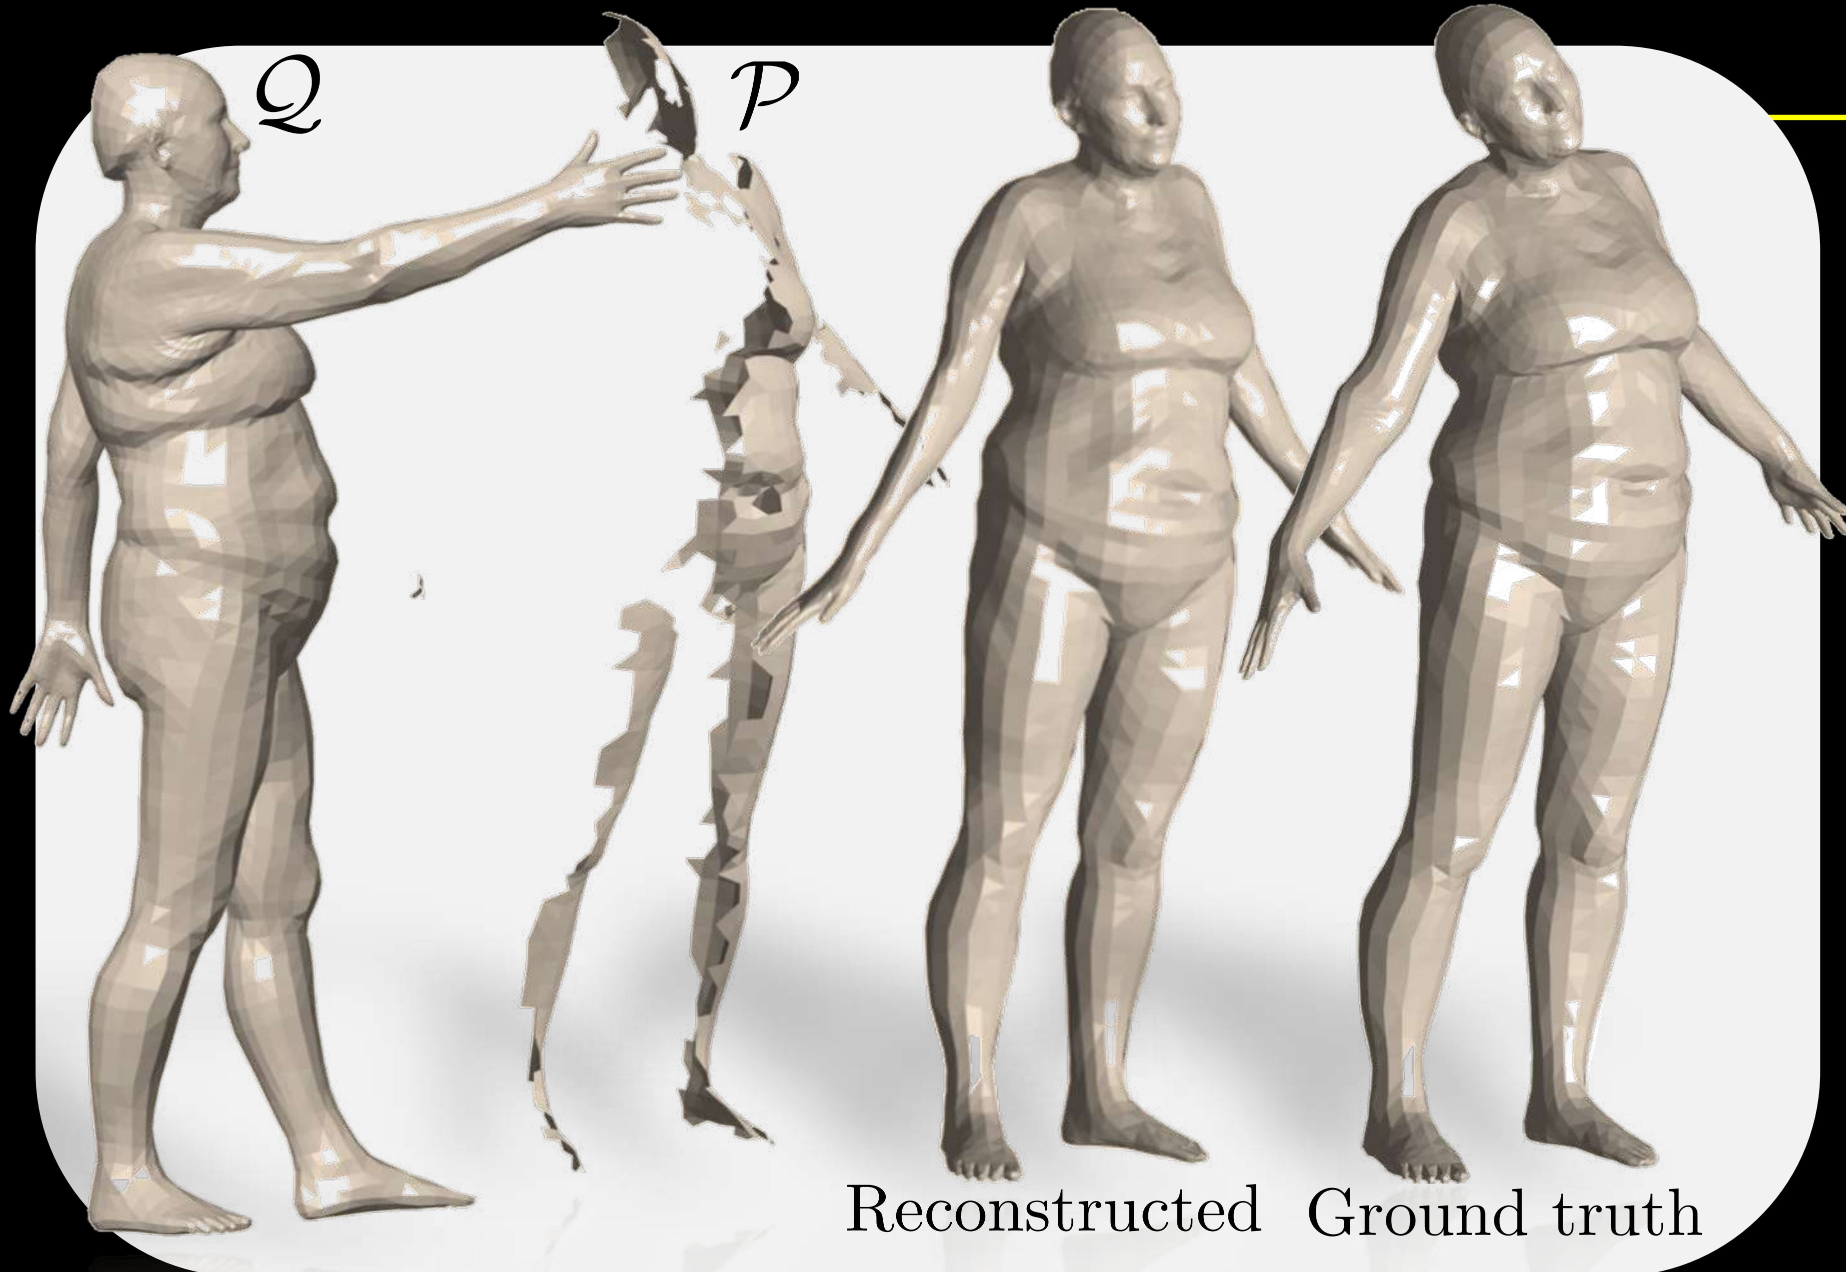

$\mathcal{Q}$
$\mathcal{P}$
Reconstructed
Ground truth

Shapes as metric spaces

$d_{GH}($  , 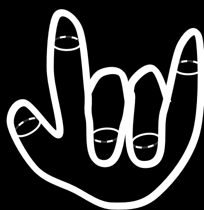 $) < d_{GH}($ 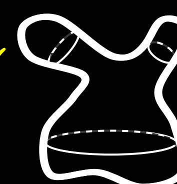 ,  $)$

$MDS, GMDS, SGMDS, GDD, PCA, RPCA,$
$F\text{-}Maps, FM\text{-}Net, GMDS\text{-}Net, SF\text{-}Maps$

Deep Eikonal solvers

$$|\nabla u(x)| = 1, \quad x \in \Omega \setminus \Gamma$$

$$u(x) = 0, \quad x \in \Gamma$$

Local numerical approximations using neural networks

- Easy to extend the support
- Approximation is learned from analytically known solutions

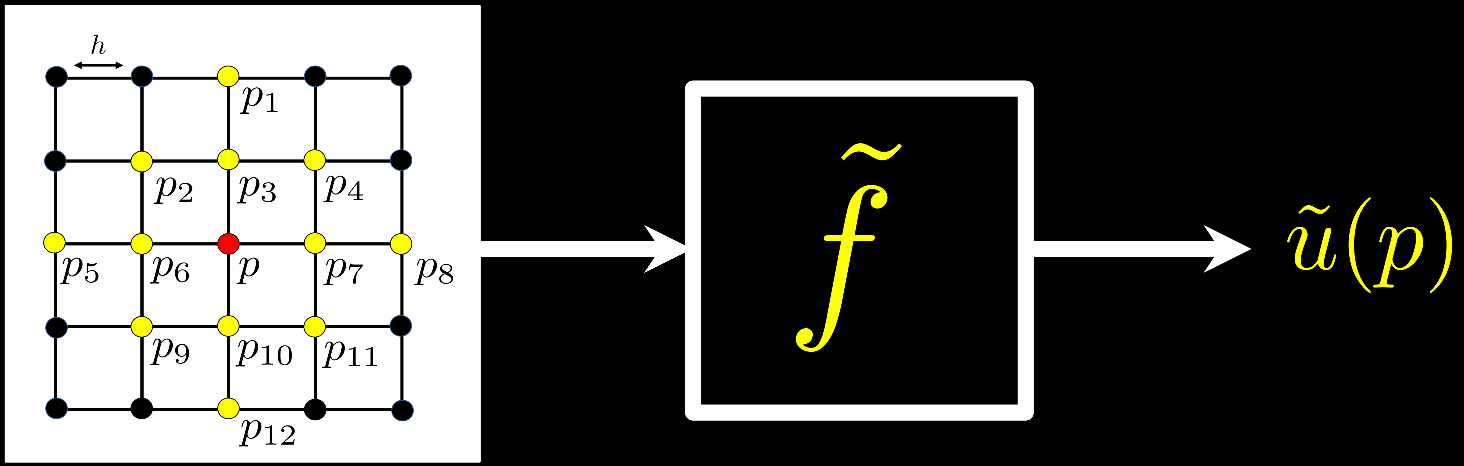

Lichtenstein, Pai, K., SSVM 2019 - Best Student Paper Award

On the evolution of accuracy/complexity

- Fast Marching (quasi-linear complexity) $\sim O(n)$

Regular grids

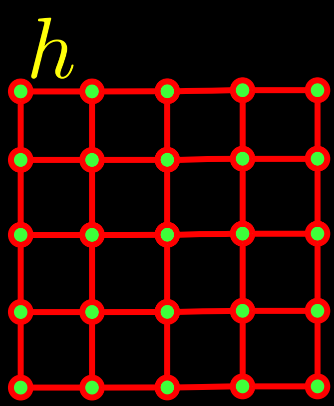

1995 $O(h)$

↓

1999 $O(h^2)$

↓

2011 $O(h^3)$

Triangulated domains

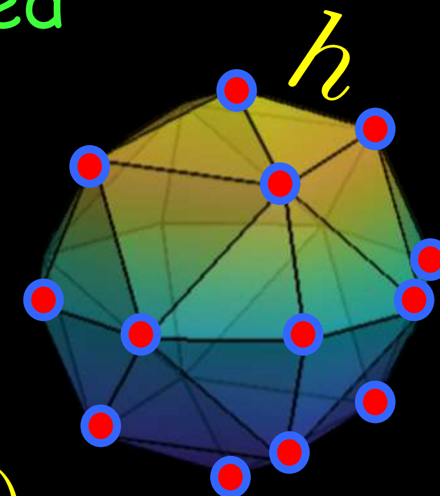

1998 $O(h)$

2005 $O(h^2)$ ← $O(n^2)$

Surazhsky et al. implementation of MMP's `exact geodesic' 1987 (~2 decades!)

