AI-Driven Human and Content Digitization

Hao Li

INSTITUTE FOR CREATIVE TECHNOLOGIES

Interaction With 3D Avatars

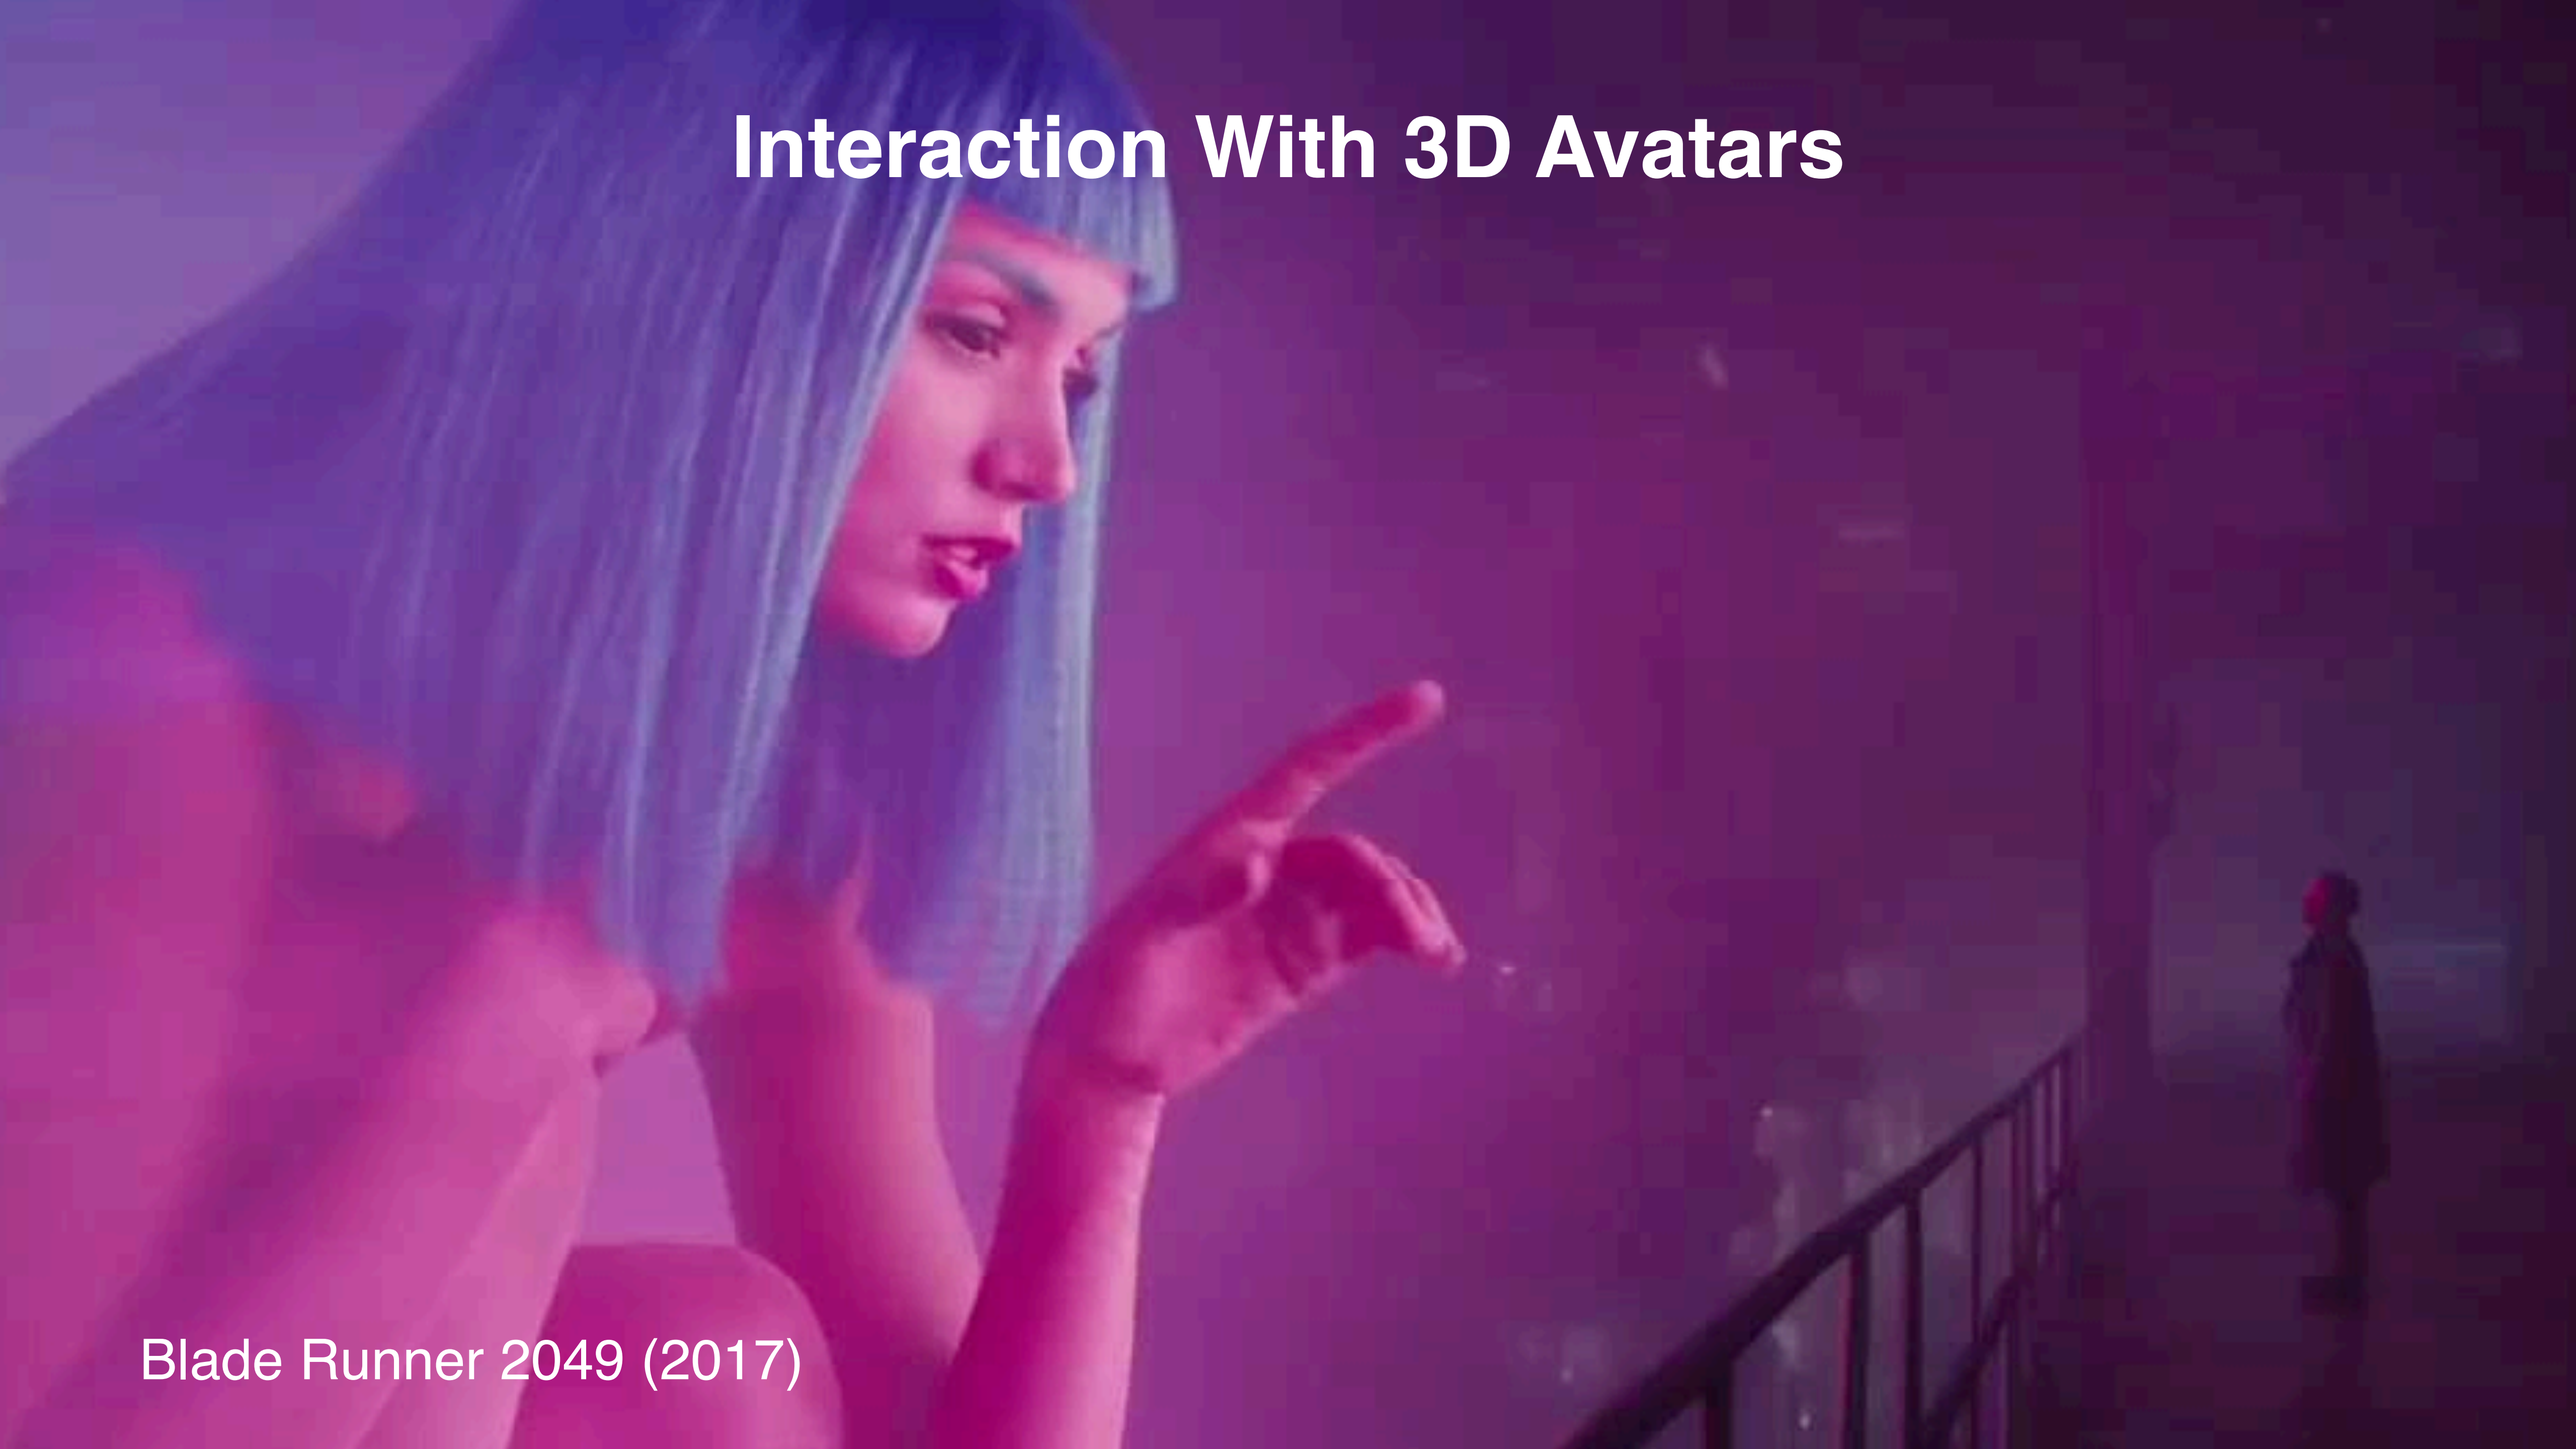

Blade Runner 2049 (2017)

Virtual Assistants / AI Chatbots

Can it get it delivered by today?

Sure, our deliveries are usually made between 7:00 AM and 6:00PM

OMG! Nice!

Virtual Shopping

CAN AI-POWERED CGI CREATIONS TAKE OVER THE INFLUENCER SPACE?

Or better yet, how can human influencers keep up?

MARIA BOBILA • APR 20, 2018

Espionage

Associated Press

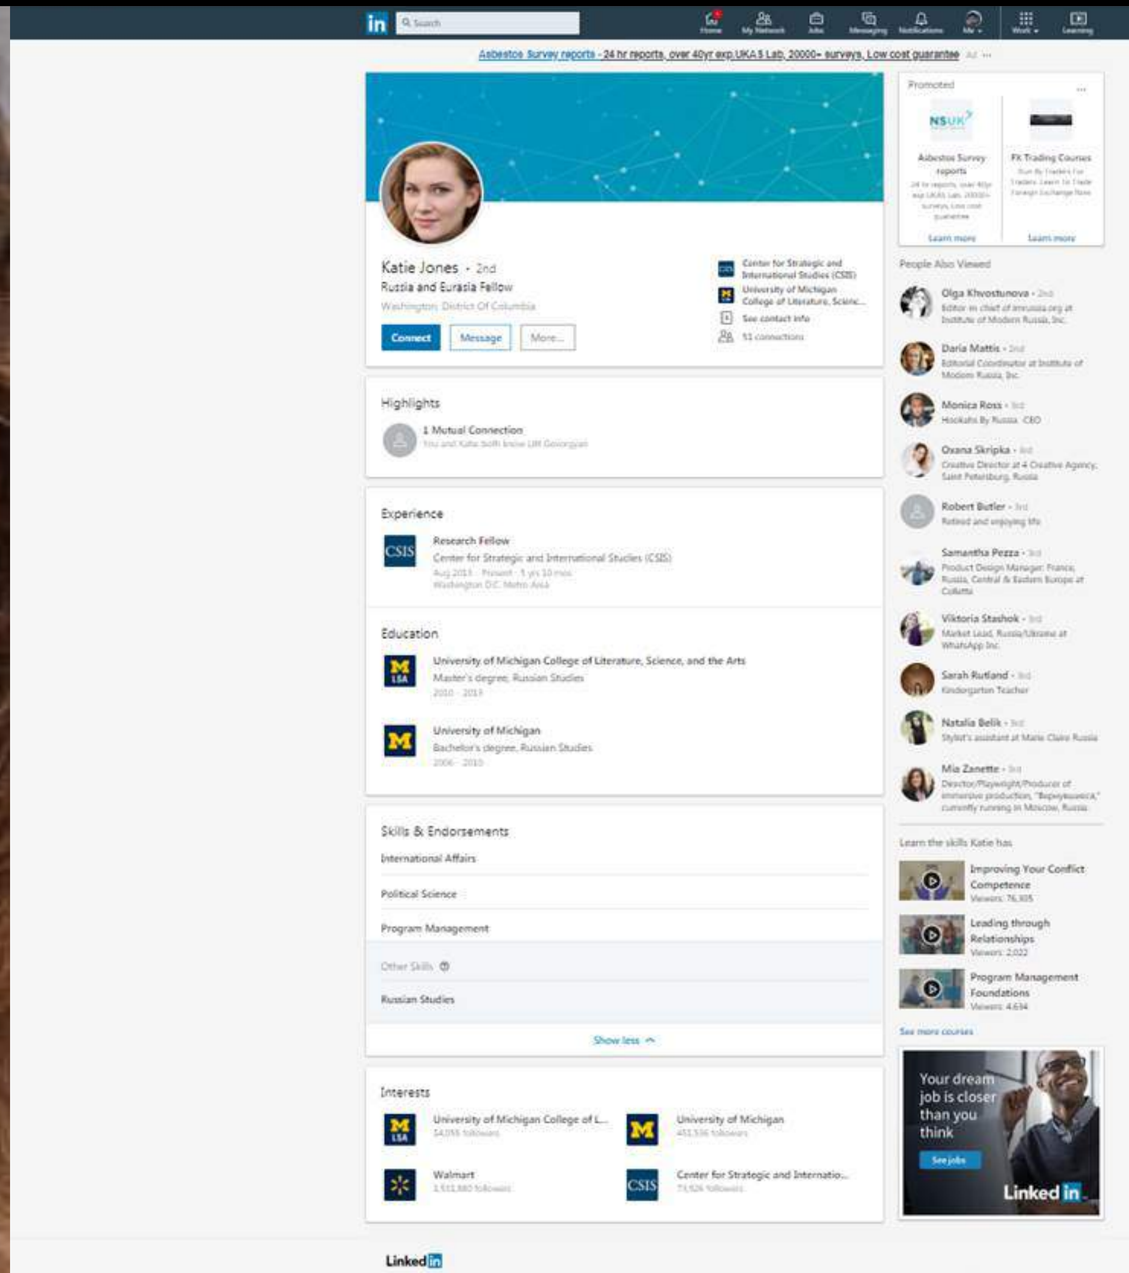

Microsoft 2019

Digital Humans

VFX Production

Weta Digital (2014)

Game Production

Epic Games / Cubic Motion / 3Lateral (2018)

Digitizing Faces

High-End Capture

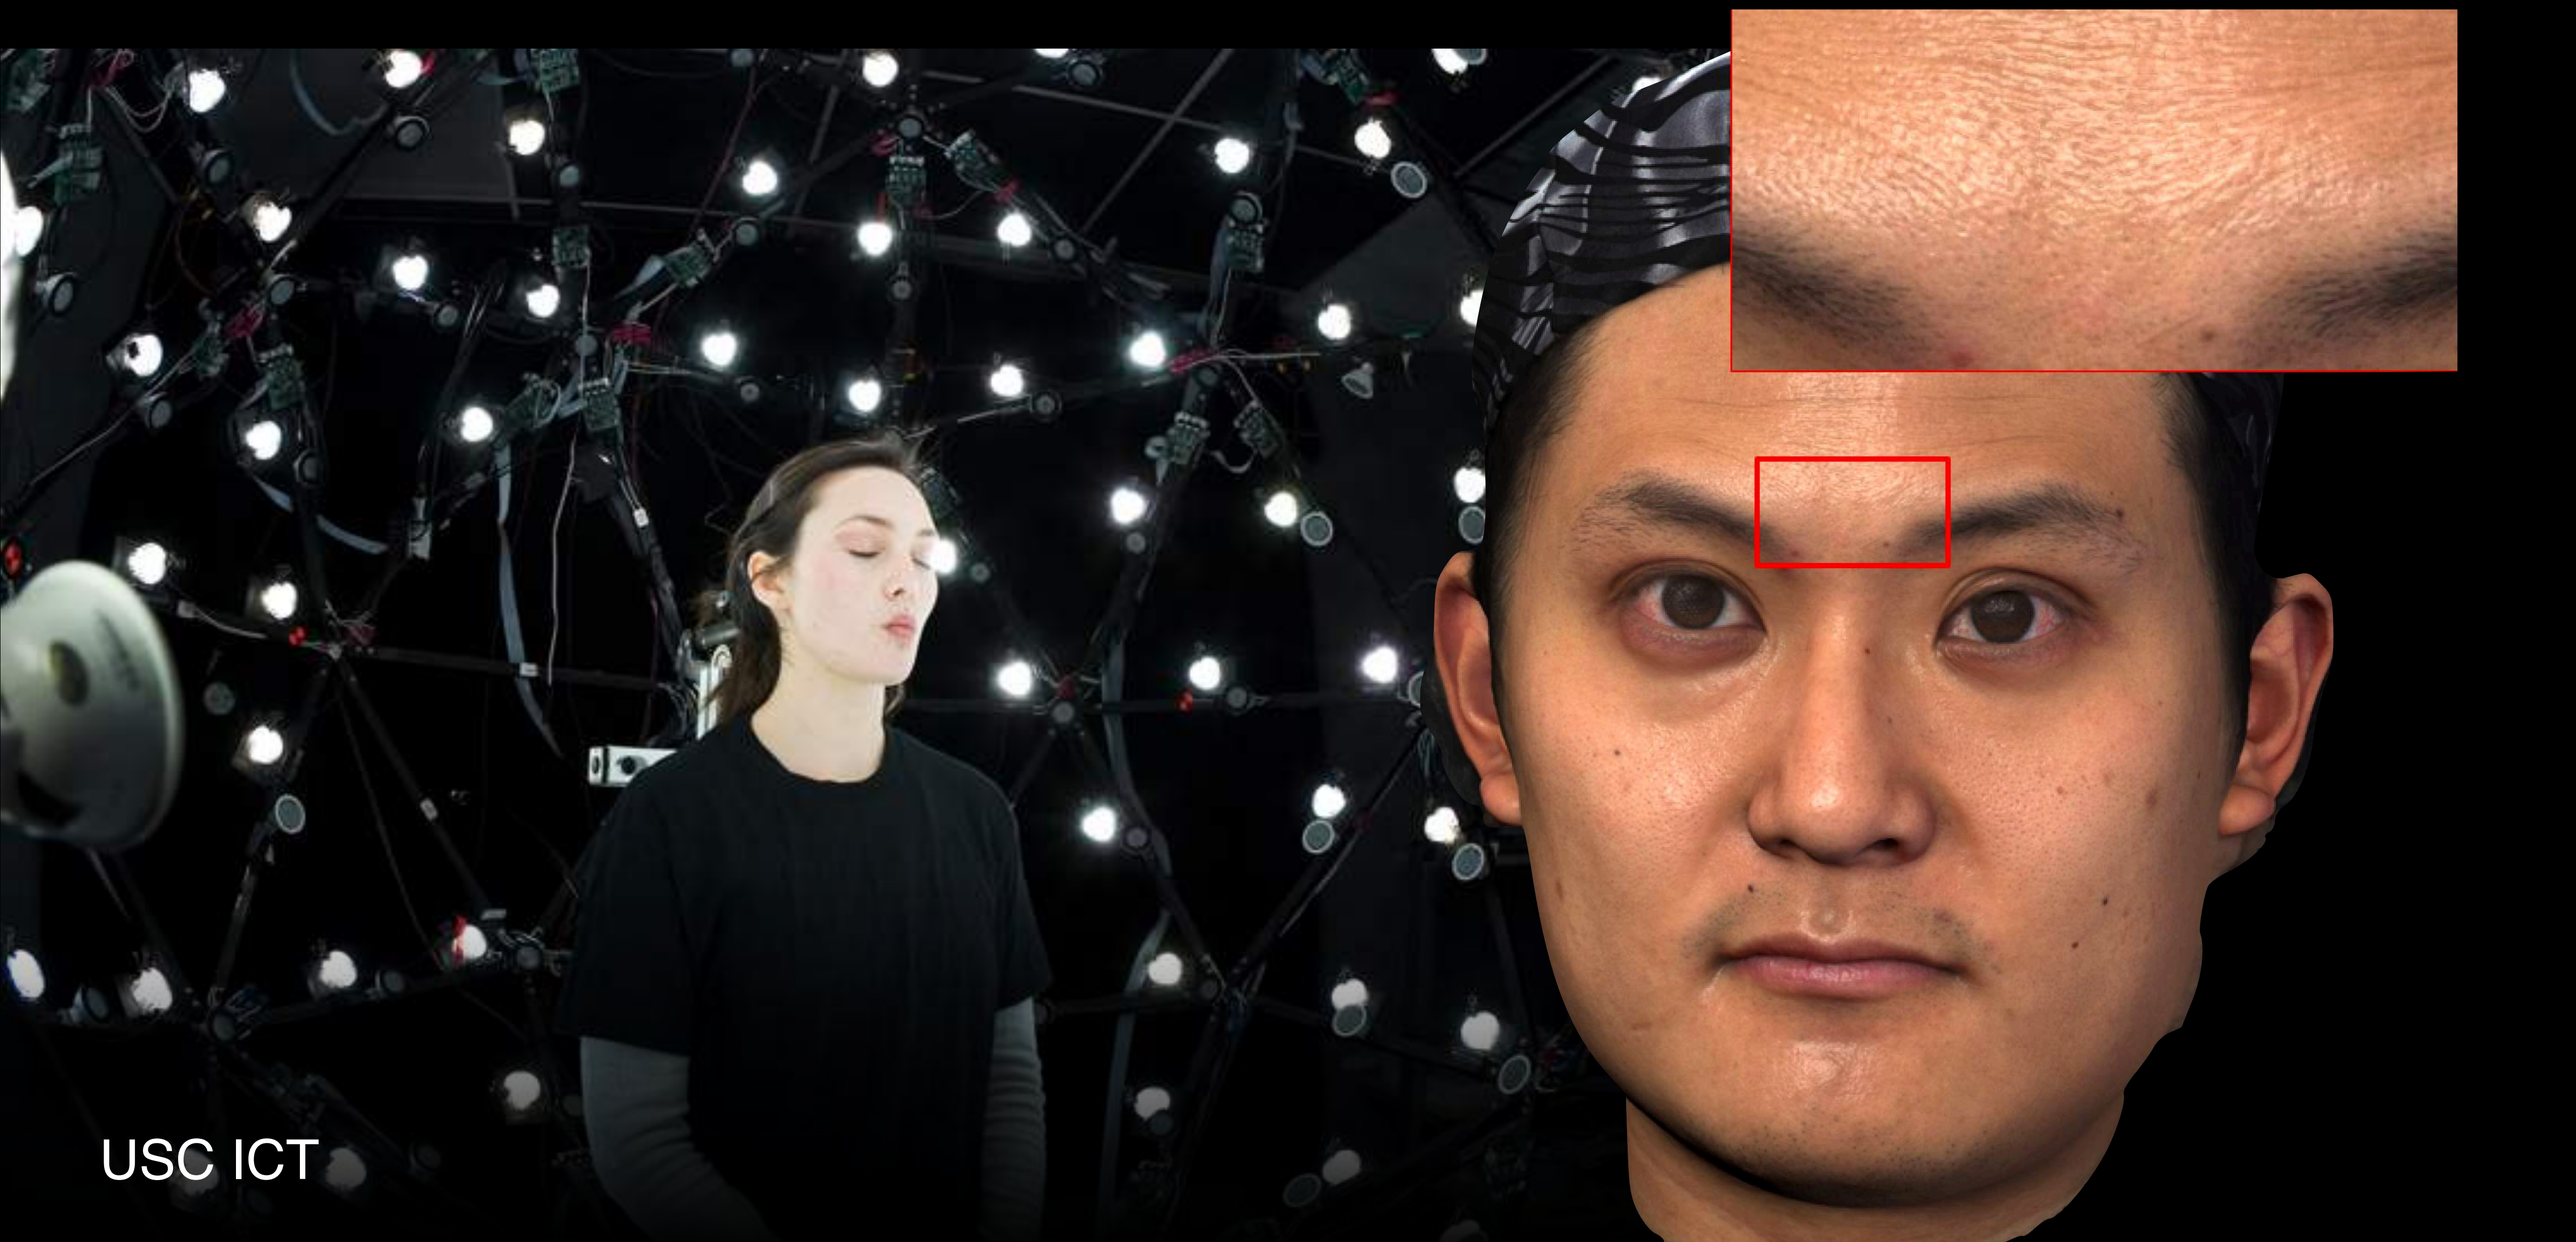

USC ICT

Full Geometry and Reflectance Inference

input image

output textured 3D face
(illumination 1)

output textured 3D face
(illumination 2)

