3D Surface Representations

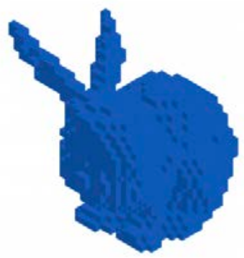


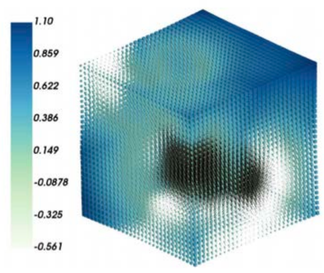

	Voxels	Explicit surface mesh	Point sets	Continuous implicit fields
High frequency details?	--	++	+	++
Arbitrary topology?	+	-	+	++
Regularity?	+	+	-	++

There are many applications at which explicit representations excel:

- High-quality rendering in computer graphics.
- Precise modeling of biological structures from biomedical data.
- Computational fluid dynamics in computer assisted design.

But:

- Their topology is fixed.
- They are not particularly deep learning friendly.

—> Implicit Surface Representations

Reminder: Level Sets

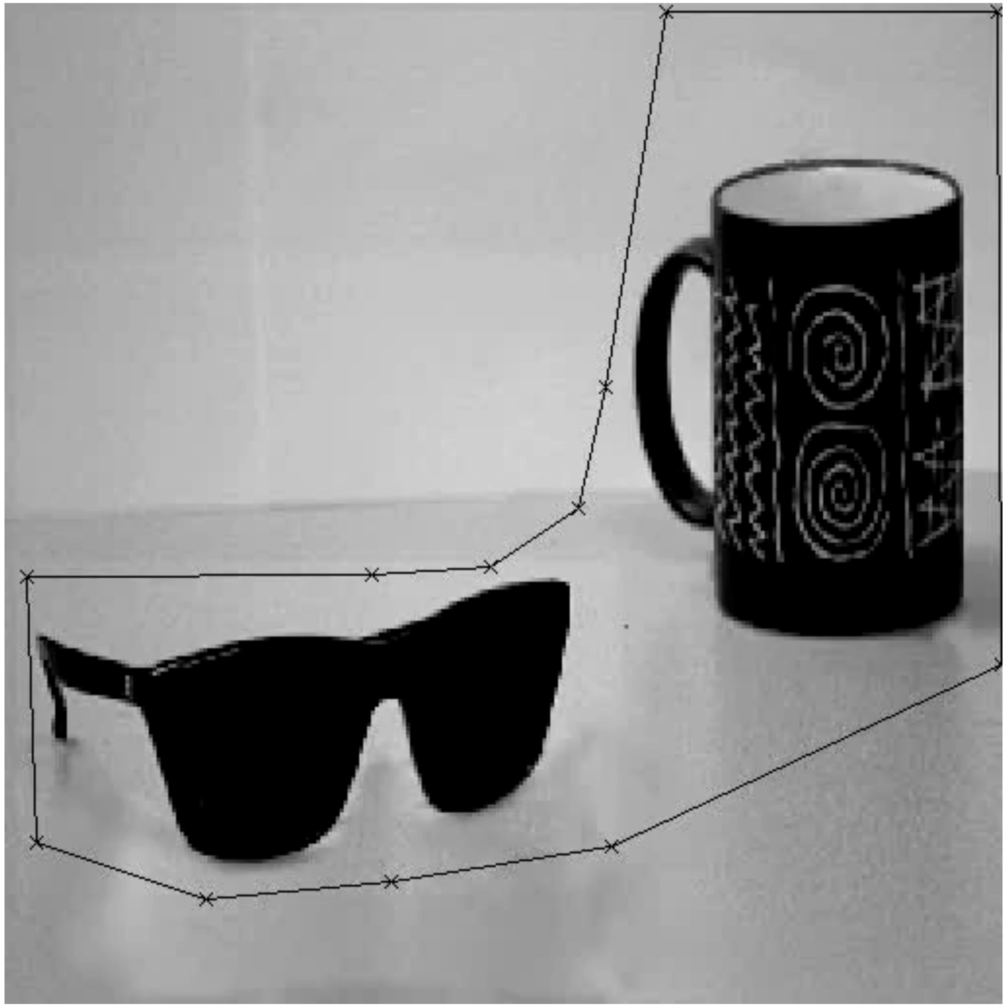

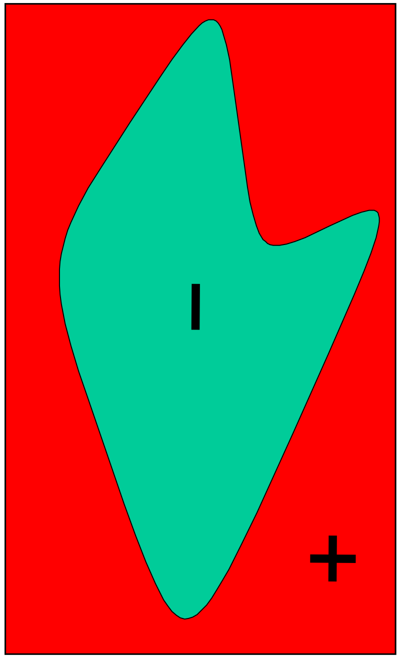

$z = \Phi(x, y, t)$,

$z > 0$ outside,

$z < 0$ inside.

Signed Distance Fields (SDF)

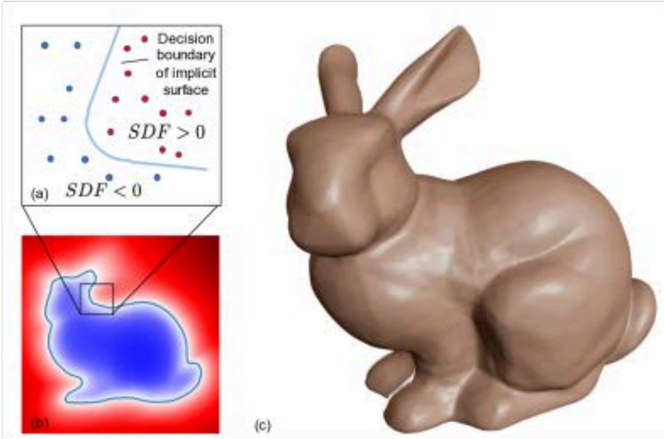

- Represent a 3D surface S by the zero crossings of a signed distance function

$$f\colon \mathbb{R}^3 \rightarrow \mathbb{R}$$

$$\forall \mathbf{x} \in \mathbb{R}^3,\ f(\mathbf{x}) \text{ is the signed distance to the surface.}$$

- Such surfaces can easily change topology, which is harder to do with explicit surface representations.

- SDFs have long been appealing in theory but hard to use in practice because it it was necessary to store the 3D values of f in a cube like structure until ….

Deep SDF

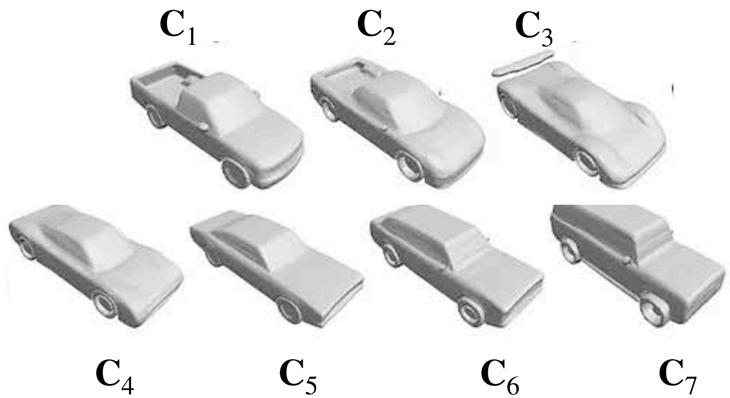

$\mathbf{x} = (x, y, z)$

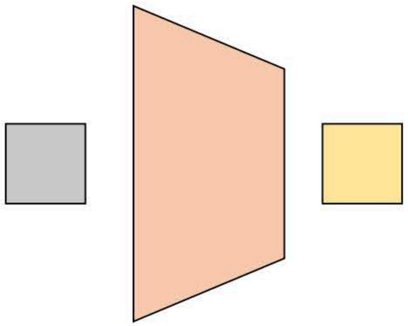

$s = f_\theta(\mathbf{x})$

C

$\mathbf{x}$

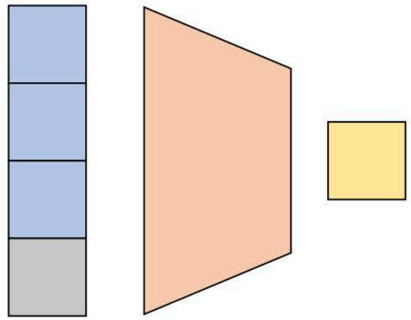

$s = f_\theta(\mathbf{x} \mid C)$

Single Shape DeepSDF

Coded Shape DeepSDF

[Park et al., CVPR'19]

Learning a Latent Vector Model

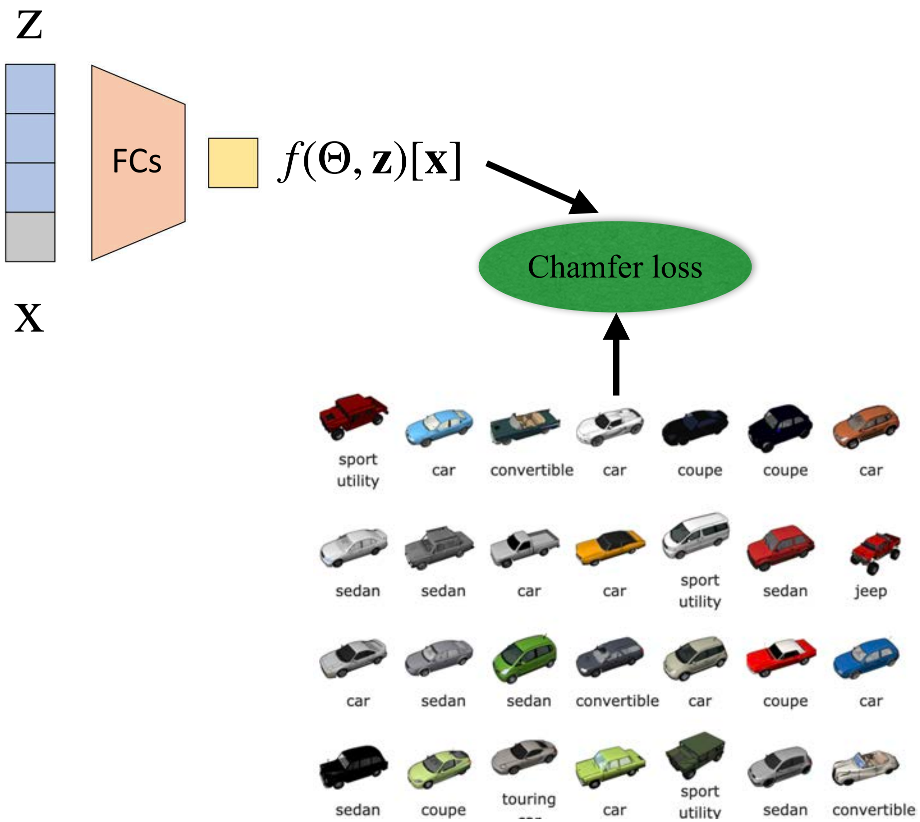

Train an auto-decoder using a set T of 3D car models:

$$\Theta^*, \mathbf{Z}^* = \operatorname{argmin}_{\Theta, \mathbf{z}, \ldots, \mathbf{z}_L} \mathscr{L}_{model}(\Theta, T) \ ,$$

$$\mathscr{L}_{model}(\Theta, T) = \sum_{l=1}^{L} \left[\sum_{j=1}^{K_j} loss(f(\Theta, \mathbf{z}_l)[\mathbf{x}_l^j], s_l^j) \right] + \lambda\, reg(\mathbf{z}_l)$$

Data Term

Regularization Term

End-to-End Differentiable Pipeline

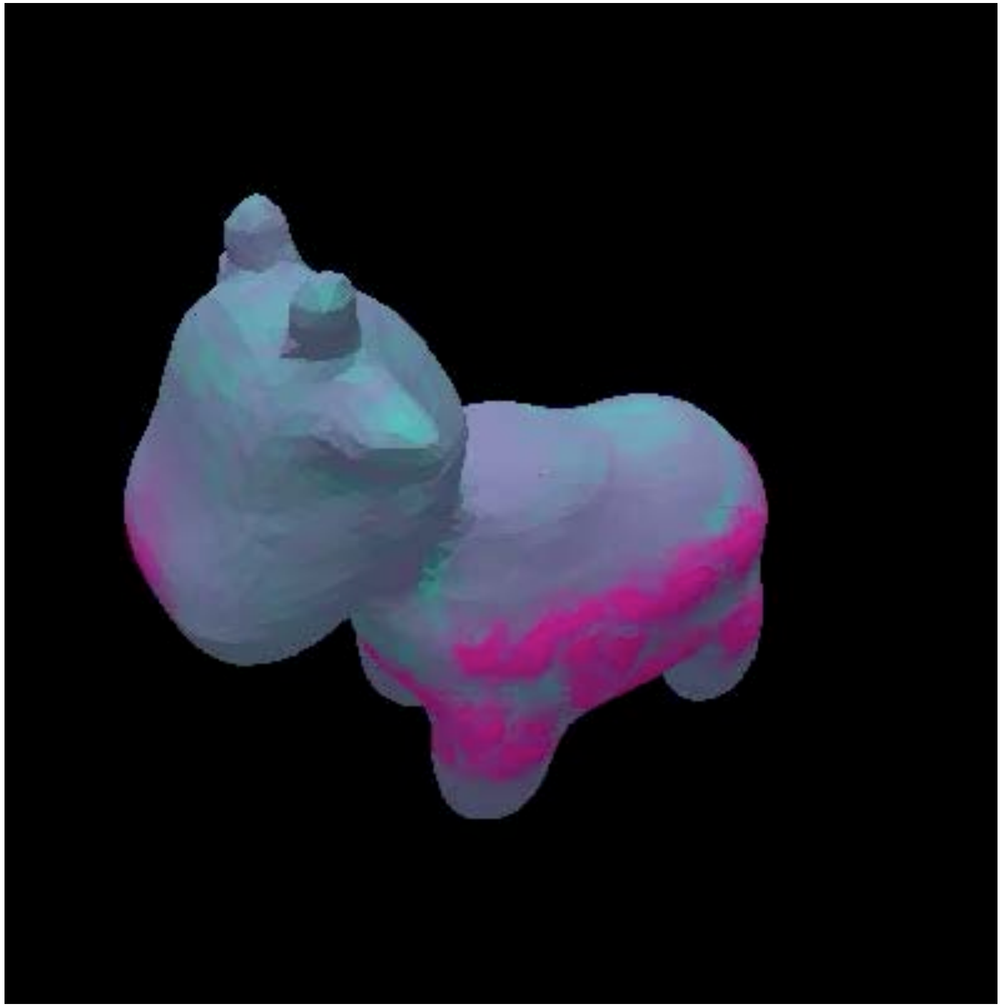

1. Start with a Deep SDF code.
2. Use marching cube to compute vertices and facets.
3. Use them for the forward pass and for backpropagation.
4. Update the SDF code and iterate.

—> We can turn a genus 0 cow into a genus 1 duck by minimizing a differentiable objection function.