Results for Cartesian grids

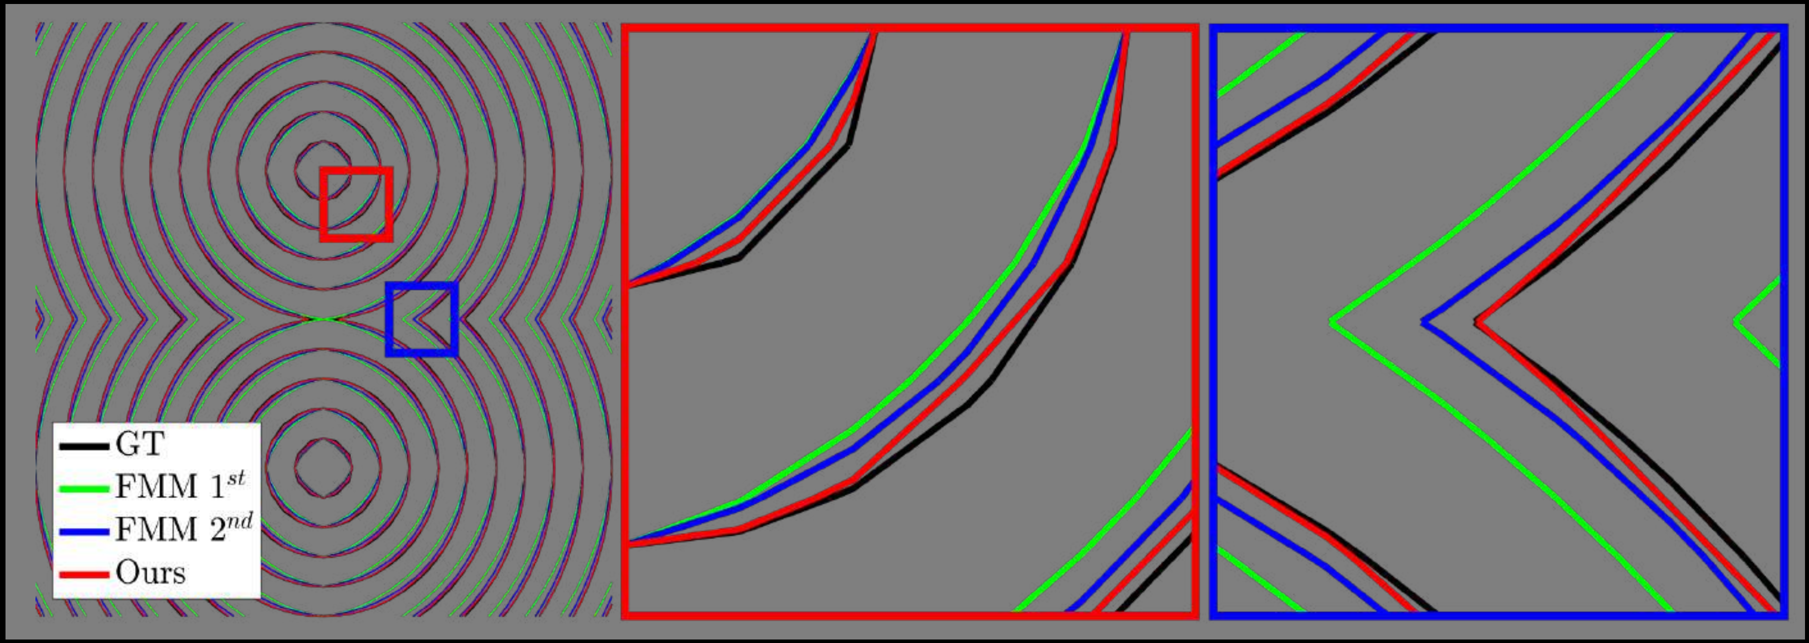

Lichtenstein, Pai, K., SSVM 2019 - Best Student Paper Award

Inter-dataset generalization

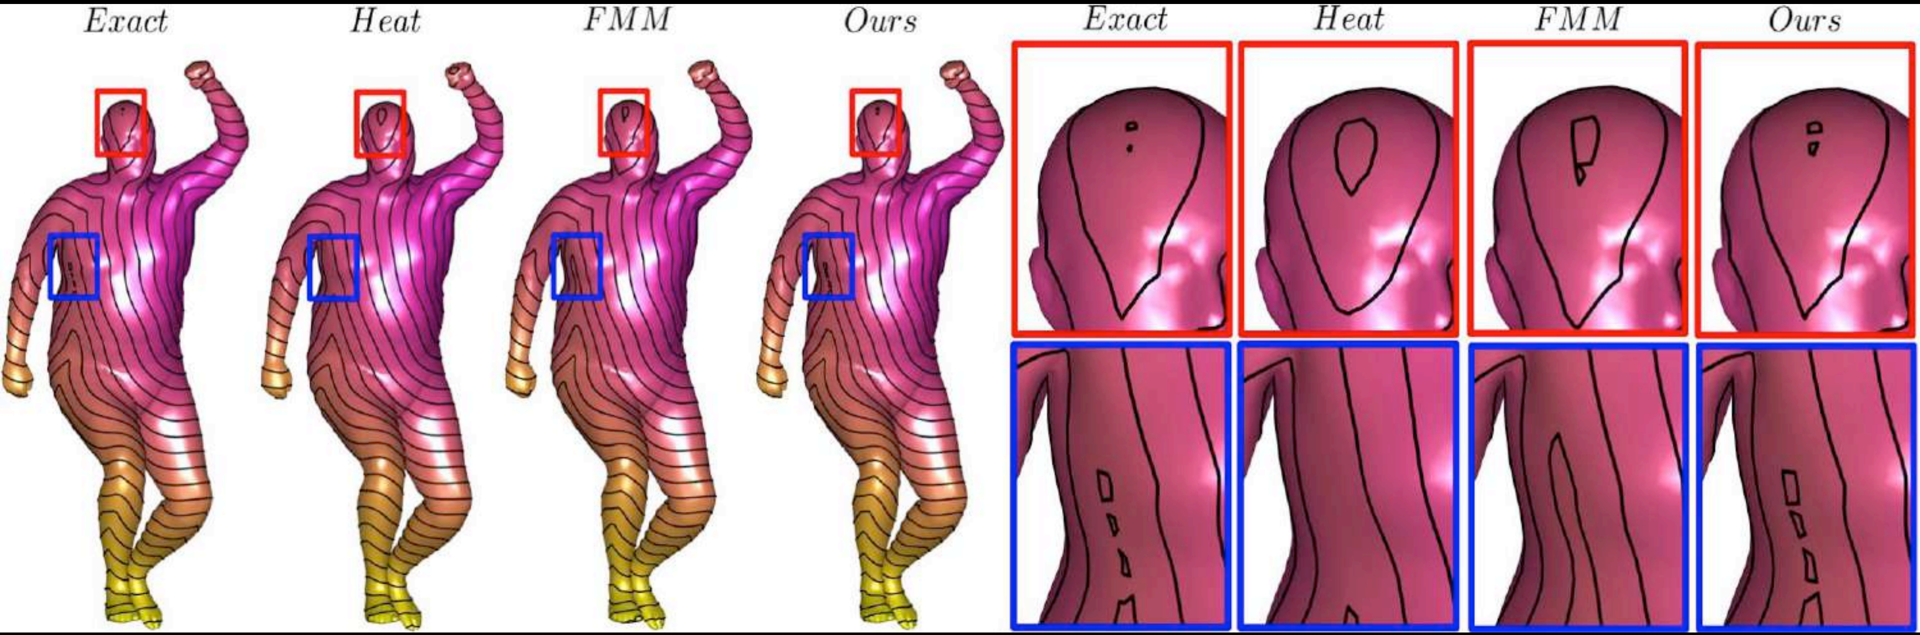

Lichtenstein, Pai, K., SSVM 2019 - Best Student Paper Award

Order of accuracy

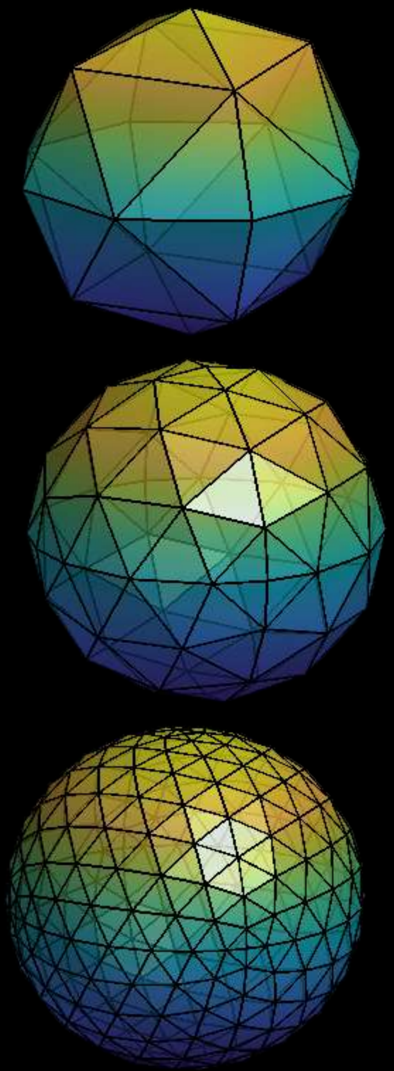

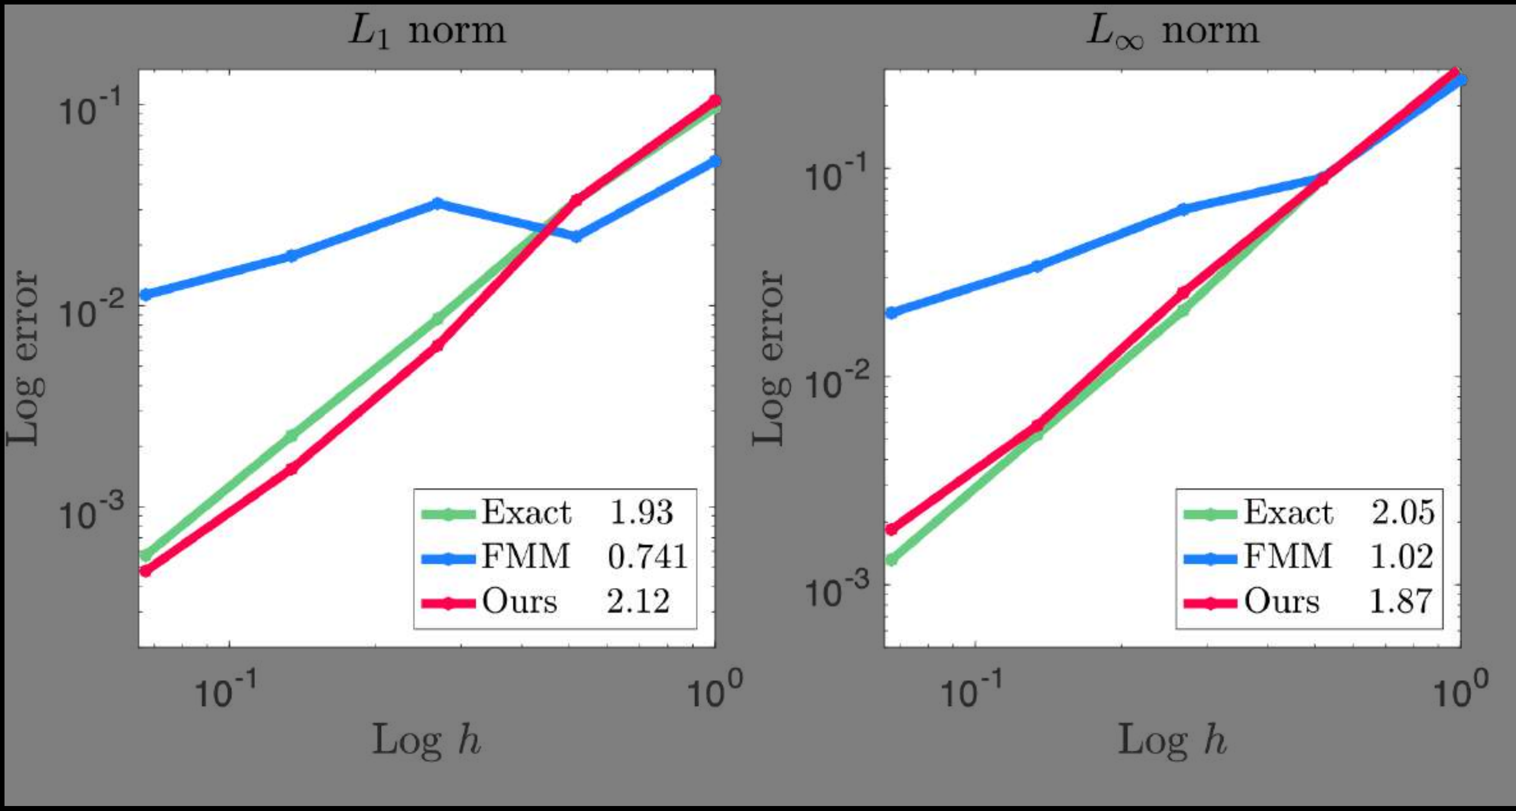

Lichtenstein, Pai, K., SSVM 2019 - Best Student Paper Award

Scale invariance?