Yamaguchi et al. 2018

Unconstrained Pictures

Deep Face Normalization

Overview

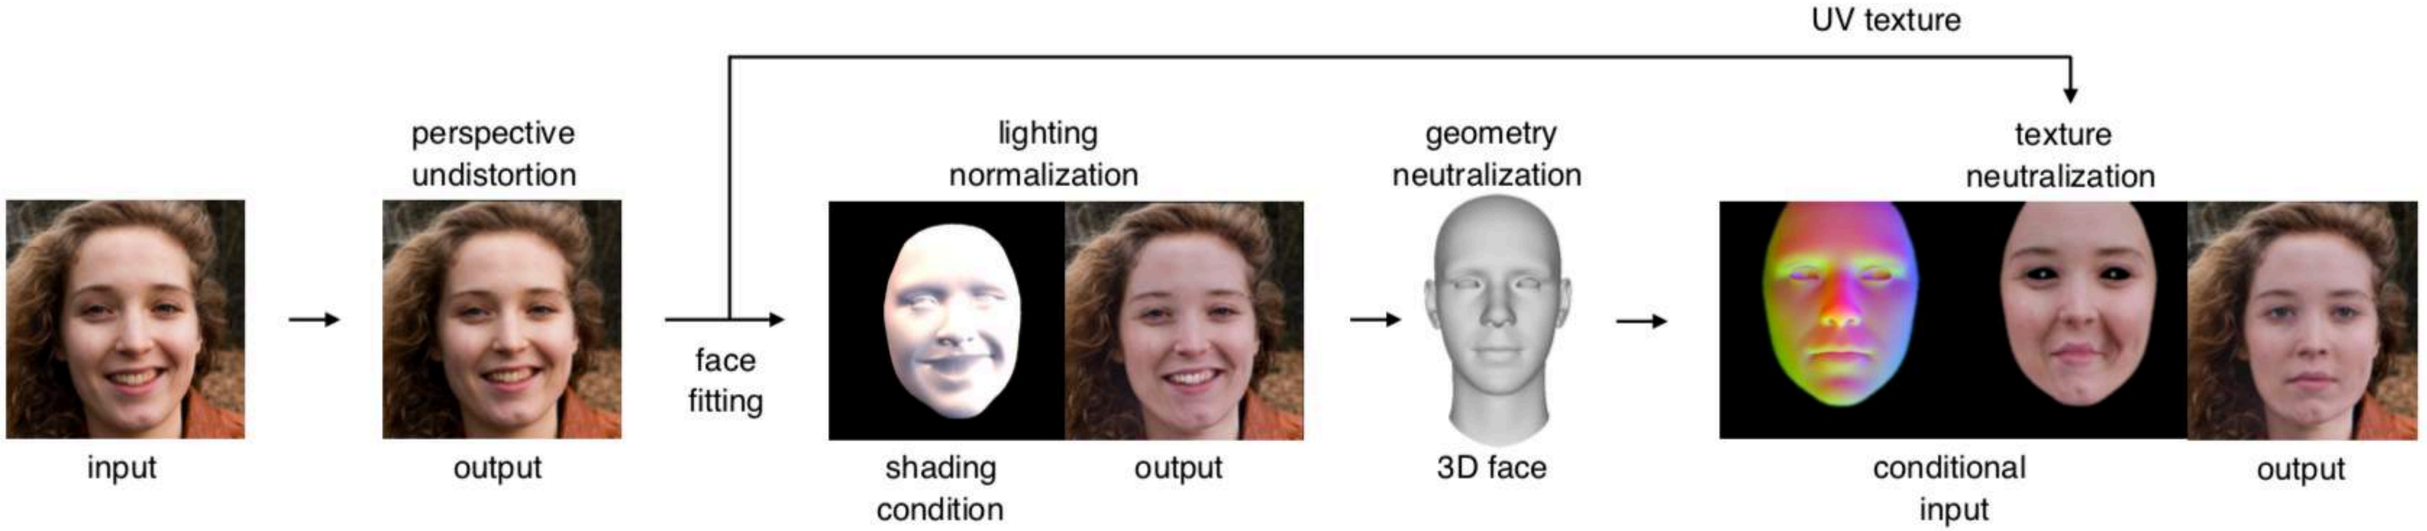
UV texture
perspective undistortion
lighting normalization
geometry neutralization
texture neutralization
face fitting
input
output
shading condition
output
3D face
conditional input
output

Illumination Normalization

Expression Normalization

Comparison

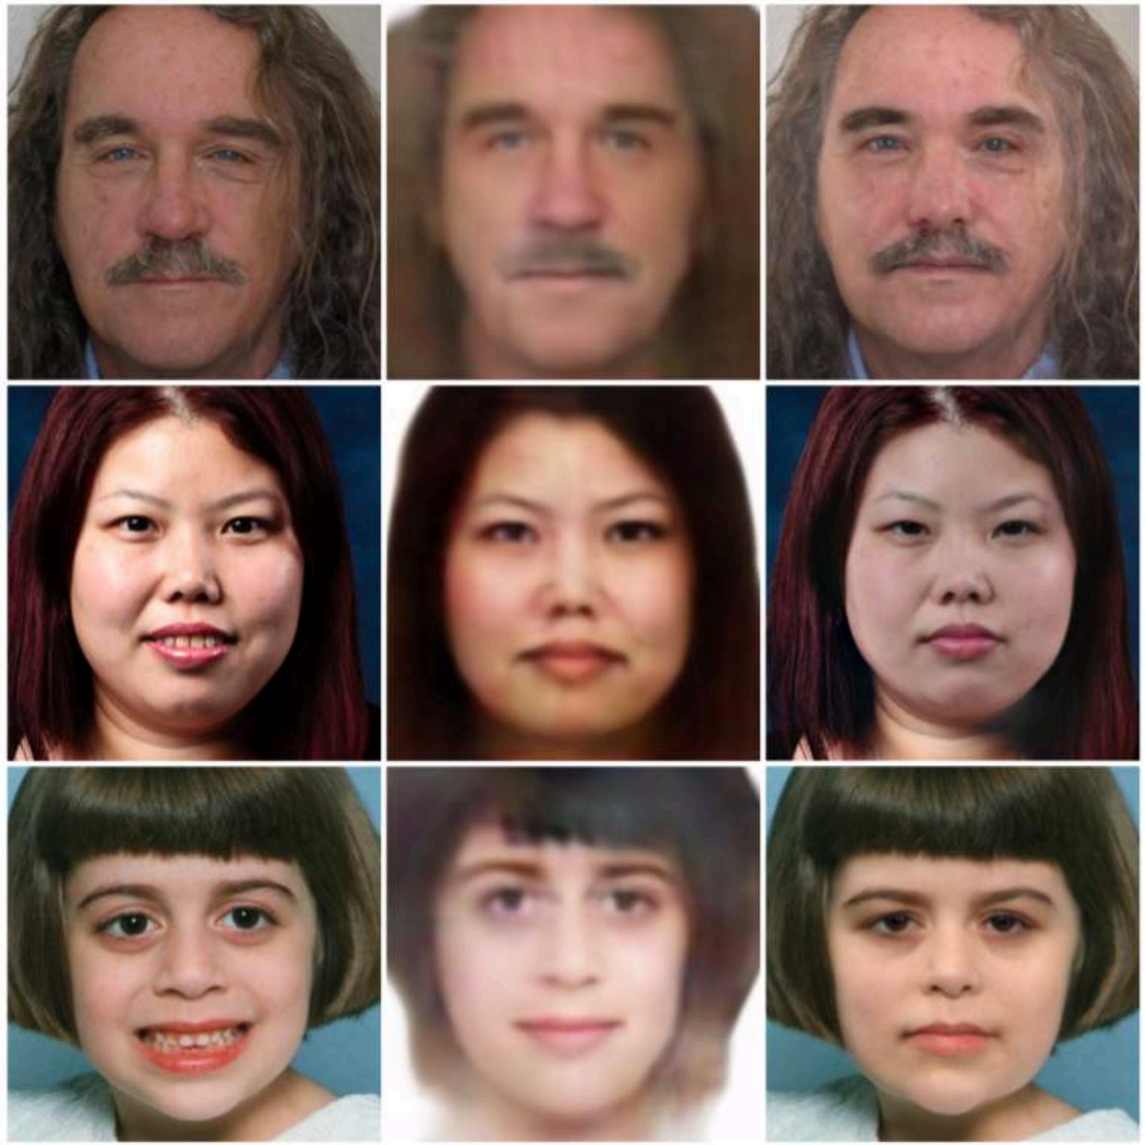

input [Cole et al.] ours

Deep Face Normalization

Live Demo [later!]