From Silhouettes to 3D Shapes

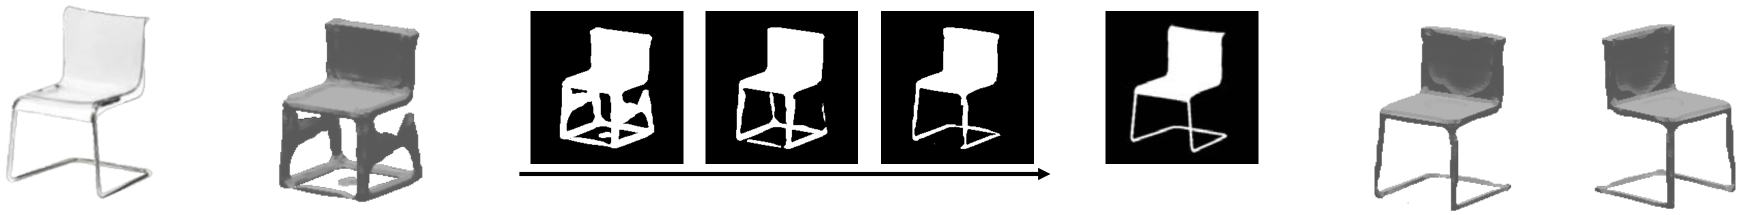

3D Model from Image

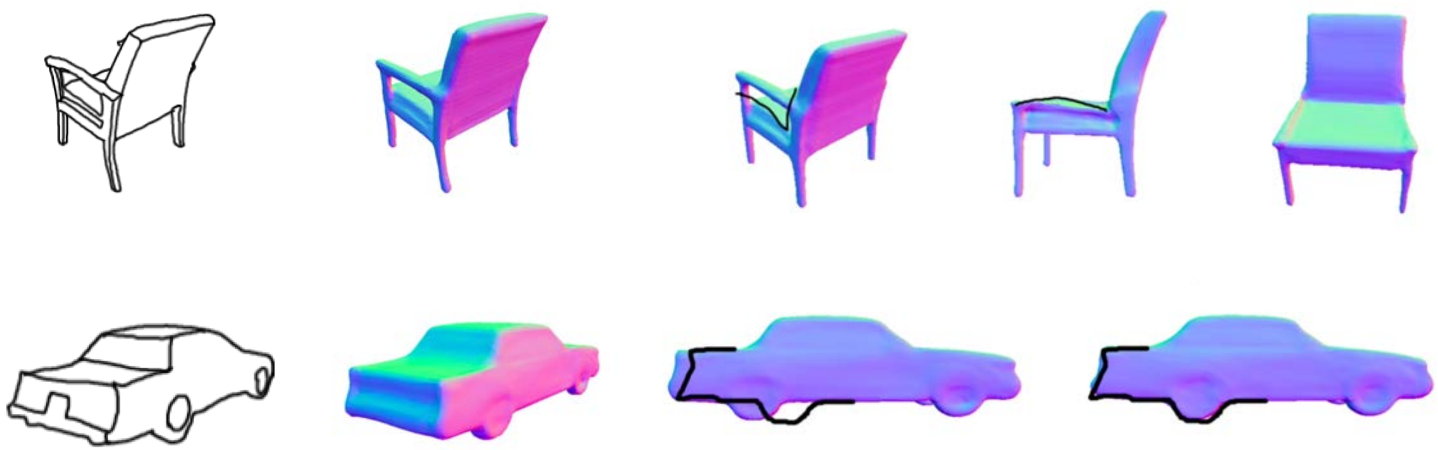

Editable 3D Model from Sketch

Non Watertight Surfaces

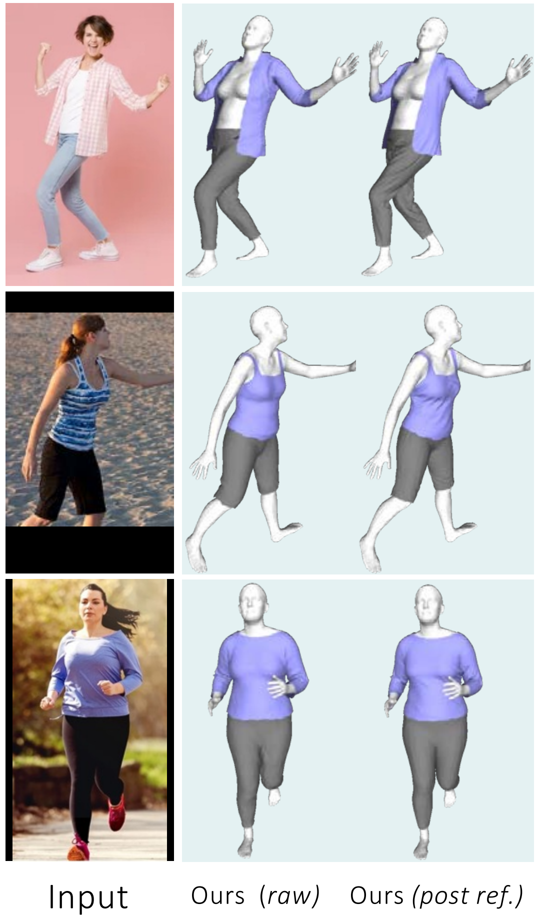

Replace SDF by UDF:

- Adapt the Marching Cube algorithm.
- The derivatives can still be computed.

➡ More accurate reconstructions.

sDF vs uDF

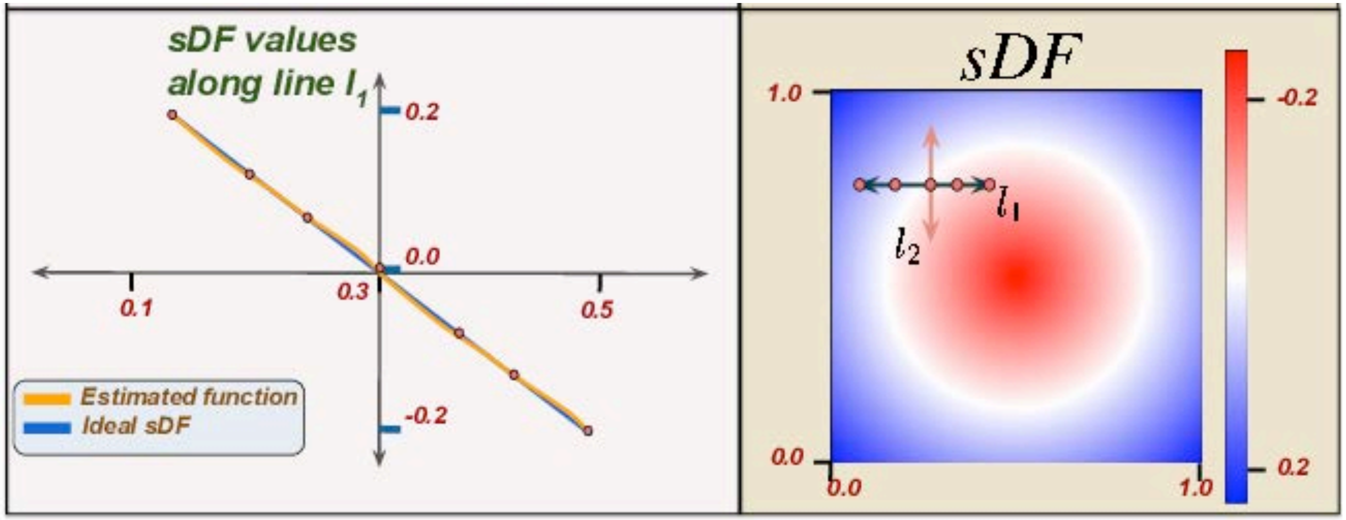

- Numerically stable
- Only watertight surfaces

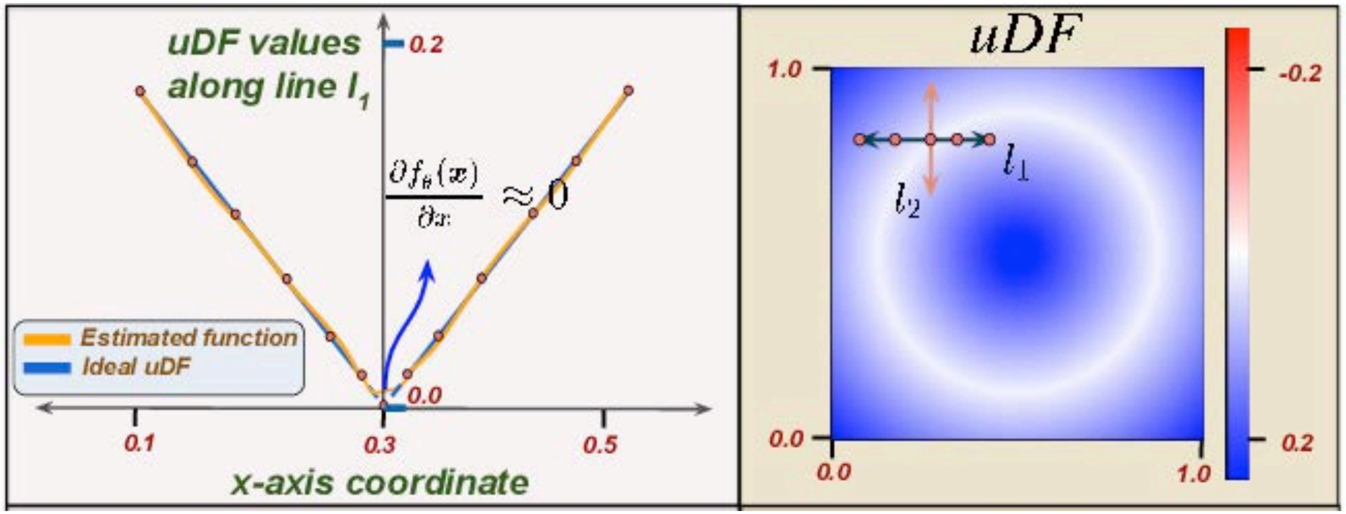

- Less stable numerically
- Can handle open surfaces

Marching Cubes

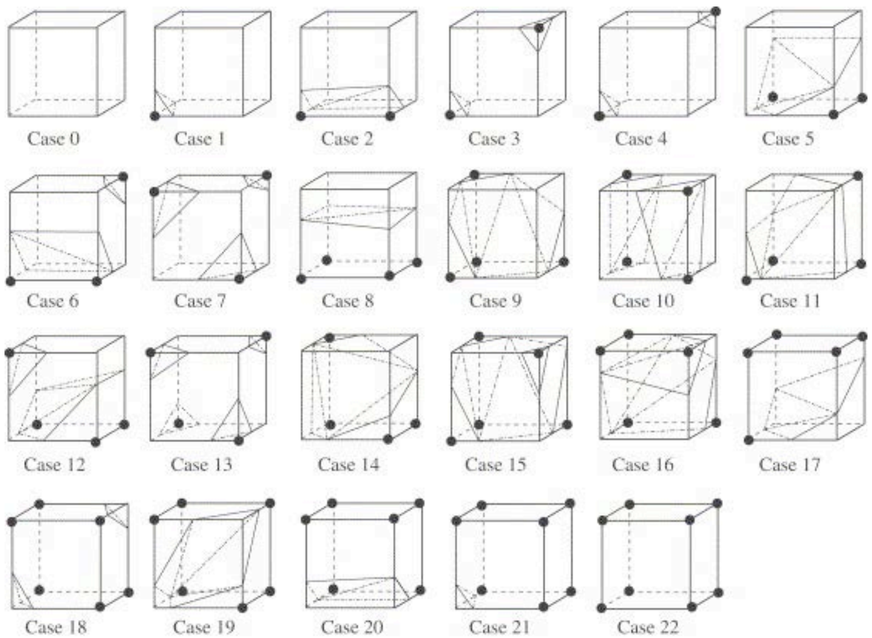

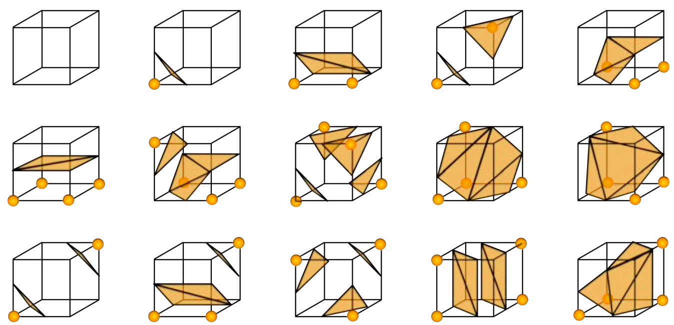

Check for sign changes at the corners

Create surface within voxels where appropriate

Will not work for UDFs!