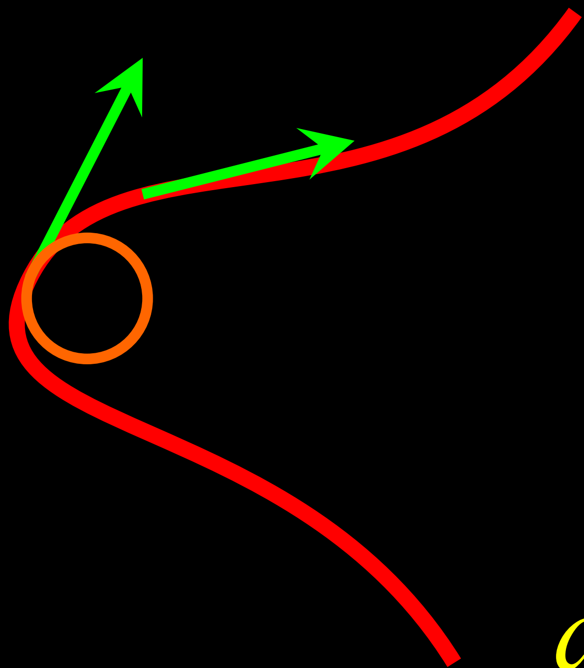

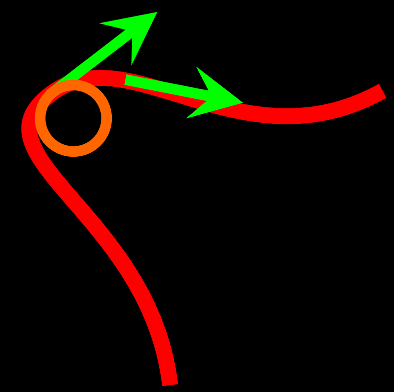

$$d\theta = |\kappa| ds$$

Manifold vs. Riemannian Manifold

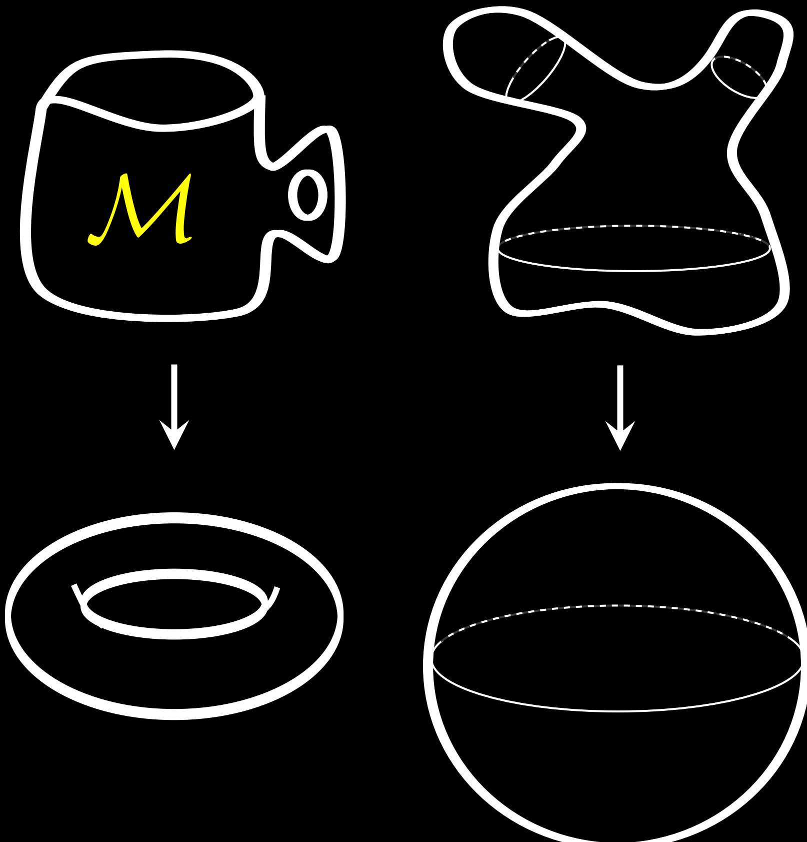

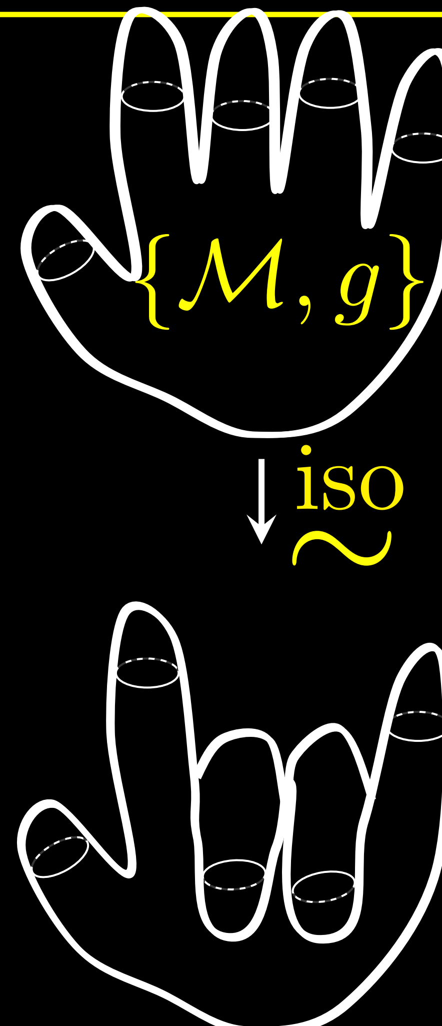

Manifold vs. Riemannian Manifolds

$$\Delta_g \psi_i = \lambda_i \psi_i$$

$$\Delta_{\tilde{g}} \tilde{\psi}_i = \tilde{\lambda}_i \tilde{\psi}_i$$

$$C_{ij} = \langle \psi_i, \tilde{\psi}_j \rangle_{\tilde{g}}$$

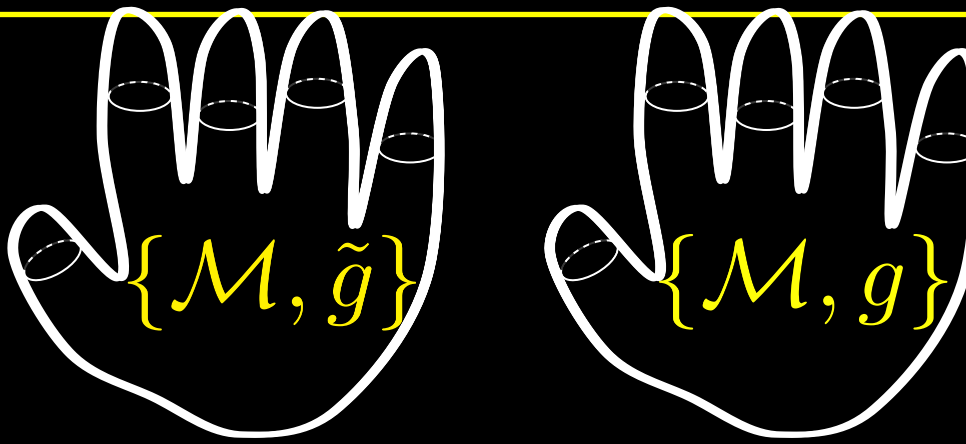

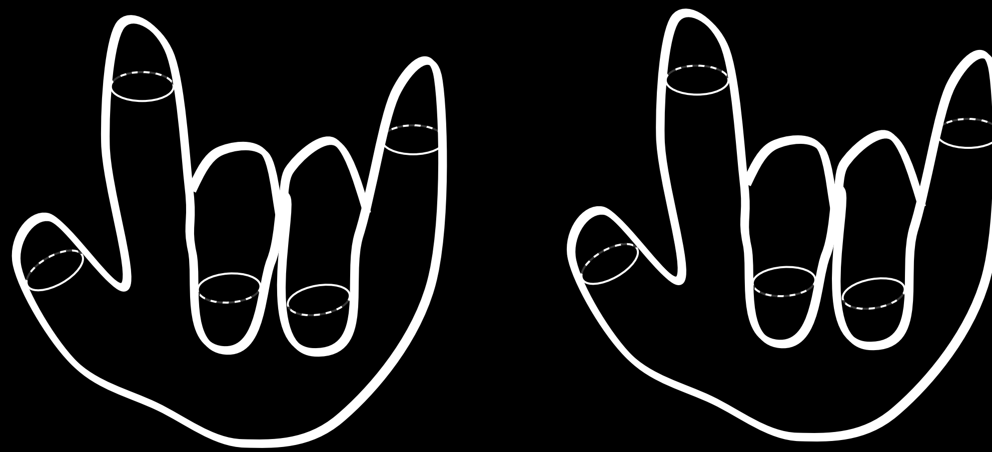

Functional Maps

$$\mathcal{Y} \approx P\mathcal{X}$$

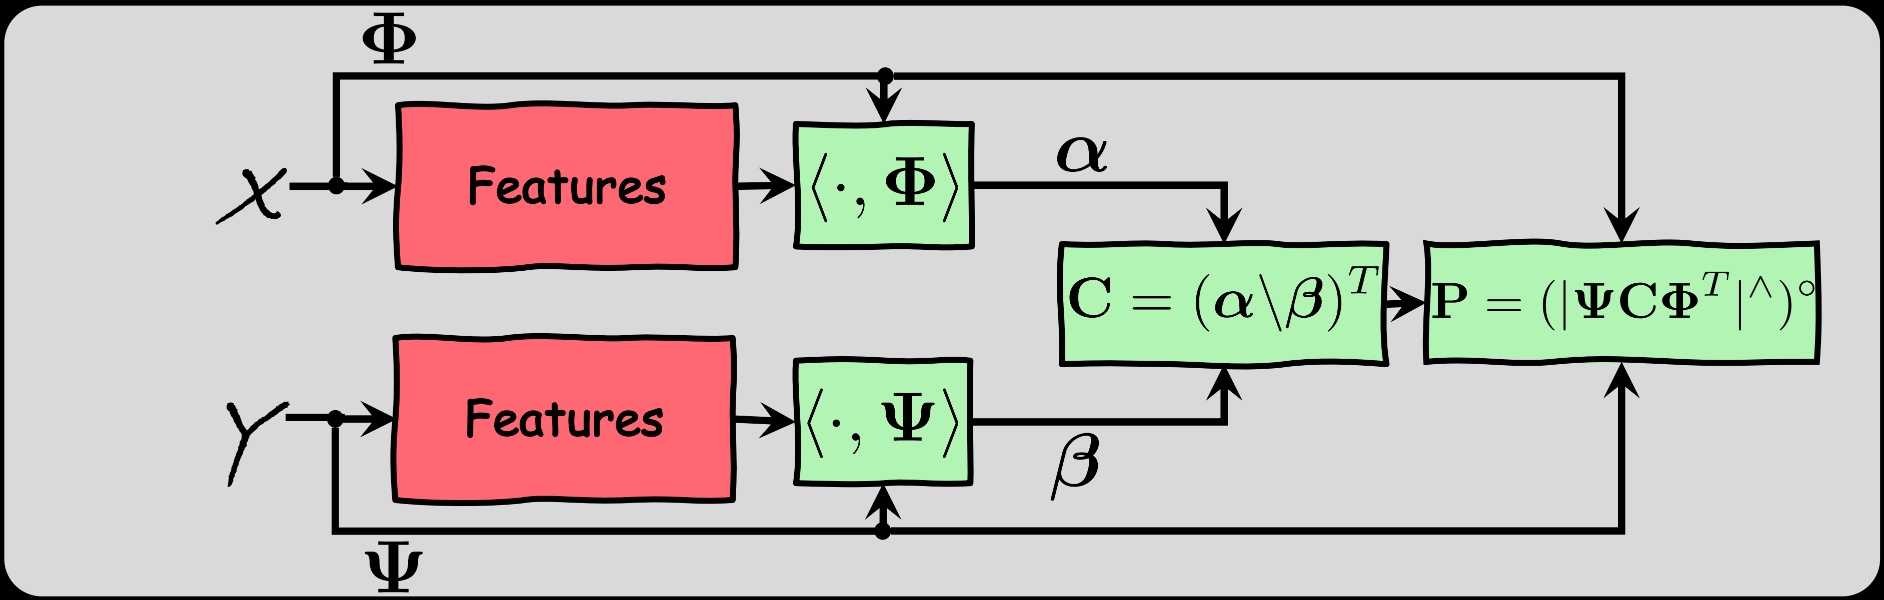

Ovsjanikov et al. 2012

Litany, Remez, Rodola, A Bronstein, M Bronstein ICCV'17

Halimi, Litany, Rodola, A Bronstein, K. CVPR'19

Roufosse, Sharma, Ovsjanikov, ICCV'19

Functional Maps Ovsjanikov et al. 2012

$$\{\psi_i\}$$

$$f = \sum_i \langle f, \psi_i \rangle \psi_i = \sum_i \alpha_i \psi_i$$

$$\{\phi_i\}$$

$$f = \sum_i \langle f, \phi_i \rangle \phi_i = \sum_i \beta_i \phi_i$$

$$\beta = C\alpha$$

$$C_{ij} = \langle T(\phi_i), \psi_j \rangle$$

$$l_{sup}(\mathcal{X},\mathcal{Y}) = \sum_{i\in\mathcal{X}}\sum_{j\in\mathcal{Y}} p_{ij} d^2_{\mathcal{Y}}(j,\pi^*(i))$$

$$l_{uns}(\mathcal{X},\mathcal{Y}) = \|D_{\mathcal{X}} - P D_{\mathcal{Y}} P^T\|$$

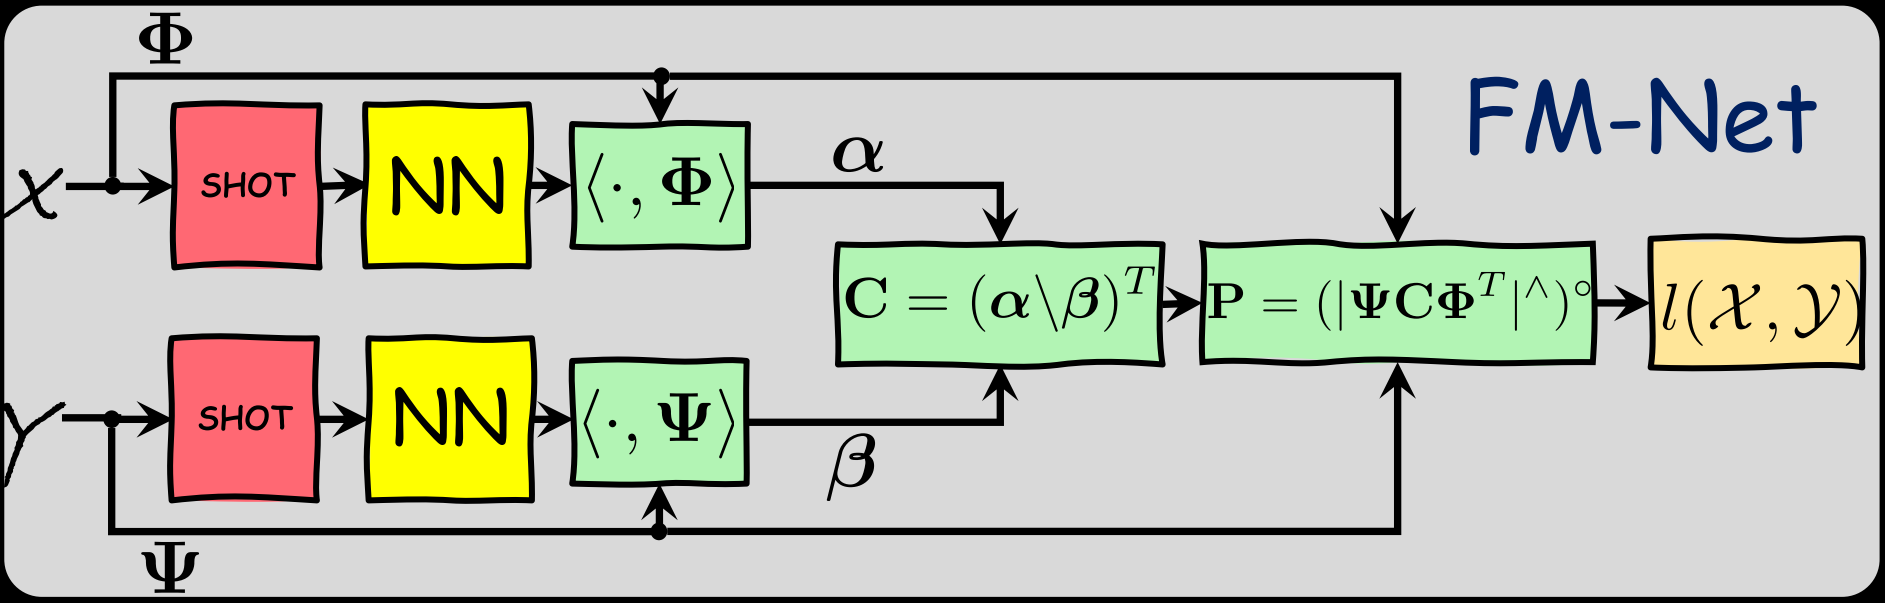

Litany, Remez, Rodola, A Bronstein, M Bronstein ICCV'17
Halimi, Litany, Rodola, A Bronstein, K. CVPR'19
Roufosse, Sharma, Ovsjanikov, ICCV'19