Pinscreen 2019

Pinscreen App

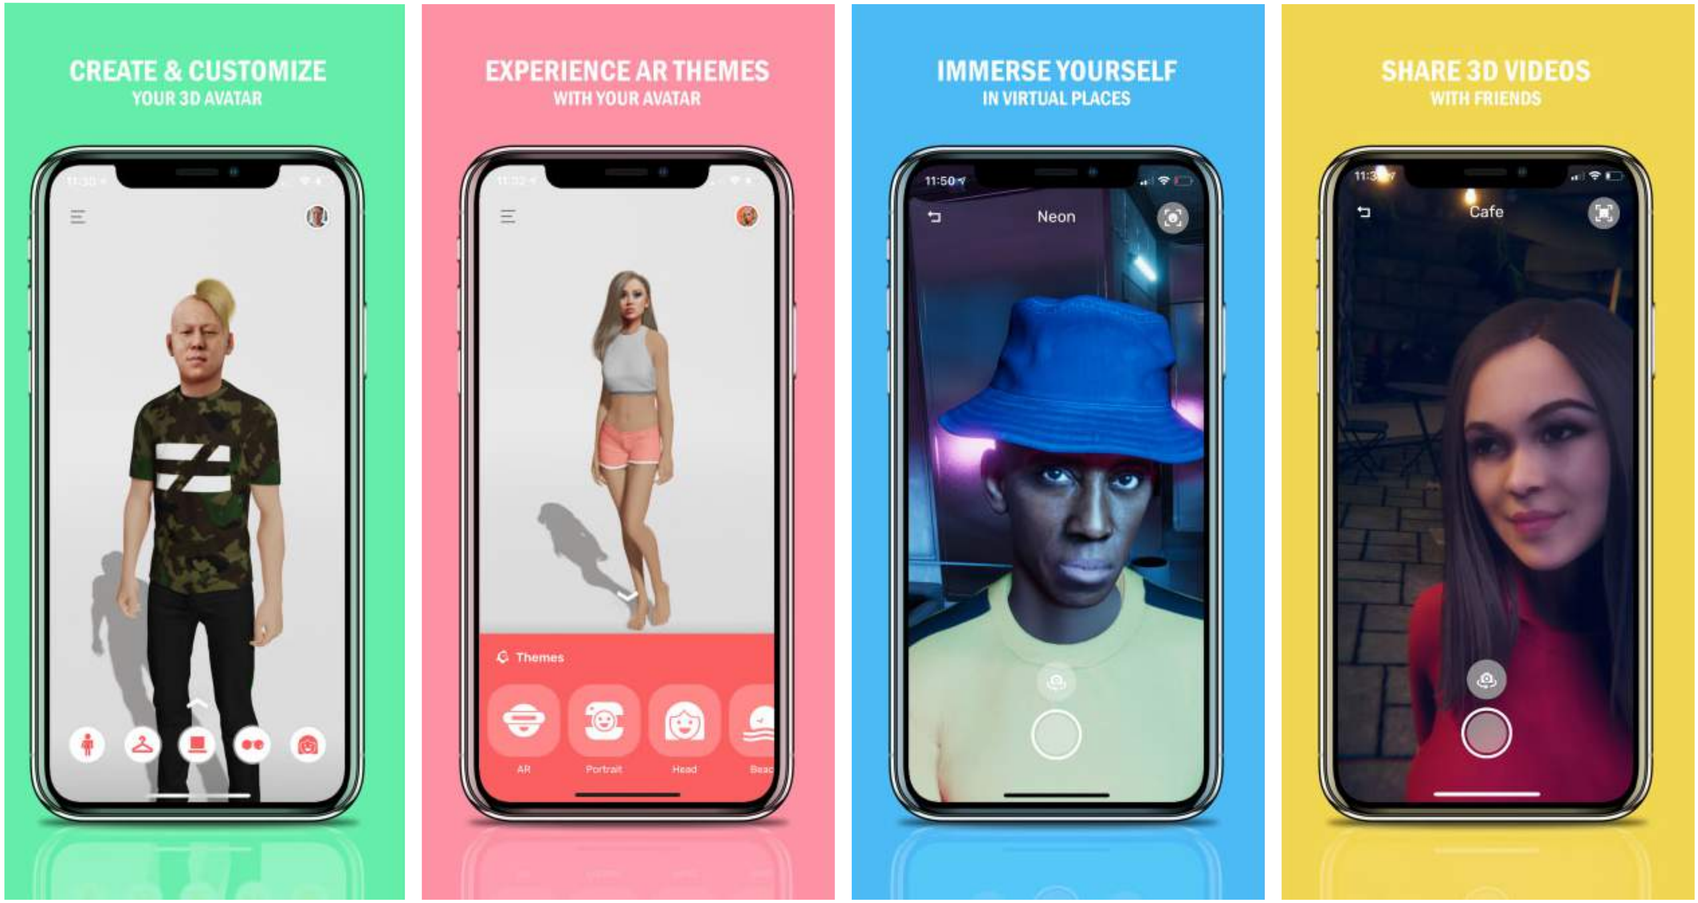

Download on the App Store

www.pinscreen.com

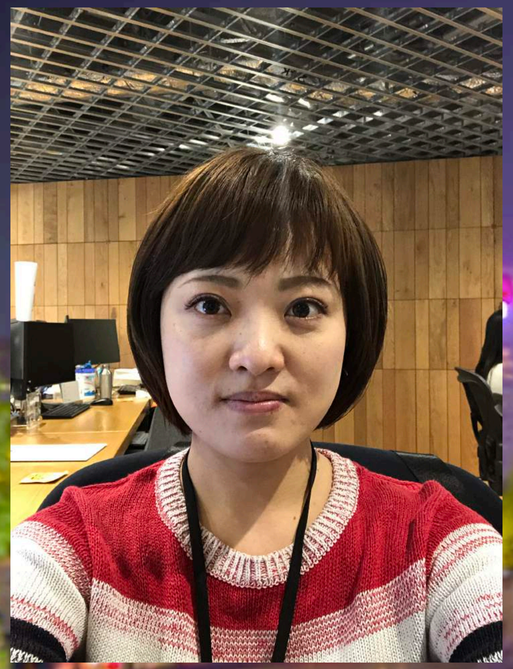

Pinscreen 2019

Facial Performances

Pinscreen 2018

Deep Learning-Based Face Synthesis

source video

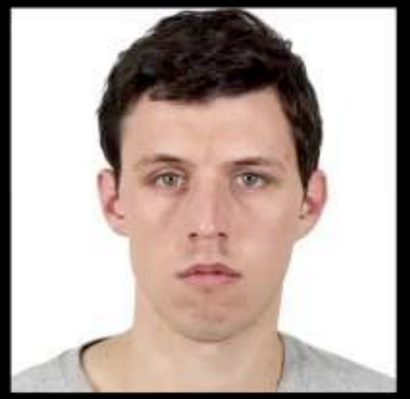

input image

subject A

subject B

subject C

our result

Nagano et al. 2018

Nagano et al. 2018

6/13/2019

US Congress holds hearing on Deepfakes

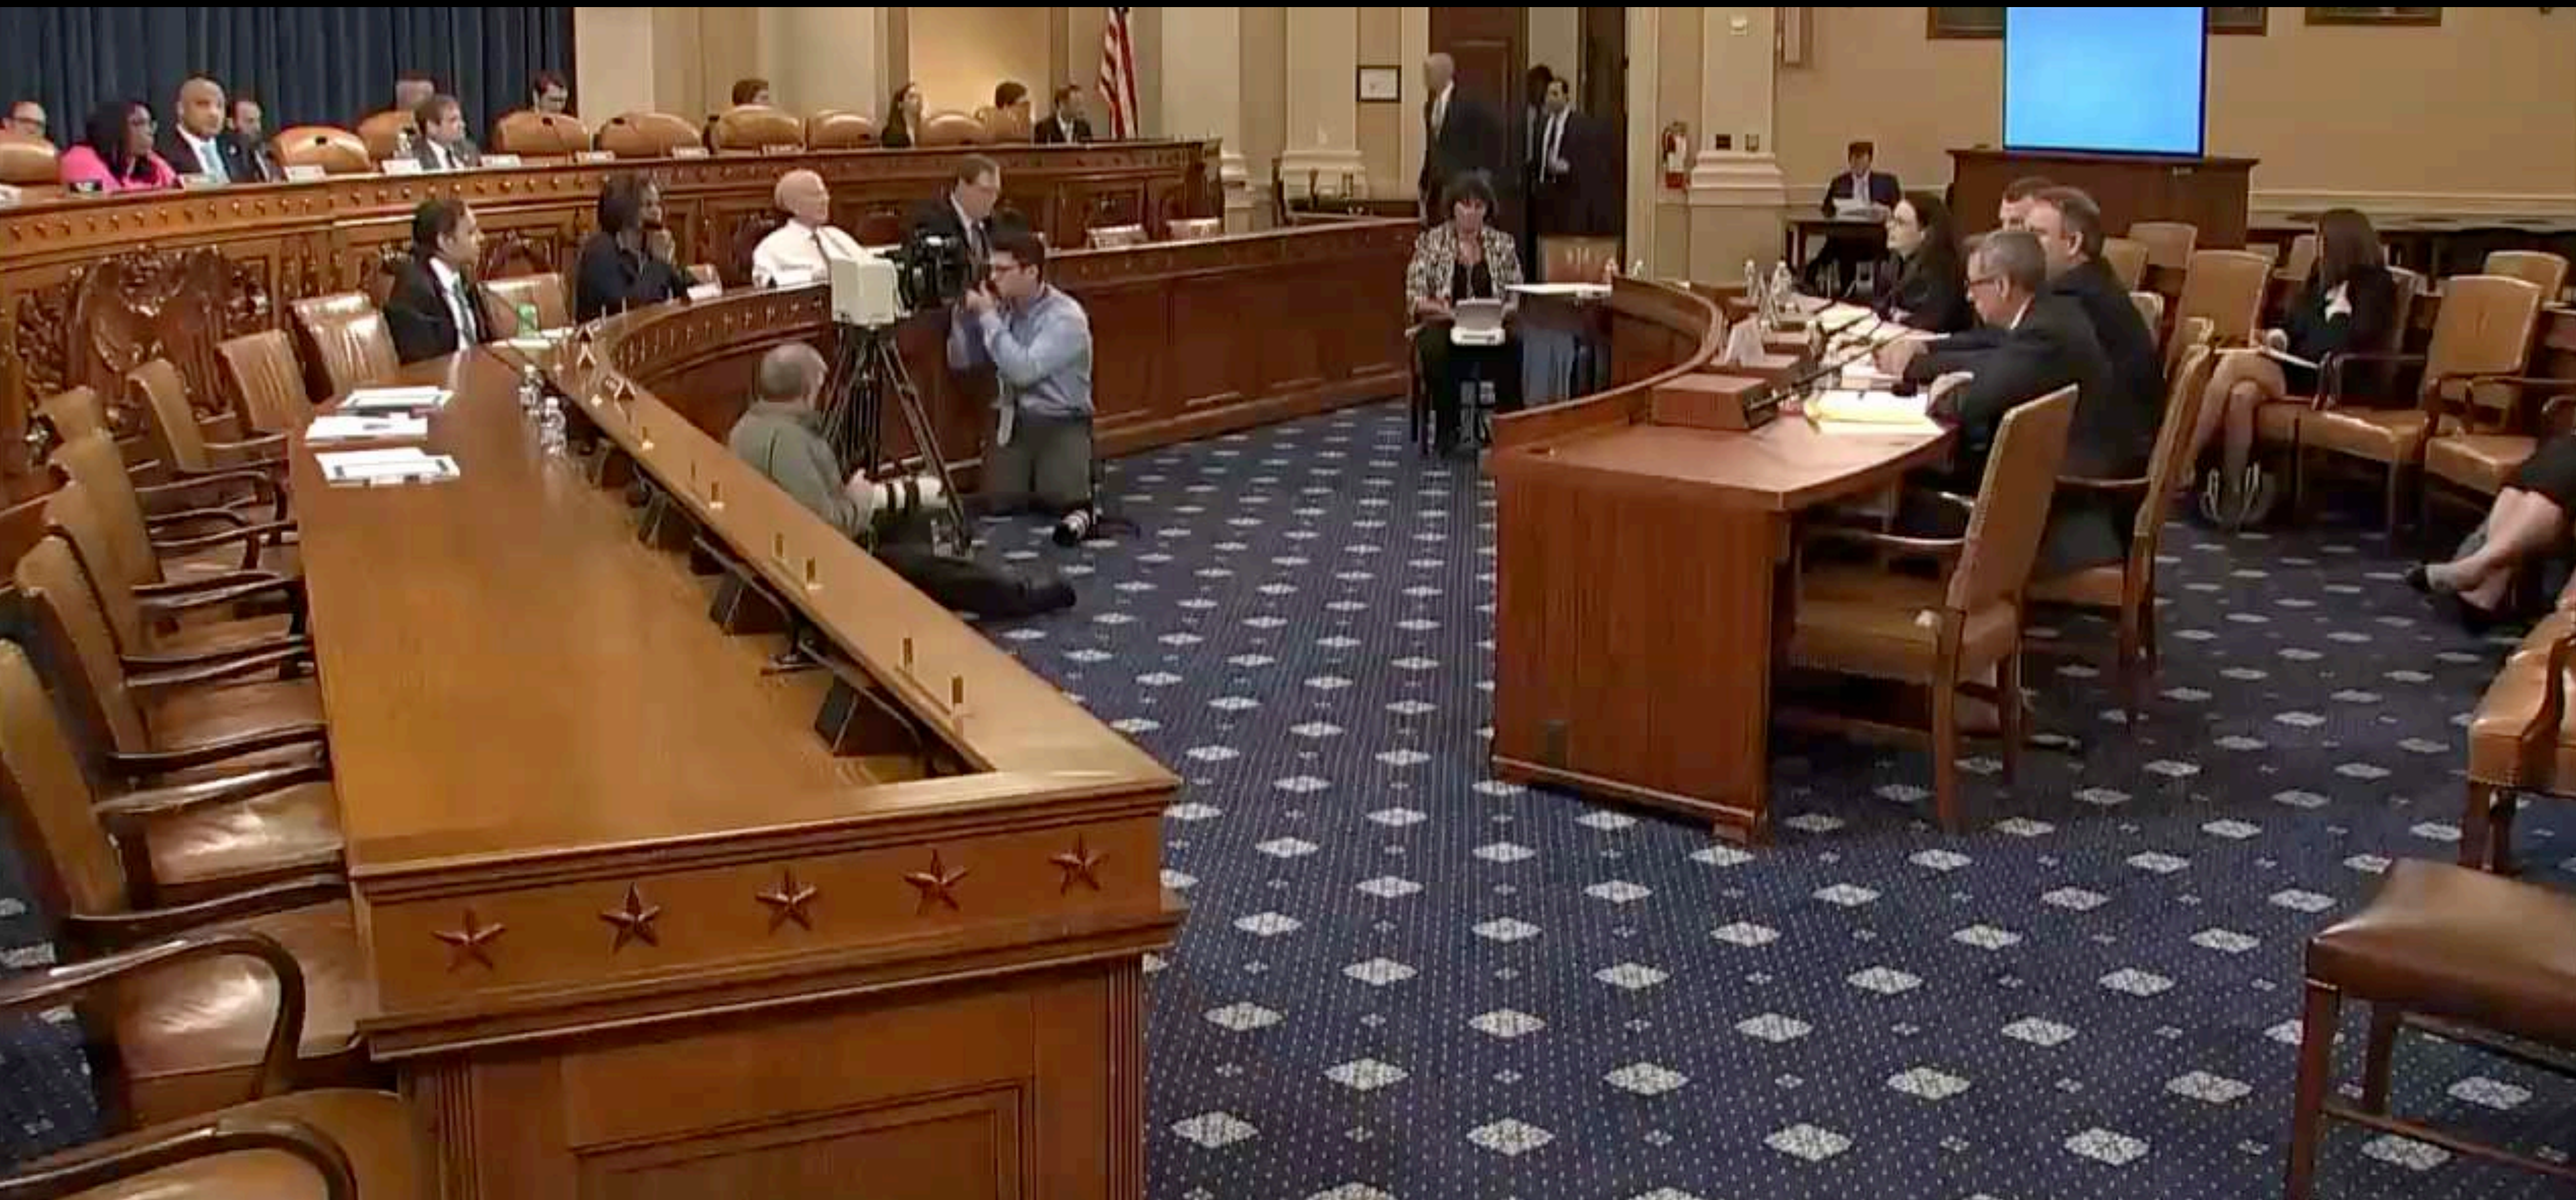

Deep Fake: FaceSwap

Deep Fake: Puppet Master

Hair Capture

Reference photo

Five-strand Dutch braid

Our result

Unstructured Data

Direct Hair Regression

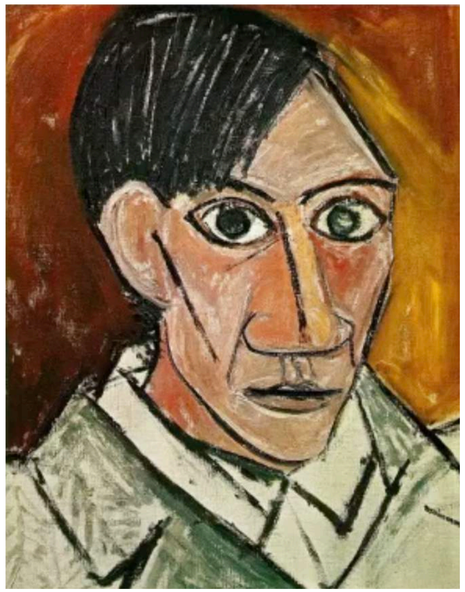

Input

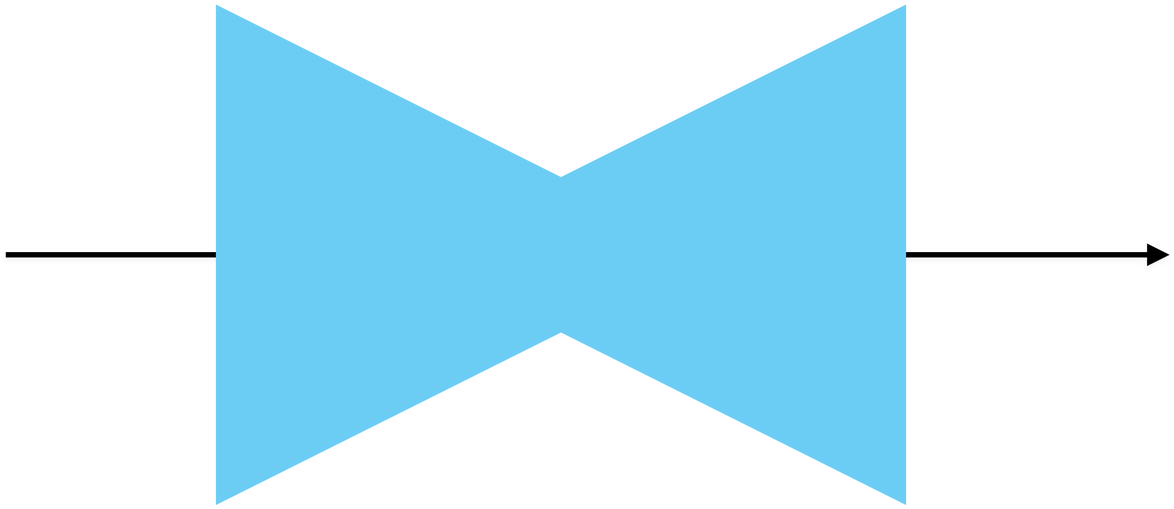

Volumetric Hair Representation

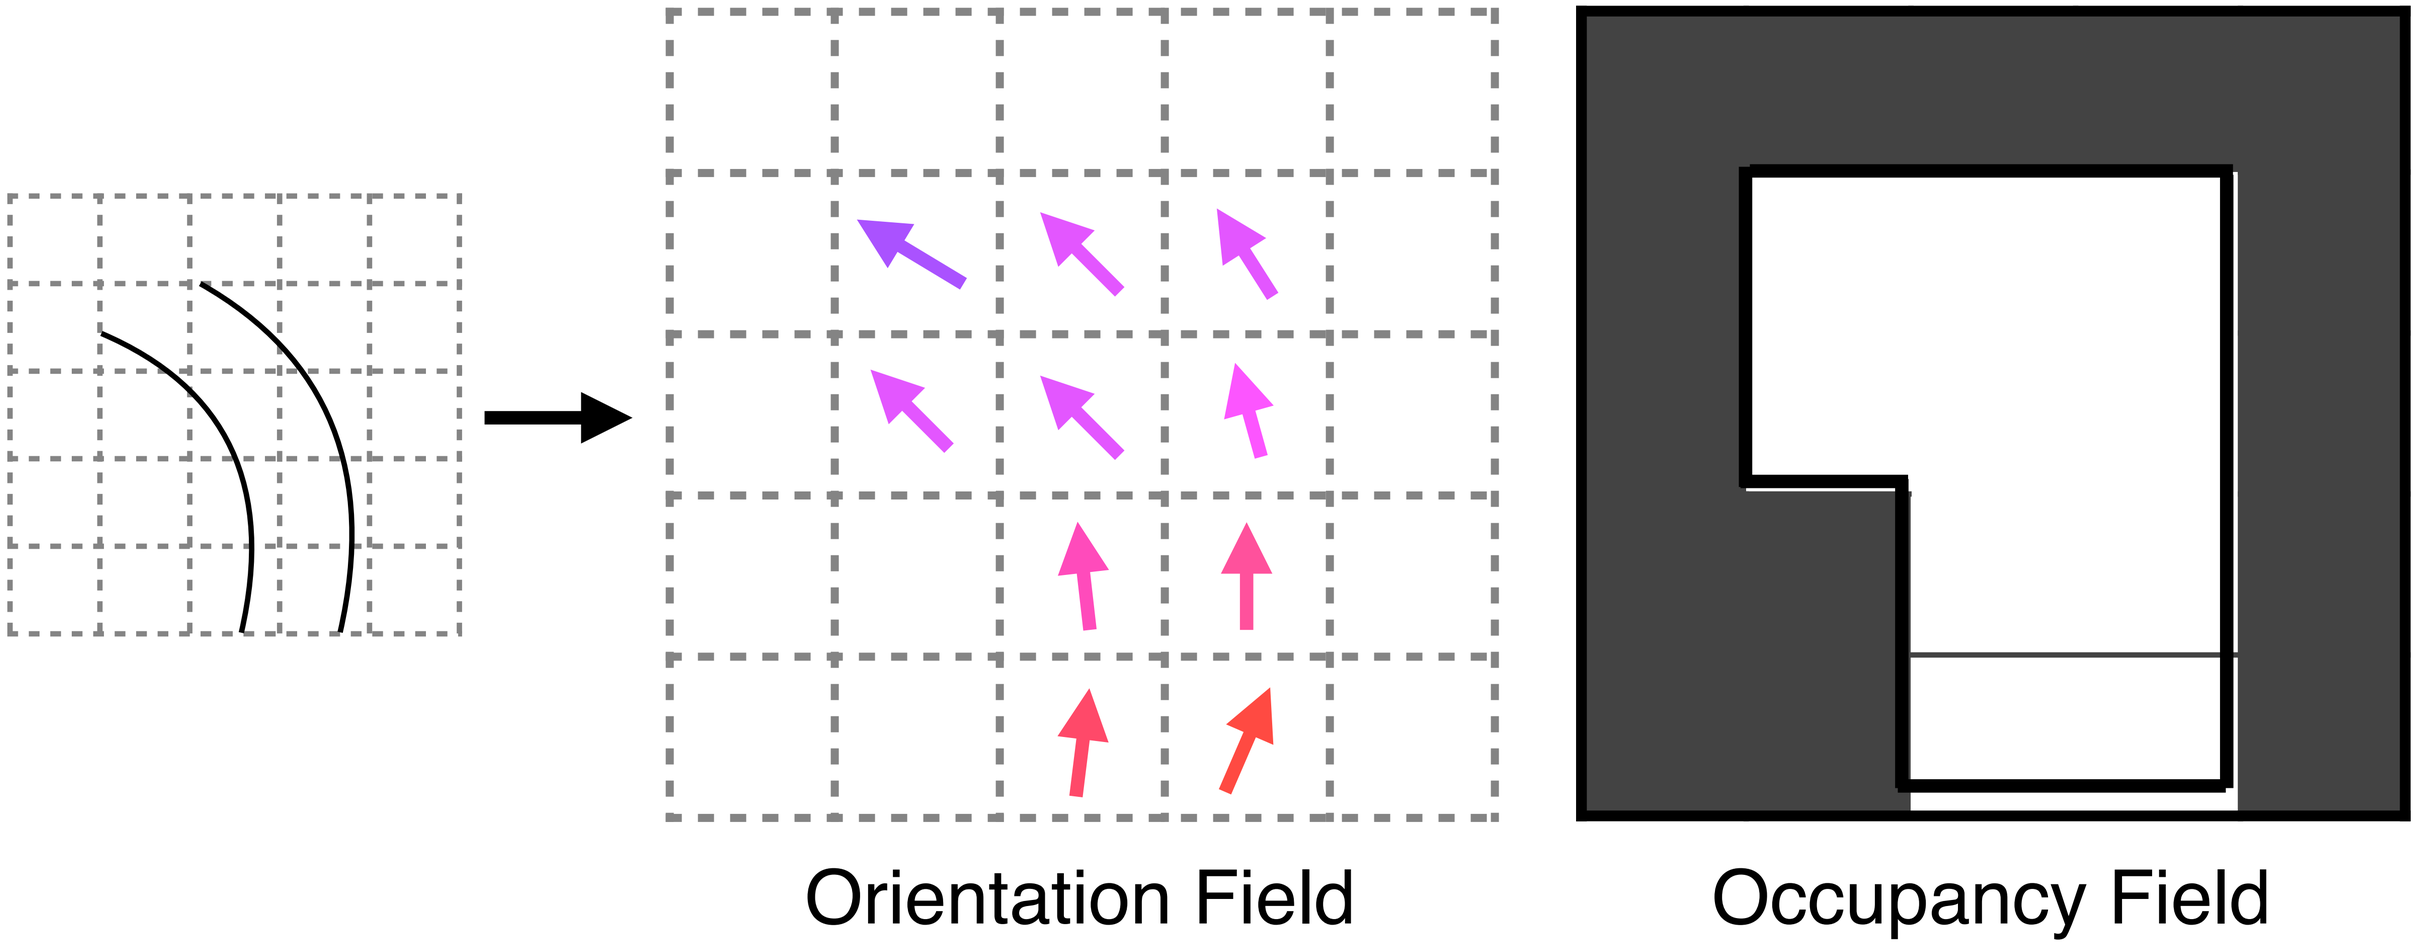

Orientation Field

Occupancy Field

Volumetric Hair Representation

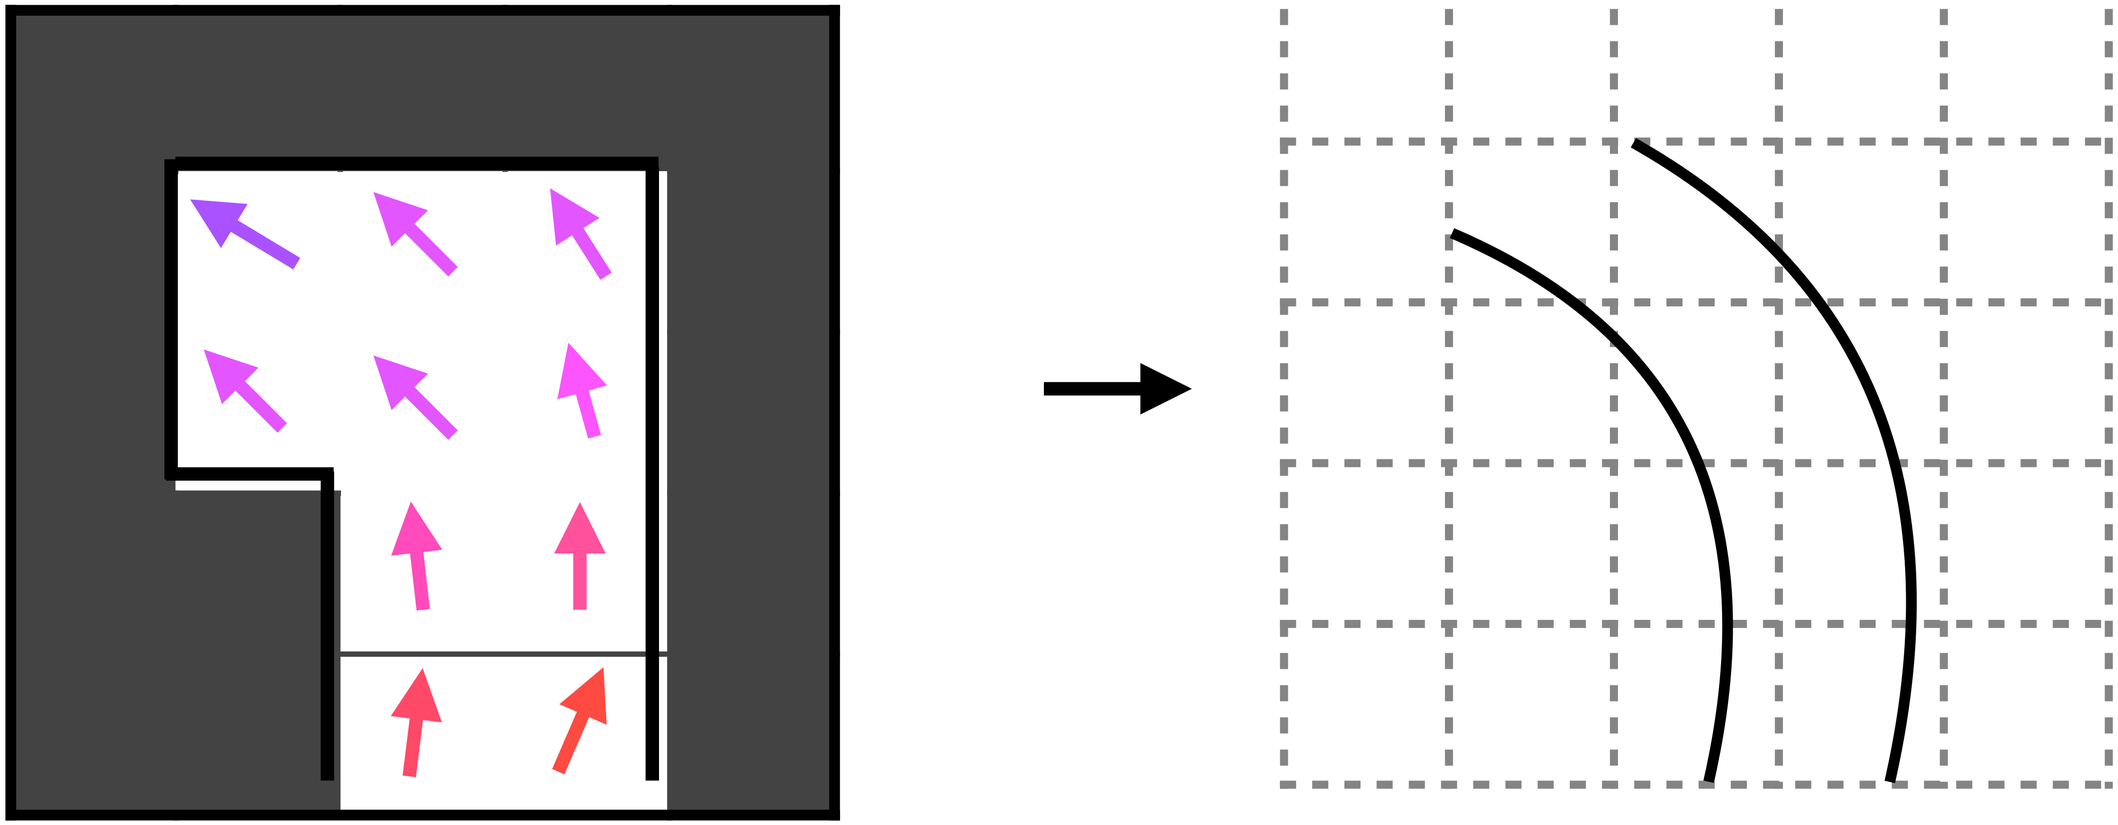

Self Reconstruction Check

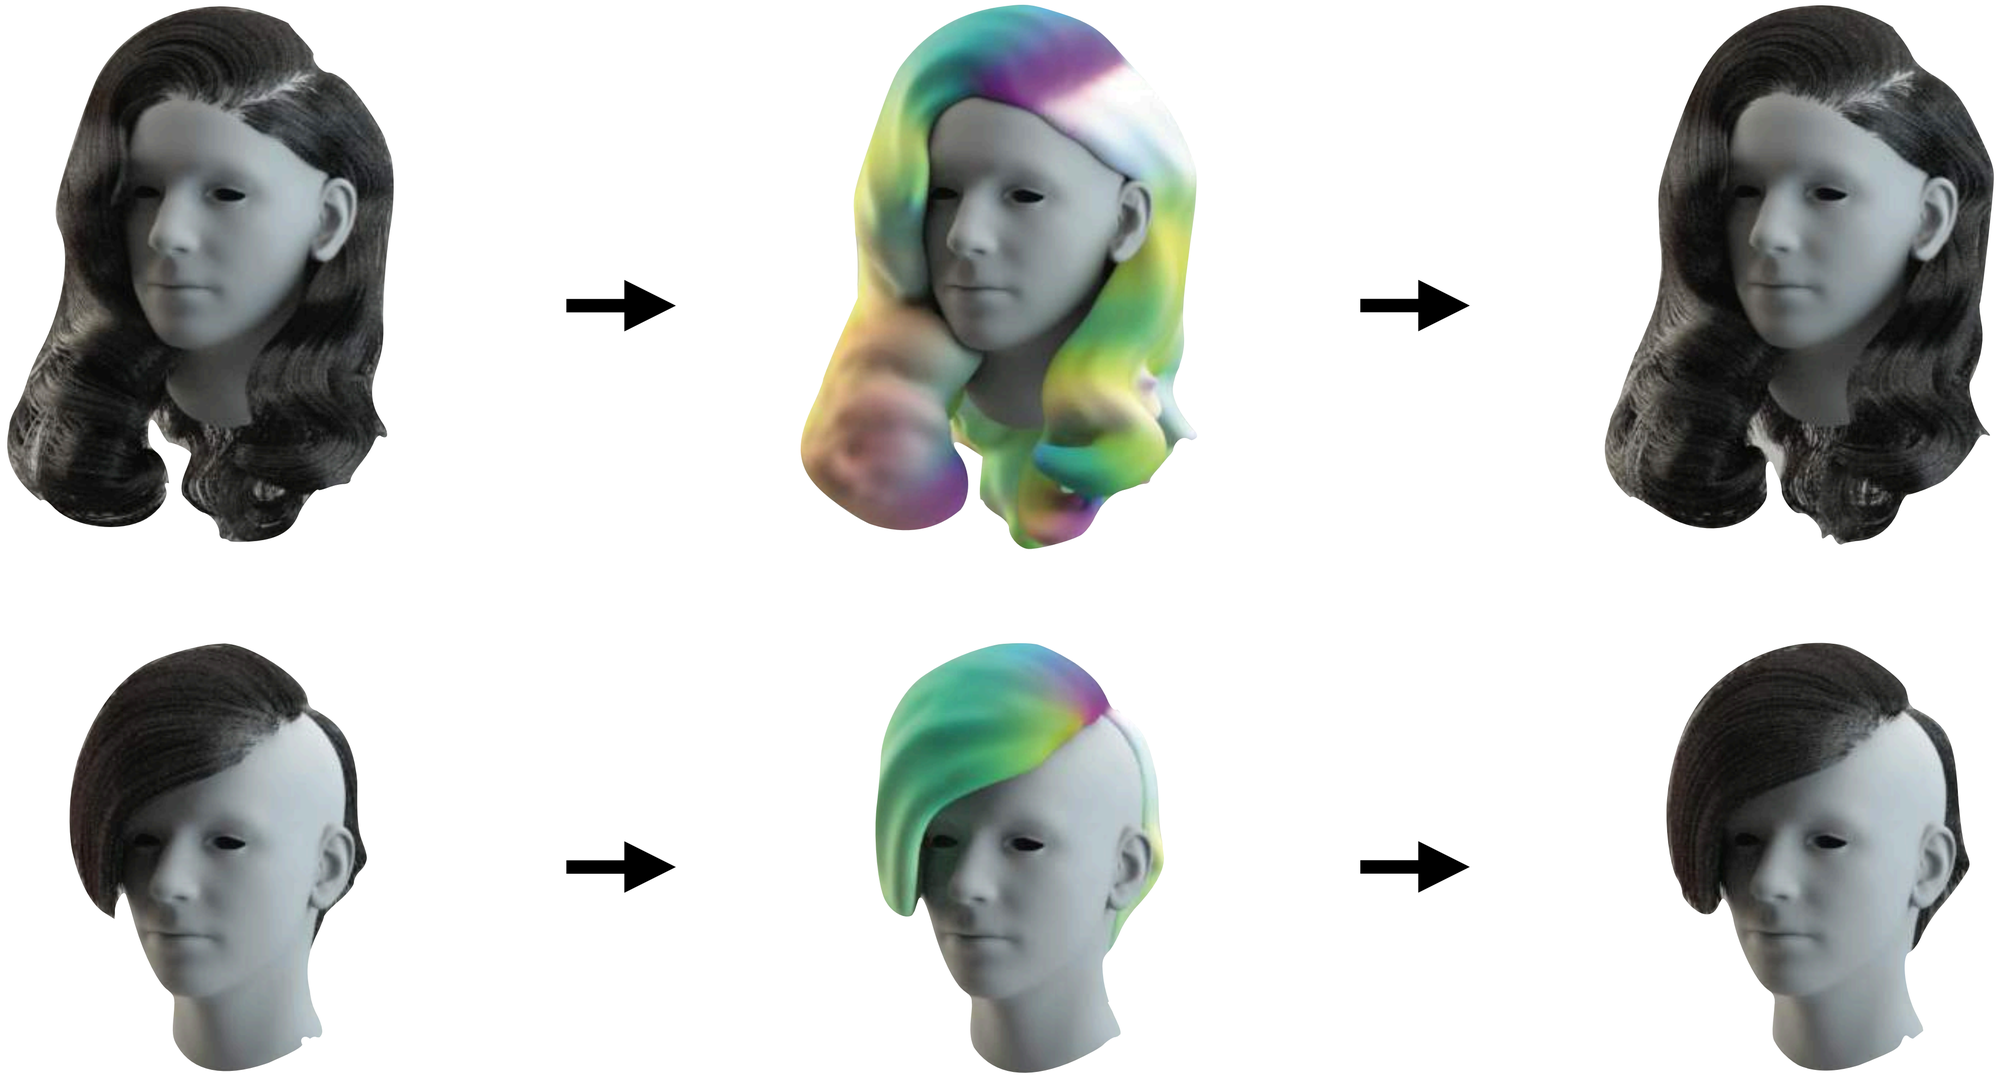

Training Data

816 portrait images [Hu et al. 2015]

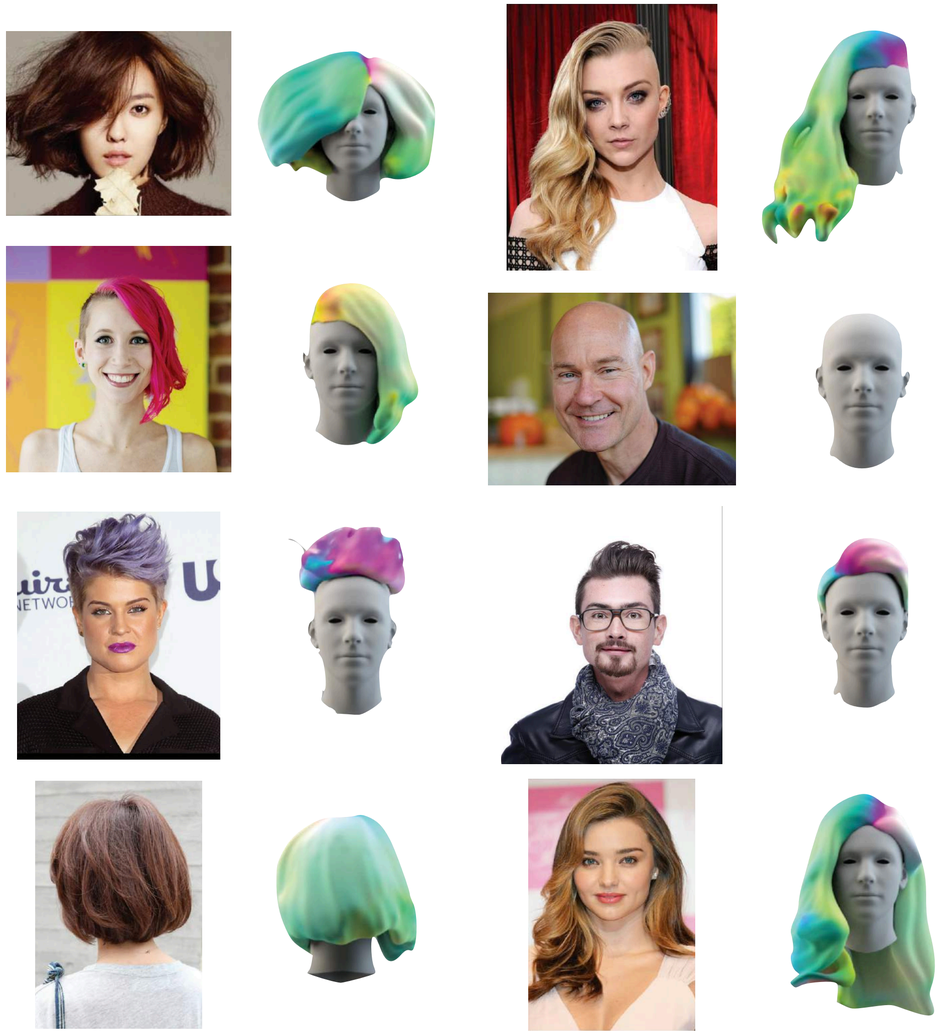

+

USC-HairSalon dataset

Deep Learning for Hair Modeling

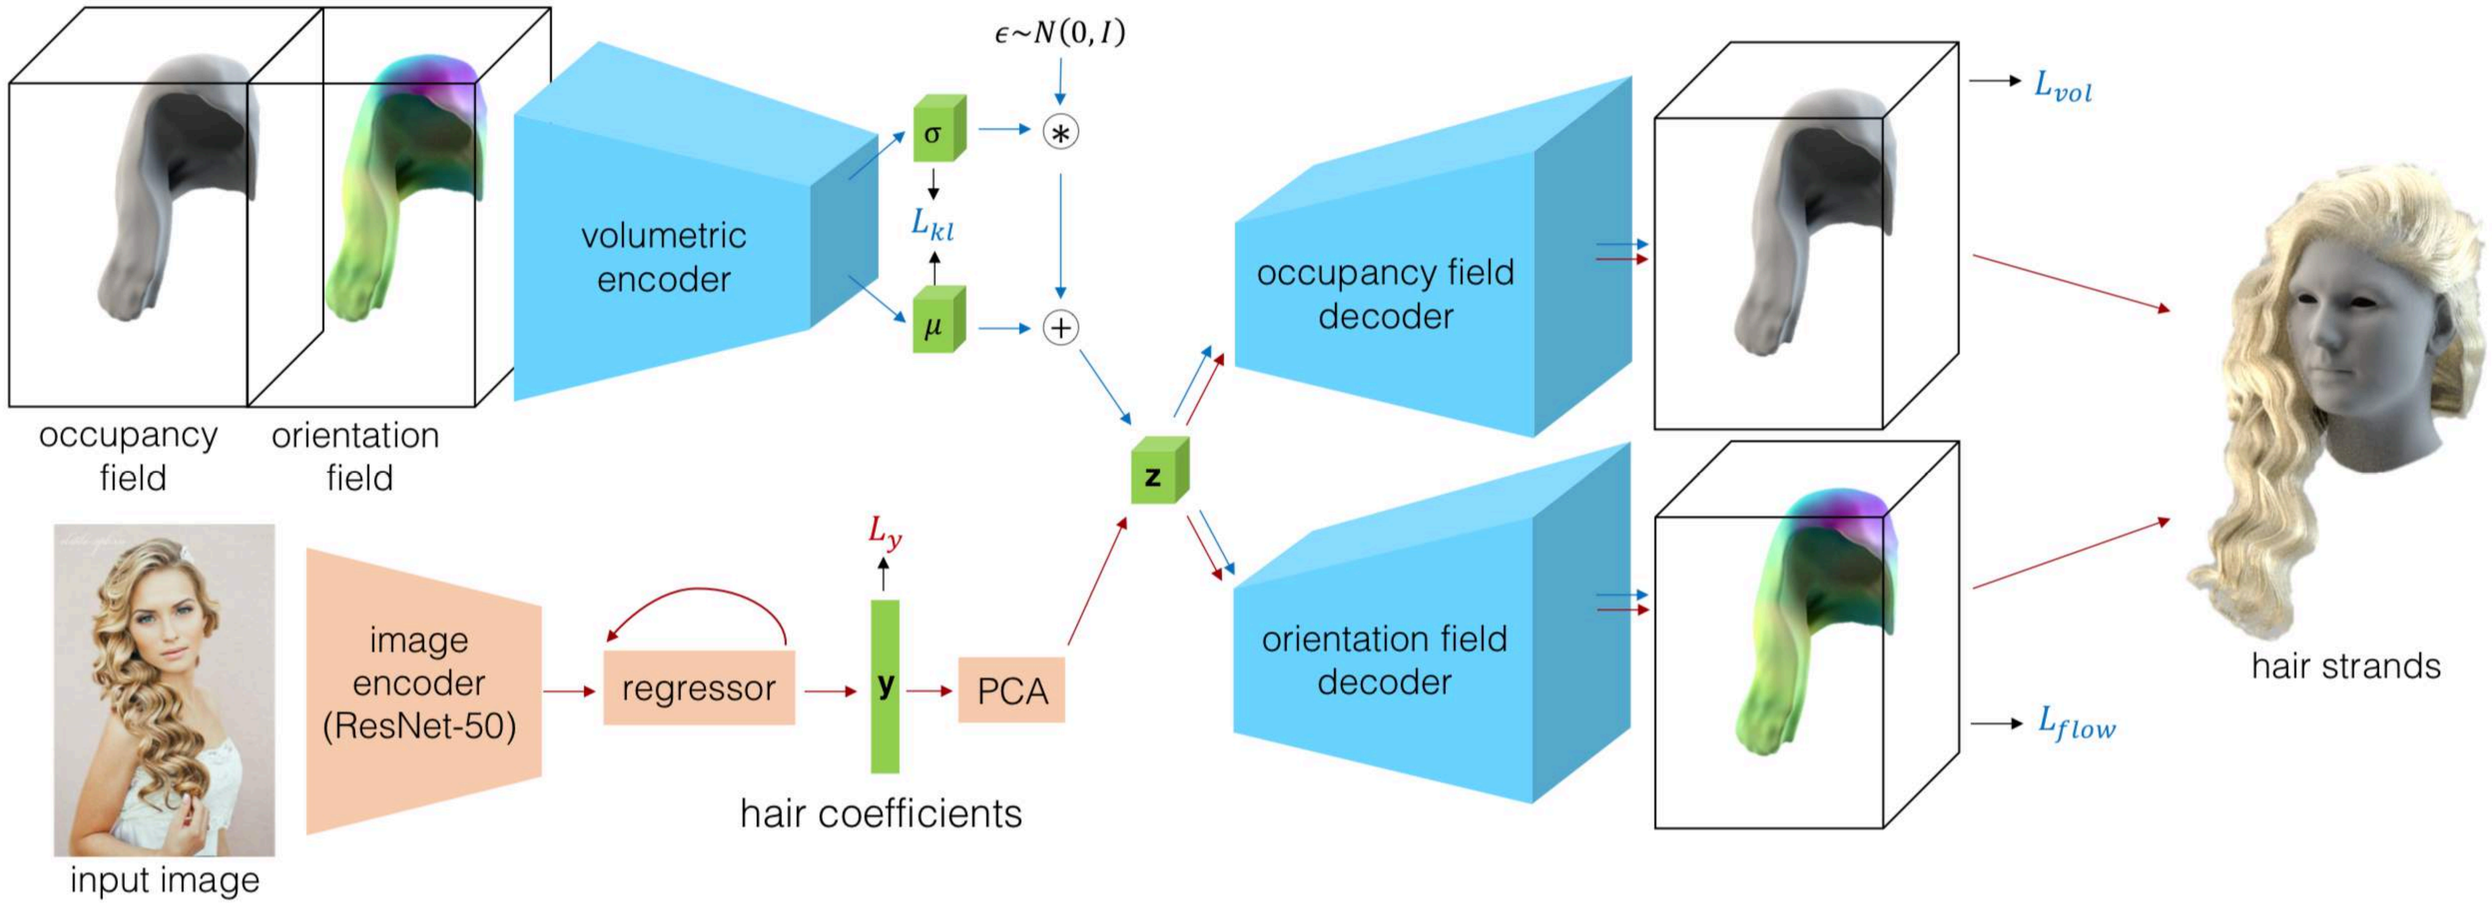

Saito et al. (2018)