From UDF to local SDF

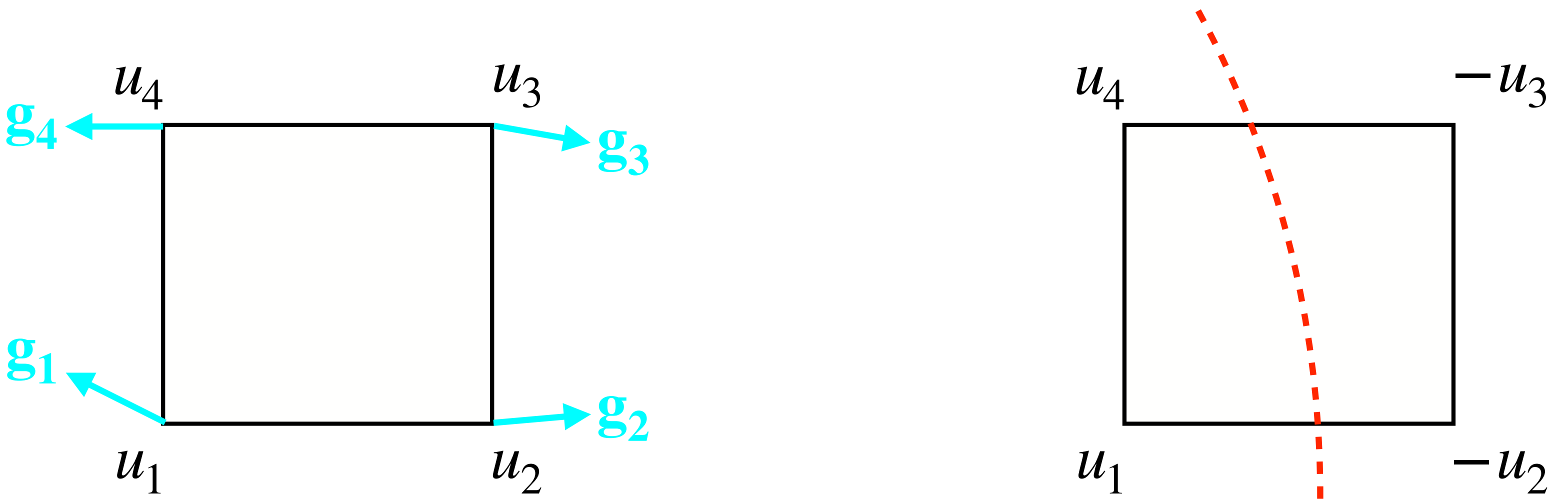

$$\mathbf{g} = \nabla u$$
$$s_i = (g_1 \cdot g_i) u_i$$

UDF vs SDF

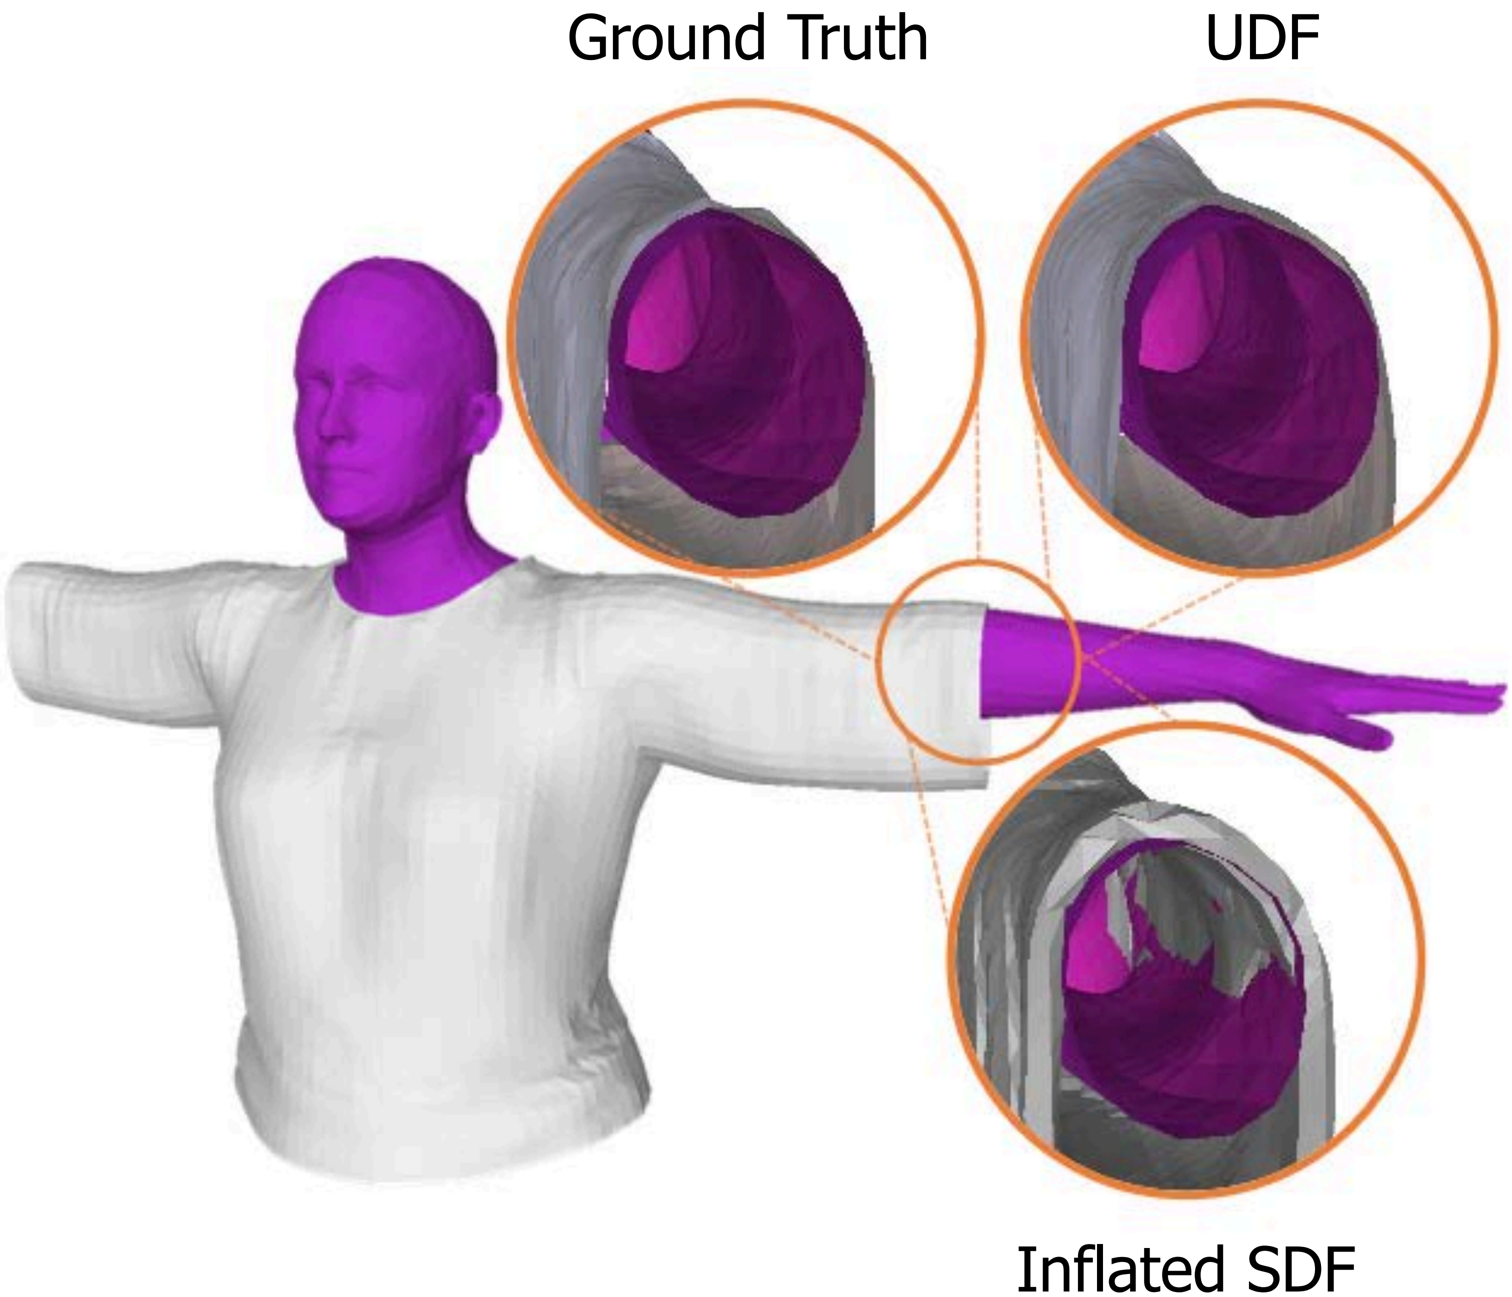

Draping the Garment on the Body

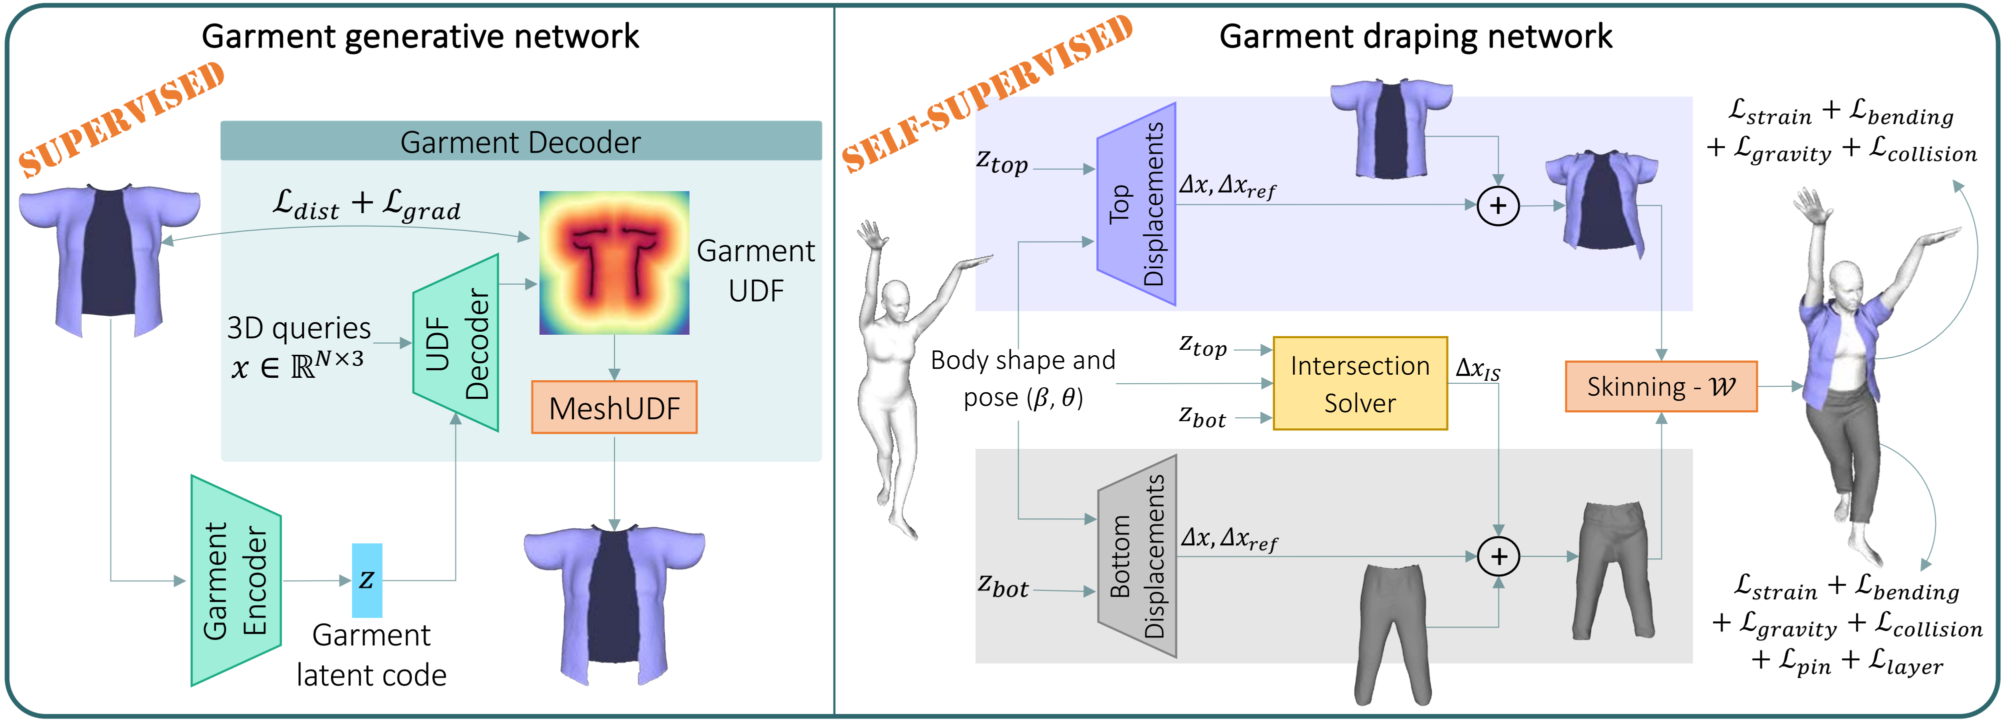

- Generate a garment in a rest pose.
- Drape the garment in a different pose.
- Use physics to provide self-supervision.

Synthesize and Fit

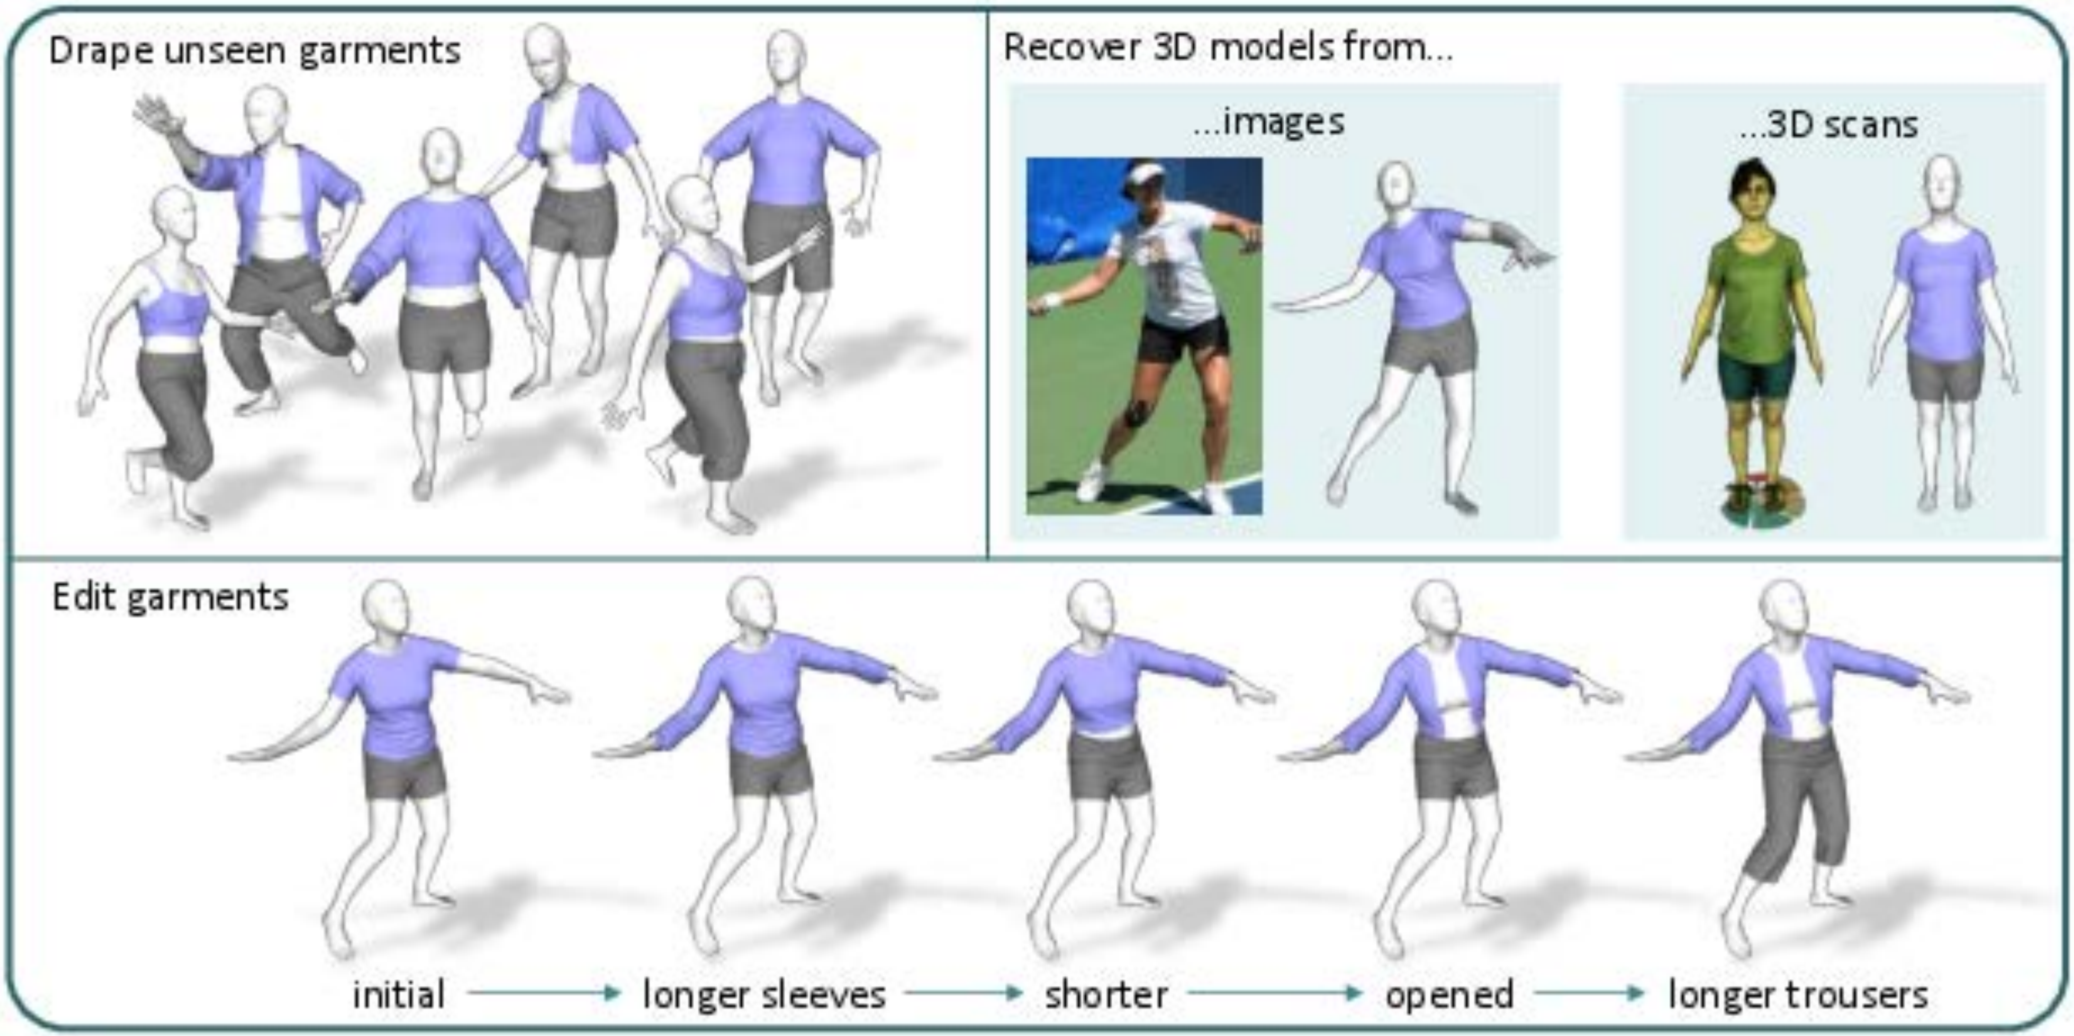

Garment Panels

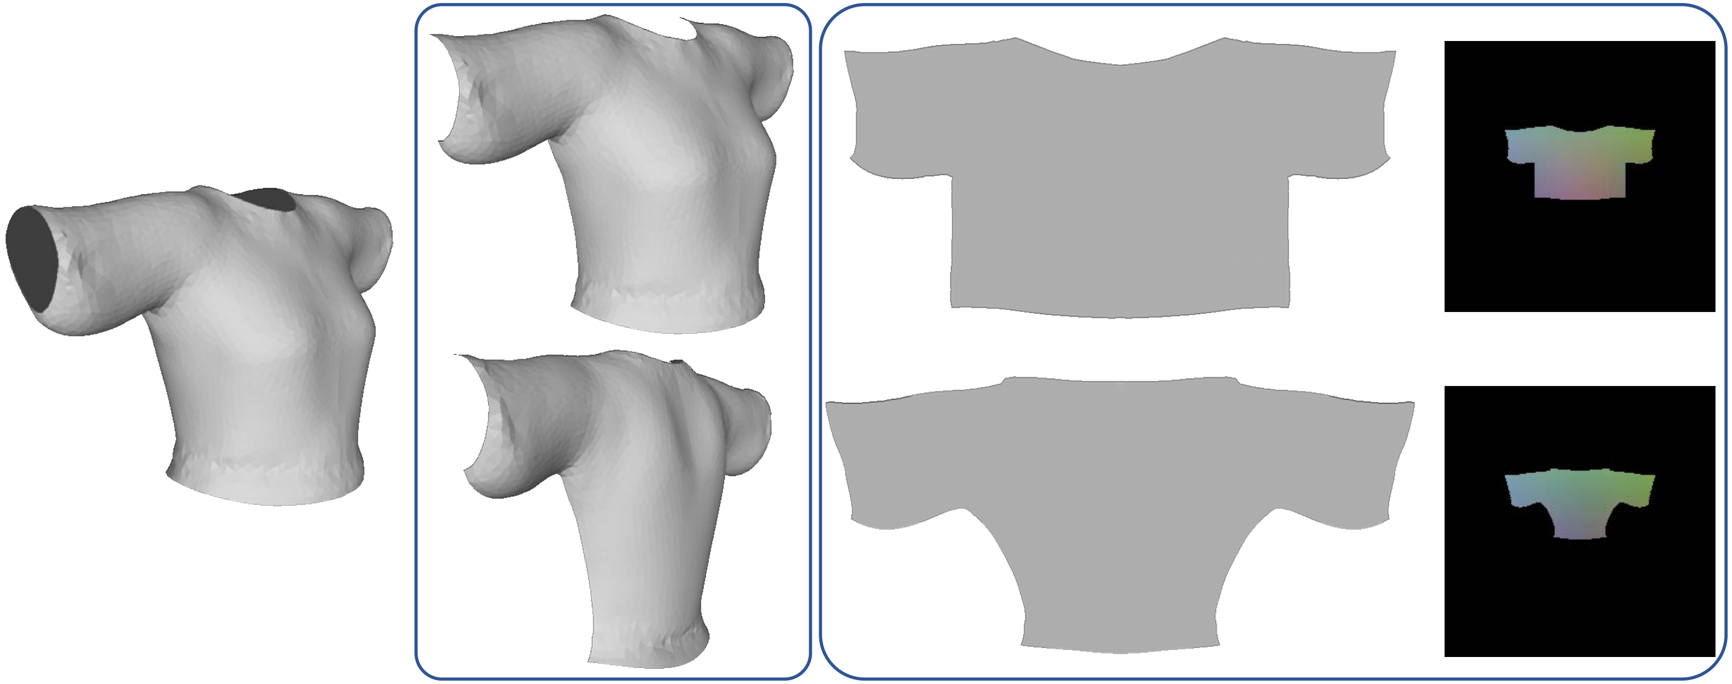

3D Garment

Front and back panels

2D shape plus UV map

Directly inspired but the way garments are models in the clothing industry.