Performance comparison

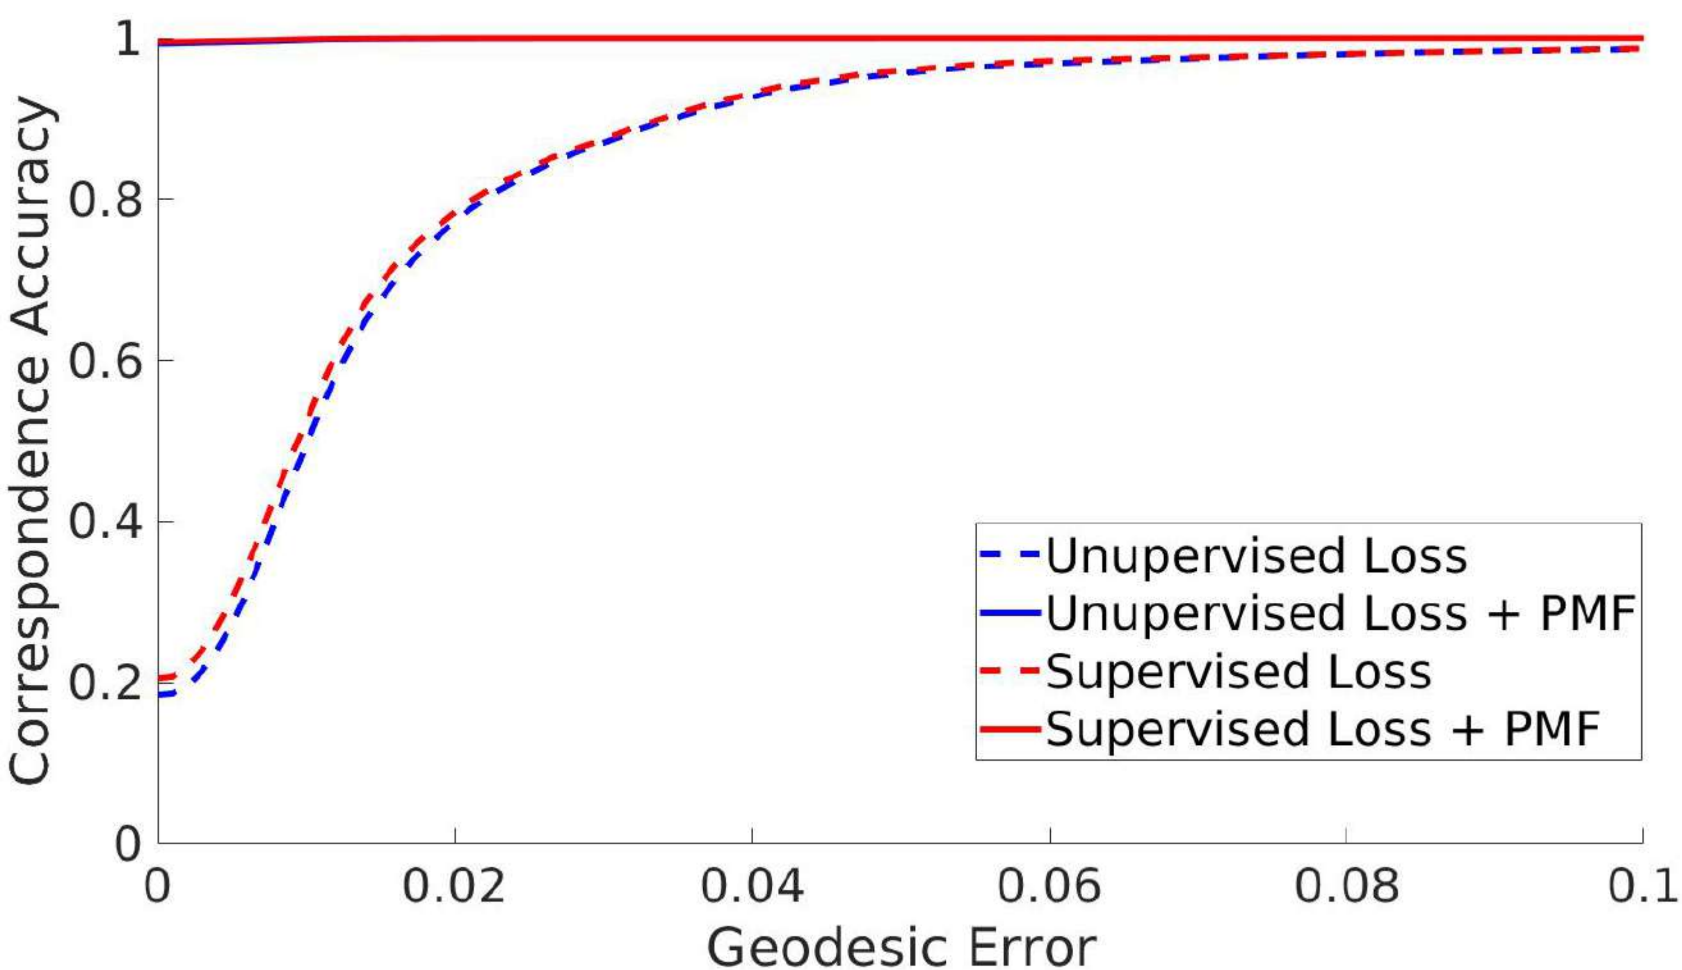

Halimi, Litany, Rodolà, Bronstein, K. CVPR'19

Partial Correspondence

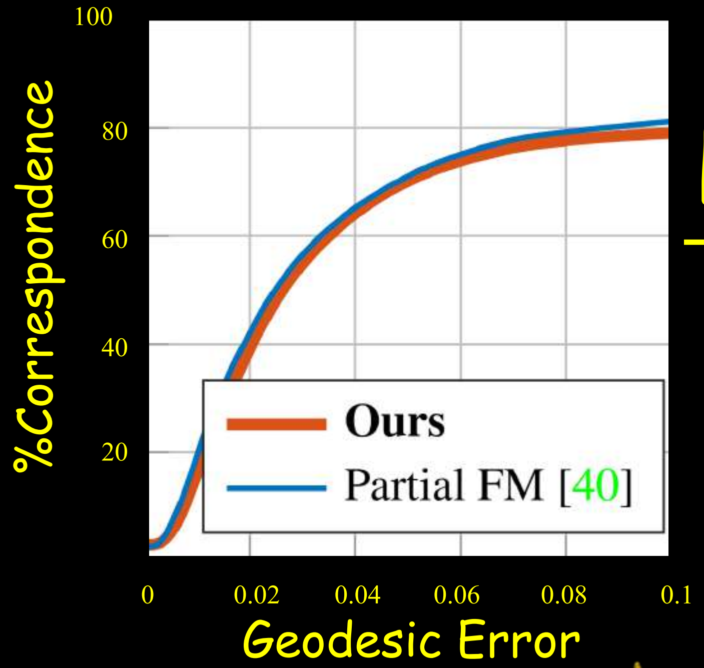

Halimi, Litany, Rodolà, Bronstein, K. CVPR'19

Single-pair correspondences

Halimi, Litany, Rodolà, Bronstein, K. CVPR'19

Self-Supervised Training

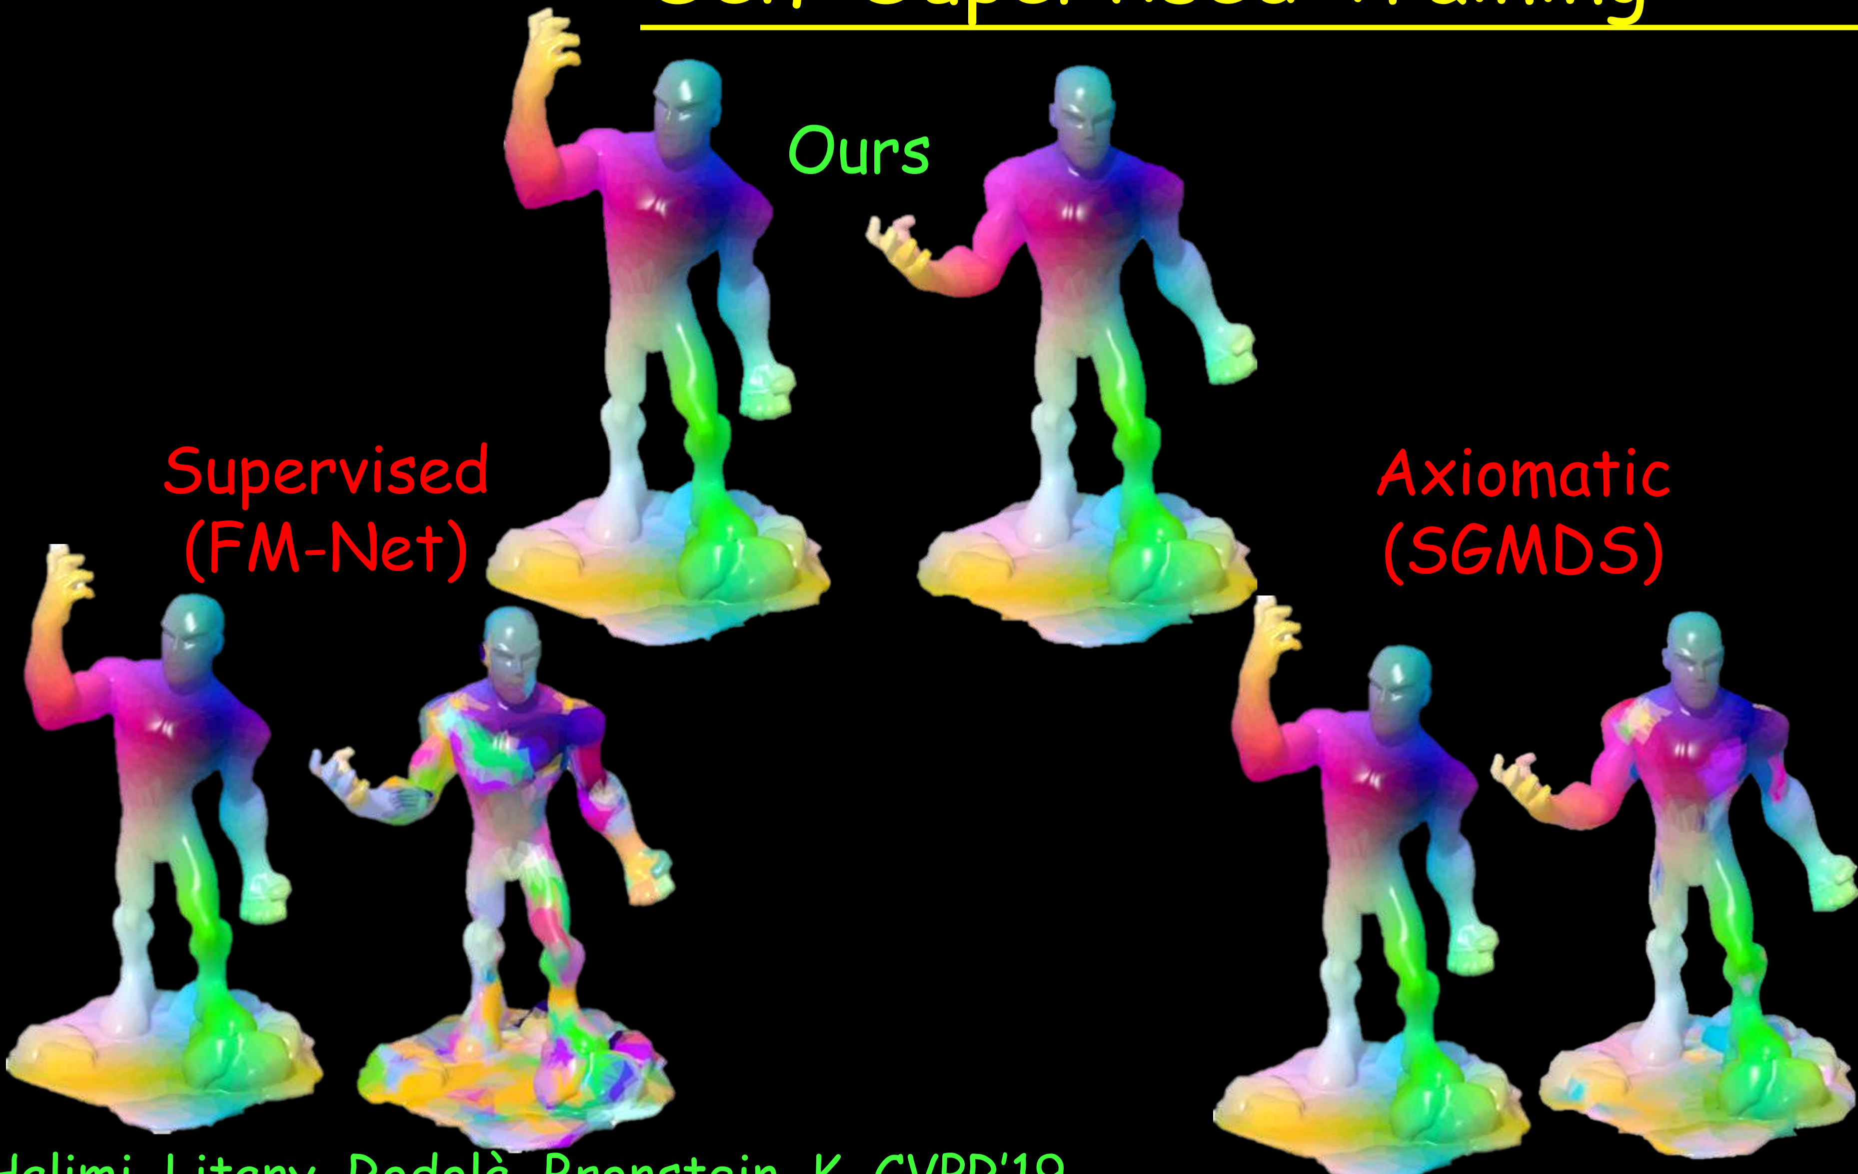

Halimi, Litany, Rodolà, Bronstein, K. CVPR'19

Surface Laplacian

$$\Delta_{\tilde{g}} \equiv -\frac{1}{\sqrt{\tilde{g}}}\partial_i \sqrt{\tilde{g}}\,\tilde{g}^{ij}\partial_j$$