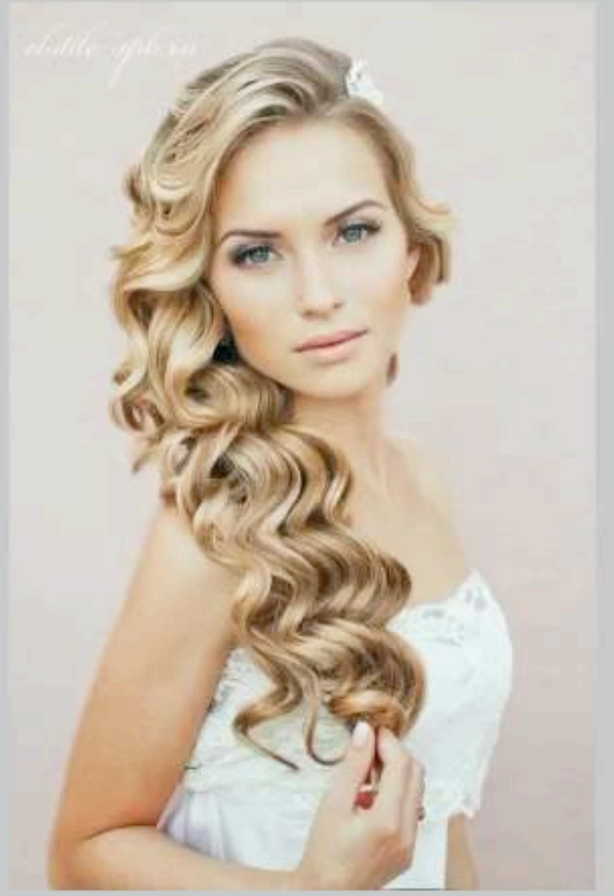

input image

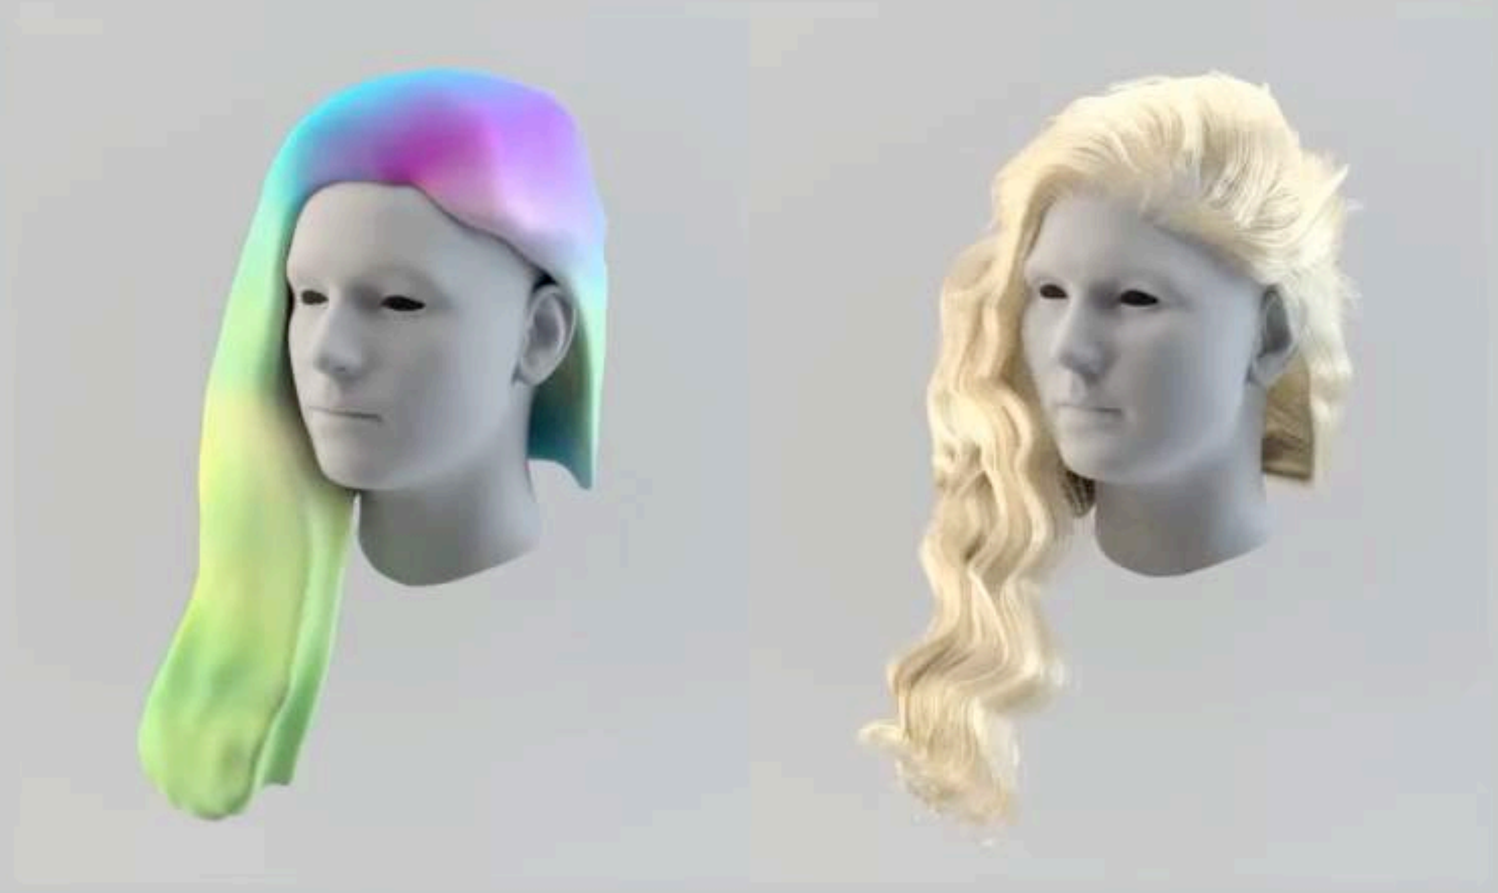

predicted volume

output strands

input image

predicted volume

output strands

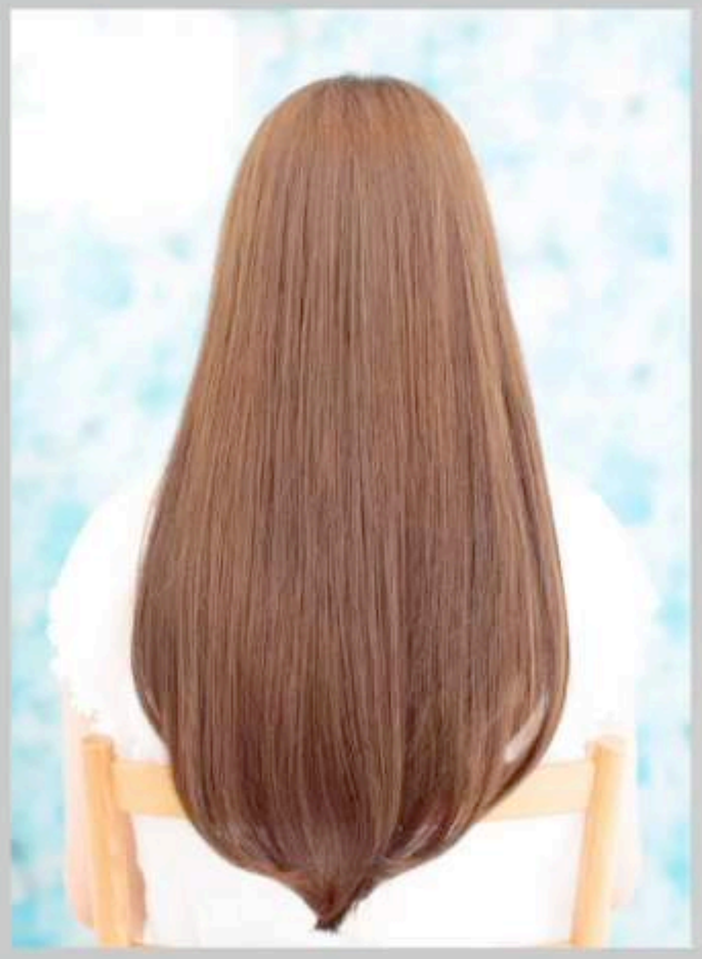

input image

predicted volume

output strands

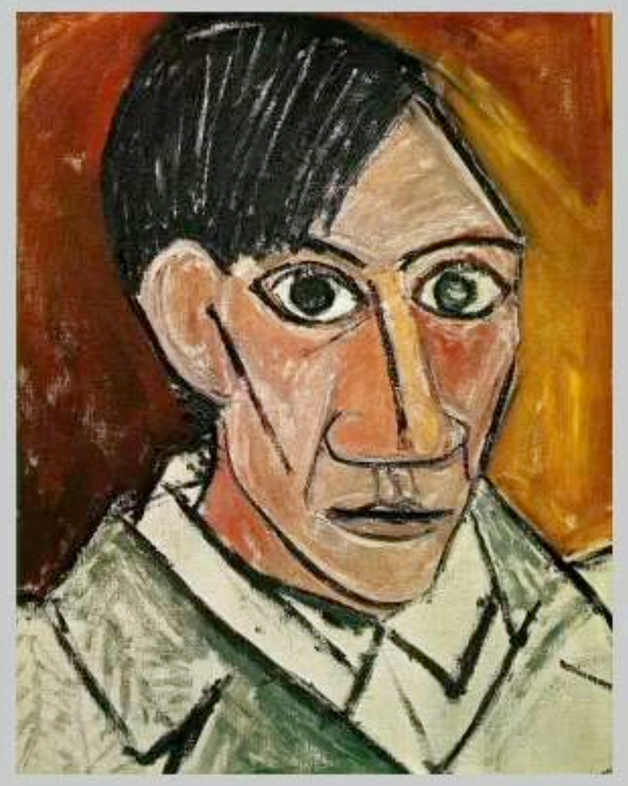

input image

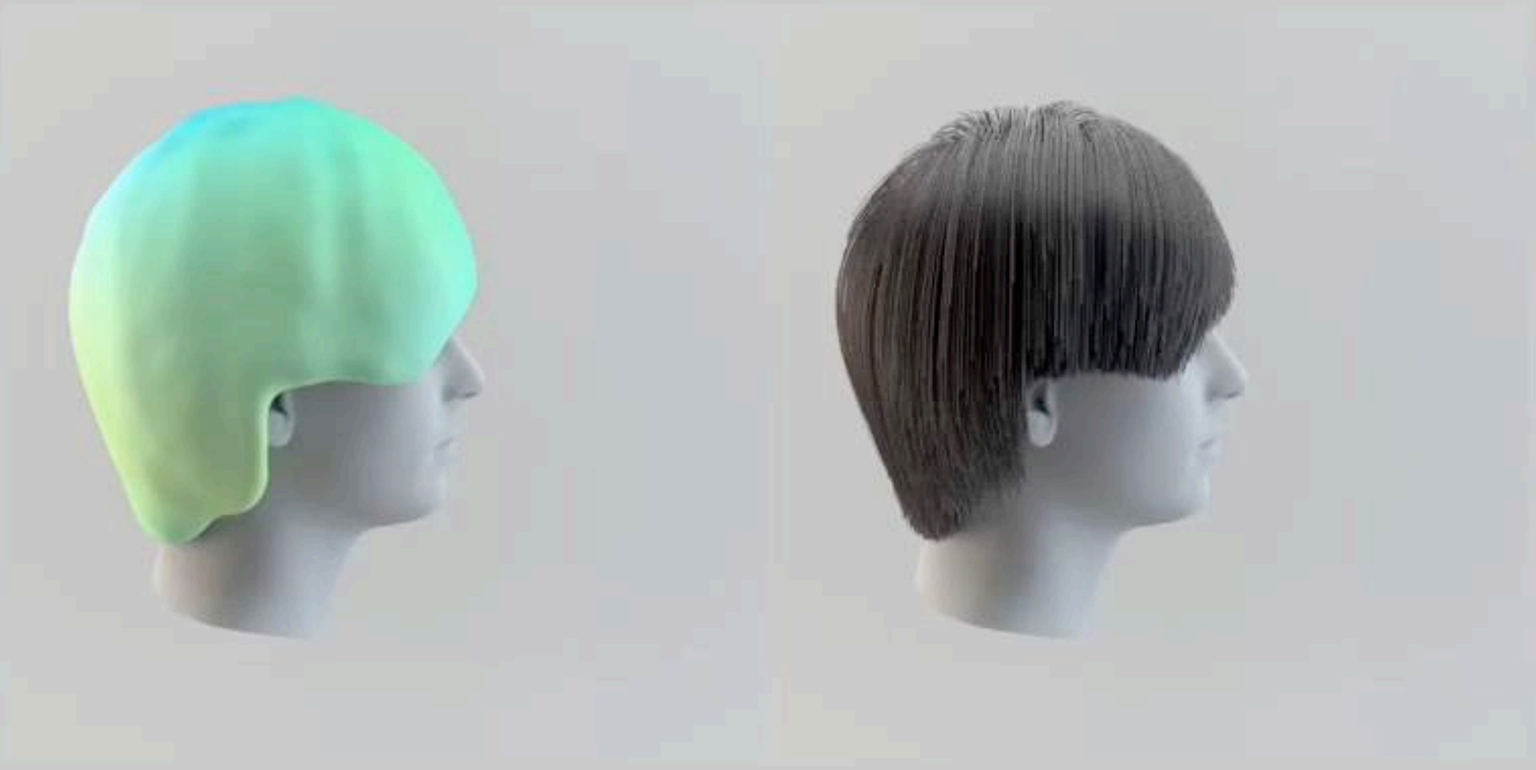

predicted volume

output strands

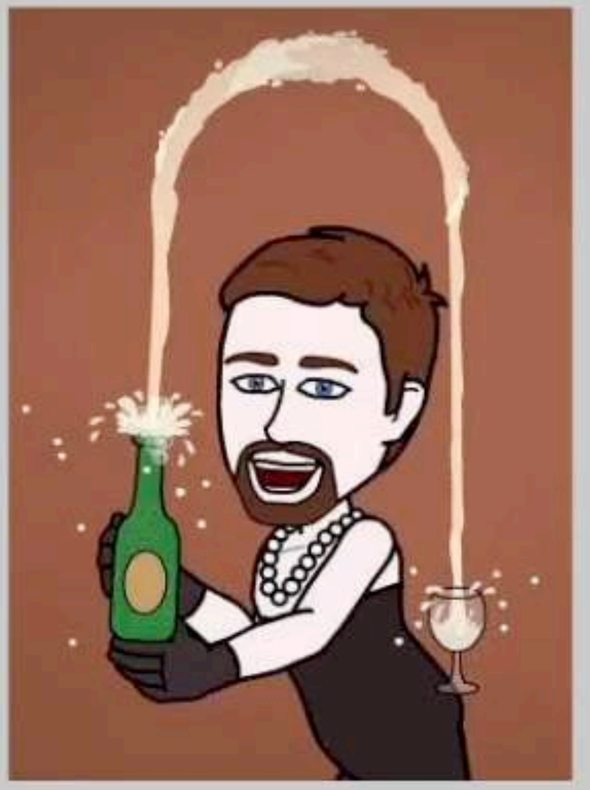

input image

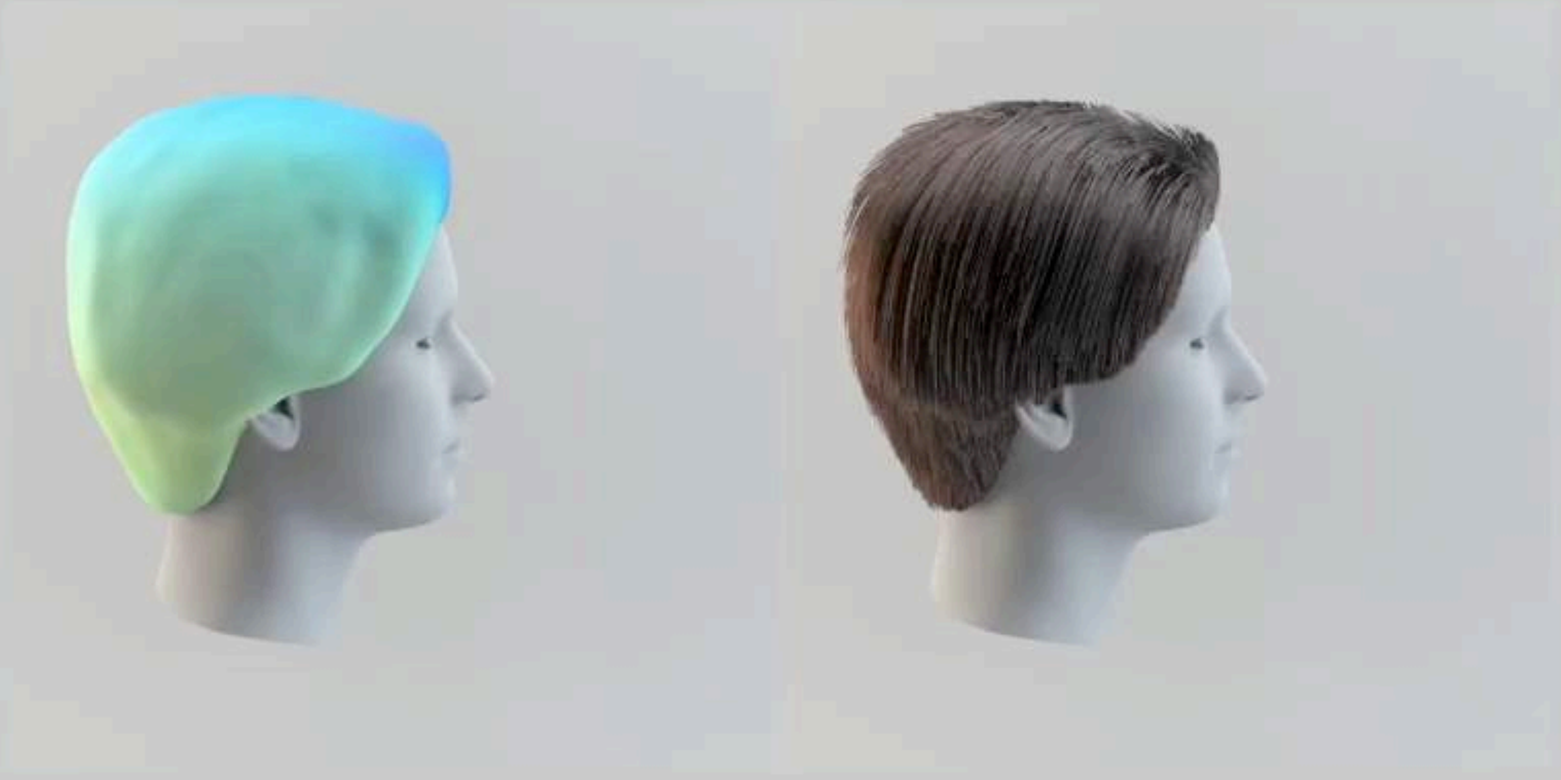

predicted volume

output strands

input image

predicted volume

output strands

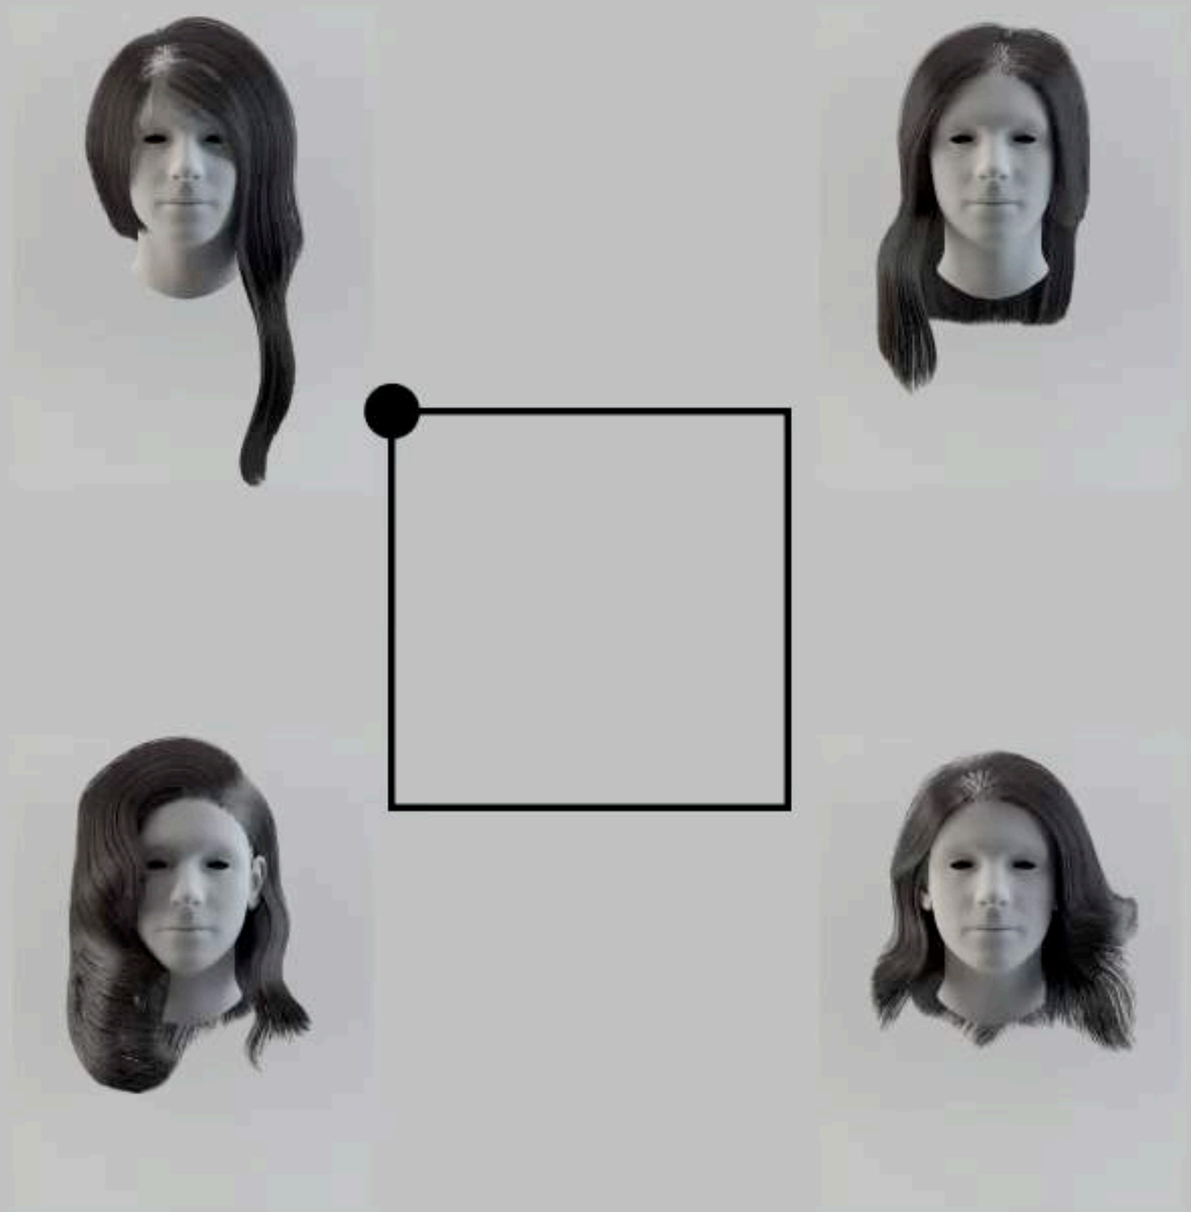

interpolation result

Human Body

Body/Performance Capture

Single-View Depth-Sensor Capture (CVPR 2018)