Image Fitting Pipeline

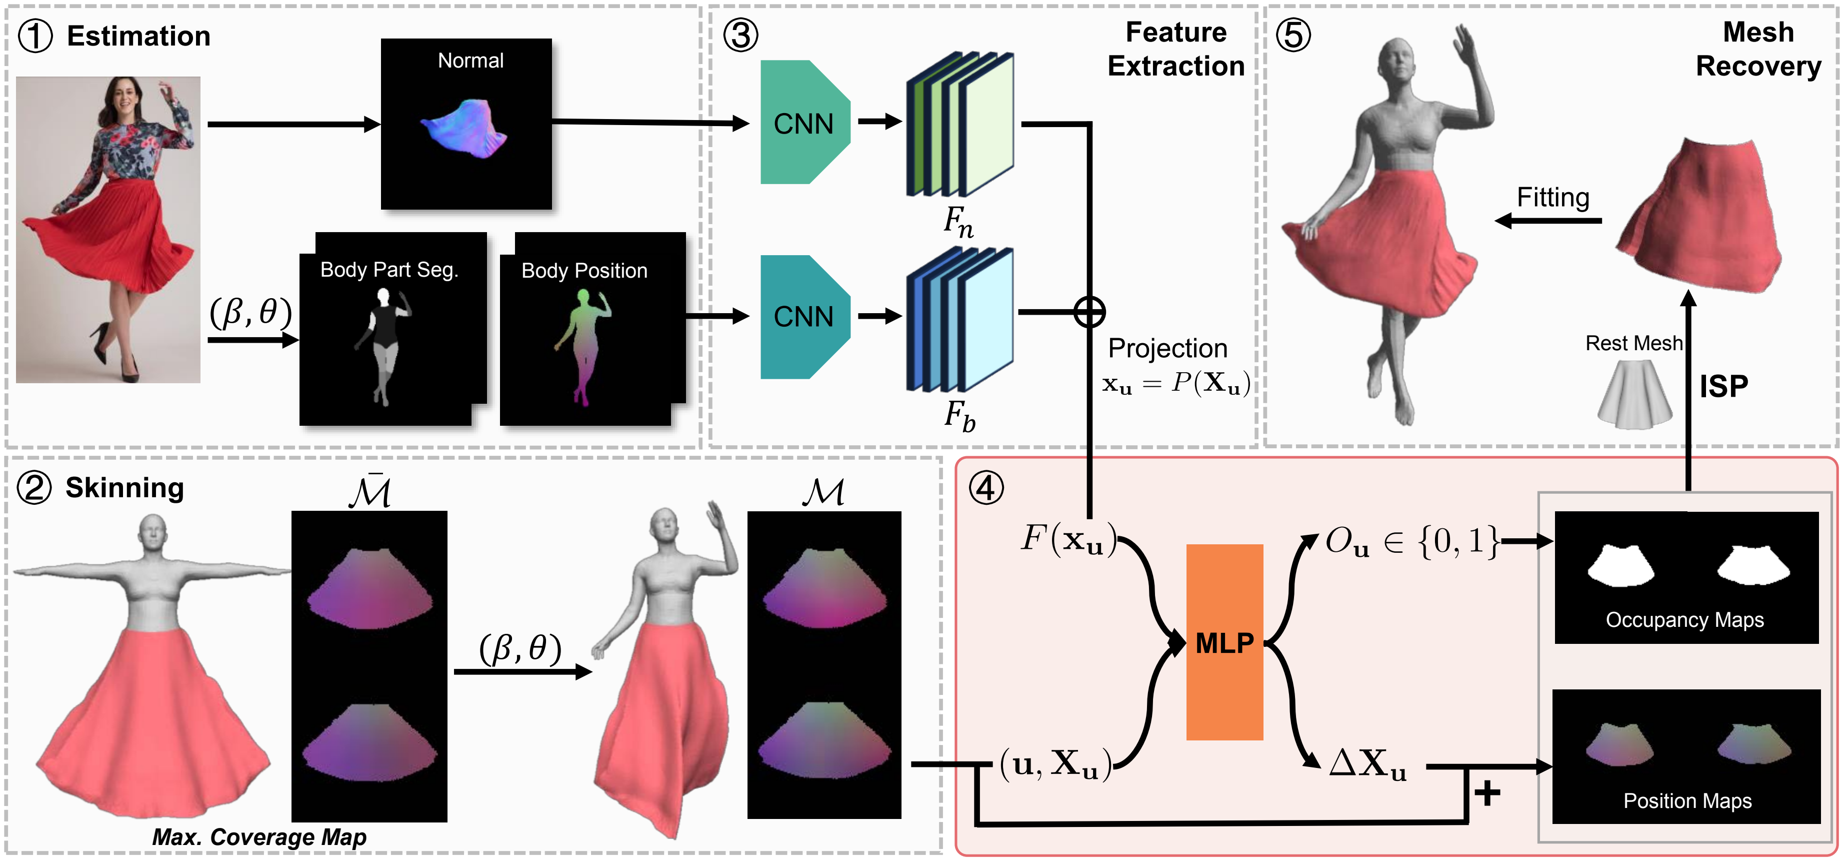

Given an image:

1. Compute image features and estimate body pose.
2. Fit the garment to closely fit the body.
3. Use pixel-aligned image features to deformation model.

Clothed People from Images

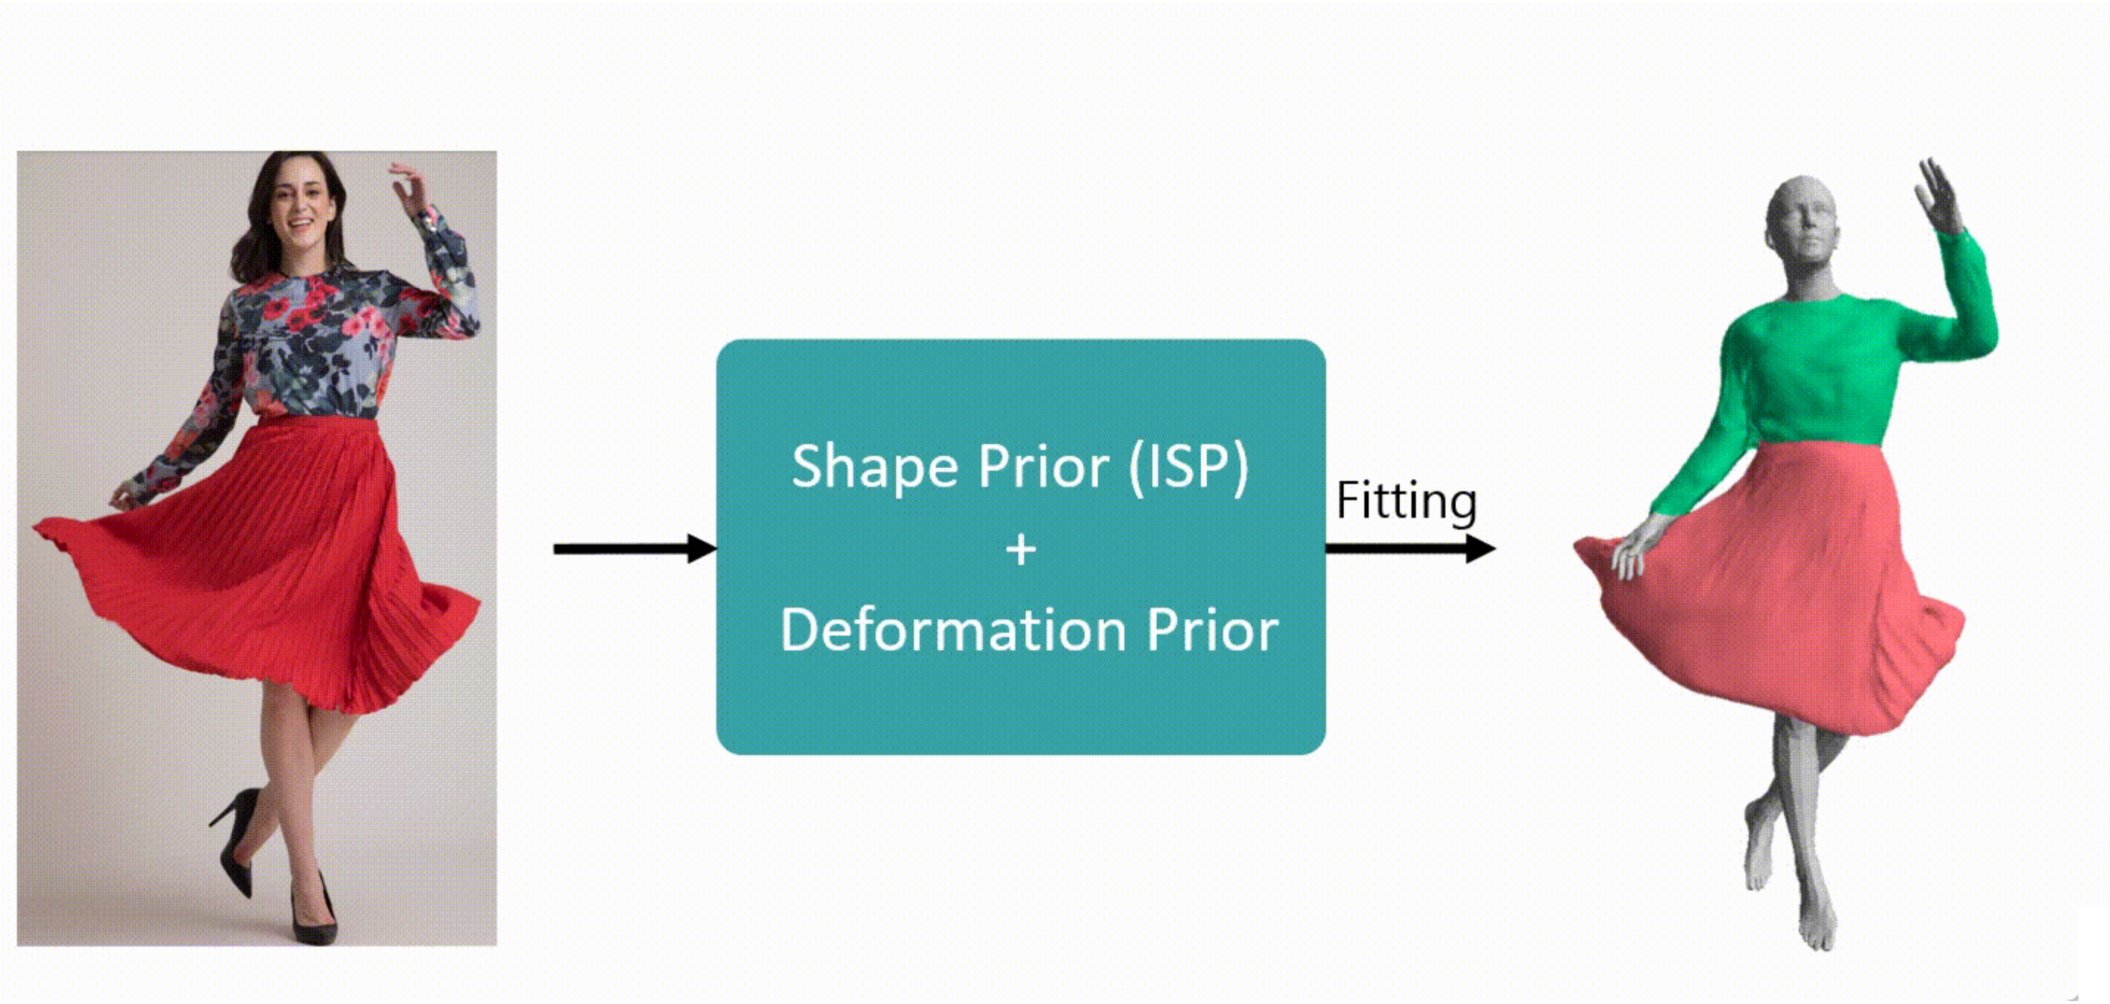

Simulation

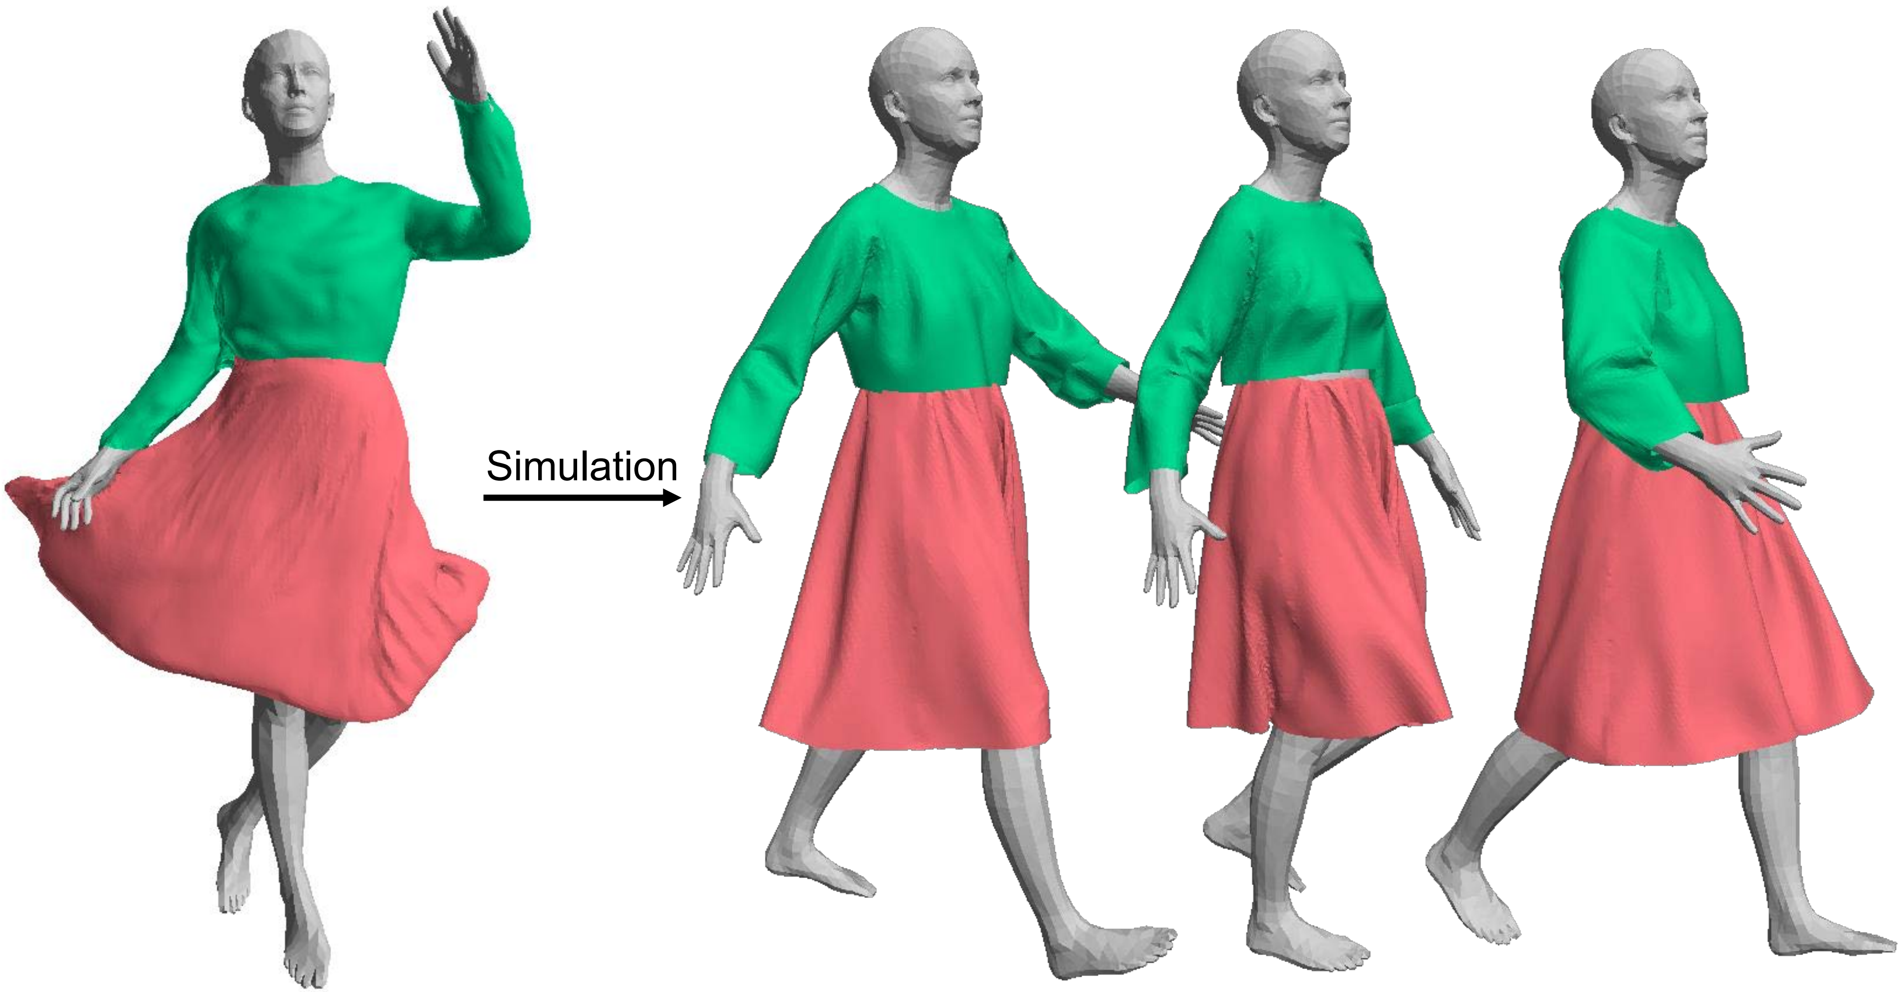

Better Clothed People from Images

- Model surfaces in terms of their distance to implicit sewing patterns.
- Add a deformation model to let garments move away from the body.
- Use a diffusion process to generate the hidden parts.