$\tilde{g}_{ij} =$  $= |\kappa_1 \kappa_2| \langle S_i, S_j \rangle$

Eigenfunctions

$$g_{ij} = \langle S_{\omega_i}, S_{\omega_j} \rangle$$

$$-\Delta_g \psi_i = \lambda_i \psi_i$$

$$\tilde{g}_{ij} = |K| g_{ij}$$

$$-\Delta_{\tilde{g}} \tilde{\psi}_i = \tilde{\lambda}_i \tilde{\psi}_i$$

Halimi & K. 2018

Self functional maps

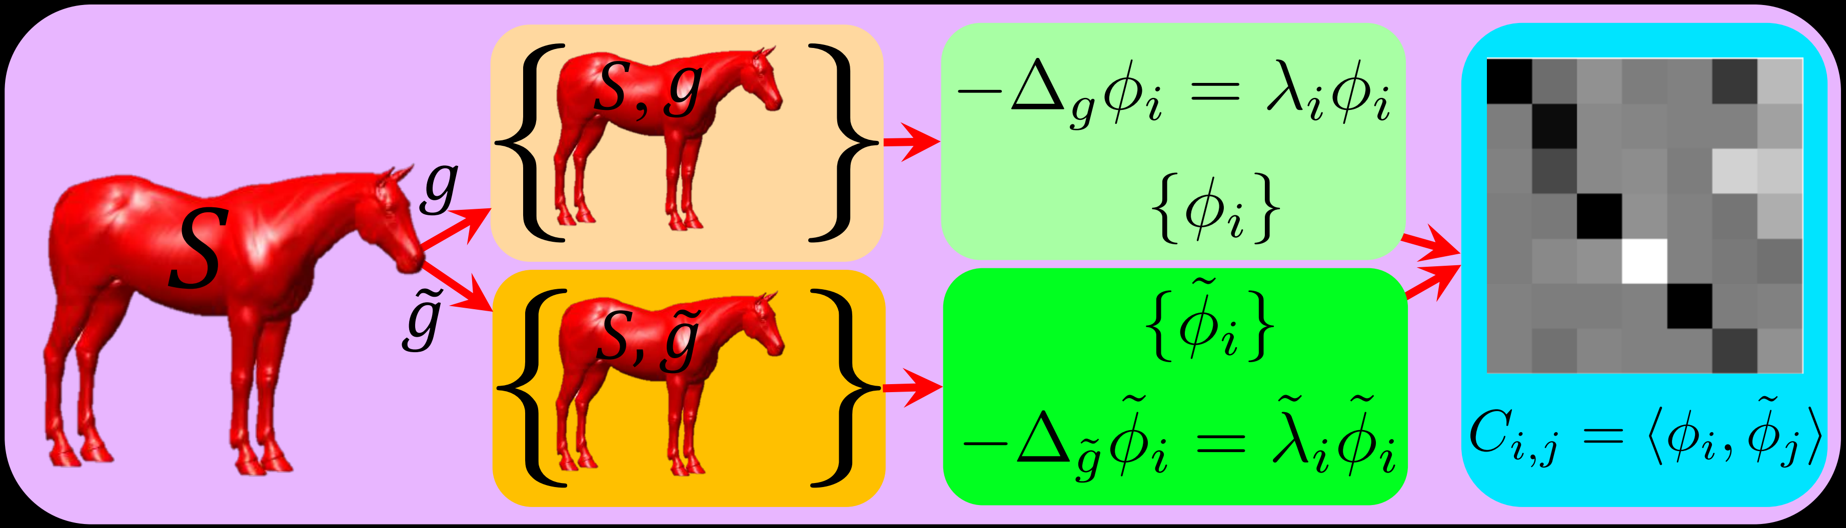

Halimi & K. 2018

Self functional maps

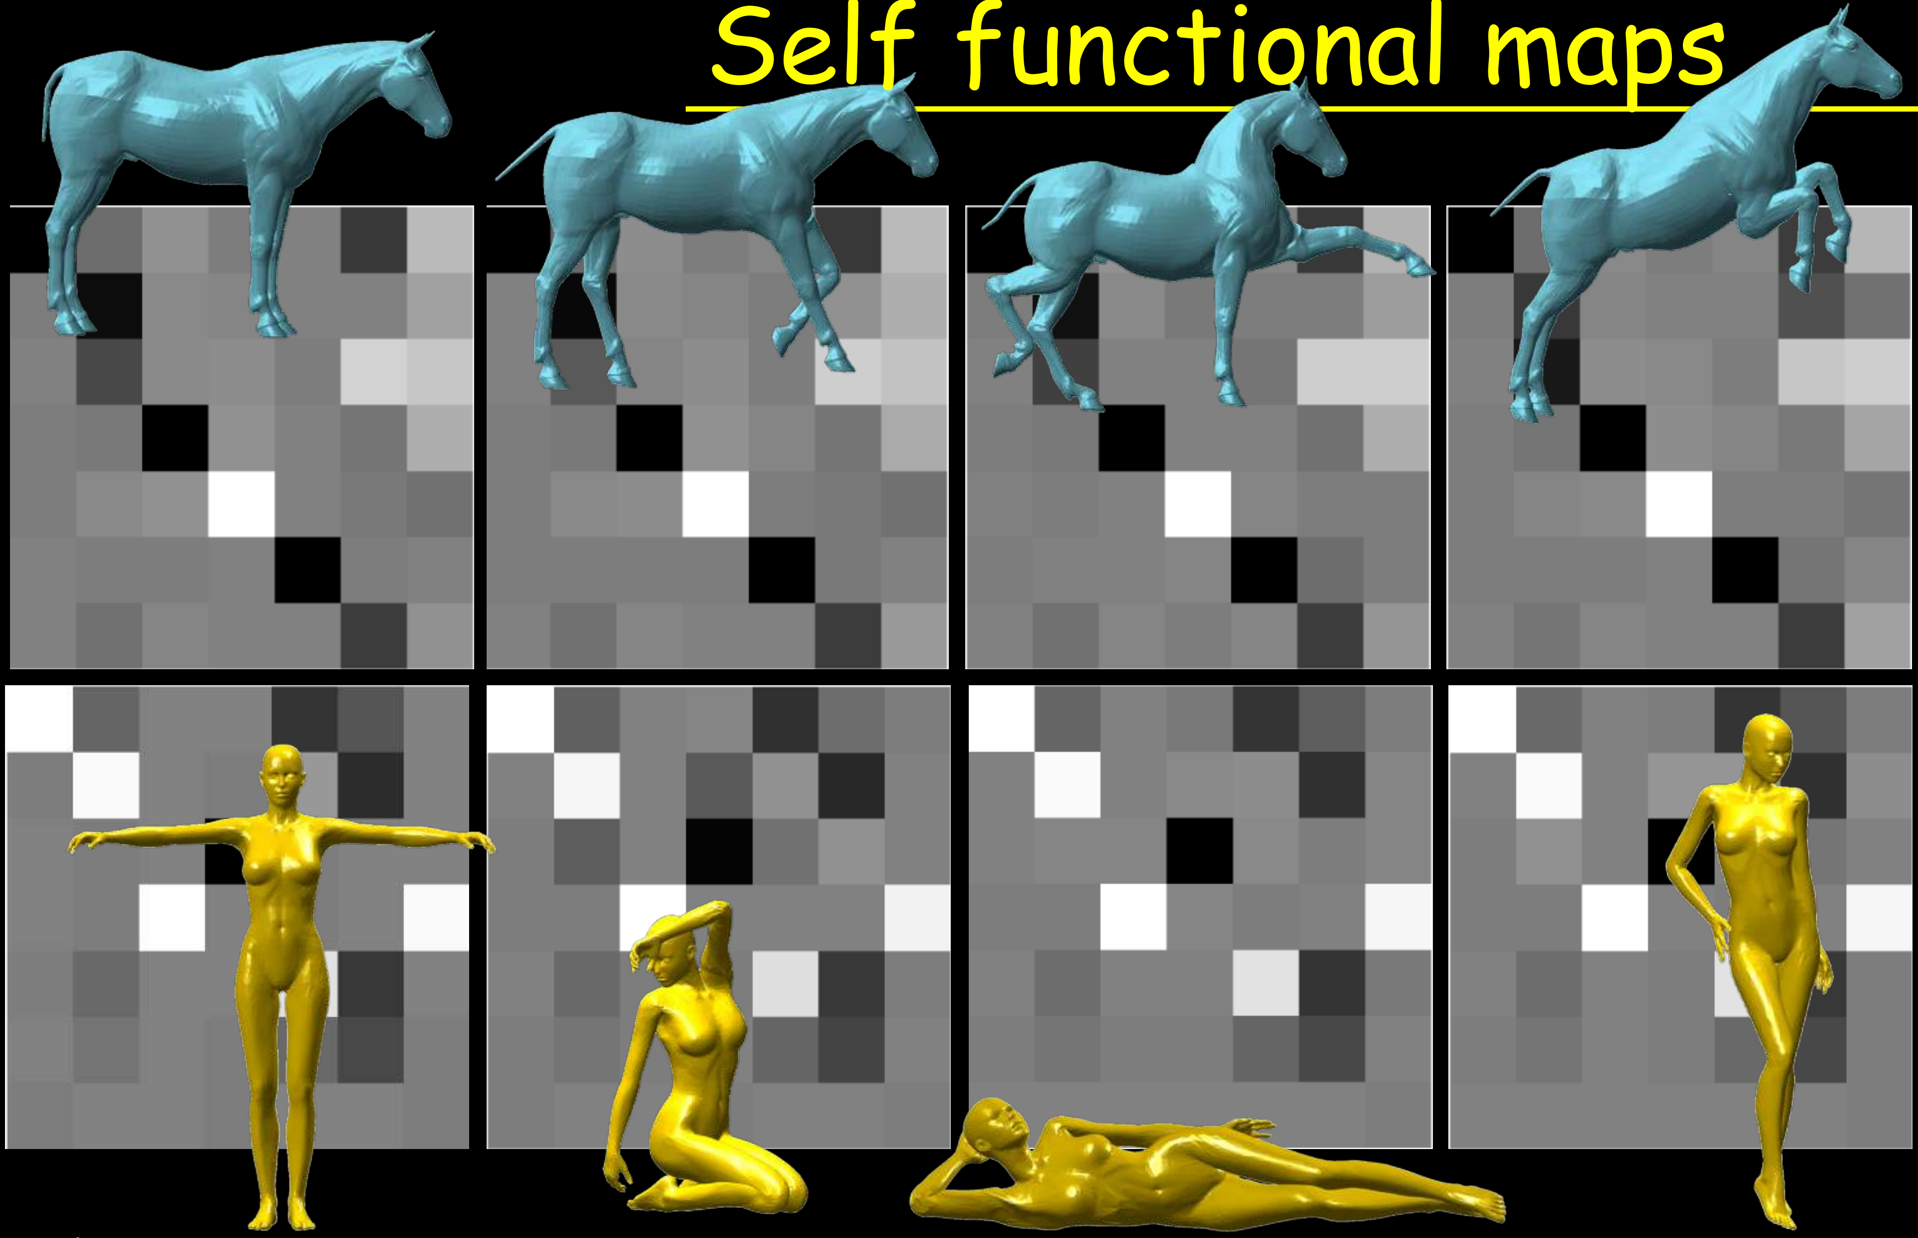

Halimi & K. 2018

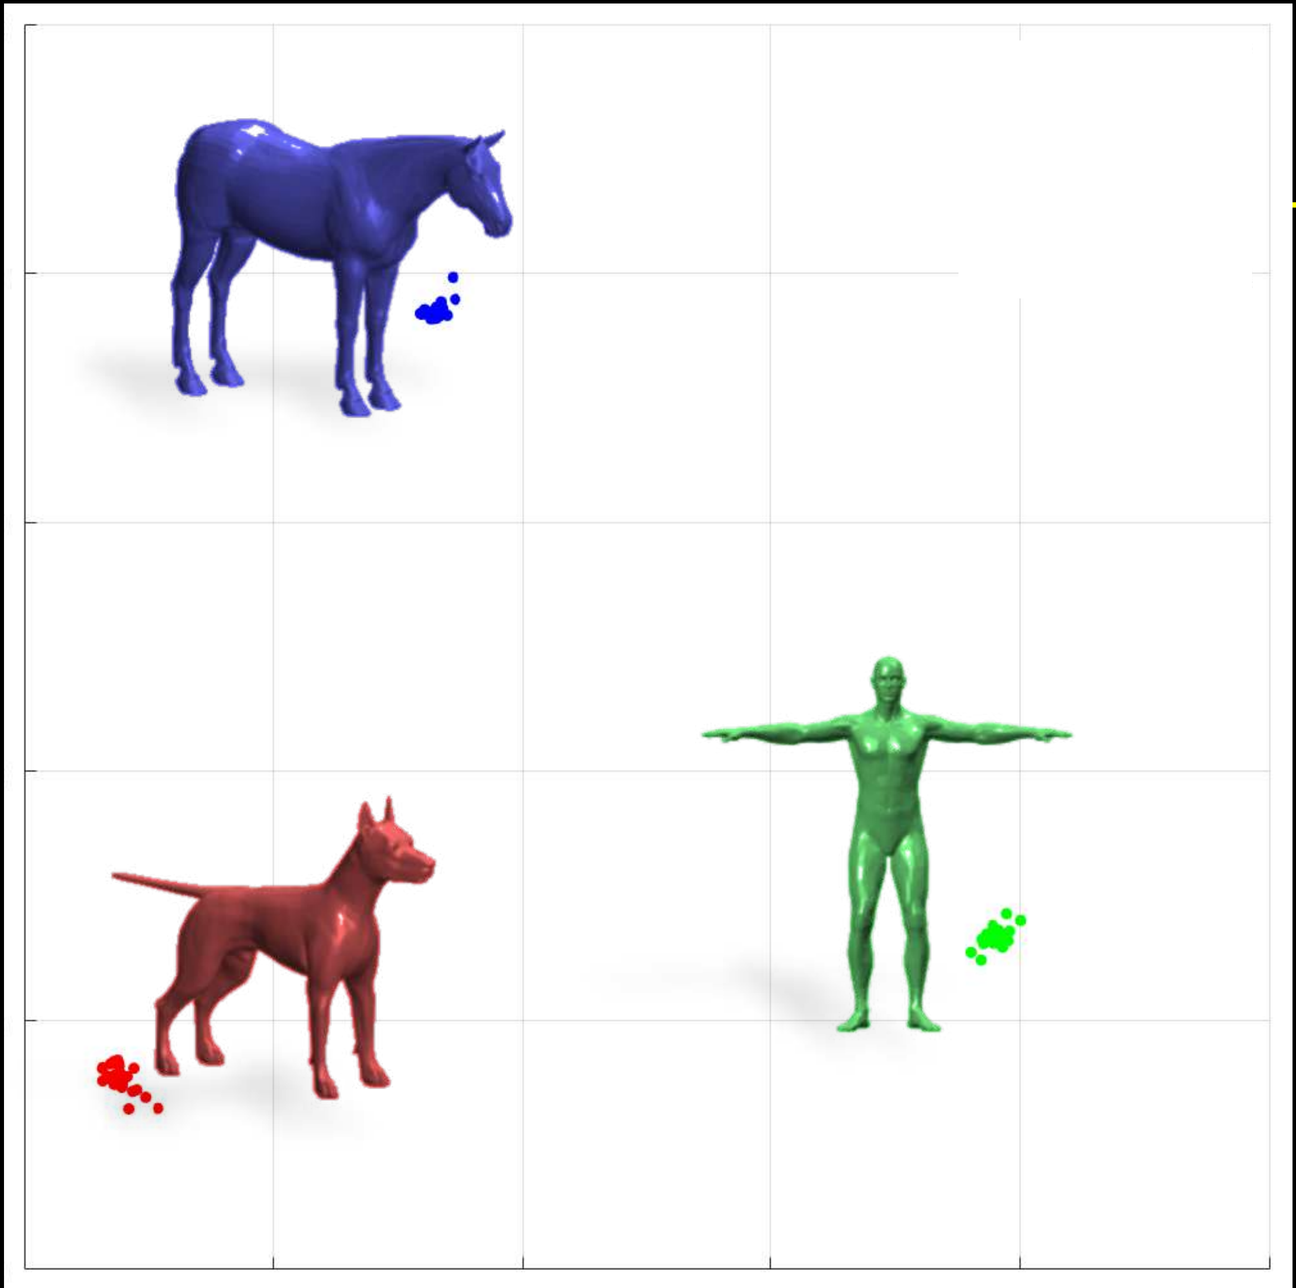

Halimi & K. 2018

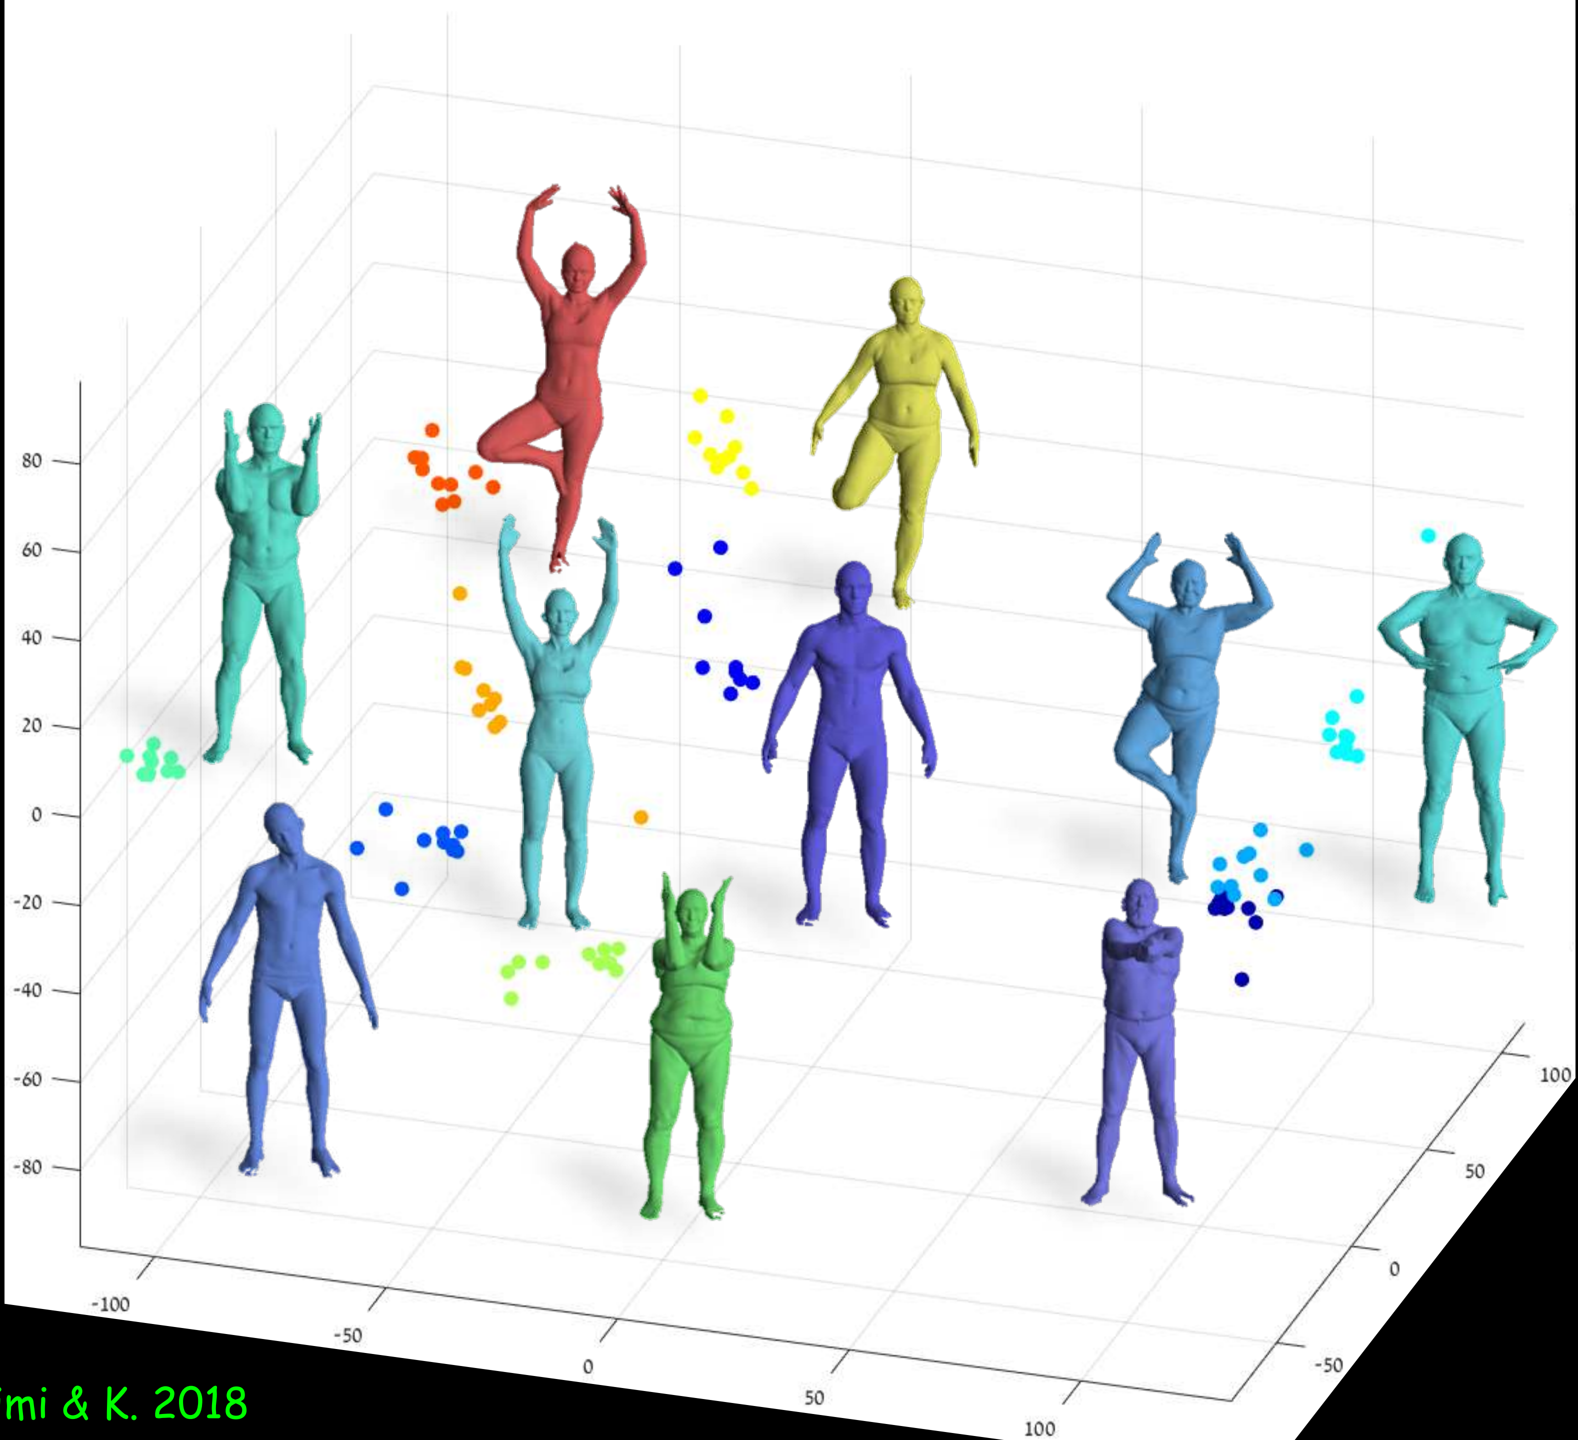

Halimi & K. 2018

$$l_{uns}(\mathcal{X},\mathcal{Y}) = \|D_{\mathcal{X}} - P D_{\mathcal{Y}} P^T\|$$