- Clothed+naked subjects modeling in real-time
- Problem: Depth Sensor, Limited Pose, Restricted Volume

3D Photogrammetry

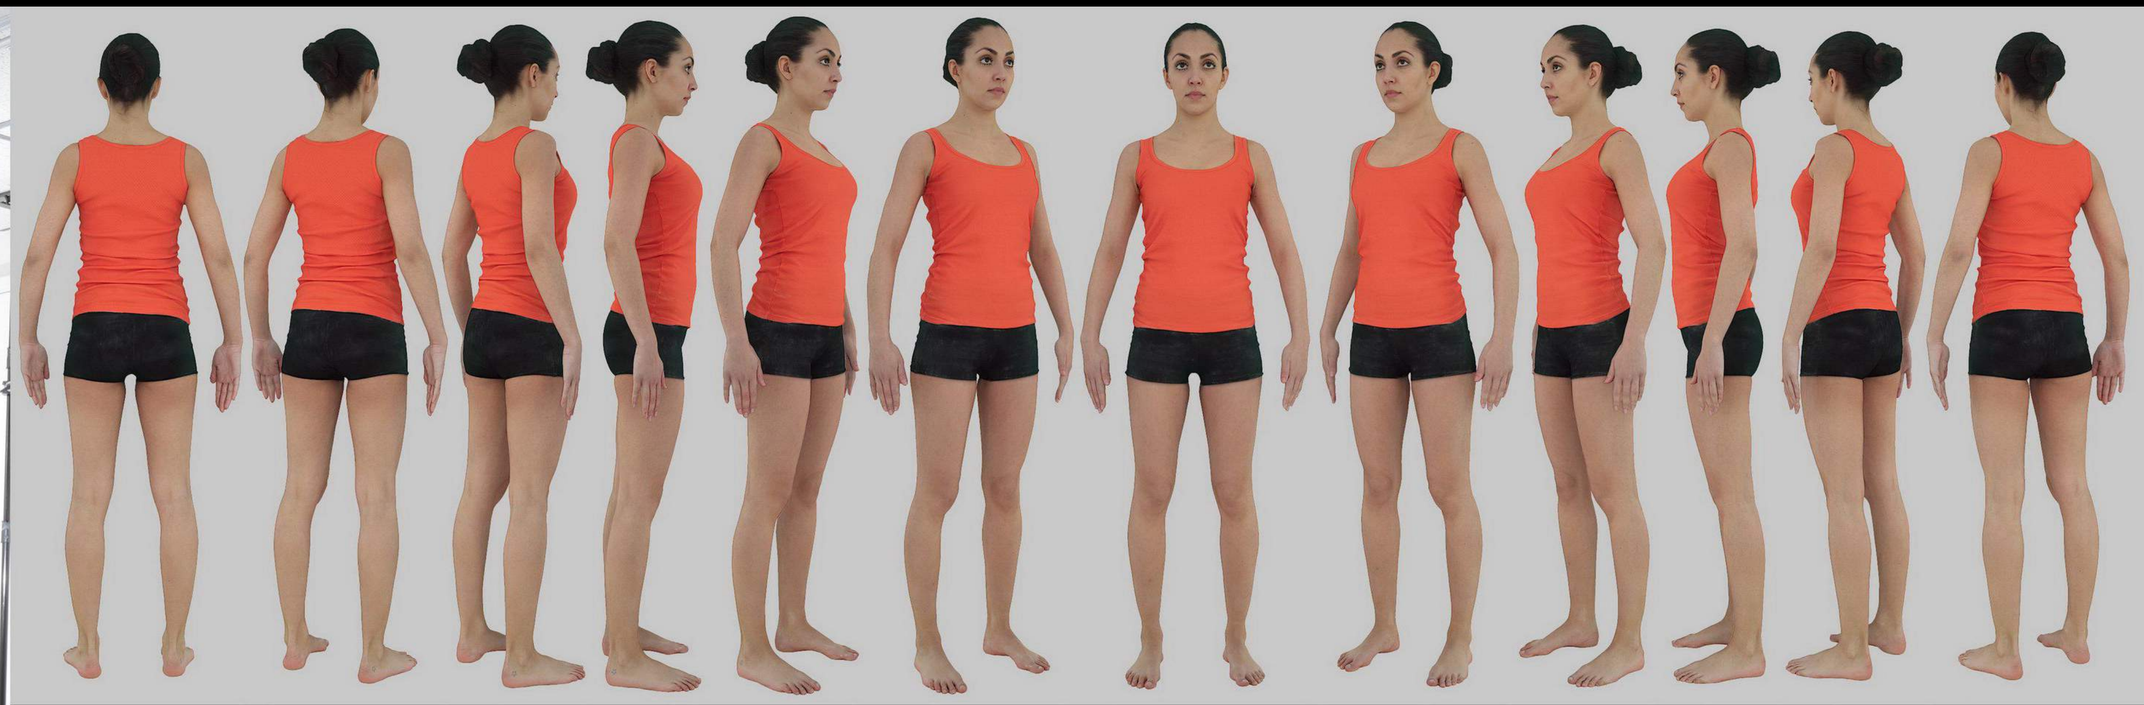

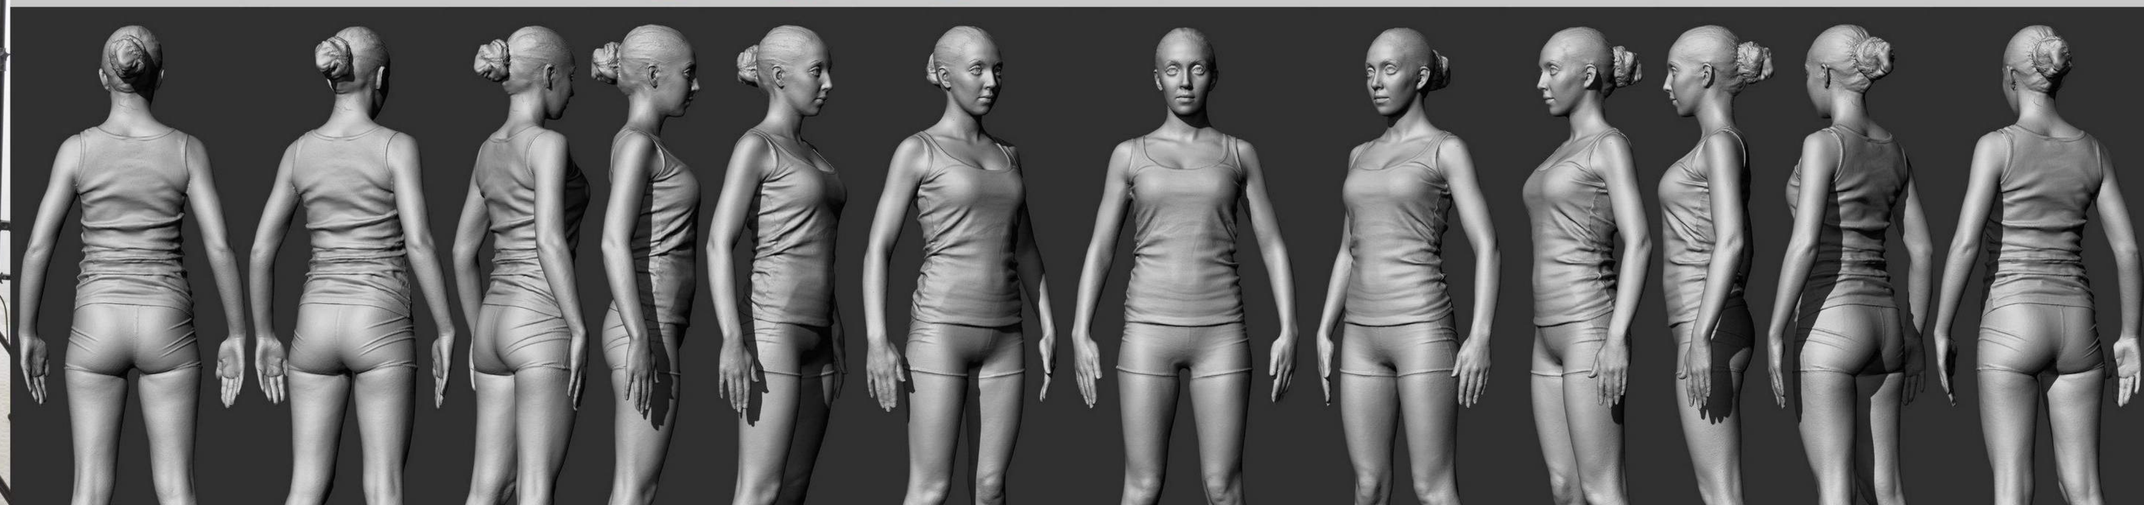

Sparse View Volumetric Capture

Sparse View Volumetric Capture

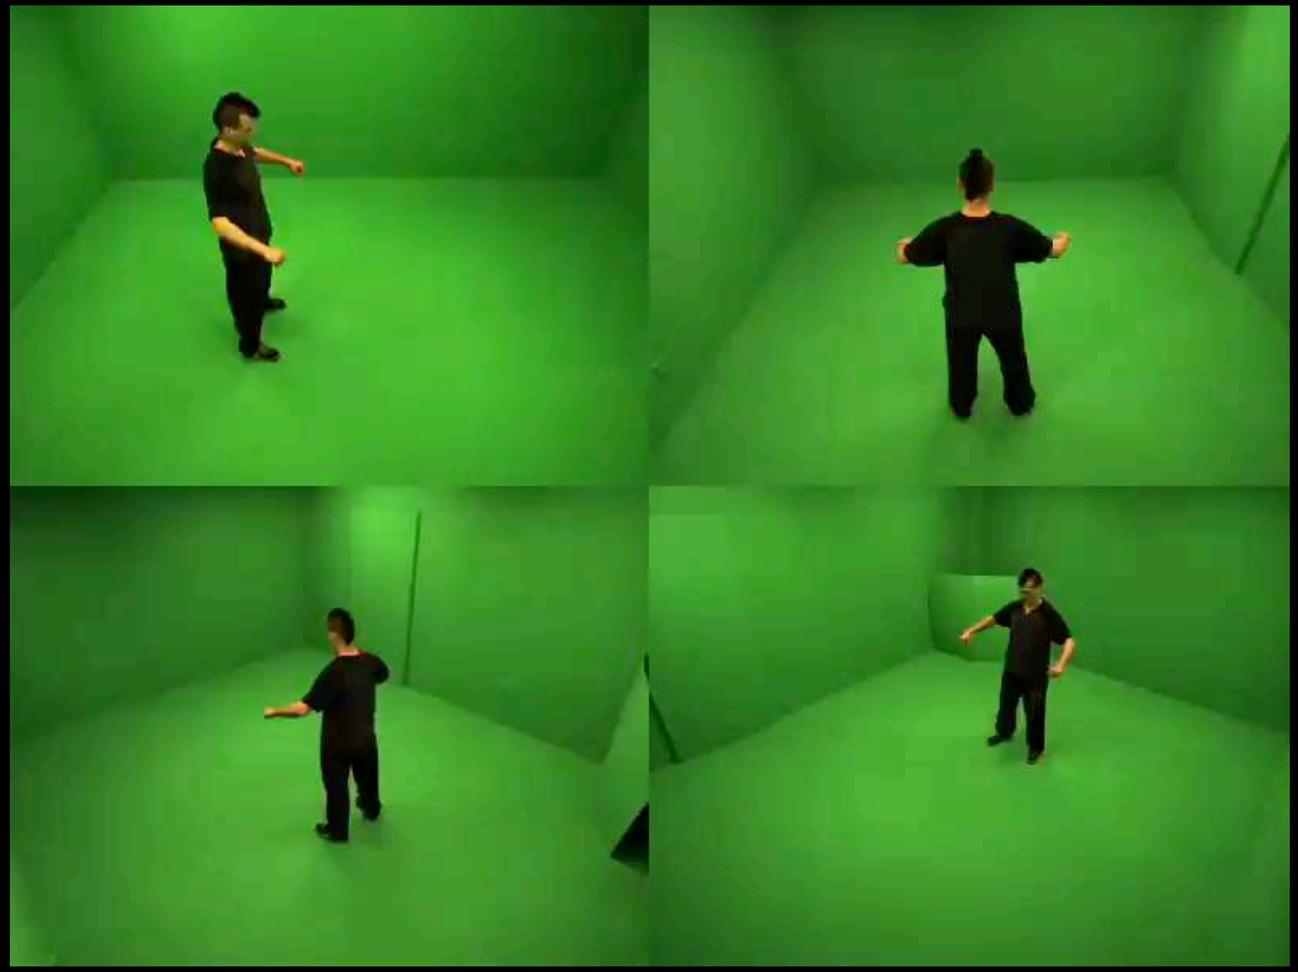

Sparse View Volumetric Capture

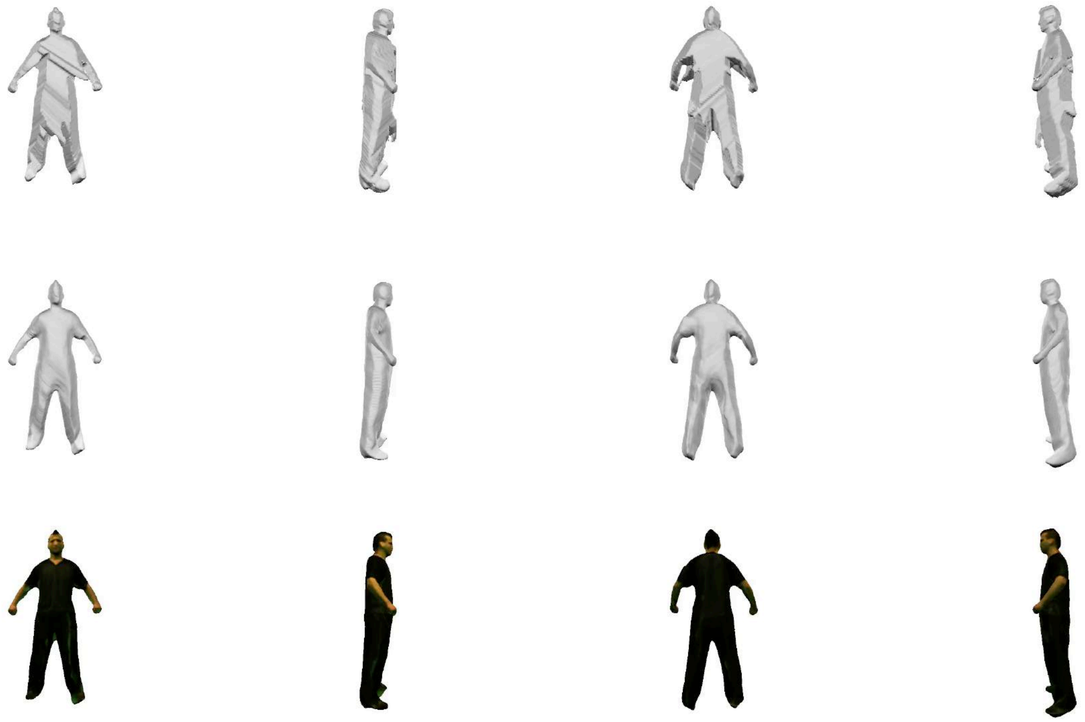

Multi-View RGB Capture (ECCV 2018)

- Clothed subjects modeling + Fast motion
- Problem: Calibrated multi-view system, no naked body estimation

Single View Prior Work

Mesh [Kanazawa et al. 18]

Voxel [Varol et al. 18]

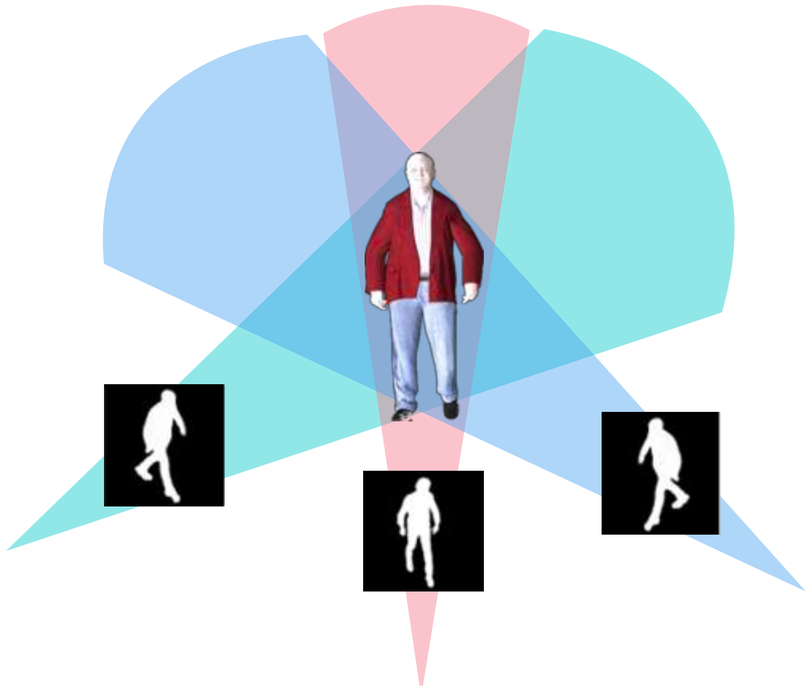

Silhouette-based
[Natsume et al. 19]

Single View Prior Work

ours

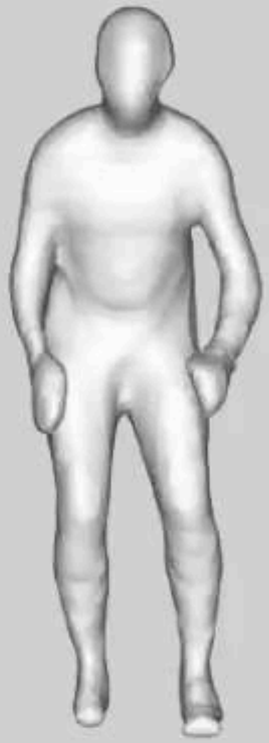

BodyNet
[Varol et al.]

HMR
[Kanazawa et al.]

Ryota Natsume*, **Shunsuke Saito***, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, Shigeo Morishima. "SiCloPe: Silhouette-Based Clothed People". CVPR 2019 (oral).

Using an Effective 3D Representation

Using an Effective 3D Representation

Inferring Texture

Key Insight: Memory Efficiency

❌ Limited resolution

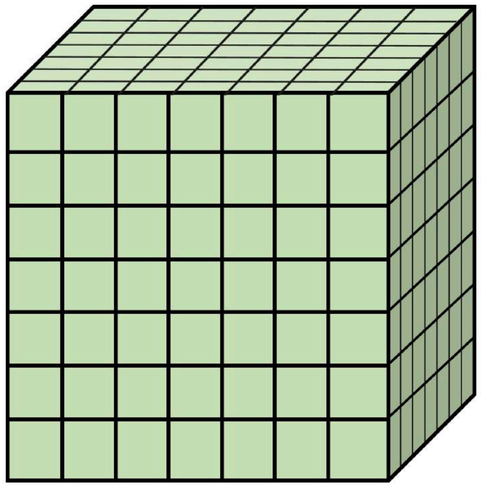

Voxel (Explicit)

✅ Continuous representation

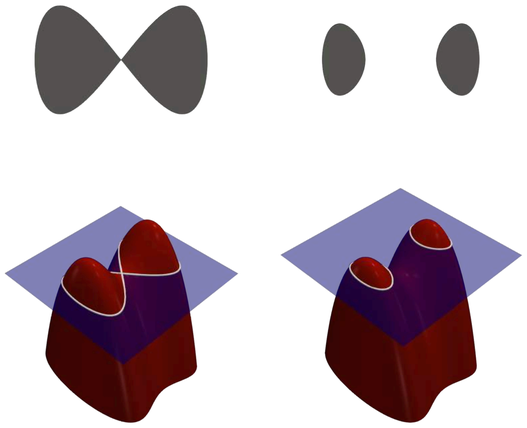

Implicit Surface

Spatial Alignment

❌ Local detail lost

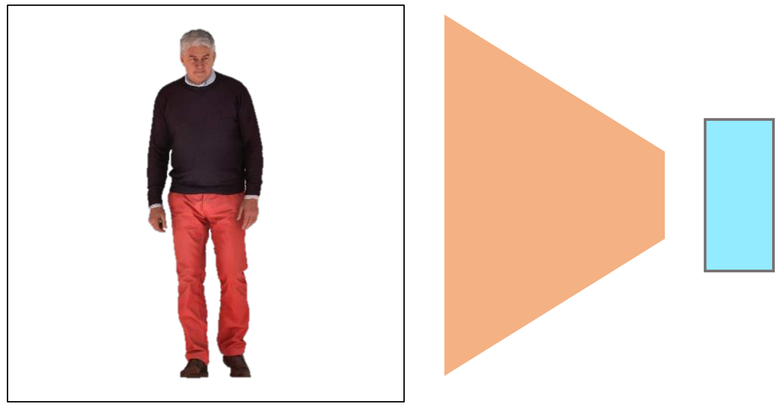

Global image representation

✅ Local detail preserved

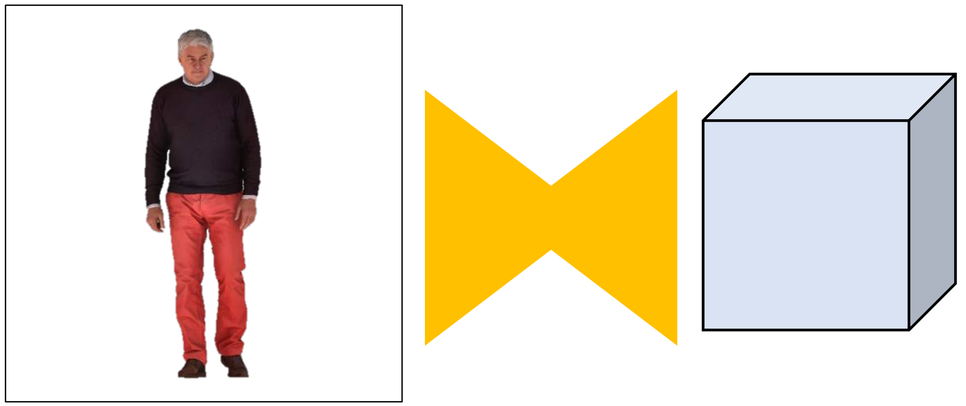

Fully-convolutional image representation

Pixel-Aligned Implicit Function (PIFu)

input

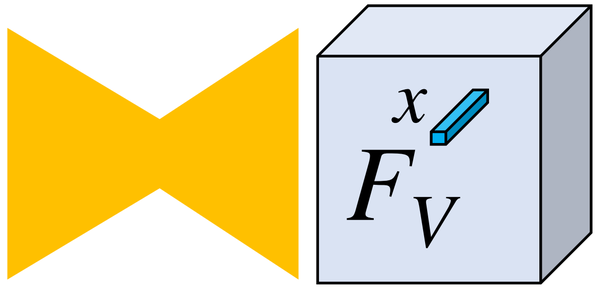

fully convolutional
image encoder

PIFu $(\forall x, \forall z)$

$$f(F_V(x), z) = \text{inside/outside}$$

- ✅ Memory Efficient
- ✅ No Discretization
- ✅ Preserves Local Detail
- ✅ Arbitrary Surface Topology

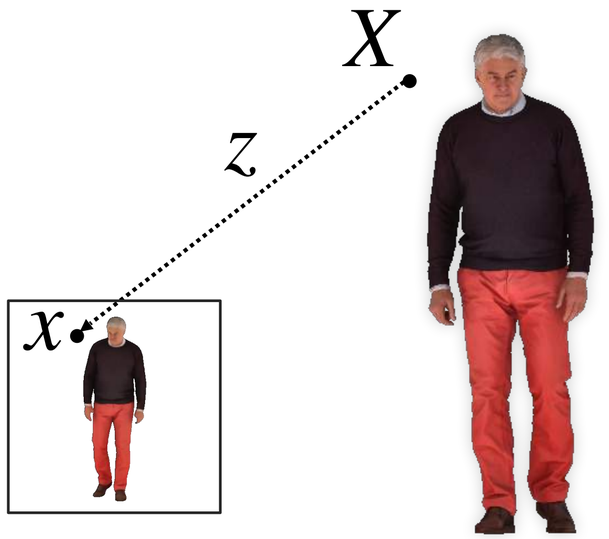

General Framework

Texture Inpainting

Tex-PIFu ($\forall x, \forall z$)

$$f_c(\text{surface } x, \text{depth } z) = \text{RGB}$$

Multi-view Incorporation

MV-PIFu ($\forall x, \forall z$)

$$f(F_i(x_i), z_i) = \Phi_i \quad \bar{\Phi} = mean(\{\Phi_i\})$$

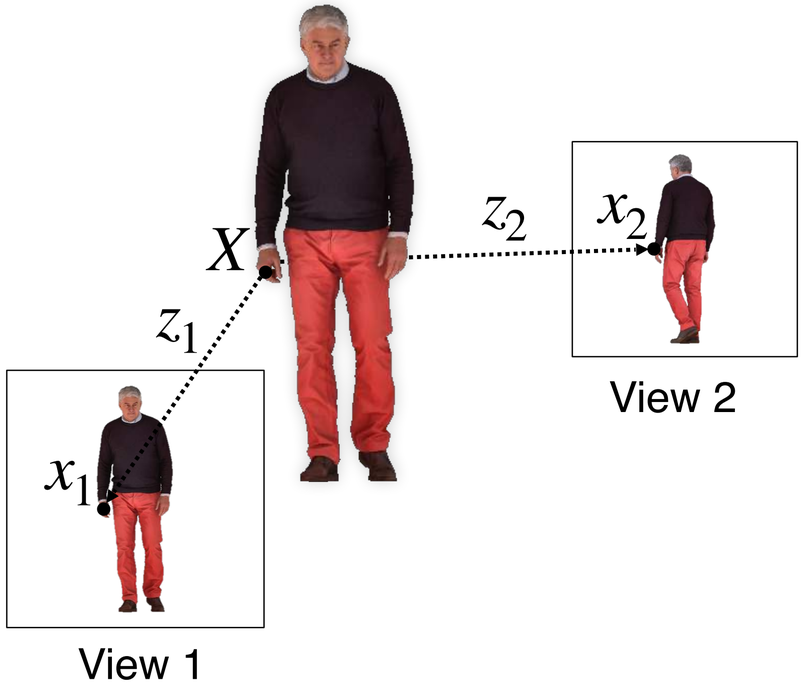

Training Data

High-resolution Photogrammetry Scans
[RenderPeople]