Li et al., SIGGRAPH'25

Clothed People from Video

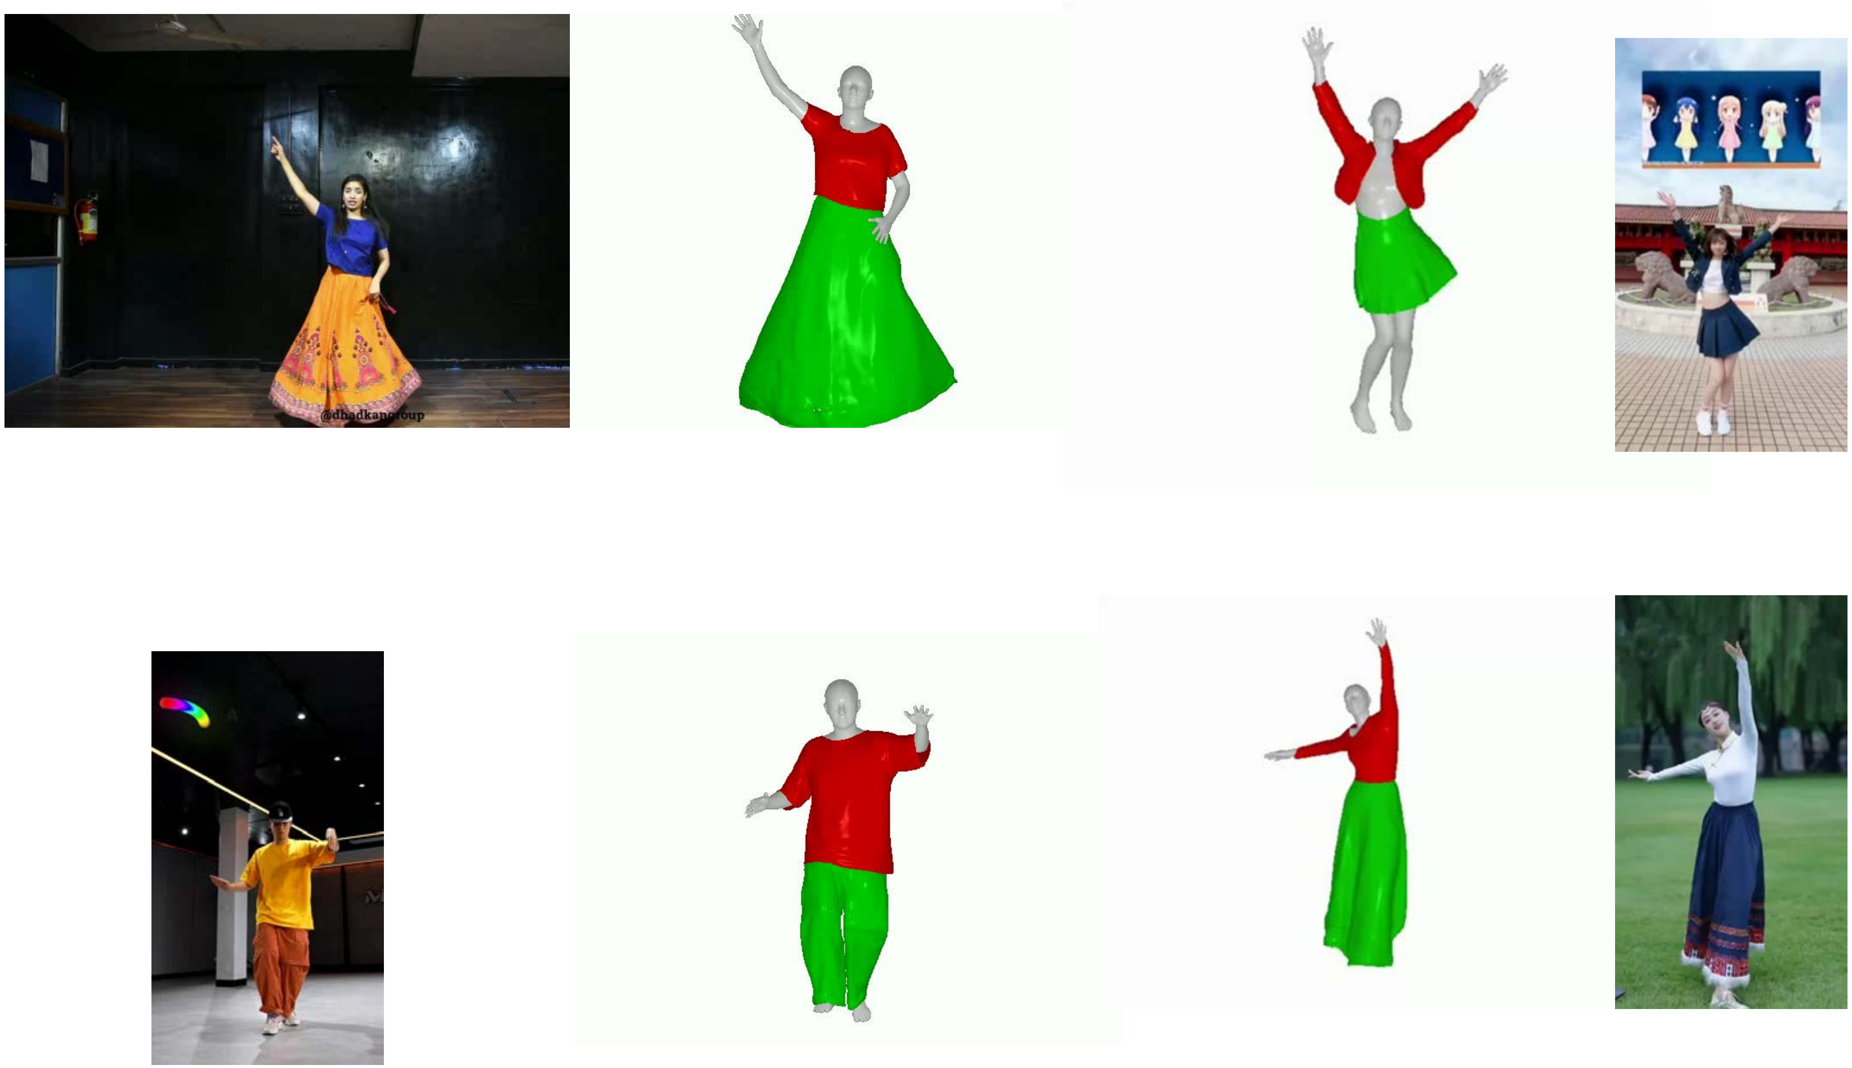

The diffusion has become temporal

3D Models for Endoscopy

Colonoscopy

Gastroscopy

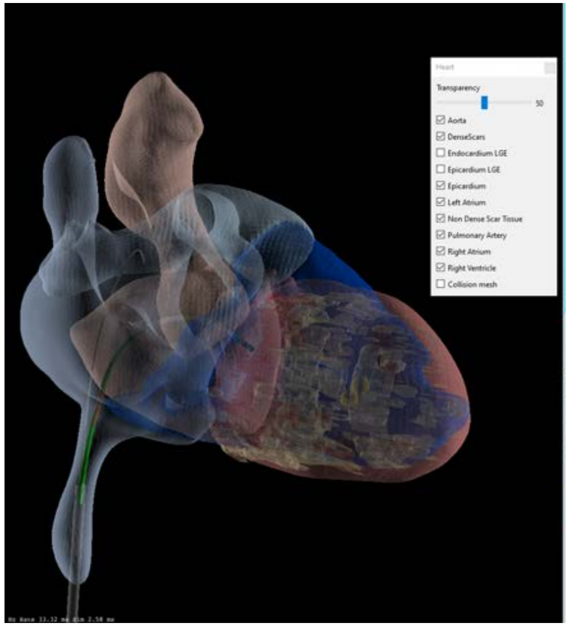

Heart surgery

Deep learning techniques dominate but:

- There rarely is enough training data.
- Anatomical and physical knowledge matter.

Surgery Simulator

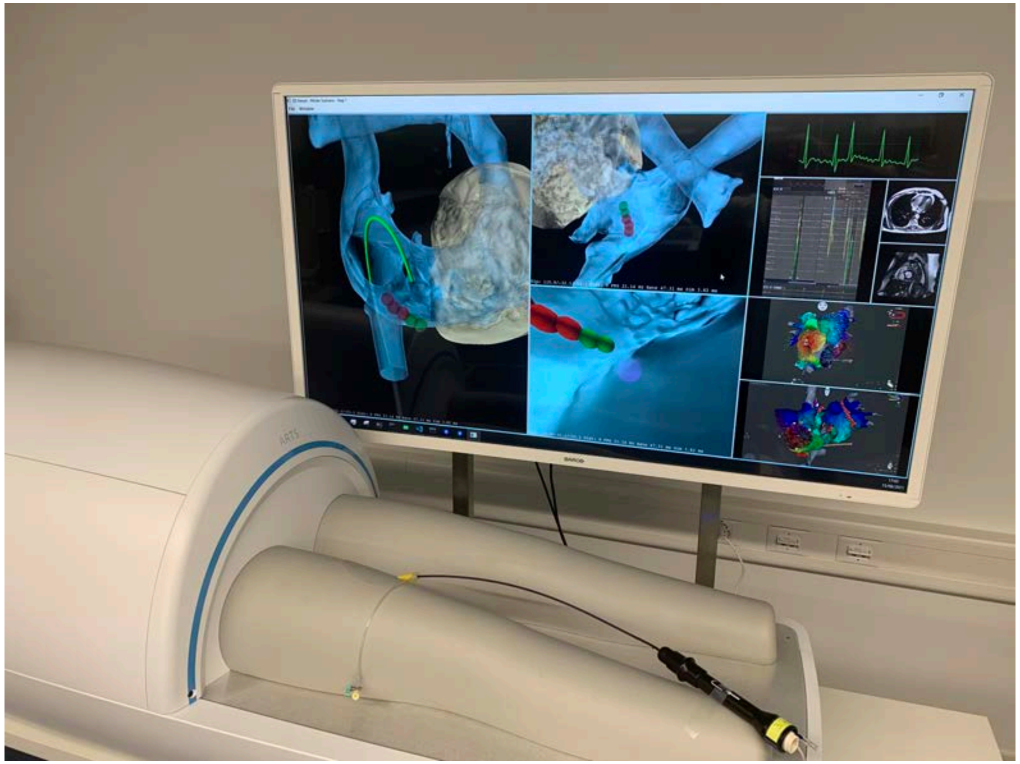

- Real catheter
- Virtual heart

➡ One is inserted into the other.

Tech Transfer

ADIS ADVANCED INTERACTIVE SYSTEMS

CHUV Centre hospitalier universitaire vaudois

➡ Heart models need to be built from MRI data.

From Simulation to Intervention

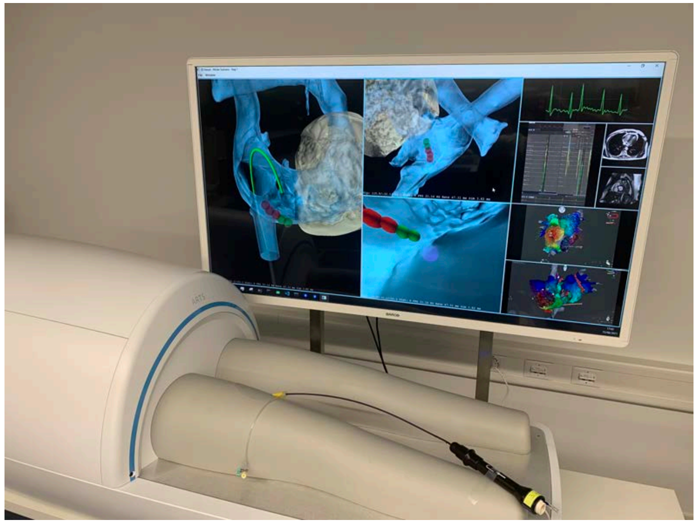

- For simulation purposes, a static model can be animated.
- For intervention purposes, dynamics must be captured in real time.
- There is no time to acquire full volumetric data.

➡ Build a static 3D model before the intervention and animate during the intervention using 2D slices.

MedTet Pipeline

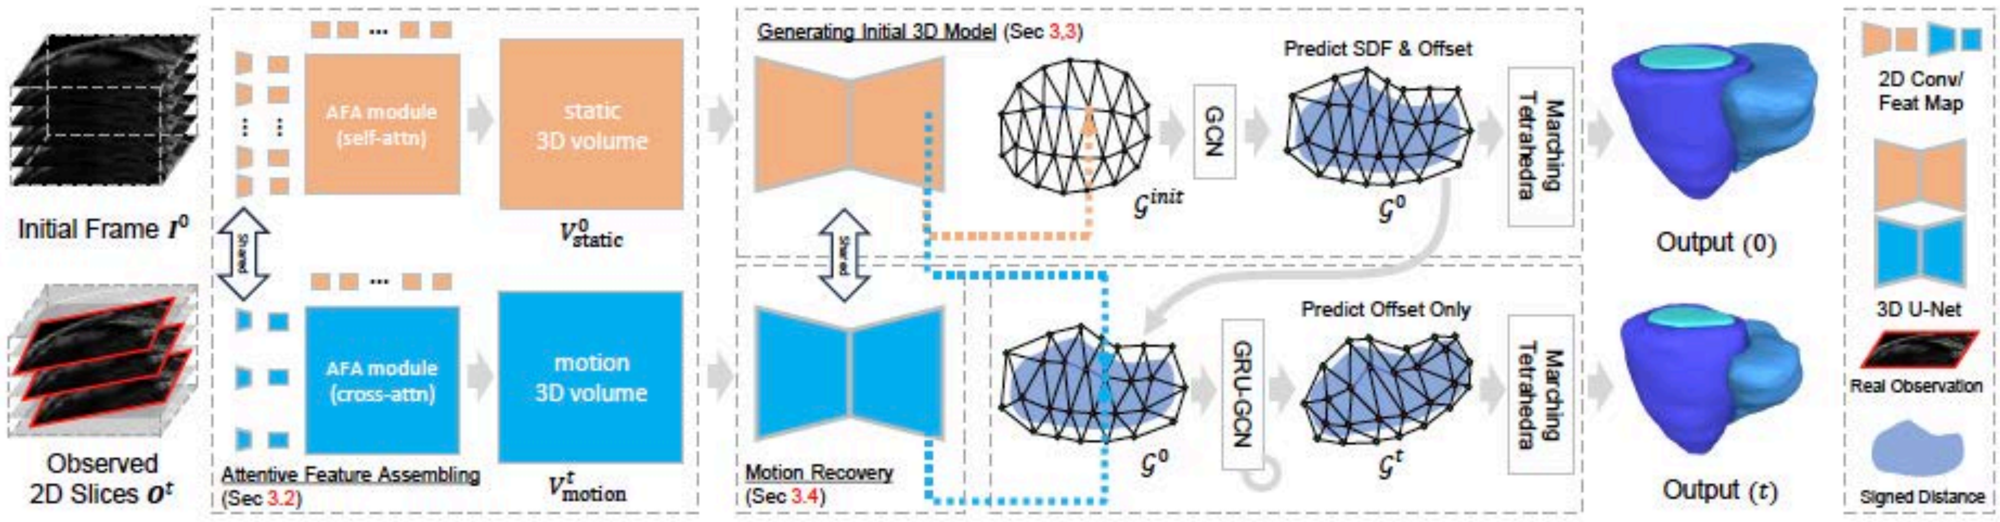

- Implicit functions defined on tetrahedral meshes to model the heart.
- A full stack is fed into the network to initialize them at time t=0.
- A time t >0 , only a few 2D slices, or even 1, are needed to estimate the deformations.