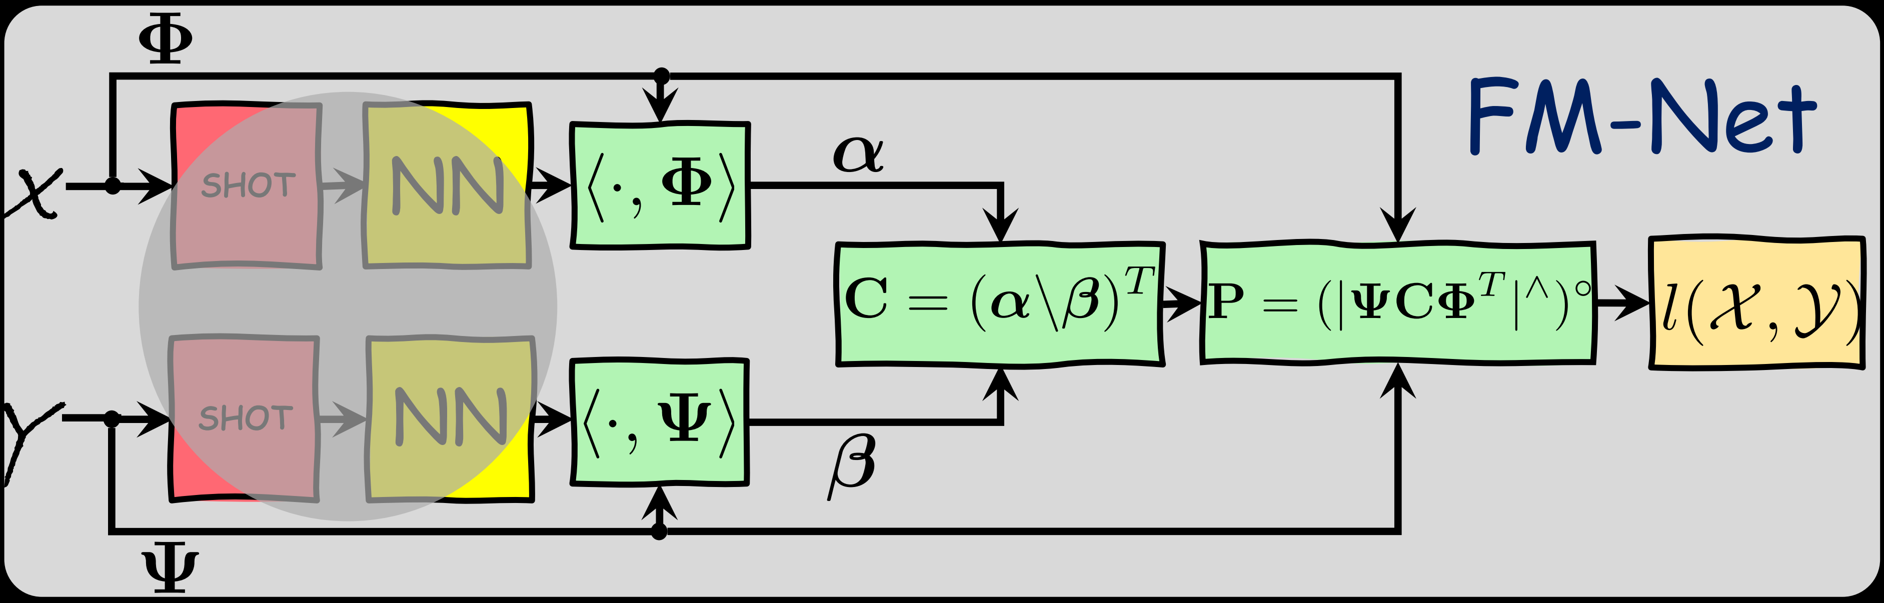

Halimi, Litany, Rodola, A Bronstein, K. CVPR'19

Aligning scale-invariant LBO eigenfunctions

$$l_{uns}(\mathcal{X},\mathcal{Y}) = \|D_{\mathcal{X}} - PD_{\mathcal{Y}}P^T\|$$

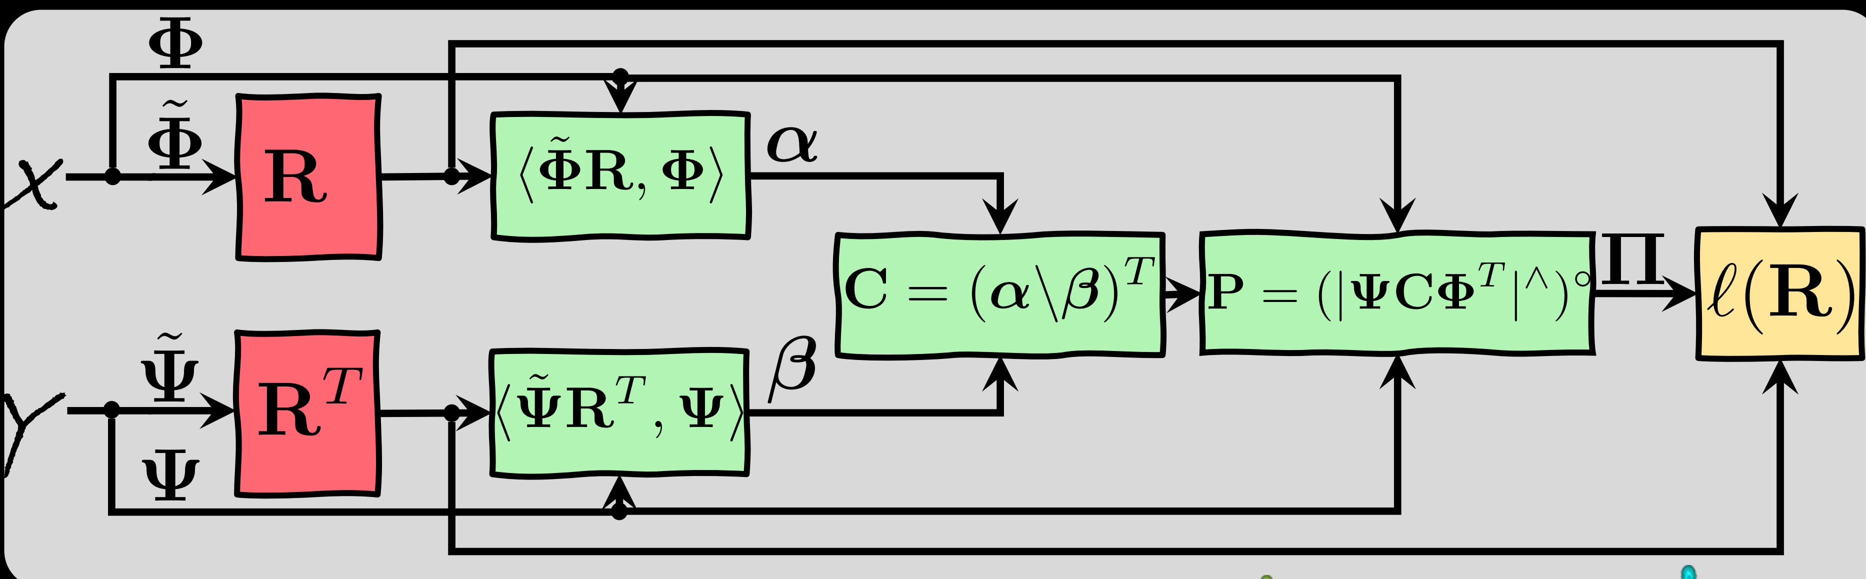

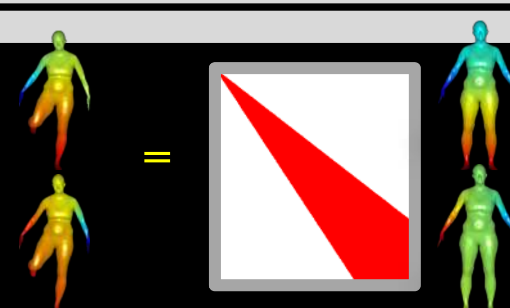

Bracha, Halimi, K. 3DOR'20

Before

After

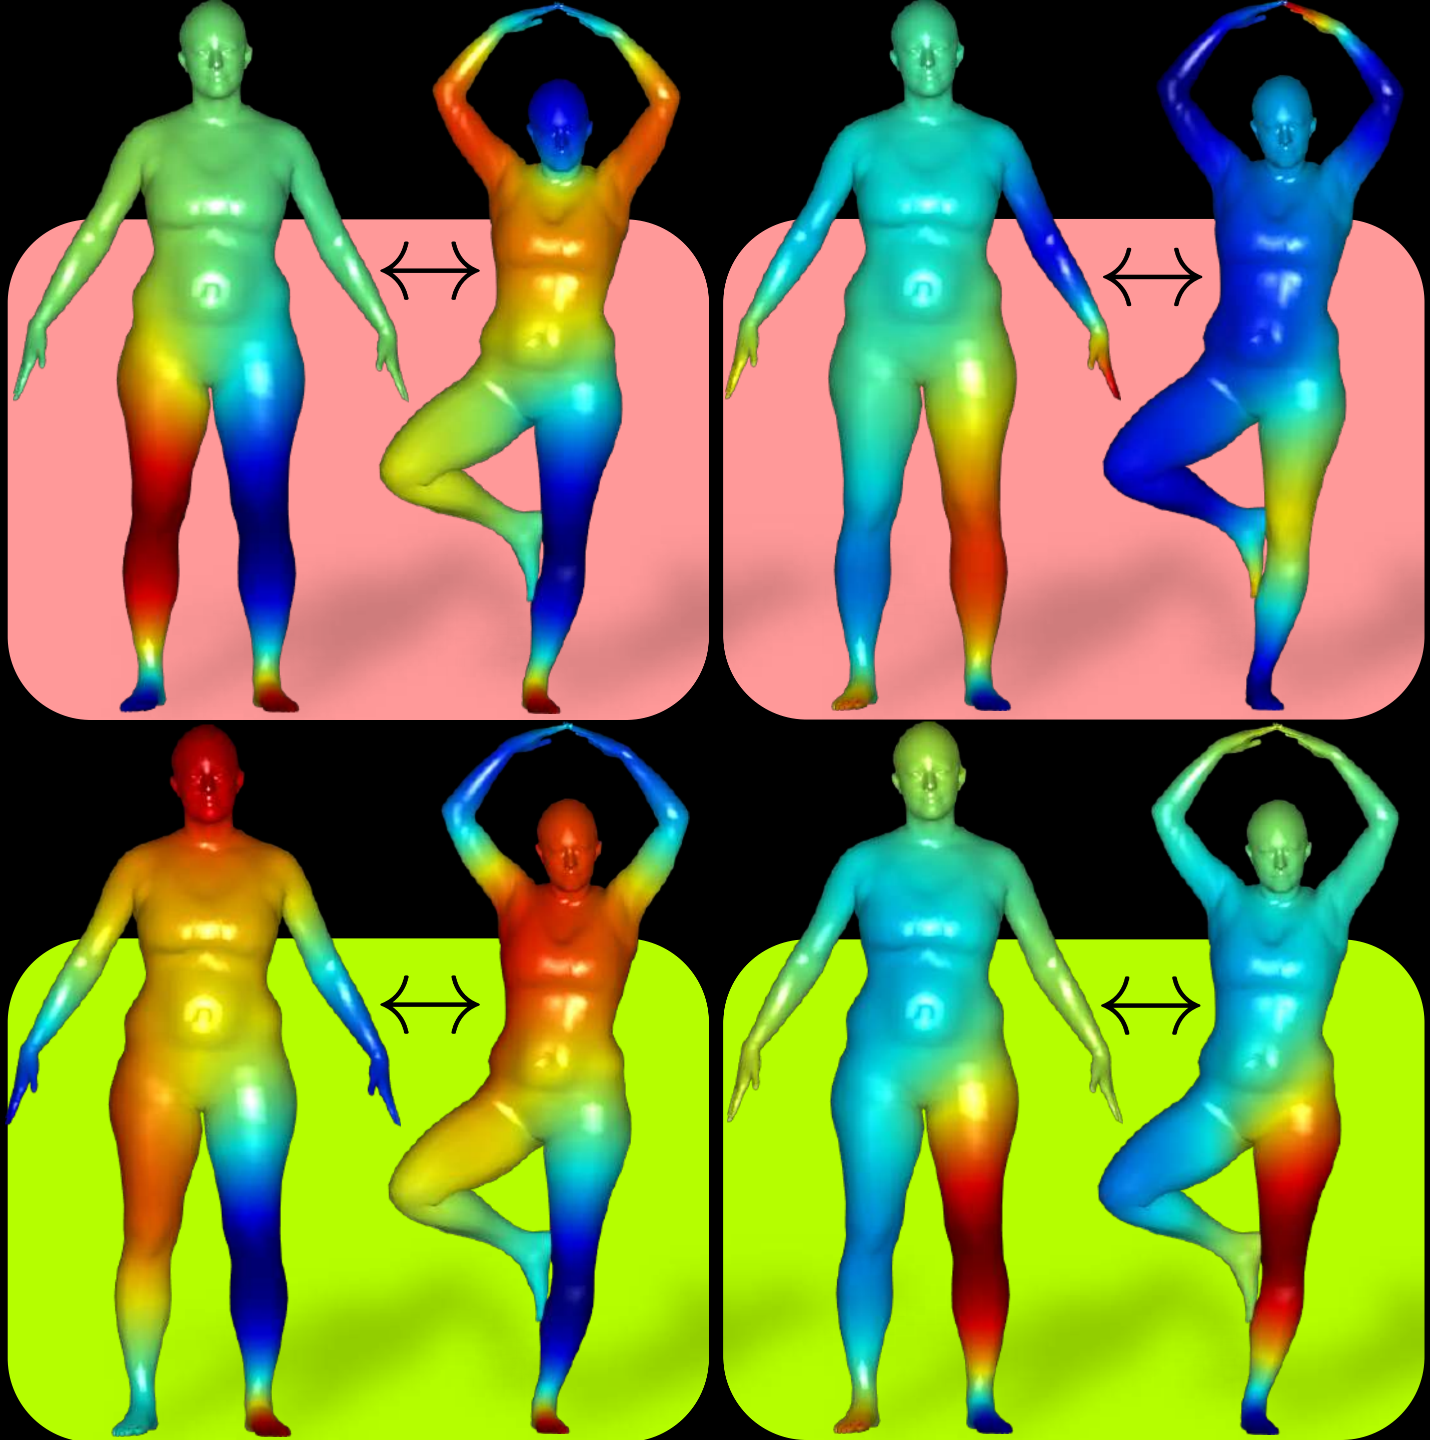

Bracha, Halimi, K. 3DOR'20

Before

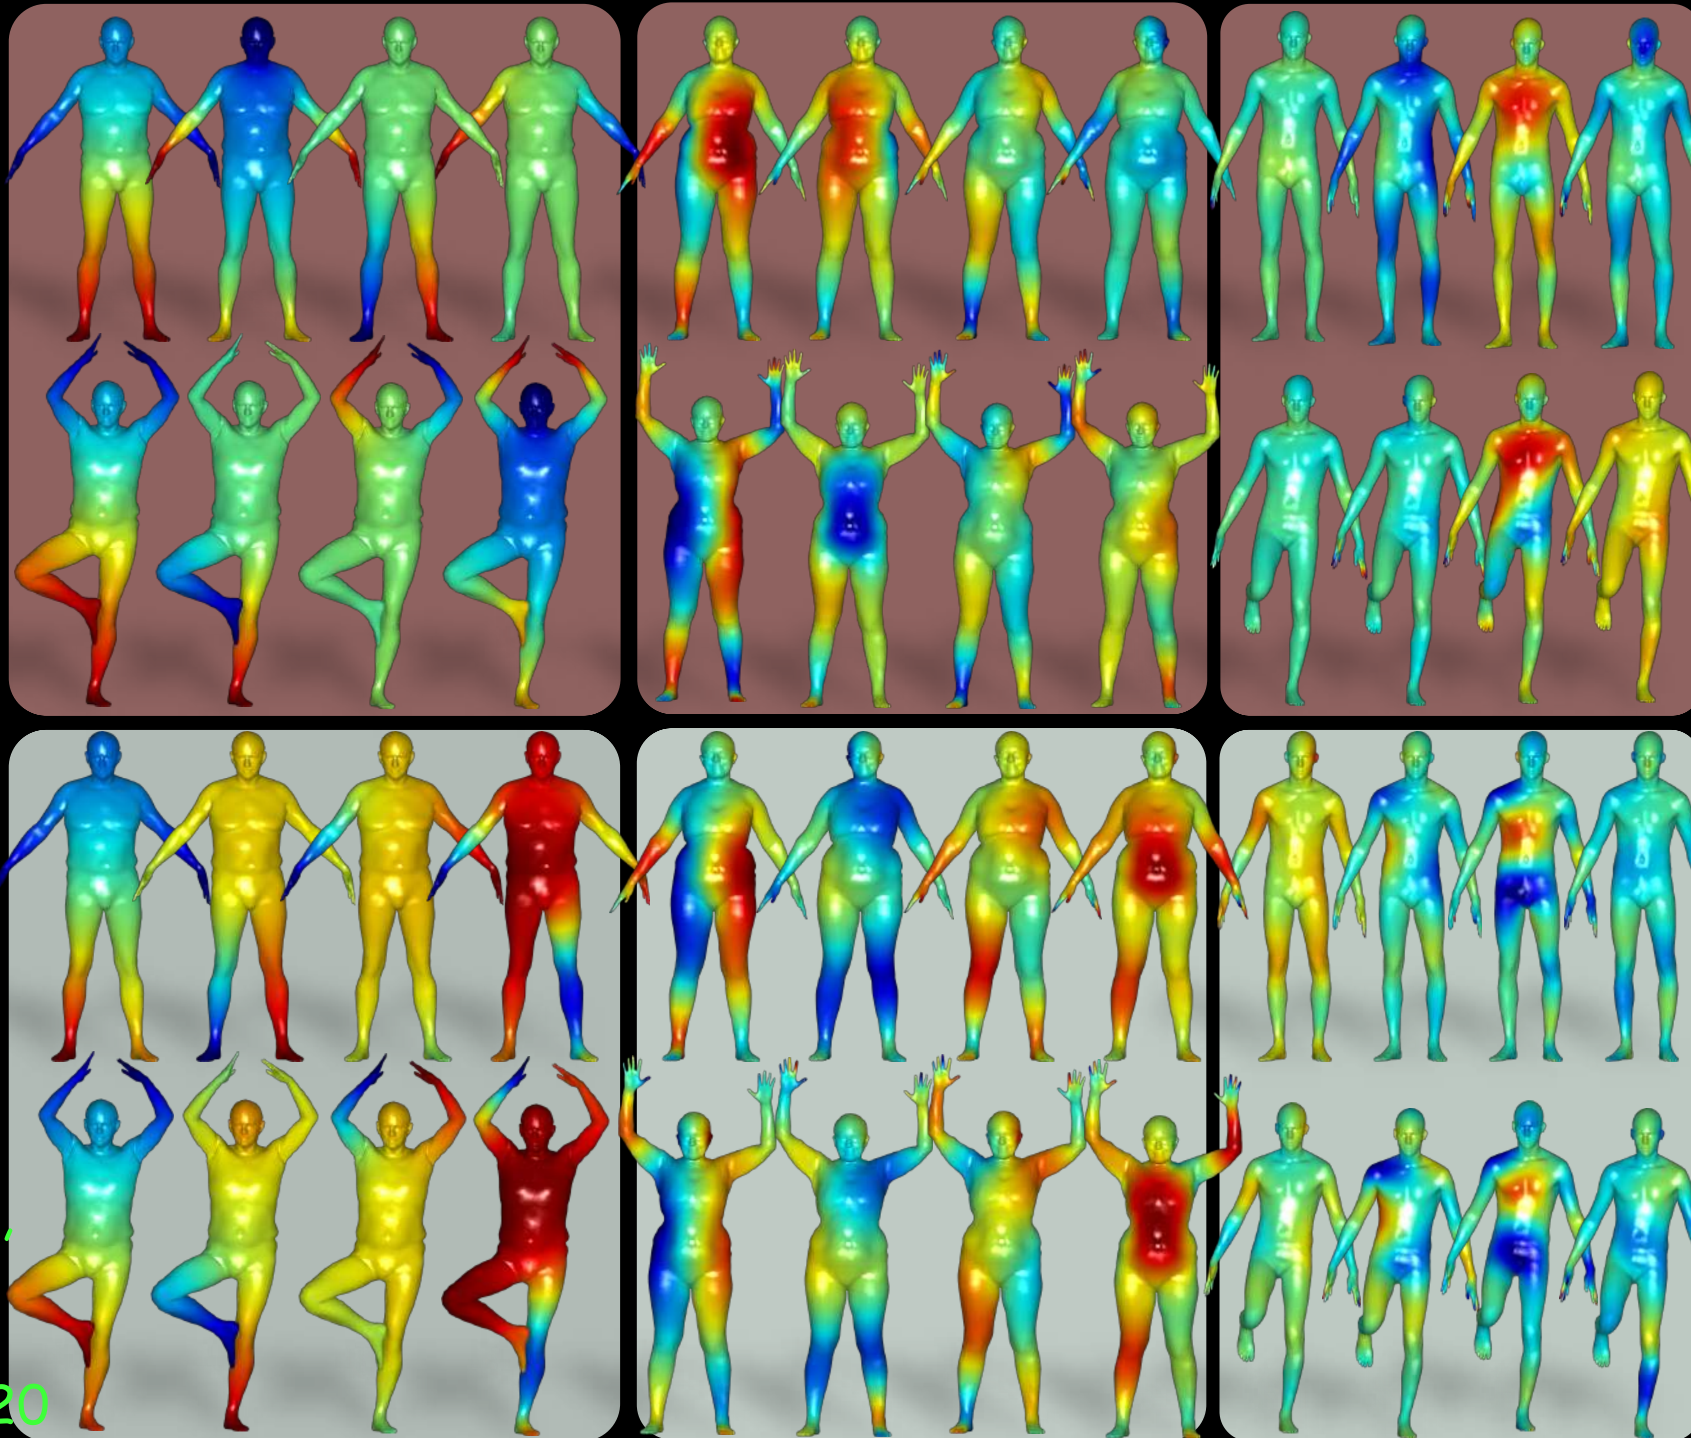

After