Photorealistic
Rendering

442 x 360 = 159,120 images

Training Method

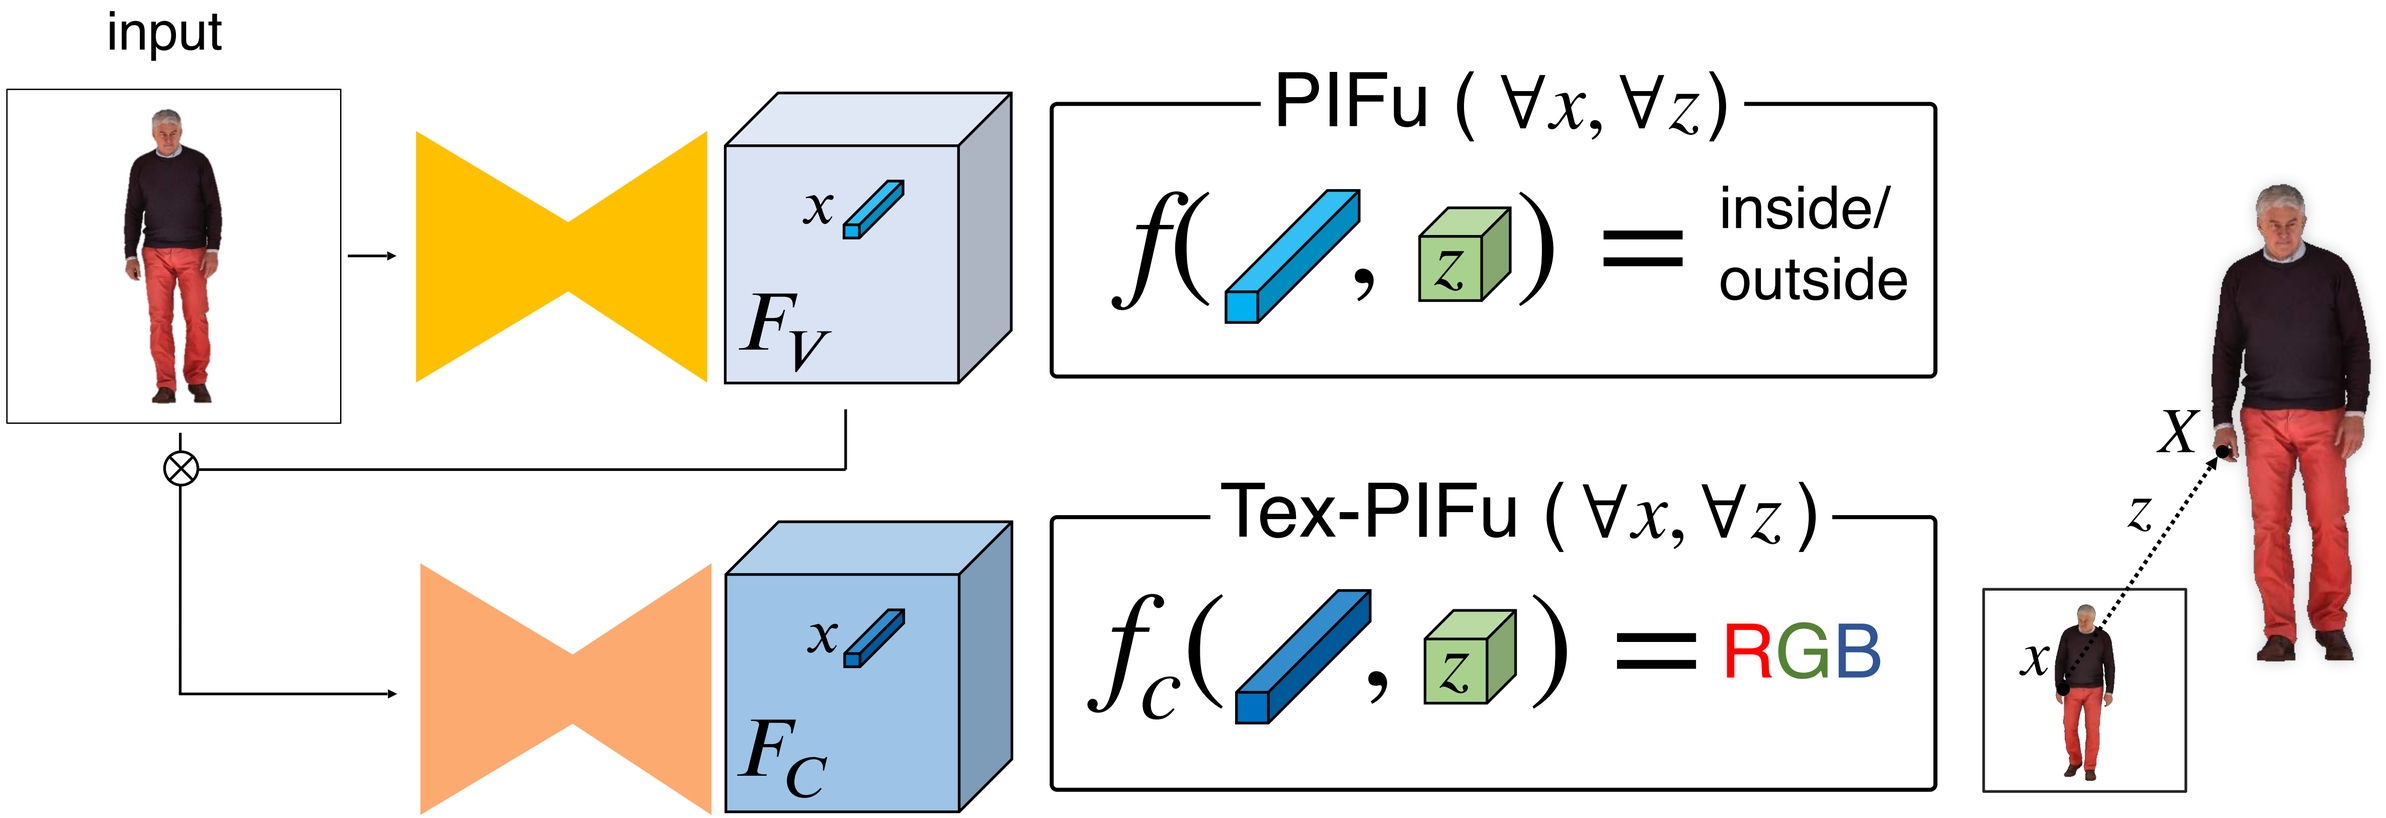

Human-Digitization Pipeline

n-view inputs
$(n \geq 1)$

PIFu

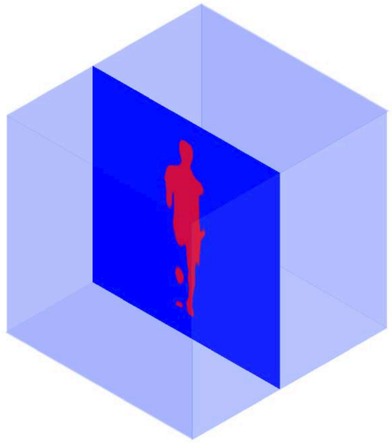

3D occupancy field

Marching Cube

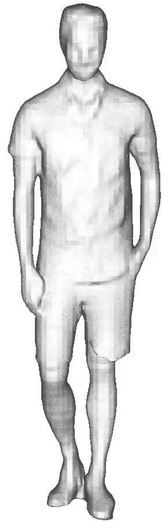

reconstructed geometry

Tex-PIFu

textured reconstruction

Results (Single Input Photo)

Results (Single Input Photo)

Results (Single Input Photo)

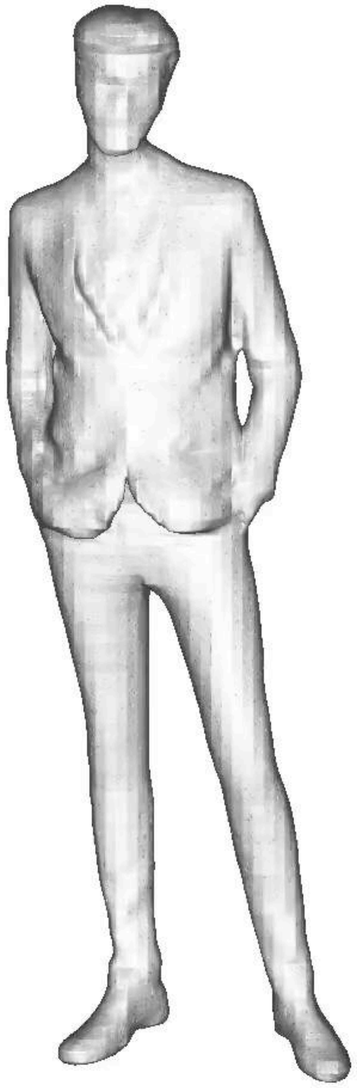

Results (Single Input Photo)

More Results




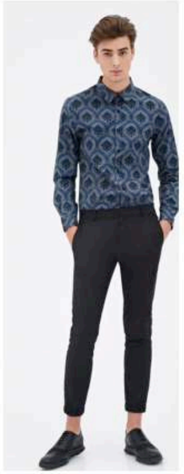
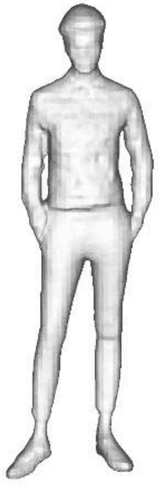





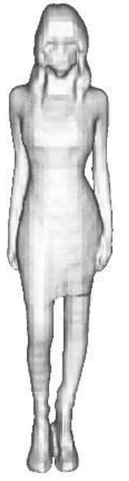


More Results











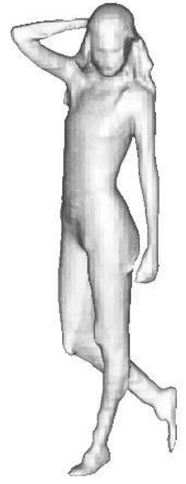


Results (Multi-View)

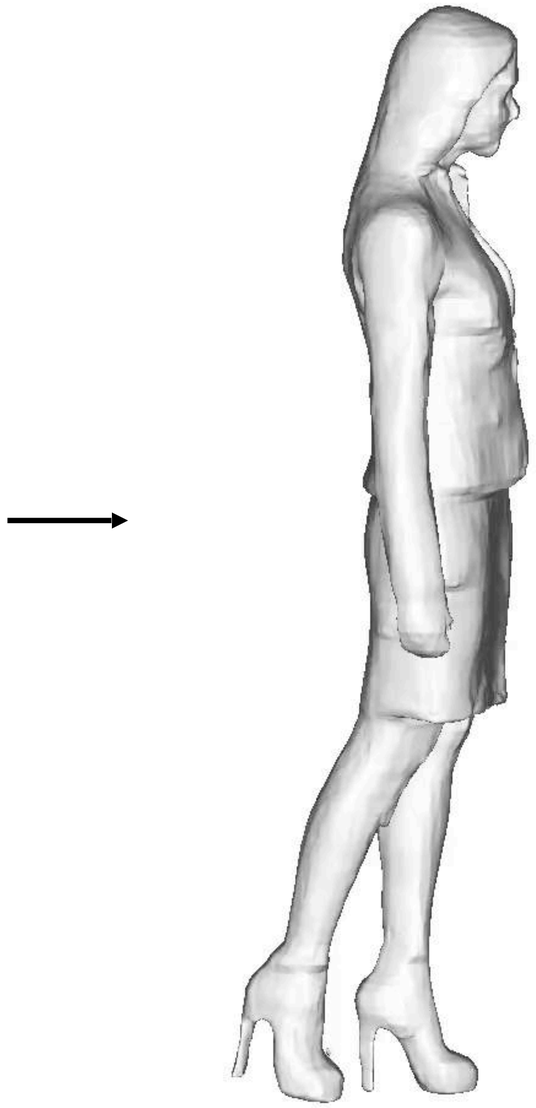

#views: 1

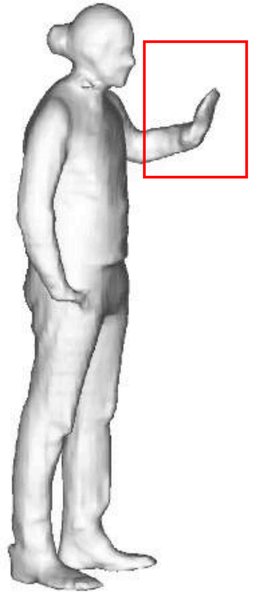

#views: 3

#views: 6

#views: 9








Results (Video Input)

Any objects?