Weak Supervision

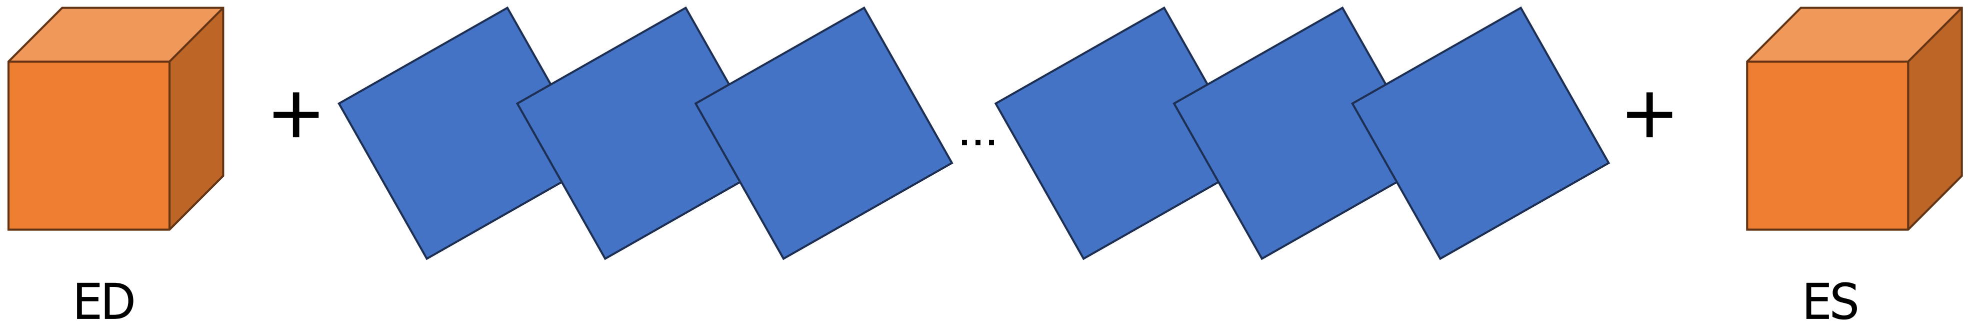

- Full 3D annotations for keyframes such as end-diastolic and end-systolic phases.
- For the intermediate frames, only unannotated 2D slices are available.

➡ Learn the 3D motion model from 2D slices in an unsupervised manner.

Moving Heart

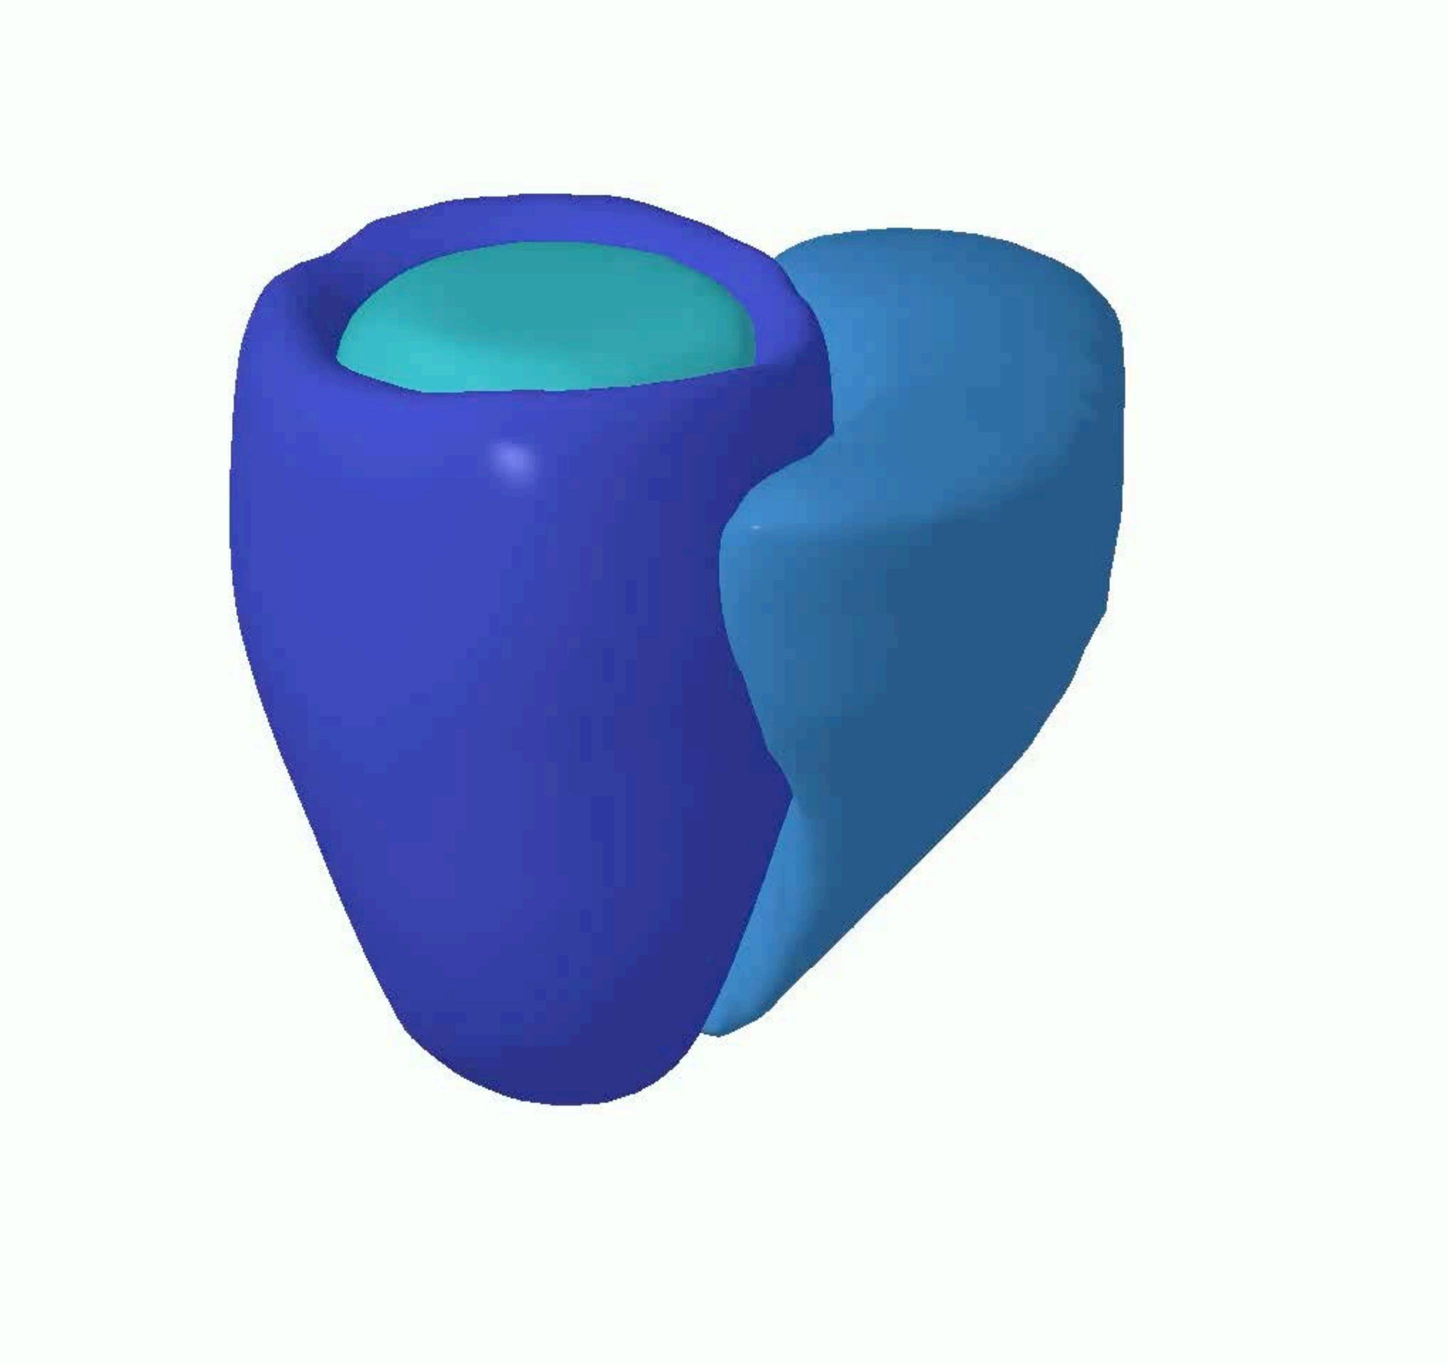

Quantitative Results

Normal heart:

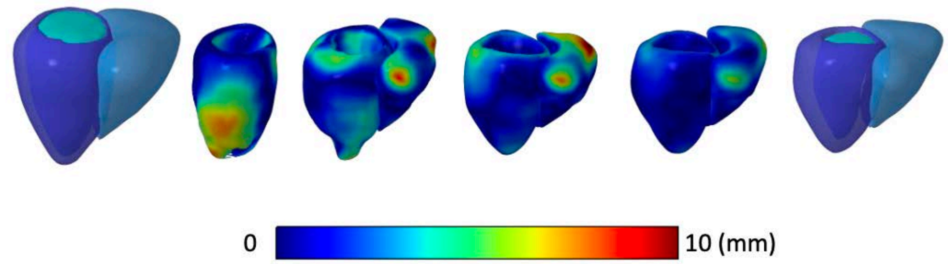

Dilated cardiomyopathy:

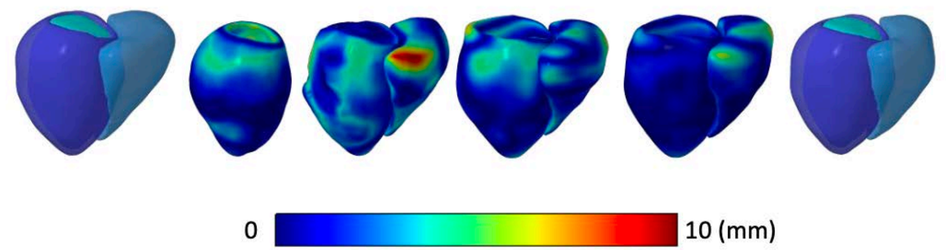

Accuracy as a function of the number of slices:

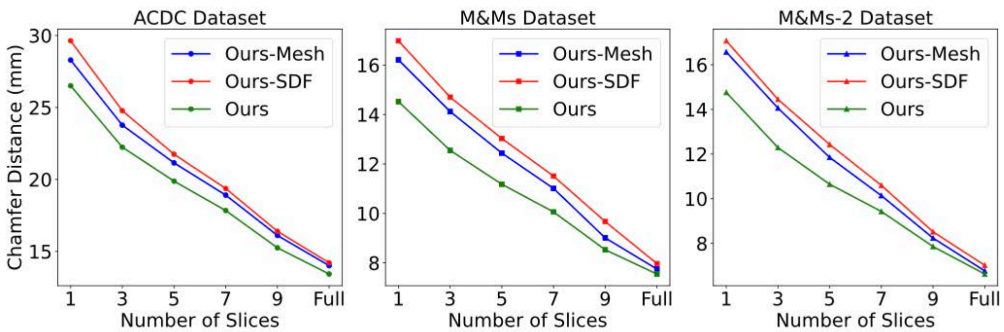

Transformative Idea

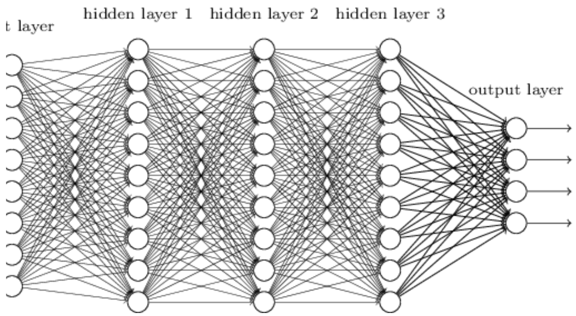

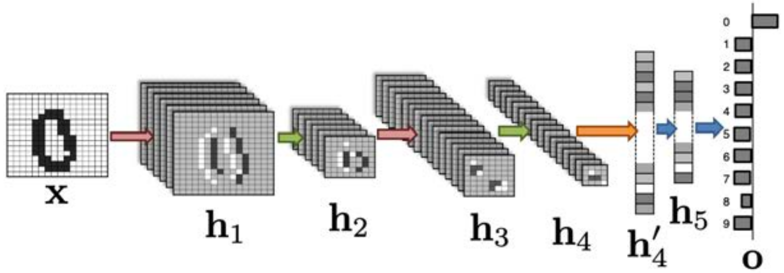

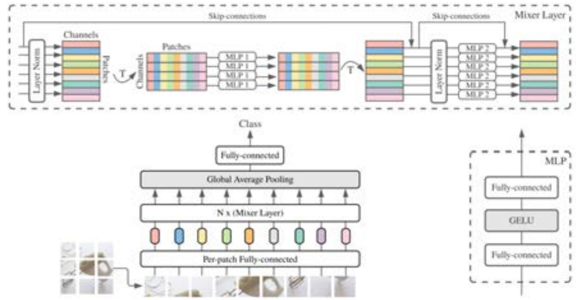

The tools we use for shape reconstruction can also be used for shape design!