Bracha, Halimi, K. 3DOR'20

3DMM

Vetter & Blanz, A morphable model for the synthesis of 3D faces, Siggraph 1999

Kemelmacher & Basri, 3D face reconstruction from a single image ... PAMI 2011

Learning using Axiomatic Knowledge

Richardson Sela Or-El K. CVPR'17

Learning using Axiomatic Knowledge

Richardson Sela Or-El K. CVPR'17

Learning using Axiomatic Knowledge

We know how to model faces

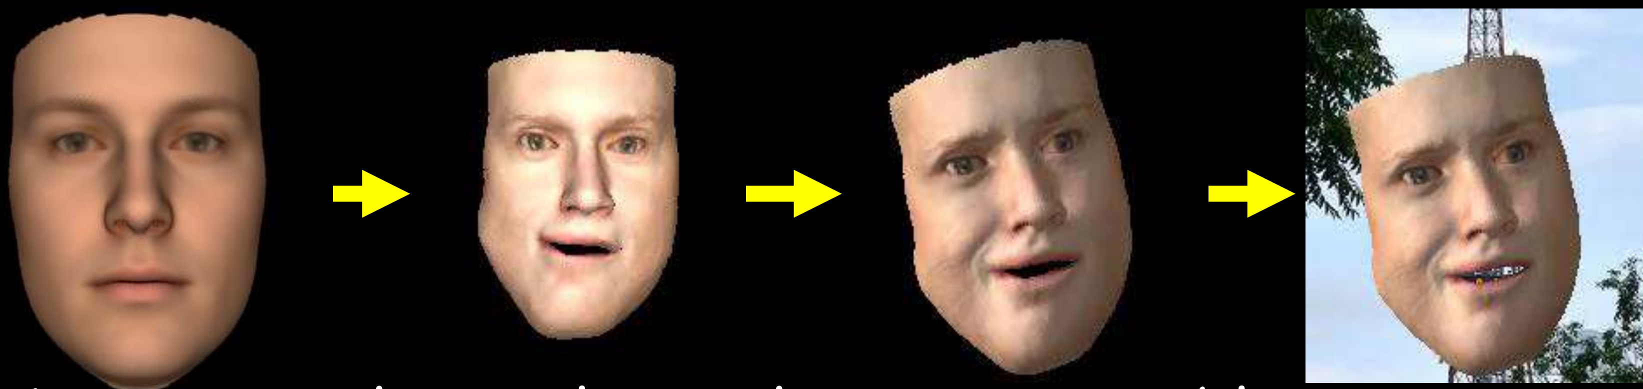

Can we use that to learn the inverse problem?