Rendering

3D Mesh

2D Image

Differentiable Rendering

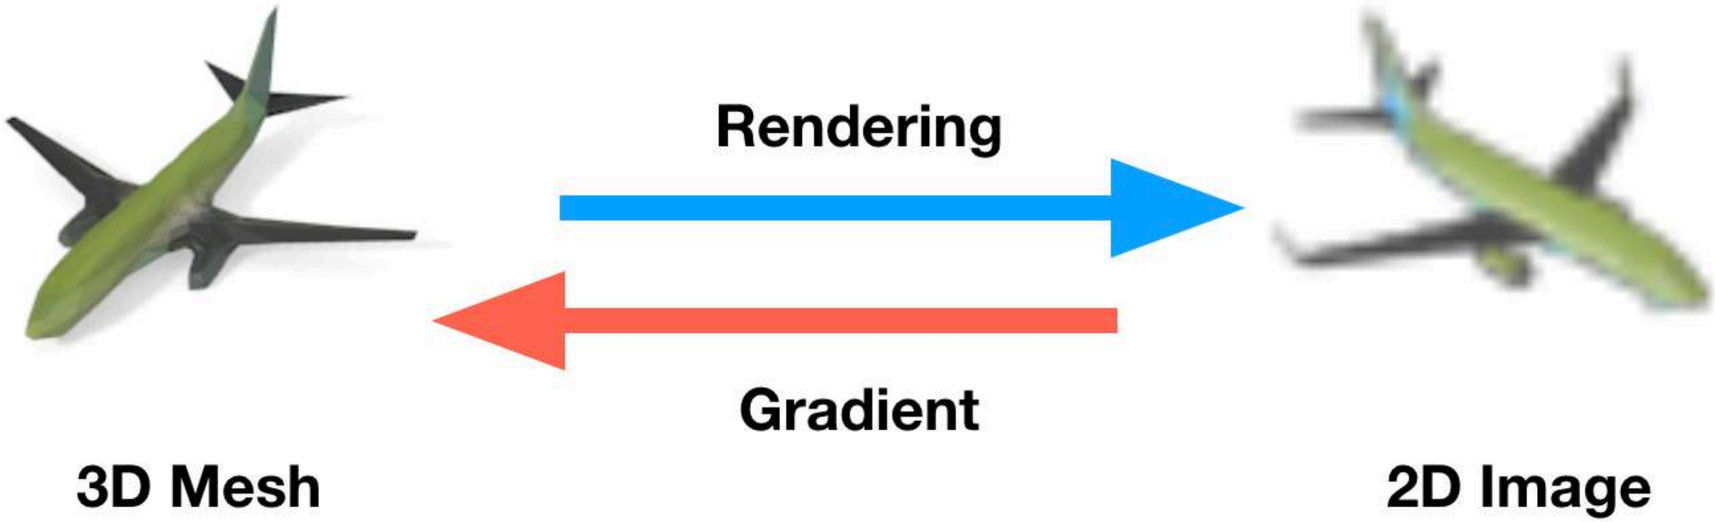

Differentiable Rendering

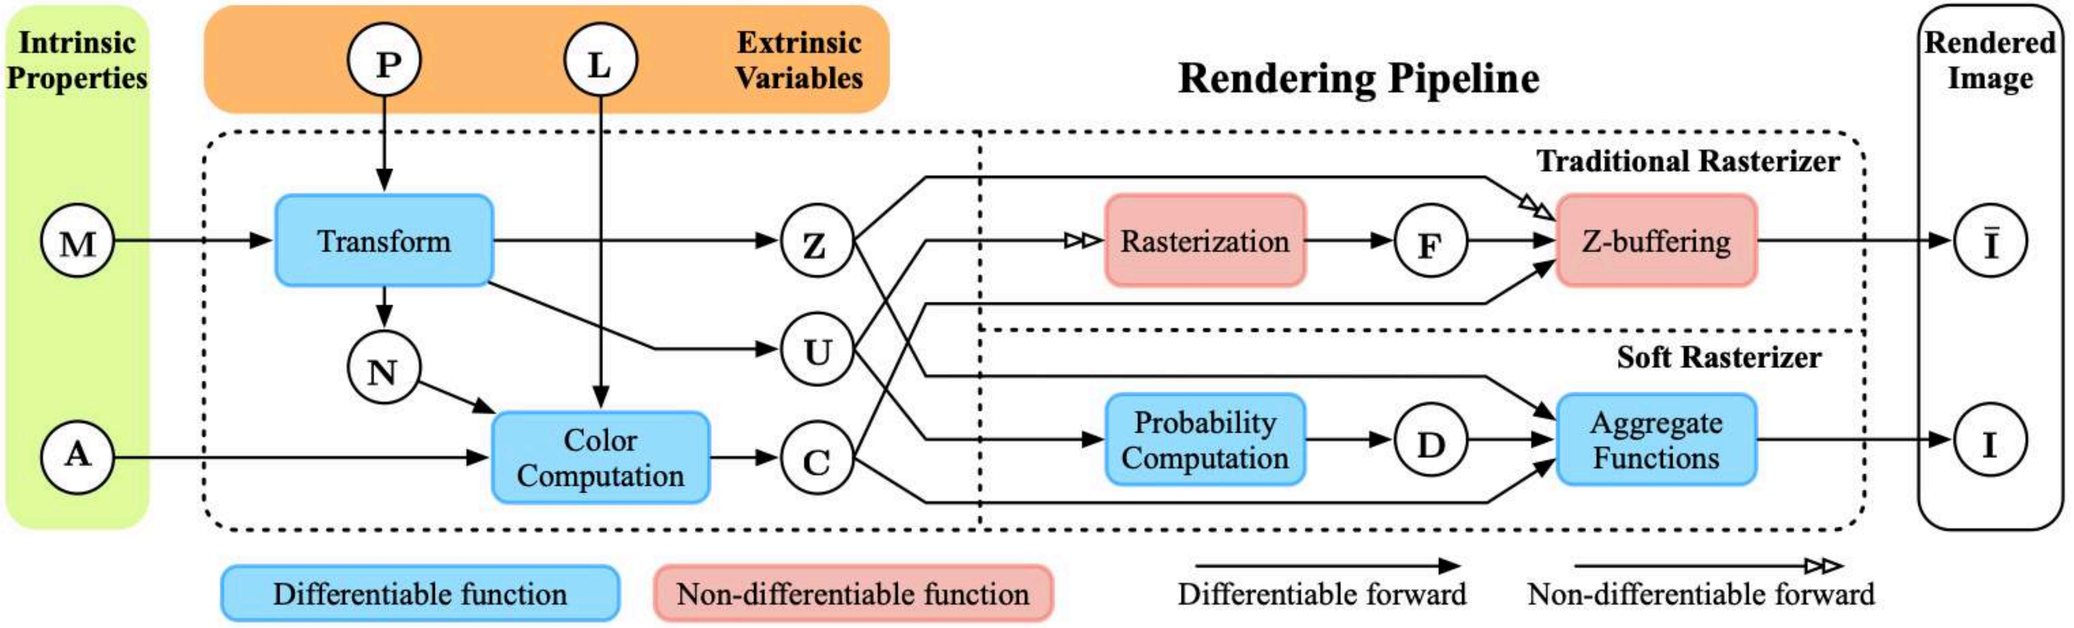

Differentiable Rendering

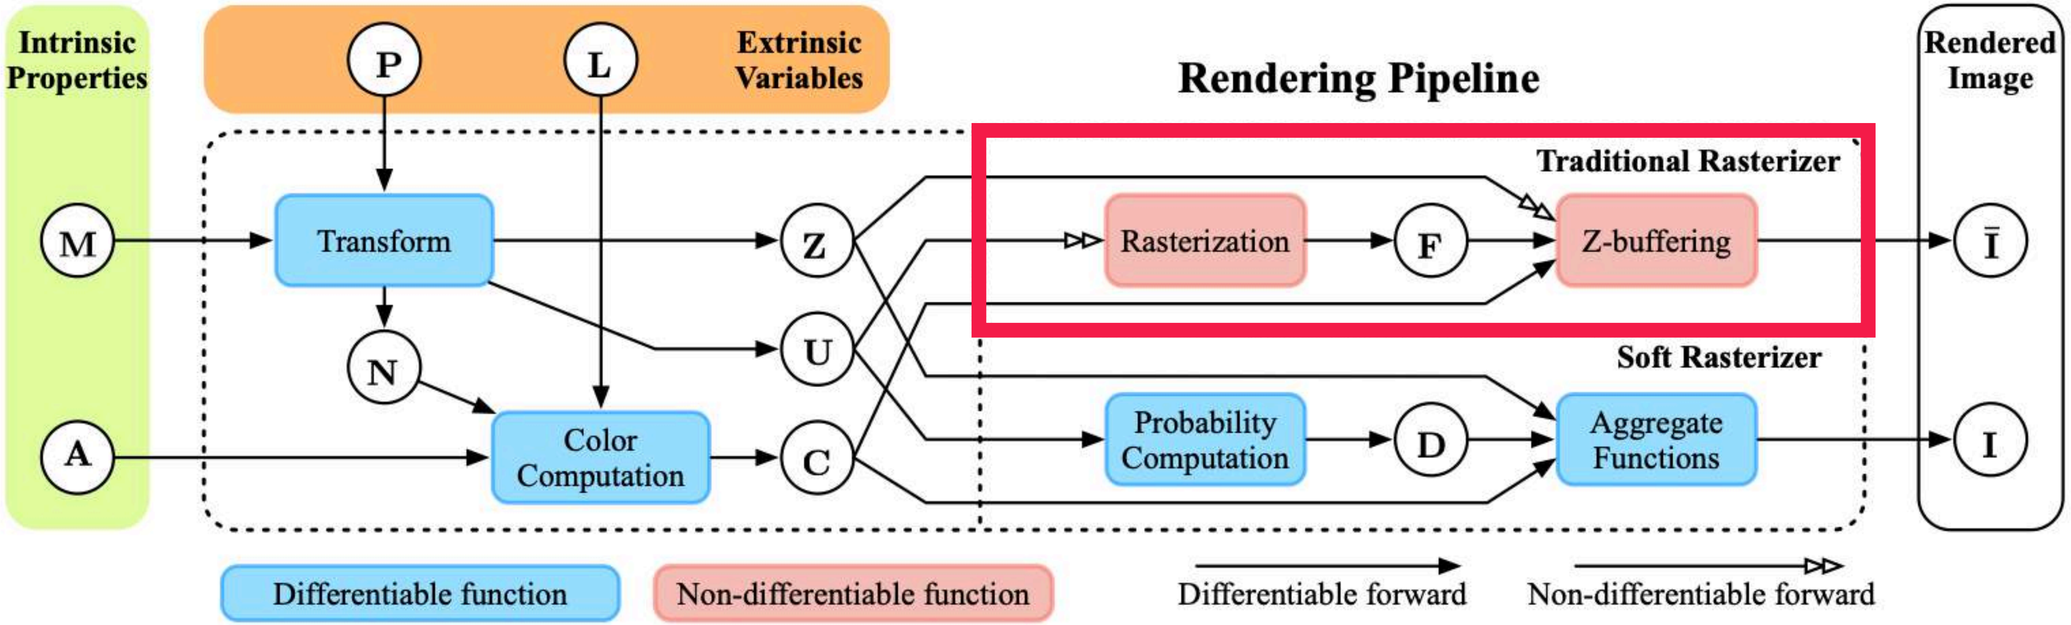

Differentiable Rendering

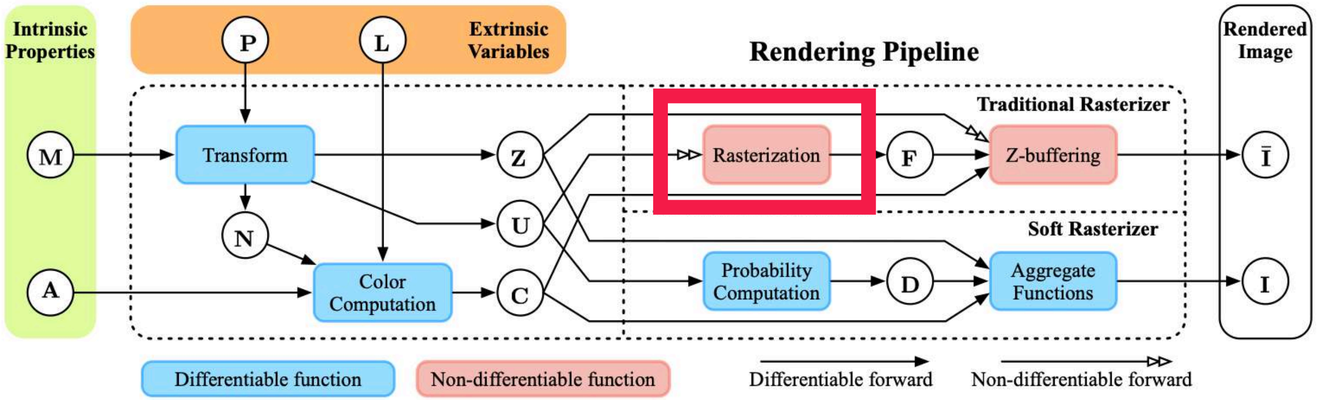

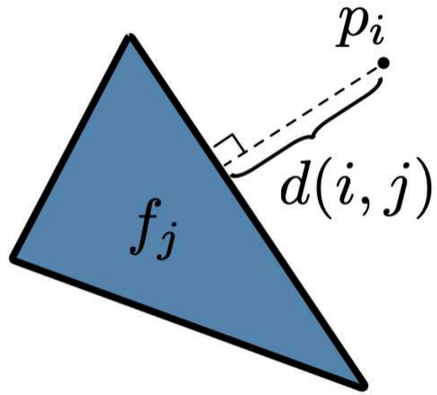

Differentiable Rendering

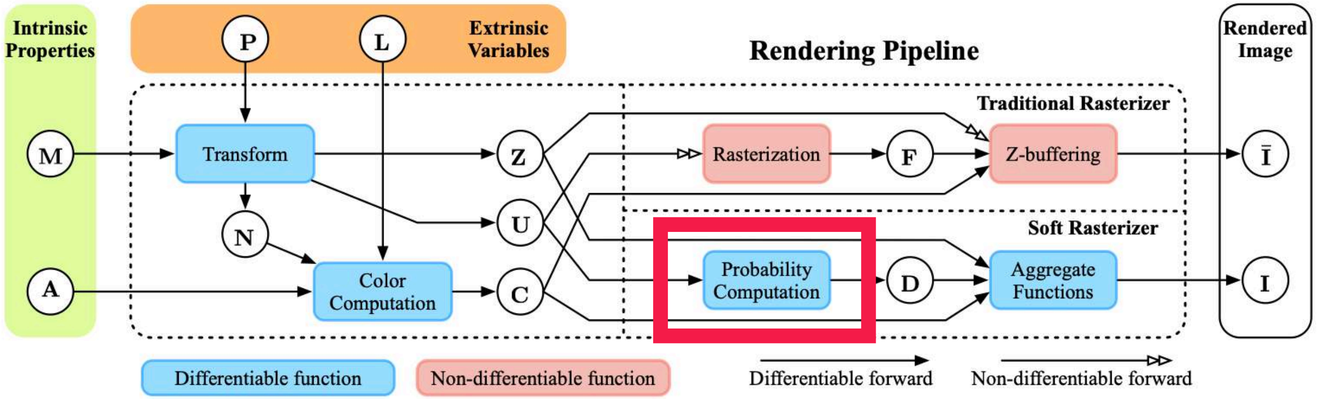

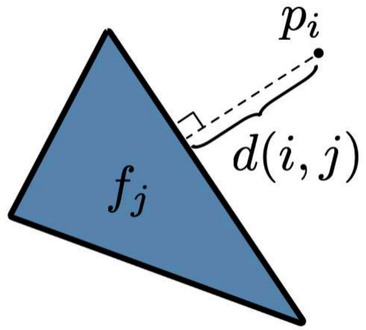

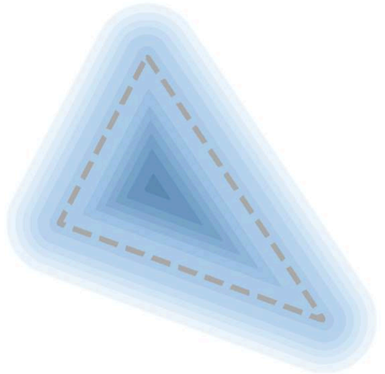

$$\mathcal{D}_j^i = sigmoid(\delta_j^i \cdot \frac{d^2(i,j)}{\sigma})$$

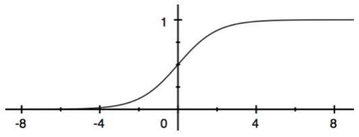

Differentiable Rendering

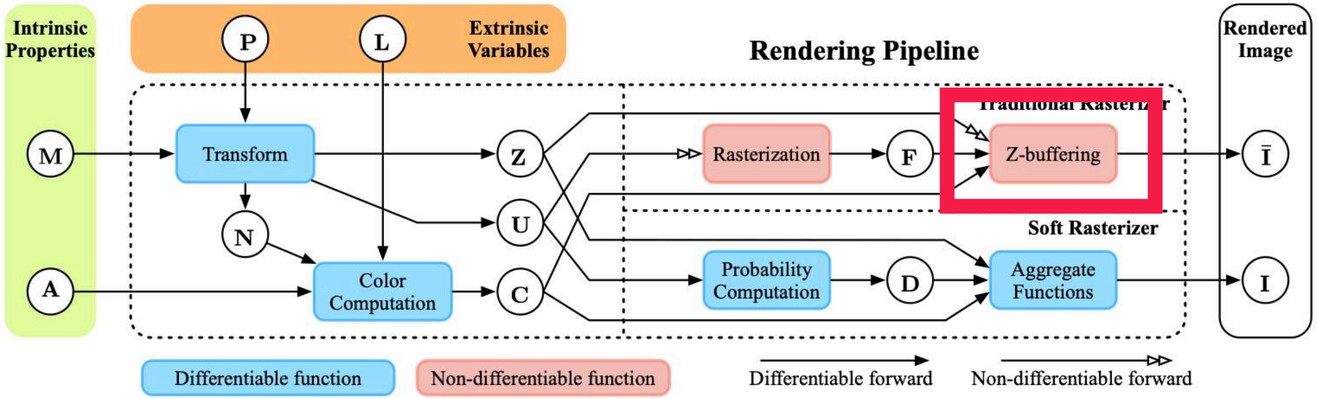

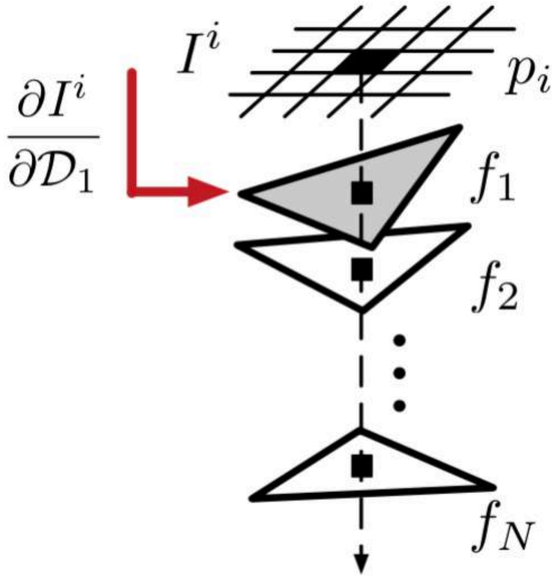

Differentiable Rendering

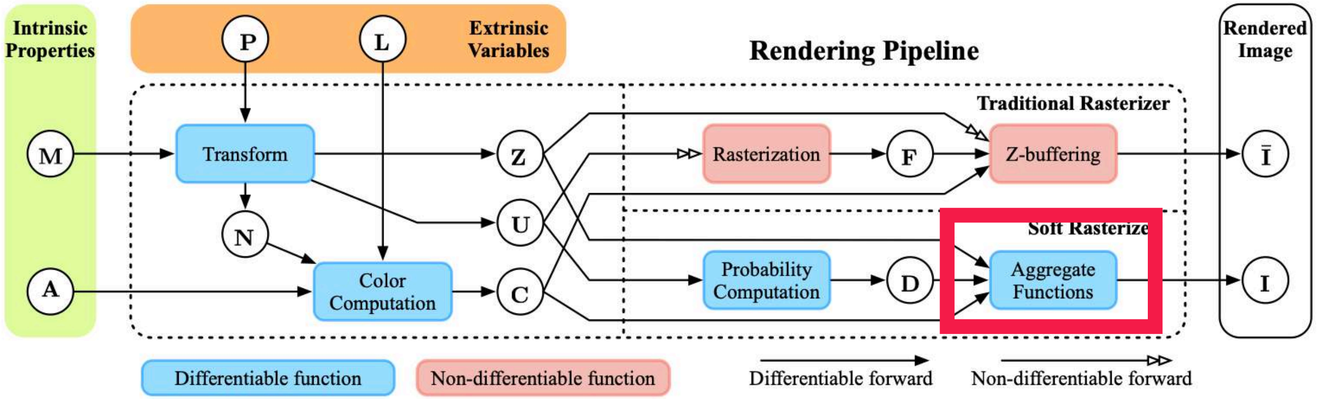

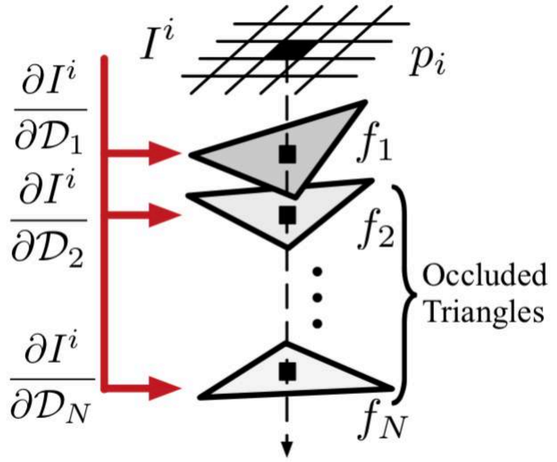

$$I^i = \mathcal{A}_S(\{C_j\}) = \sum_j w_j^i C_j^i + w_b^i C_b$$

$$w_j^i = \frac{\mathcal{D}_j^i \exp(z_j^i/\gamma)}{\sum_k \mathcal{D}_k^i \exp(z_k^i/\gamma) + \exp(\epsilon/\gamma)}$$

Differentiable Rendering

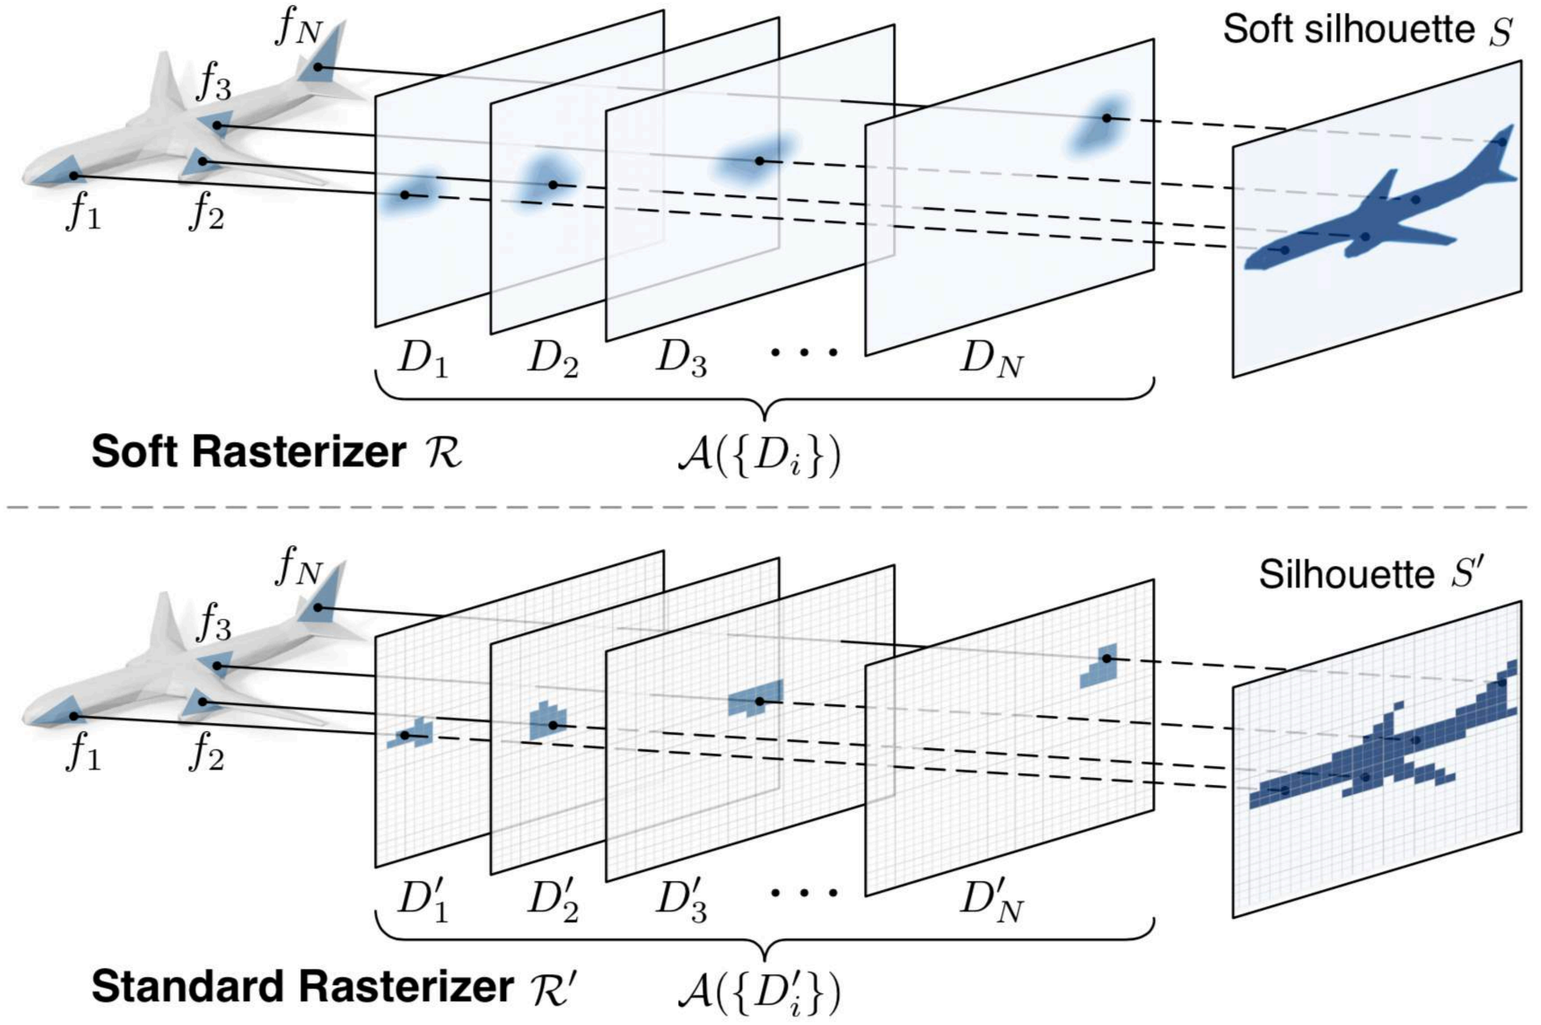

Liu et al. 2019

Shape Optimization

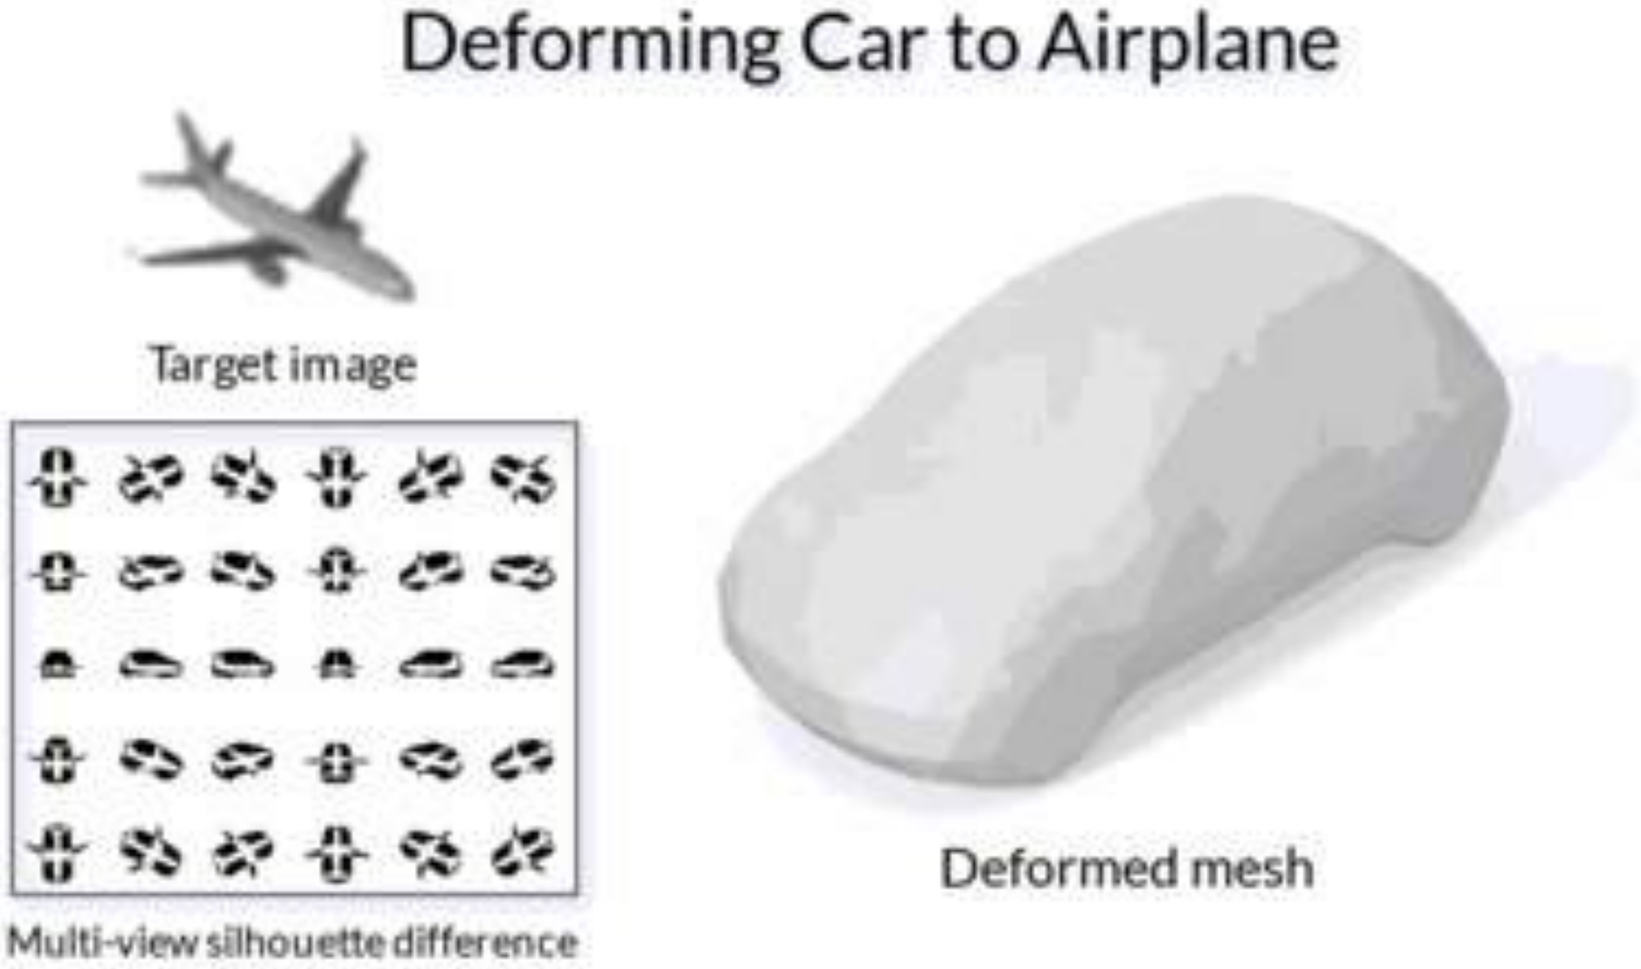

3D Pose Estimation

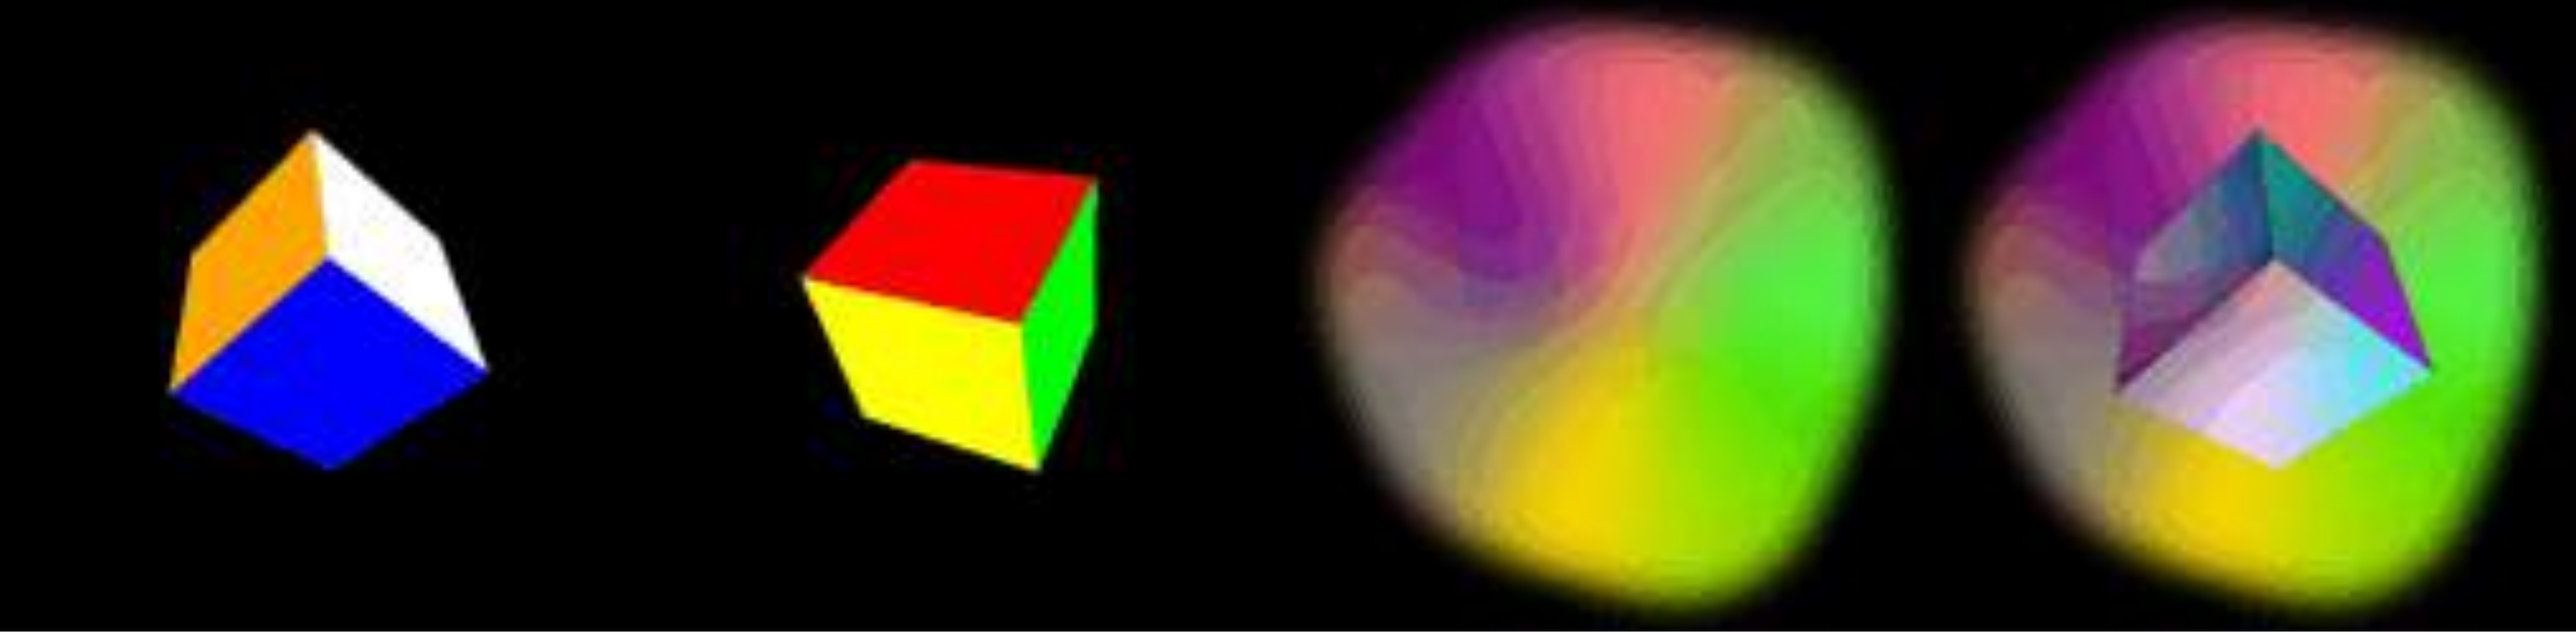

Target | Initial | Soft Rendering | Difference

3D Pose Estimation

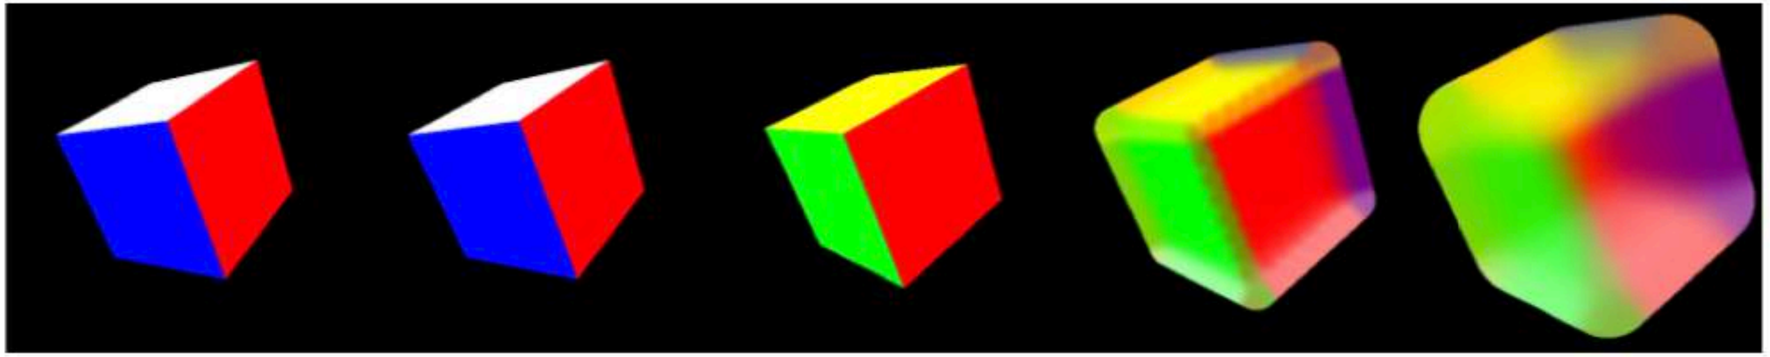

(a) target (b) our result (c) NMR result (d) smooth rendering (e) smoother rendering

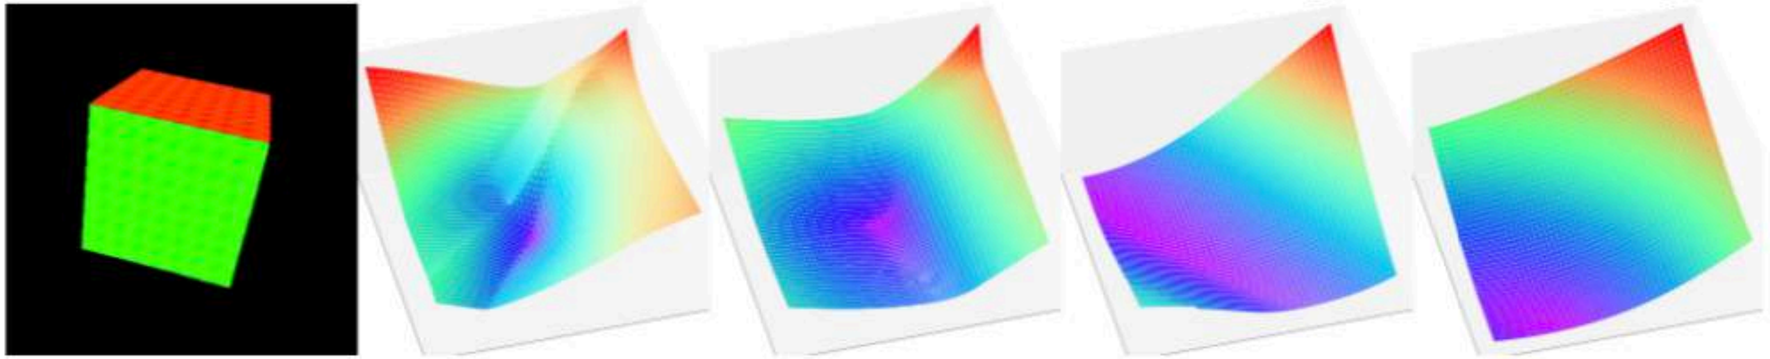

(f) initialization (g) global minimum (h) local minimum (i) $\sigma = 0.01$ $\gamma = 0.1$ (j) $\sigma = 0.03$ $\gamma = 0.3$