Drag Prediction

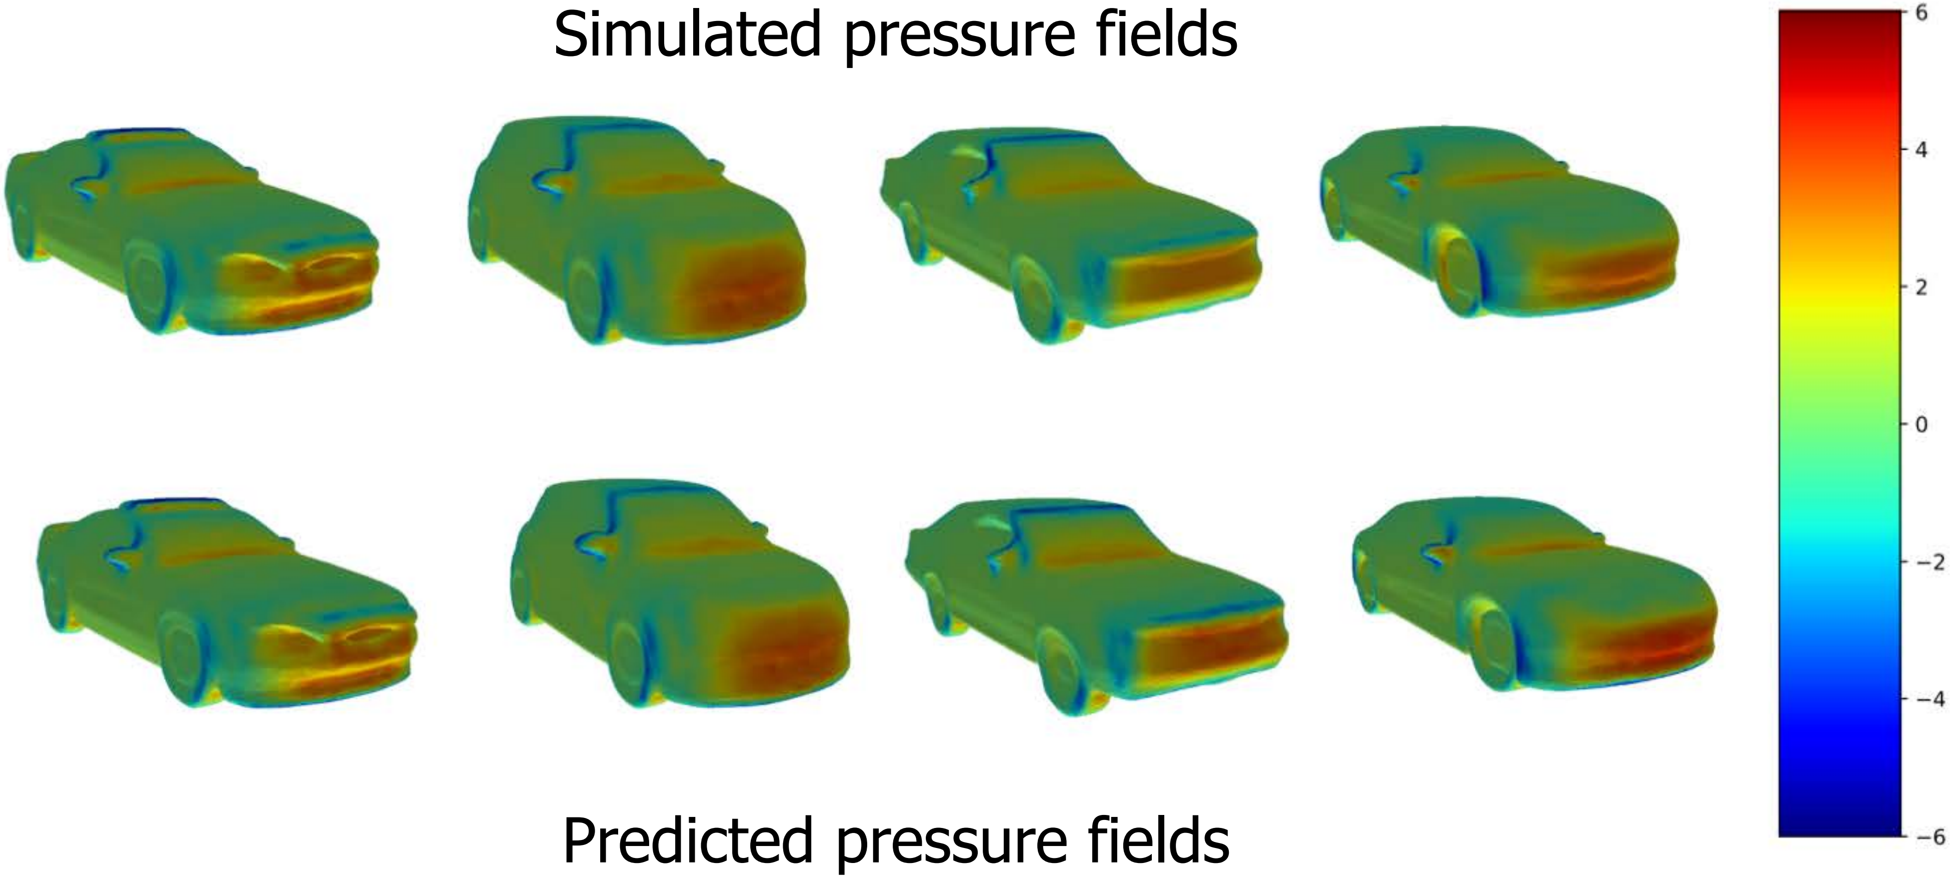

- A GCNN can be trained to predict the aerodynamic drag of a 3D shape.
- The drag becomes a differentiable function of the surface mesh vertices.

Drag Minimization

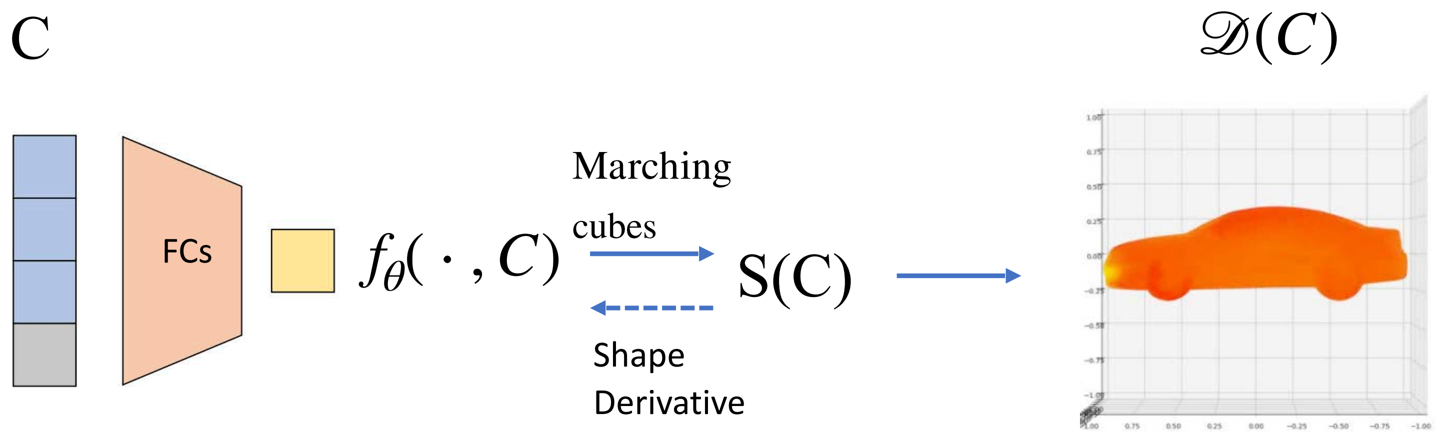

Minimize $\mathscr{D}(C)$ with respect to C under constraint.

From SUV to Sports Car

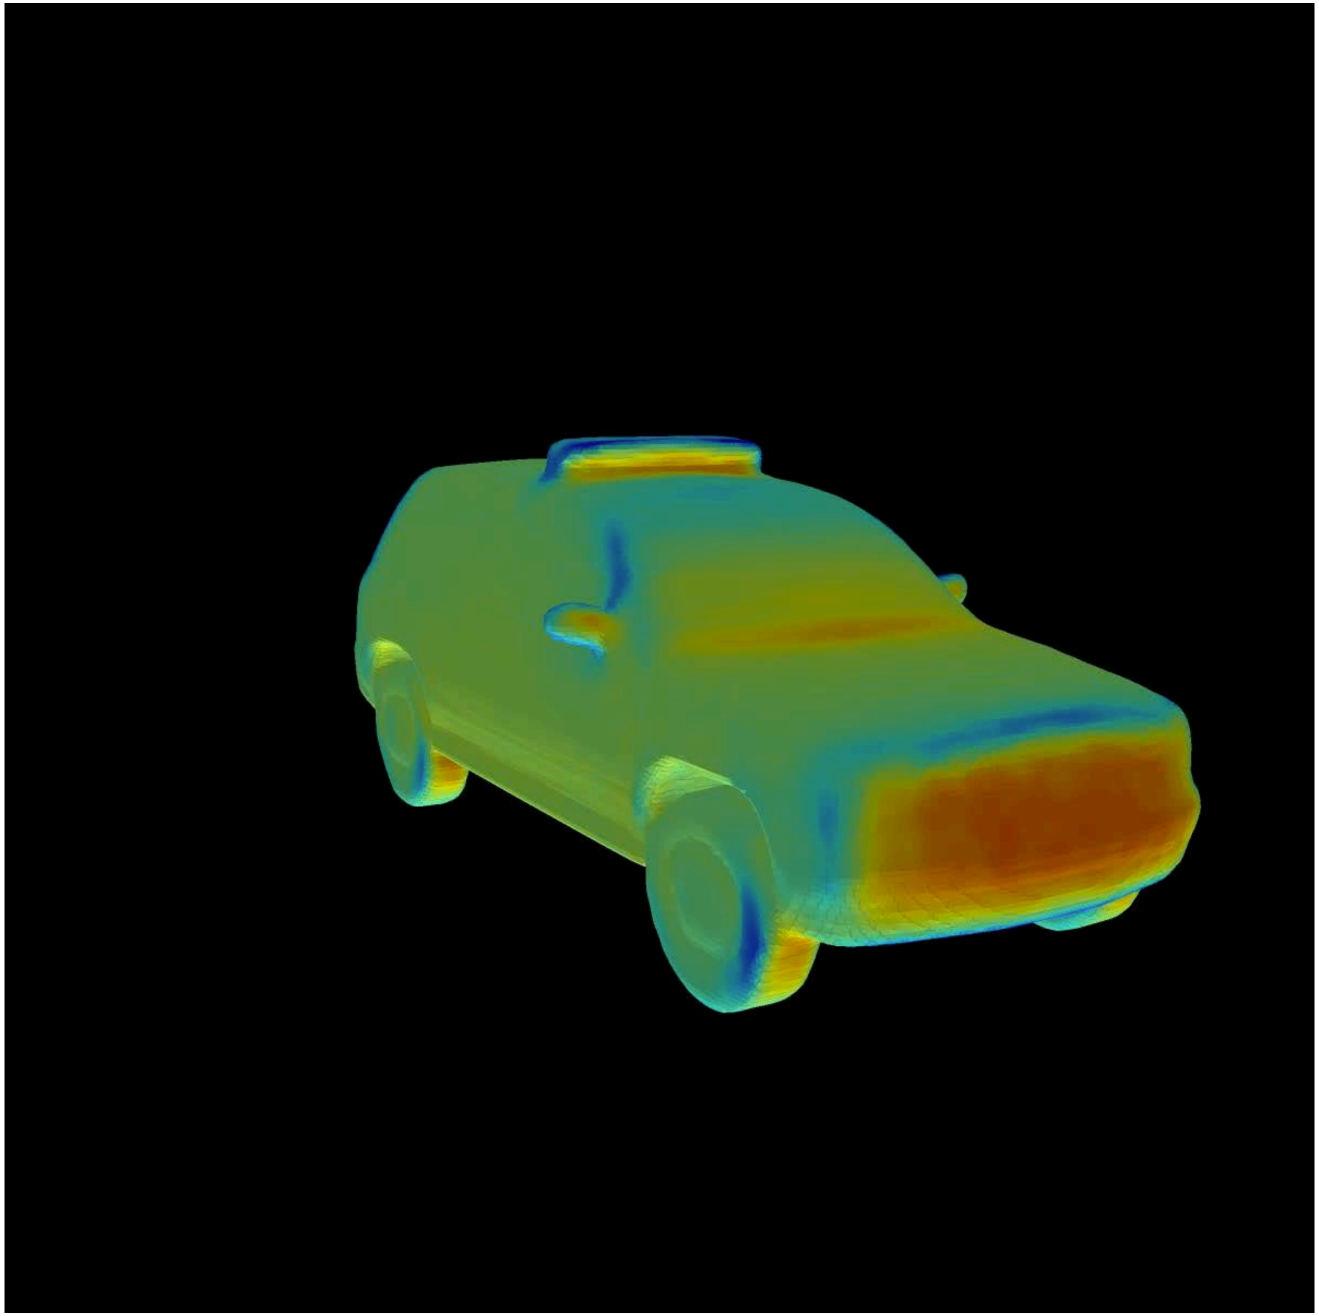

Eliminating Shock Waves

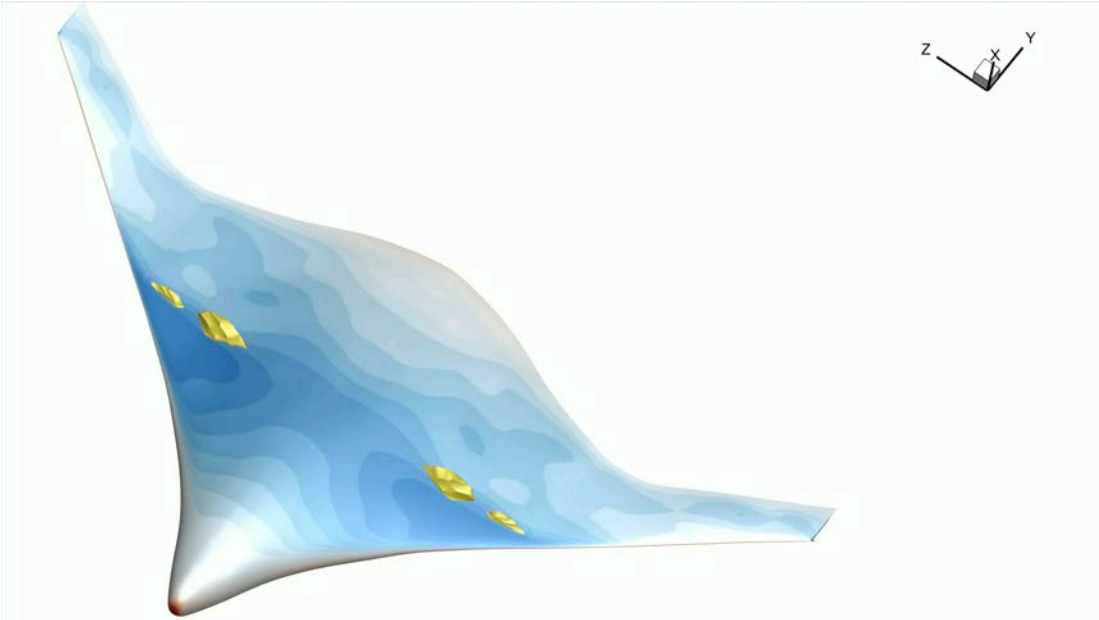

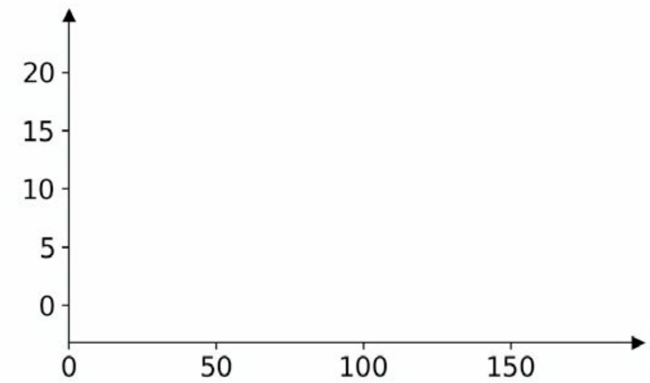

Surface Pressure Distribution

Mach = 1 Iso-Surface (shock region)