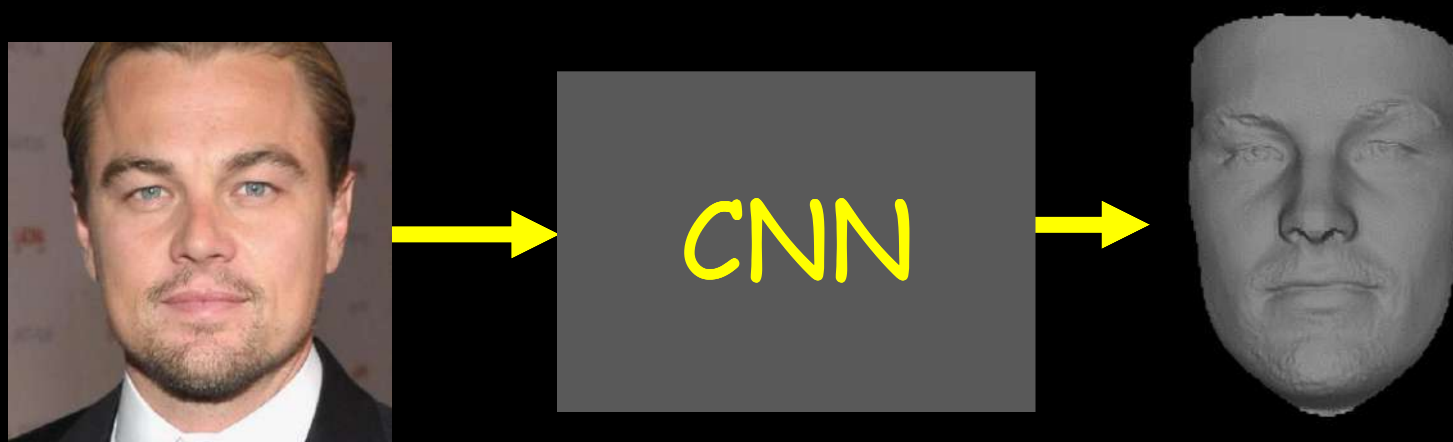

Richardson Sela Or-El K. CVPR'17

Face reconst n evolution

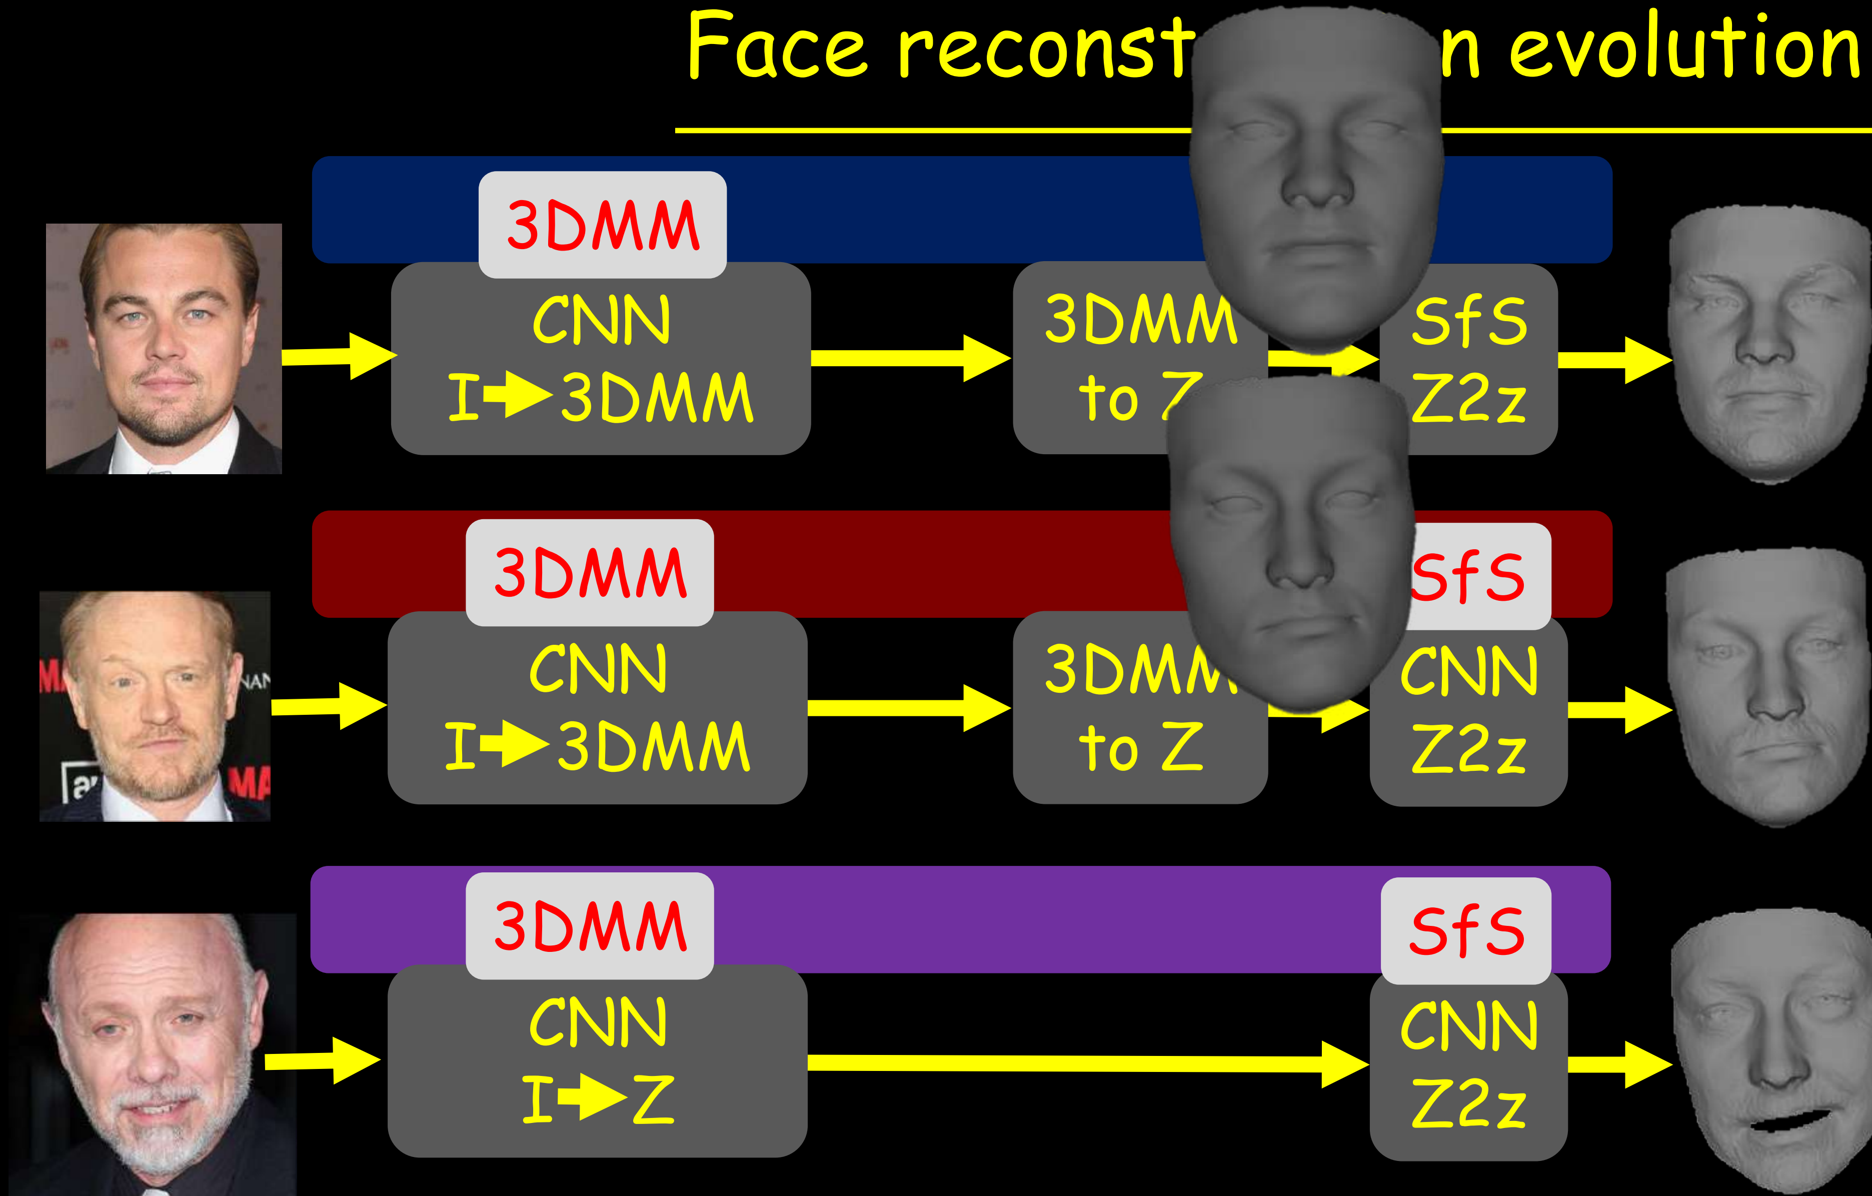

Richardson Sela K. 3DV'16/ Richardson Sela Or-El K. CVPR'17/ Richardson Sela K. ICCV'17

Shamai, Slossberg, K. 2018/2019 Synthesizing photometries and geometries with GANs &

Synthesizing Expressions

Shamai, Slossberg, K. 2018/2019 Synthesizing photometries and geometries with GANs &

Shamai, Slossberg, K. 2018/2019 Synthesizing photometries and geometries with GANs &

2D embedding identities

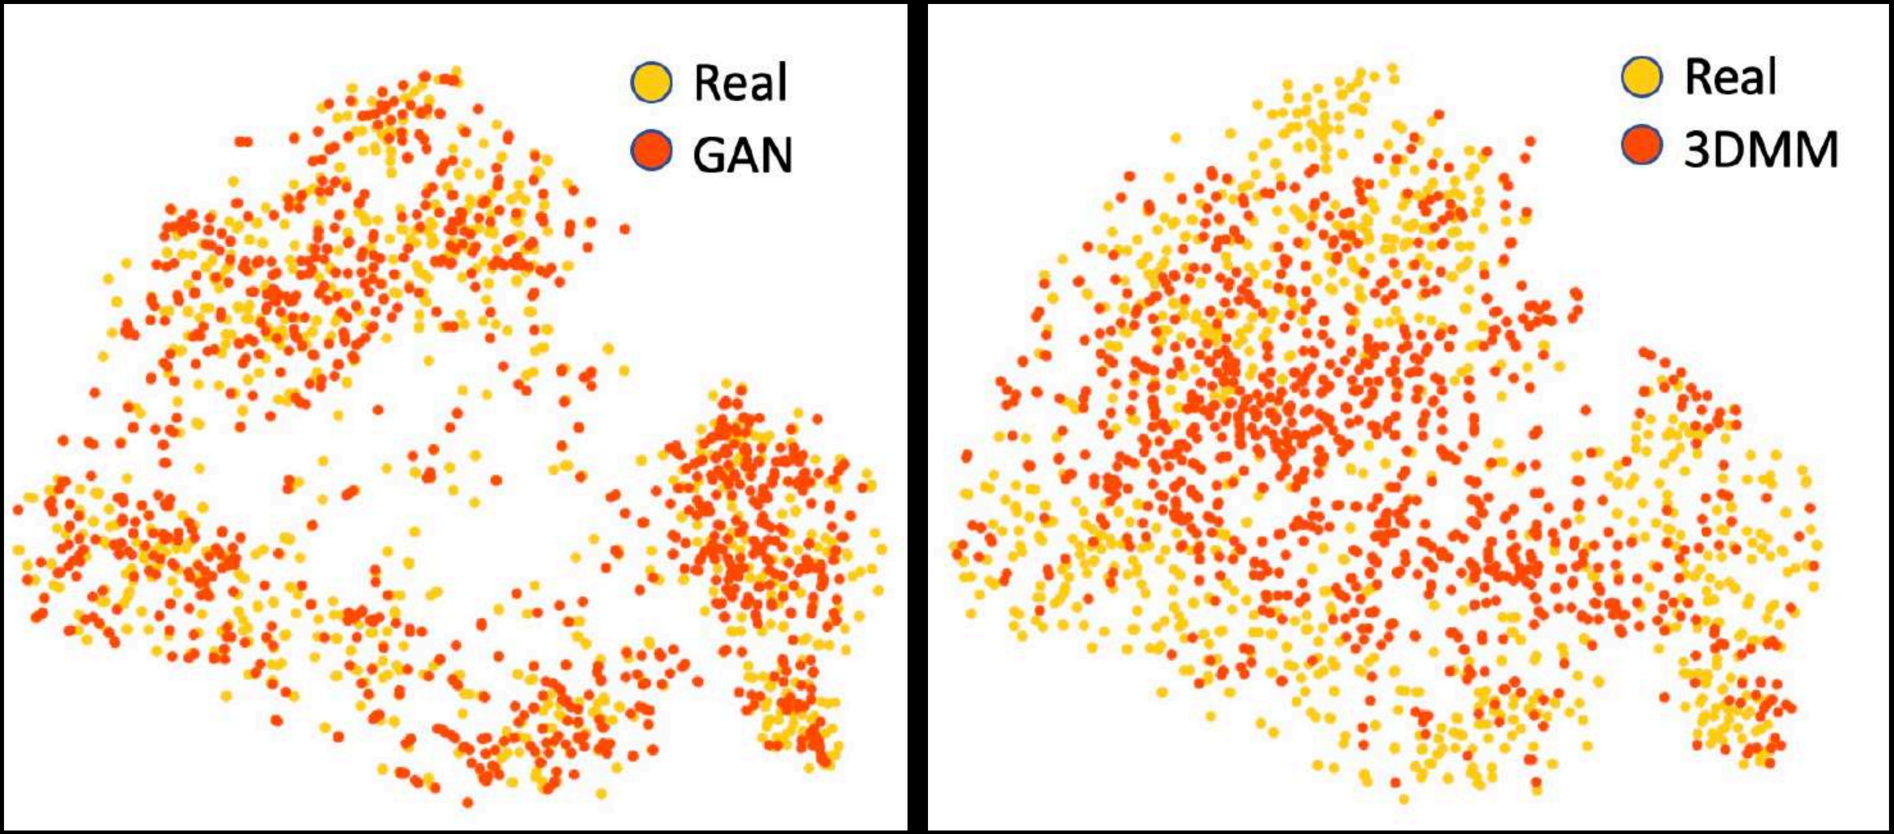

Training iterations (GAN)

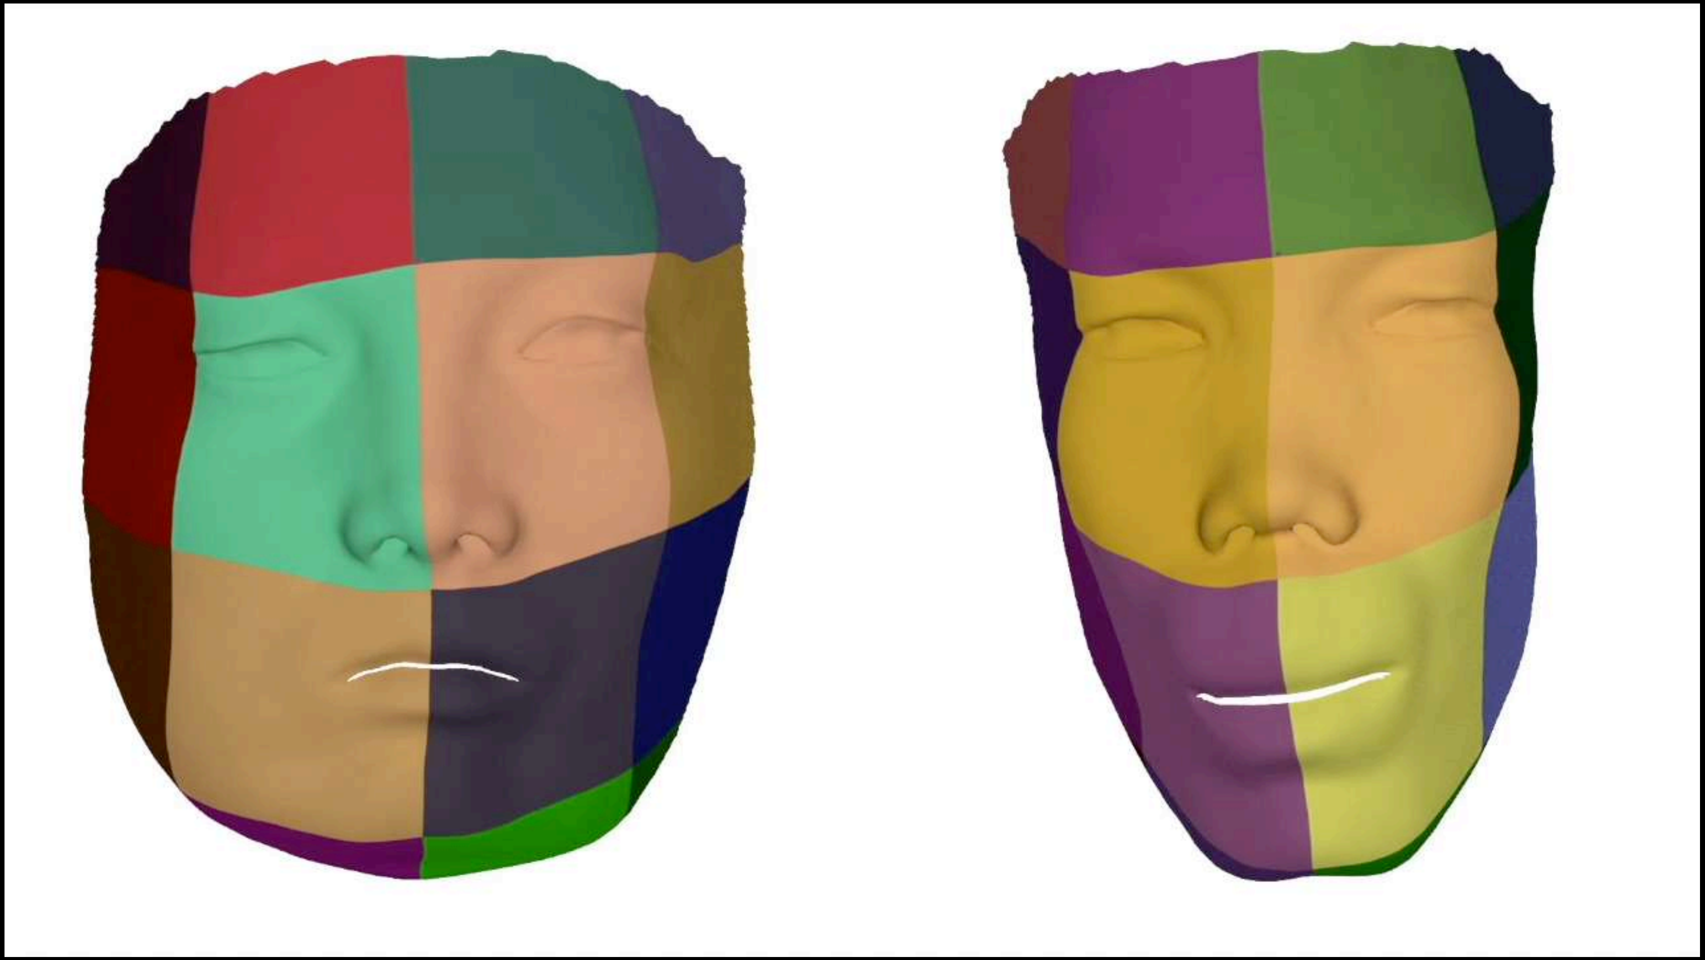

Shamai, Slossberg, K. 2018/2019 Synthesizing photometries and geometries with GANs &


Thank you for your attention