Single-View Digitization

(a) Synthetic Image

(b) Our reconstruction

(c) Real image

(d) Our reconstruction

Single-View Digitization

Comparisons

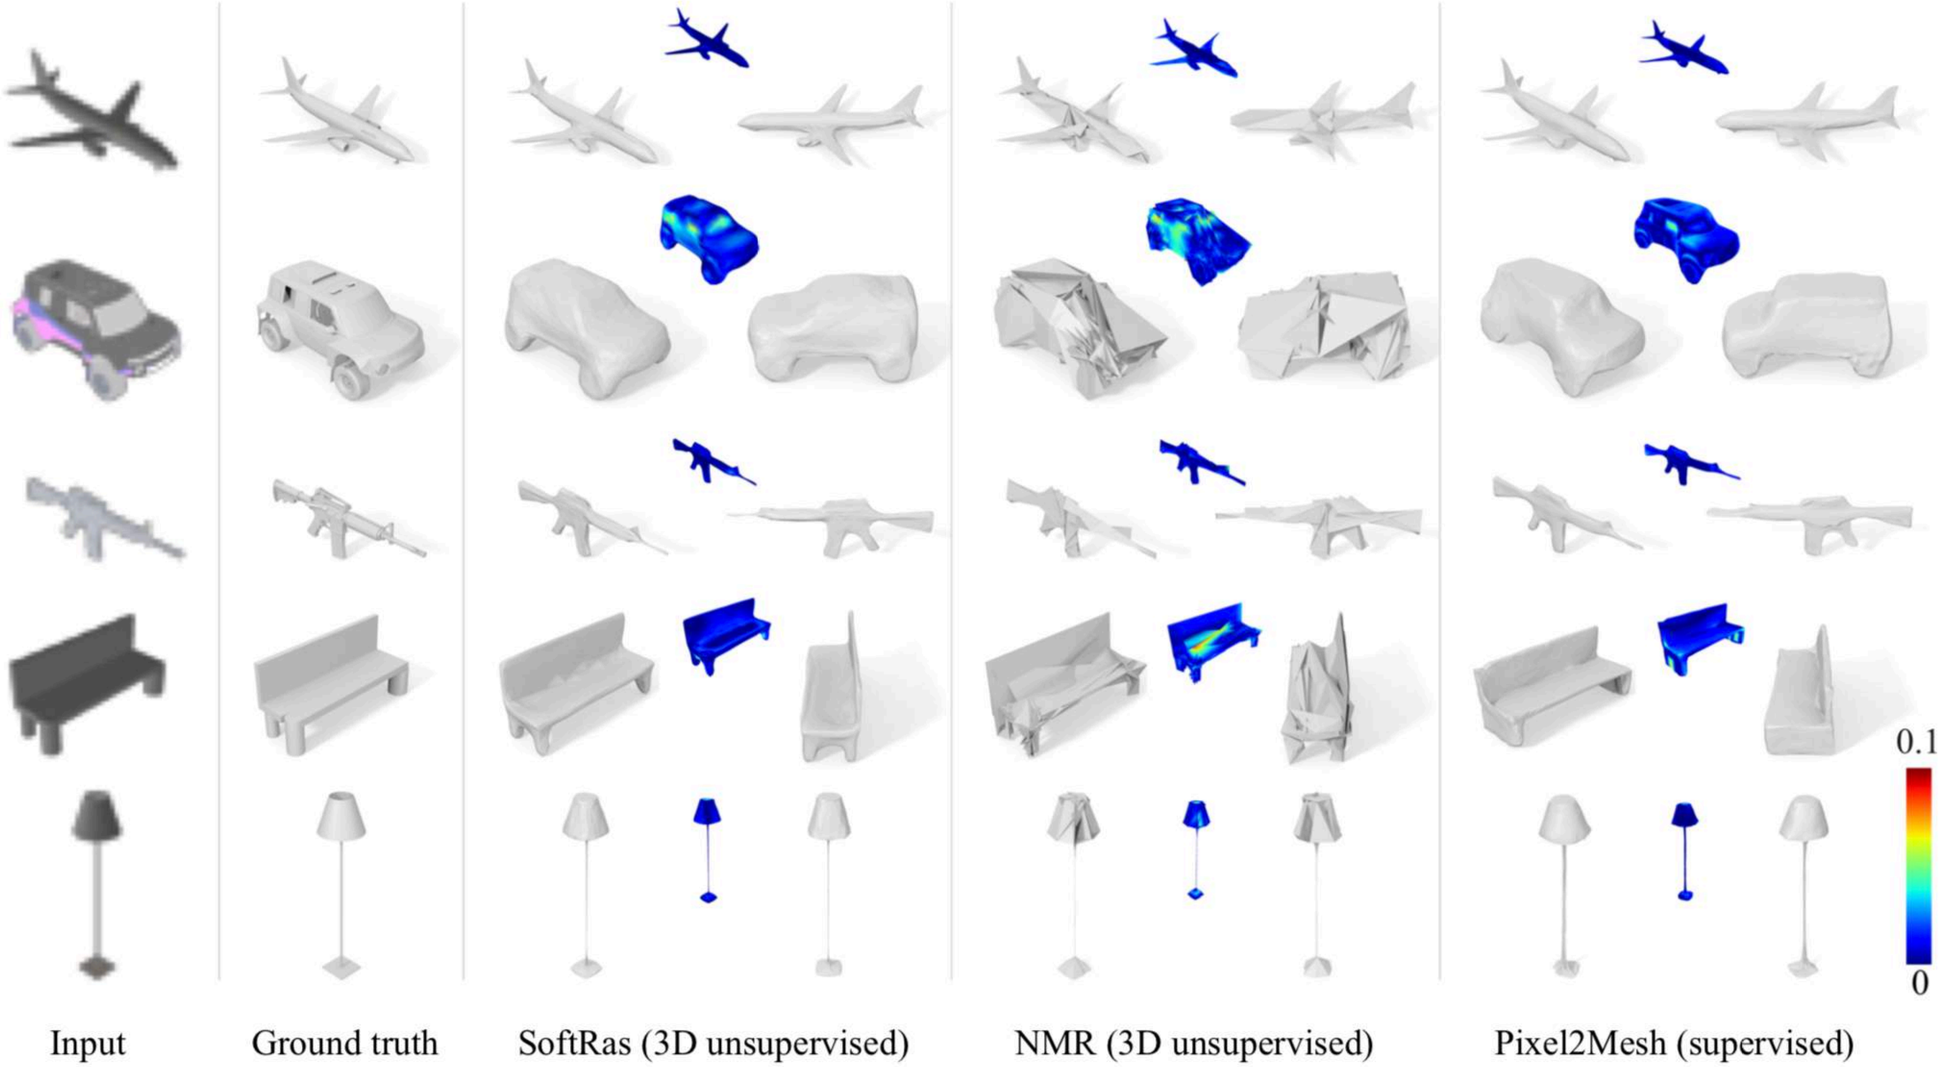

github.com/ShichenLiu/SoftRas

README.md

Soft Rasterizer (SoftRas)

This repository contains the code (in PyTorch) for "Soft Rasterizer: A Differentiable Renderer for Image-based 3D Reasoning" (ICCV'2019 Oral) by Shichen Liu, Tianye Li, Weikai Chen and Hao Li.

Contents

Introduction

Soft Rasterizer (SoftRas) is a truly differentiable renderer framework with a novel formulation that views rendering as a **differentiable aggregating process** that fuses **probabilistic contributions** of all mesh triangles with respect to the rendered pixels. Thanks to such *"soft"* formulation, our framework is able to (1) directly render colorized mesh using differentiable functions and (2) back-propagate efficient supervision signals to mesh vertices and their attributes (color, normal, etc.) from various forms of image representations, including silhouette, shading and color images.

Arbitrary Topologies?

Learning Implicit Functions via 2D Supervision

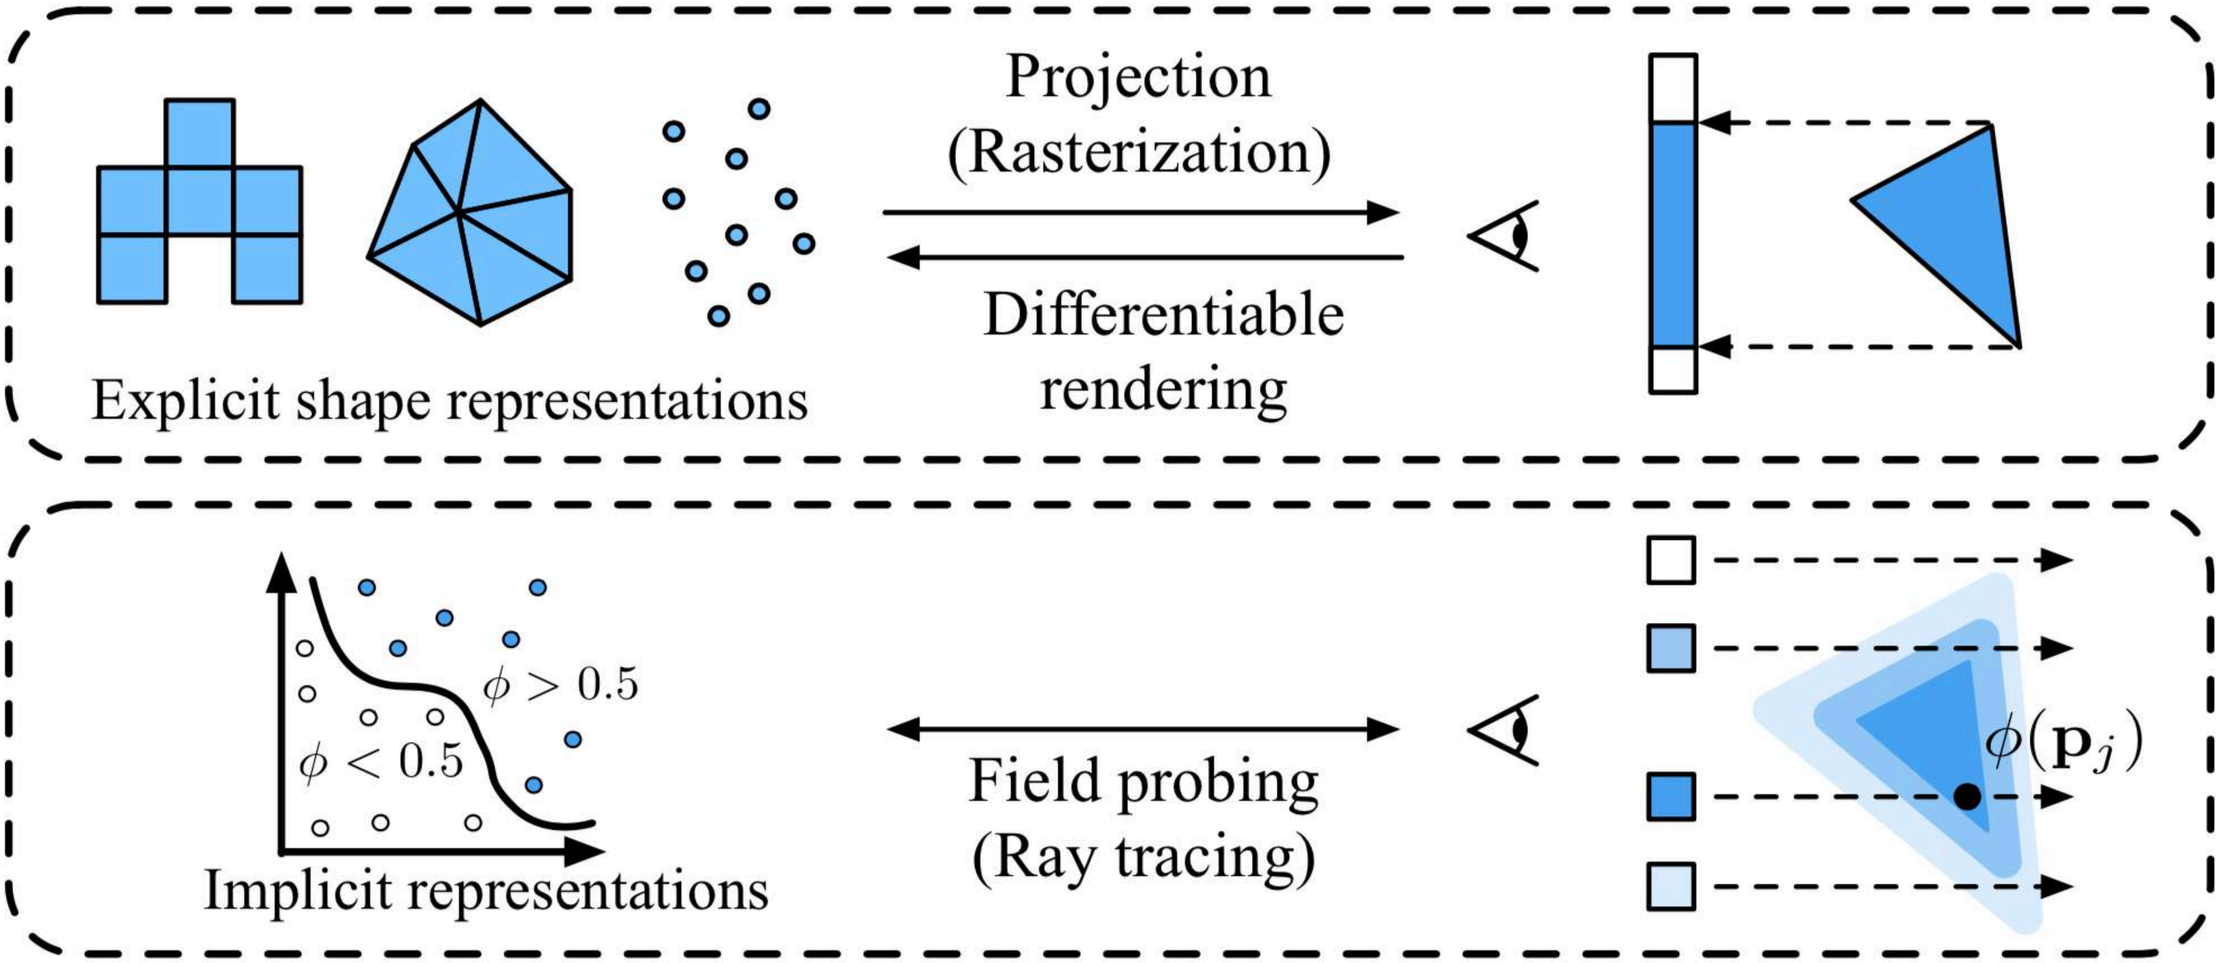

Differentiable Rendering (Tables)

Input Image

Reconstruction

Input Image

Reconstruction

Liu et al. 2019

Differentiable Rendering (Chairs)

Input Image

Reconstruction

Input Image

Reconstruction

Liu et al. 2019

Differentiable Rendering (Airplanes)

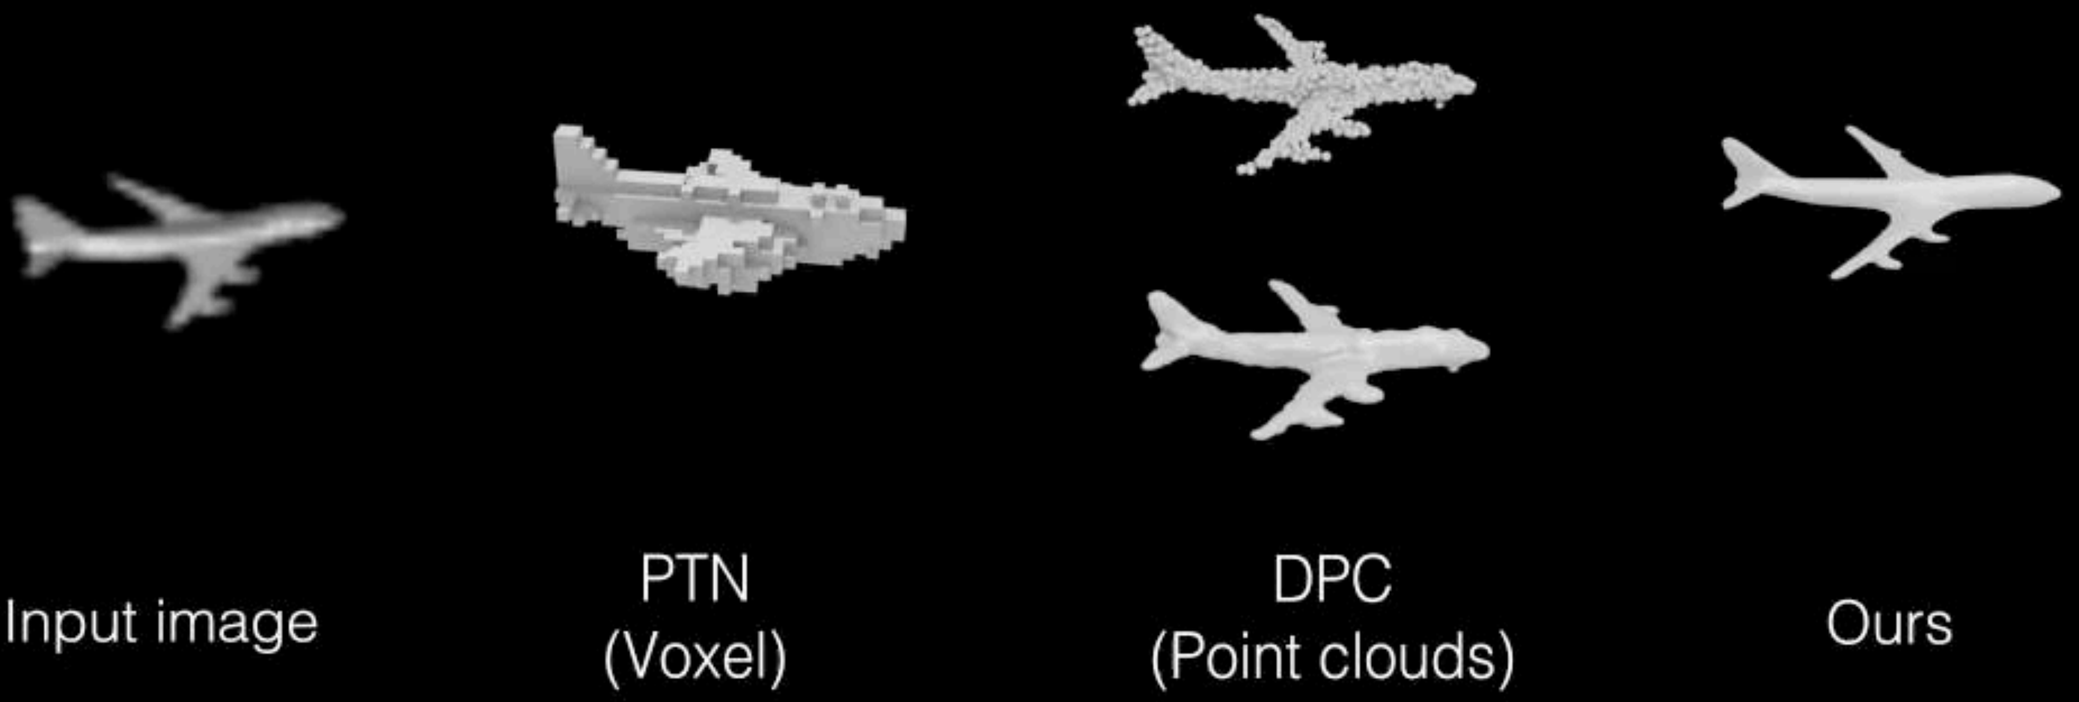

Liu et al. 2019

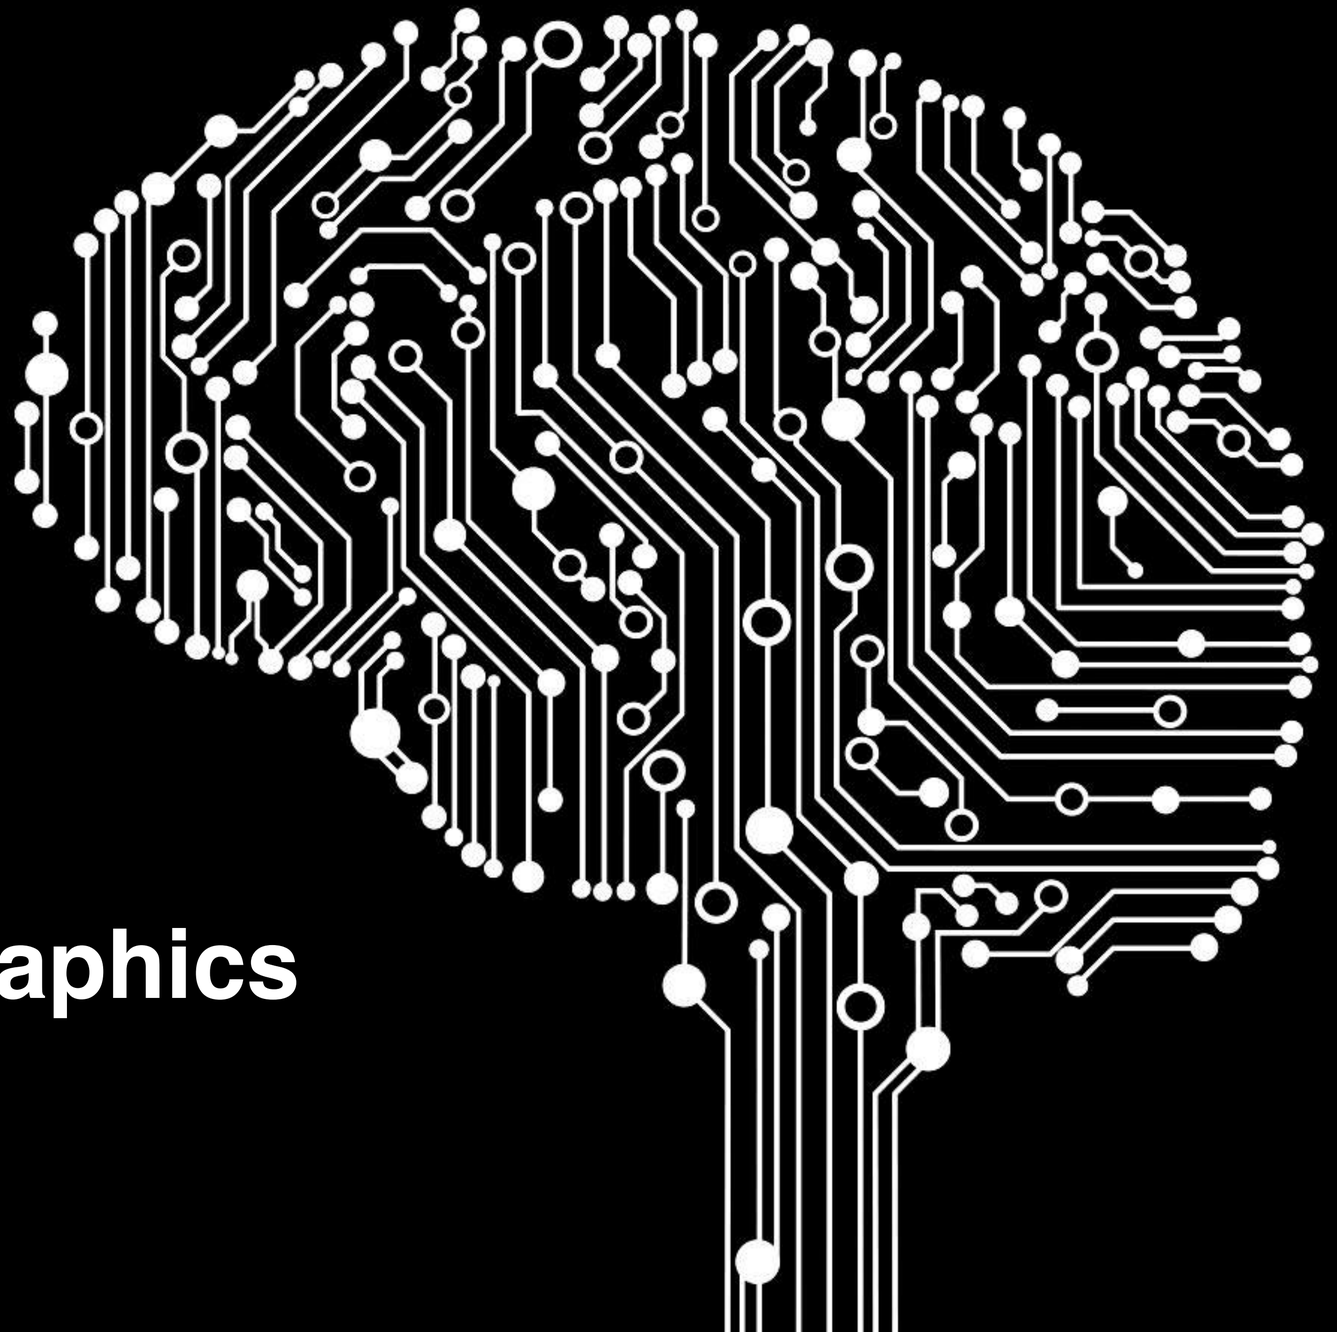

AI-Driven Graphics


YouTube
Generate

Disinformation & Privacy Concerns are Real

EDITOR'S PICK | 804,508 views | Jul 17, 2019, 07:20am

FaceApp: Is The Russian Face-Agi App A Danger To Your Privacy?

Thomas Brewster Forbes Staff

Cybersecurity

I cover crime, privacy and security in digital and physical forms.

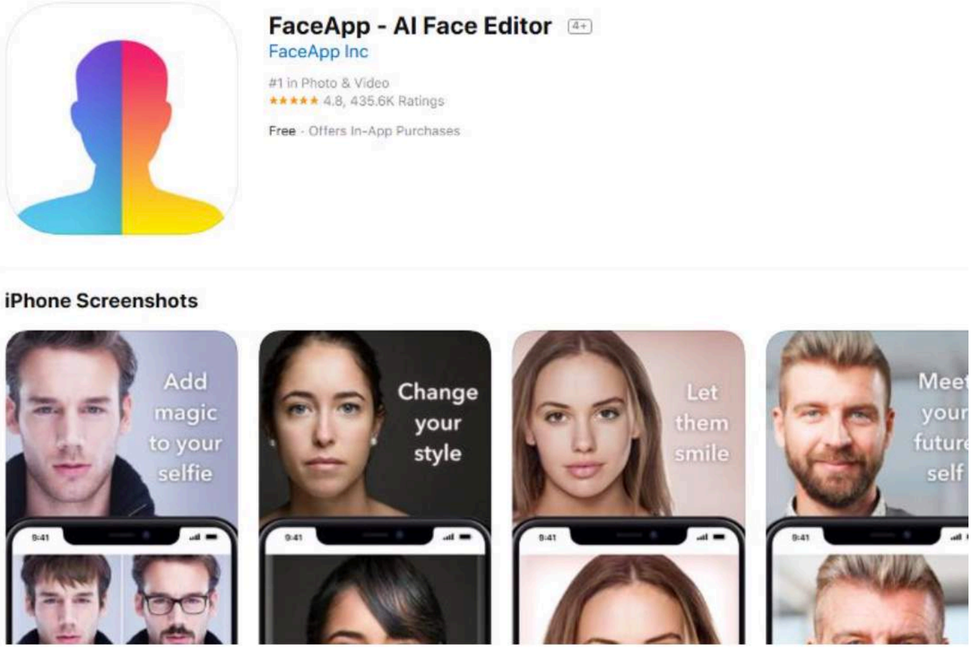

FaceApp is a massively popular face-altering app for Android and iOS, but there are privacy concerns a tech. FORBES

Another convincing deepfake app goes vira prompting immediate privacy backlash

Insert your face into TV shows or movies with just a single photograph

By Jon Porter | @JonPorty | Sep 2, 2019, 6:32am EDT

SHARE

Image: @AllanXia

Virtualization / Personalization

From Virtual to Physical

Thanks!