Bicycle Shell

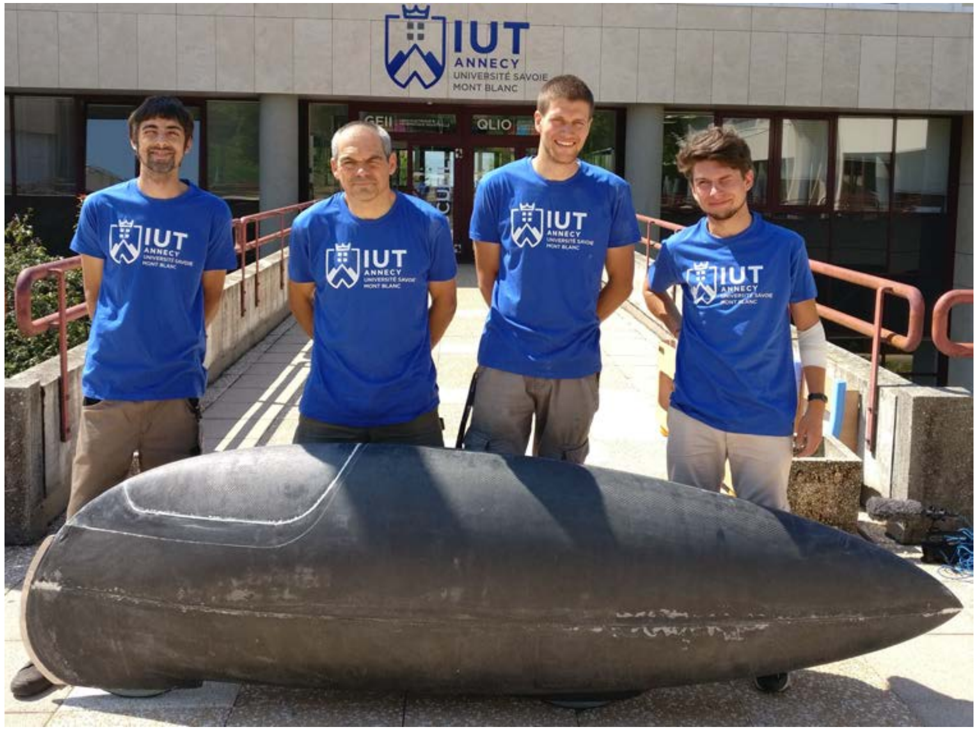

Altair 6, IUT Annecy, 2018

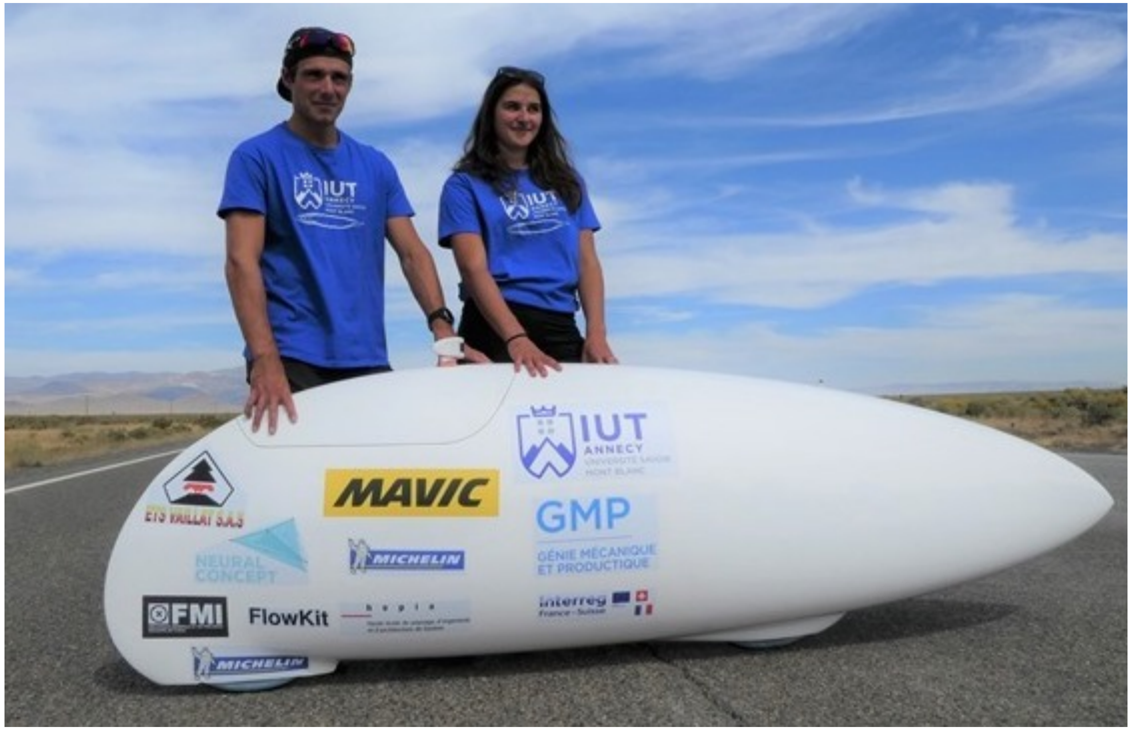

World Human Powered Speed Challenge
Battle Mountain Nevada, 2019

Women world record: 126,48 km/h
Men student world record: 136.74 km/h

Chasing another World Record

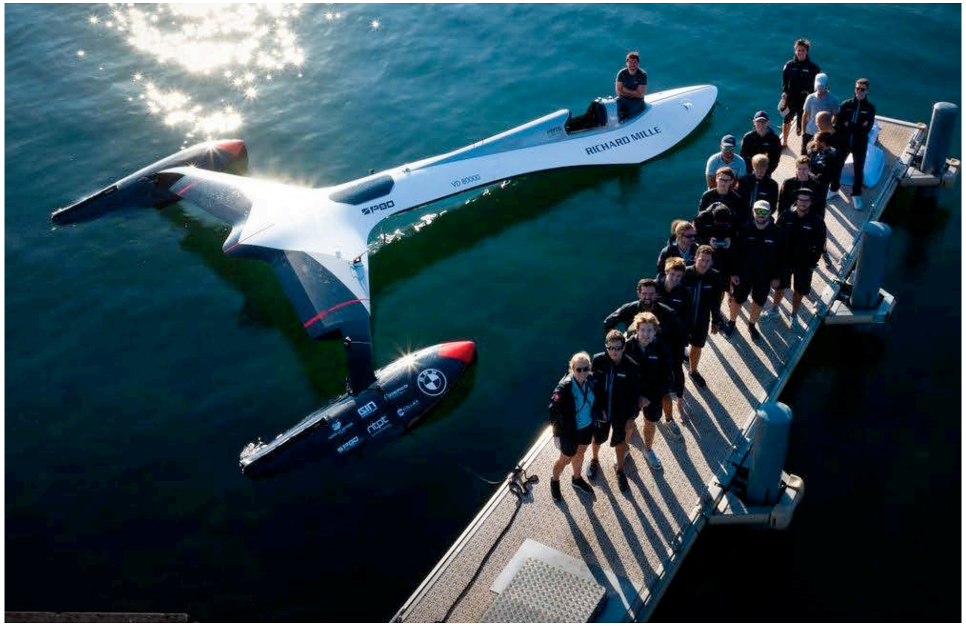

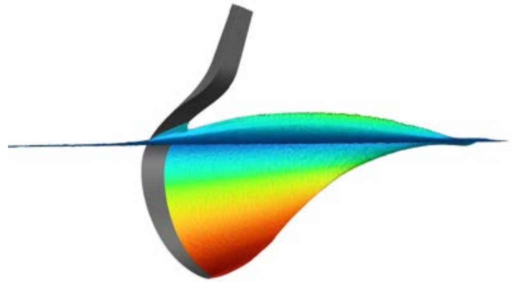

- Boat powered by a kite
- Design speed 80 knots

- Superventilating foil
- Inspired by the X15

- 2025: 58 knots in Leucate, France
- 2027: Next attempt

Designing Cities

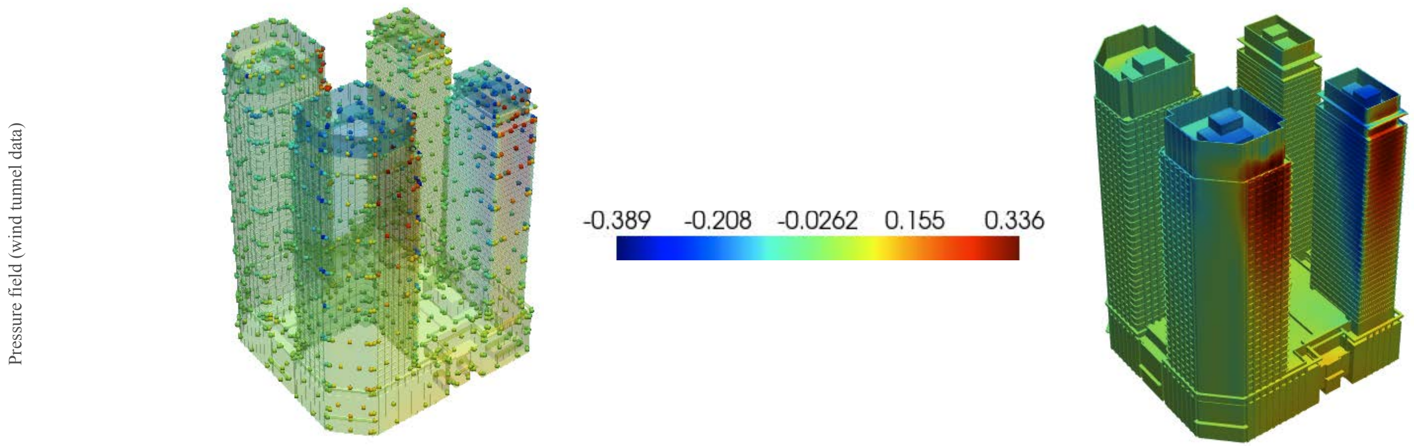

- Wind pressure fields can be predicted.
- The building layout can be optimized accordingly.

Heat Exchanger

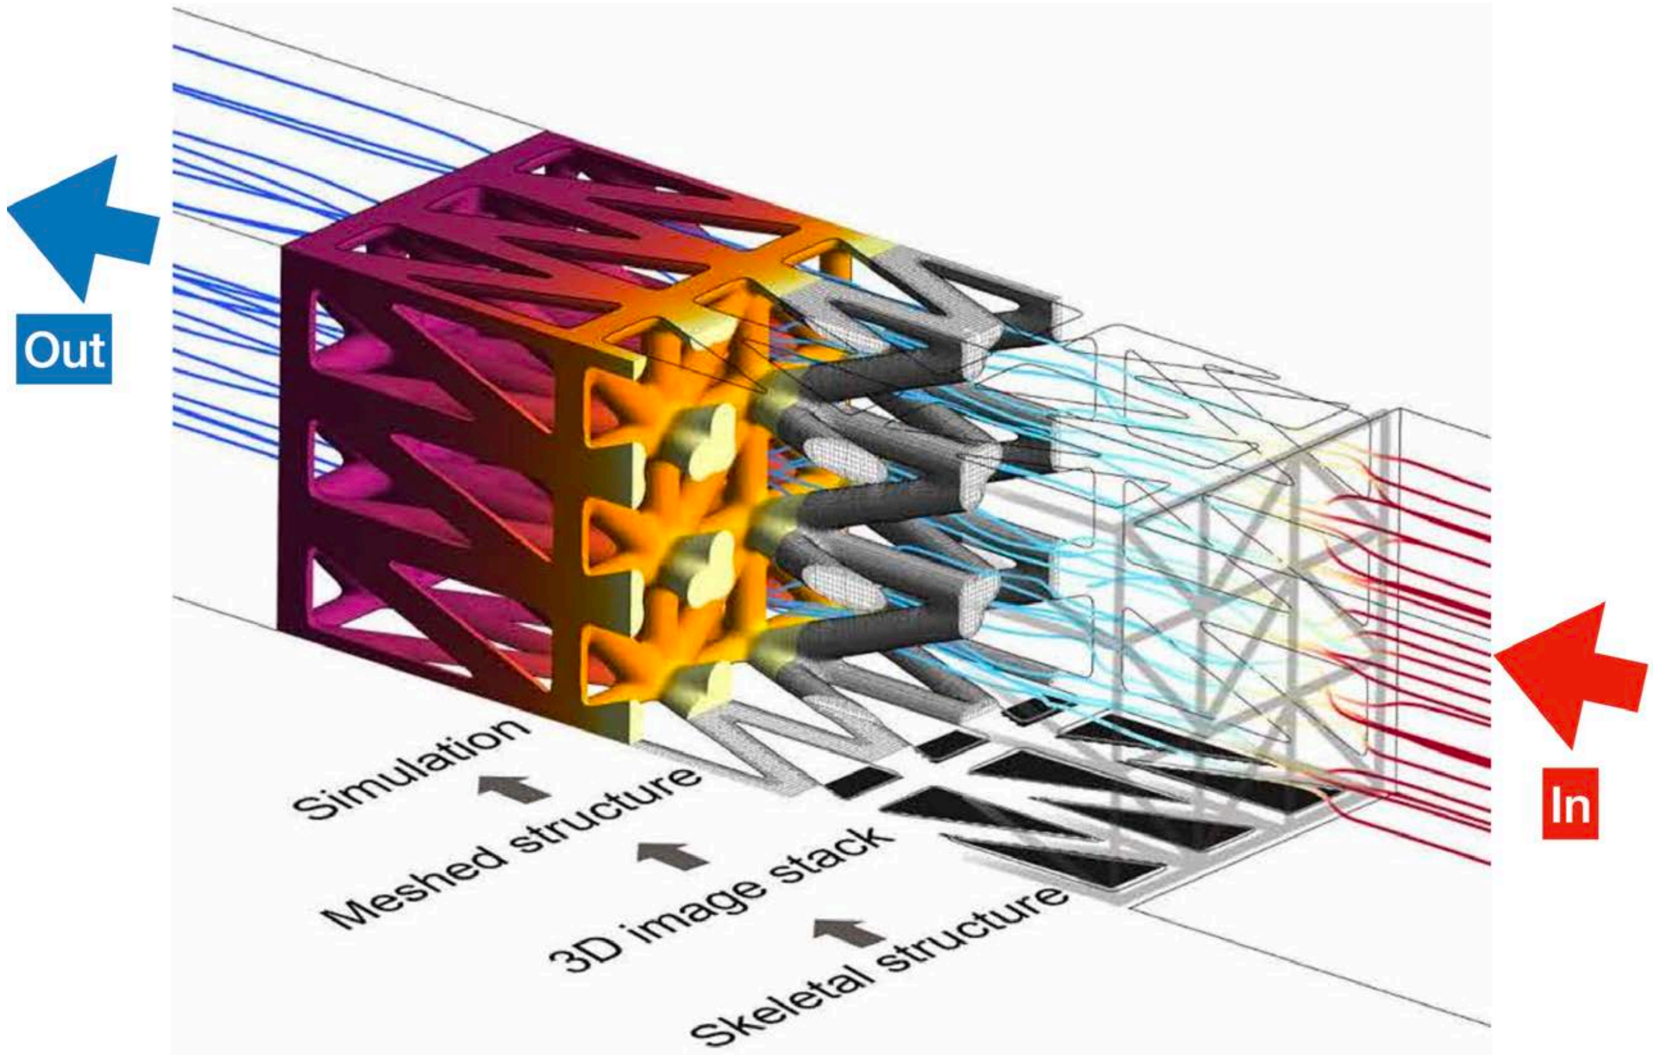

Predict and optimize the heat-exchange performance of 3D monolithic macro-porous structures

Structural Analysis

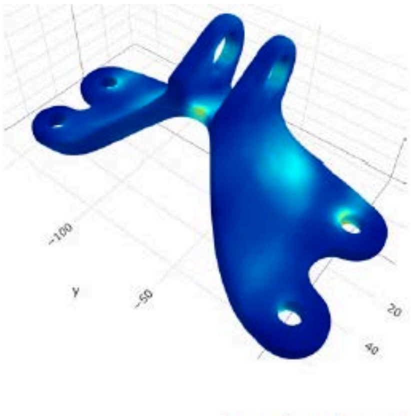

Simulation

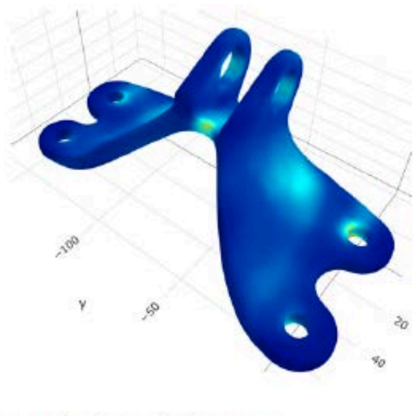

Emulation

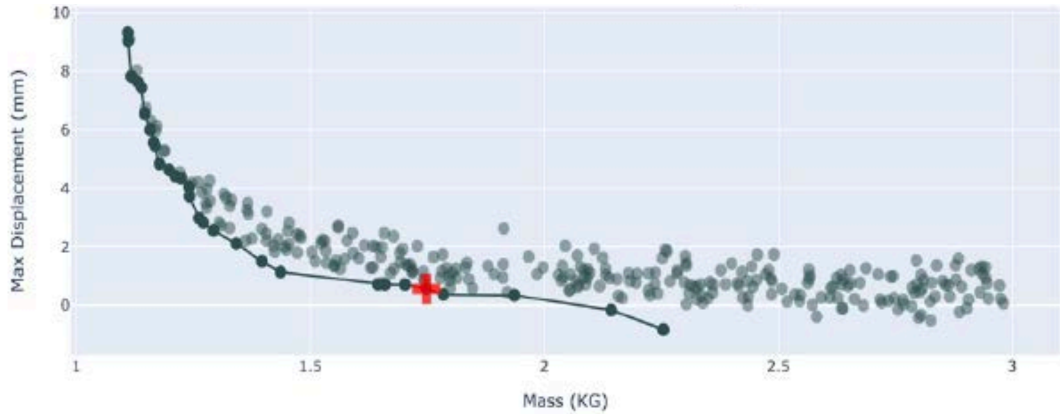

Pareto Front

- Accurate stress and displacement predictions on new geometries.
- The best compromise between mass and rigidity can be found.

From Simple to Composite Objects

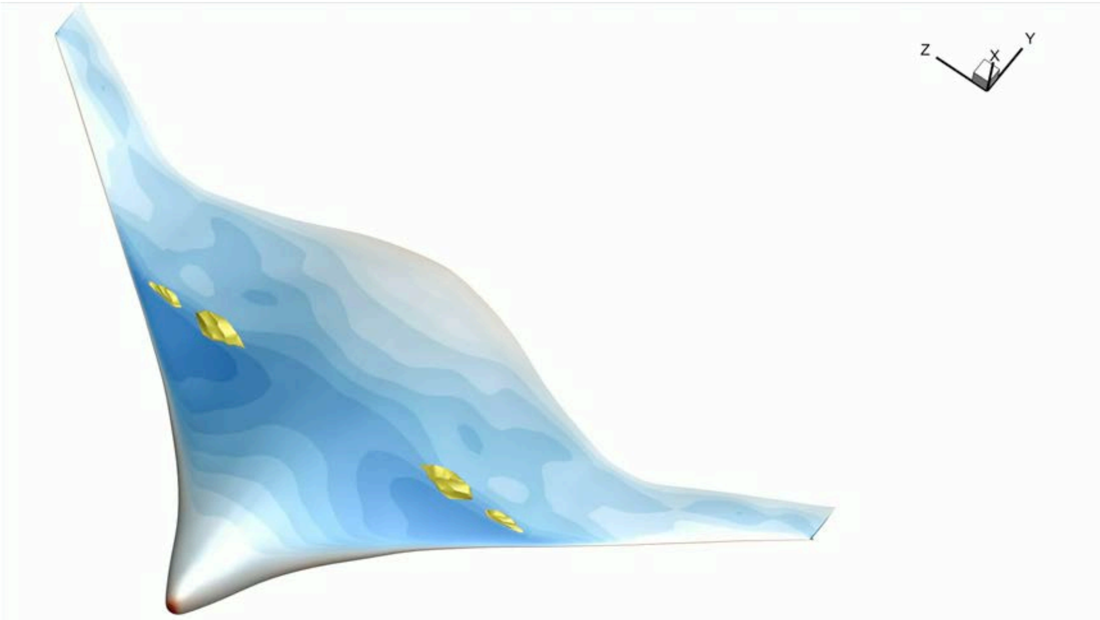

We can do this ….

… but not that.

ParfSDF

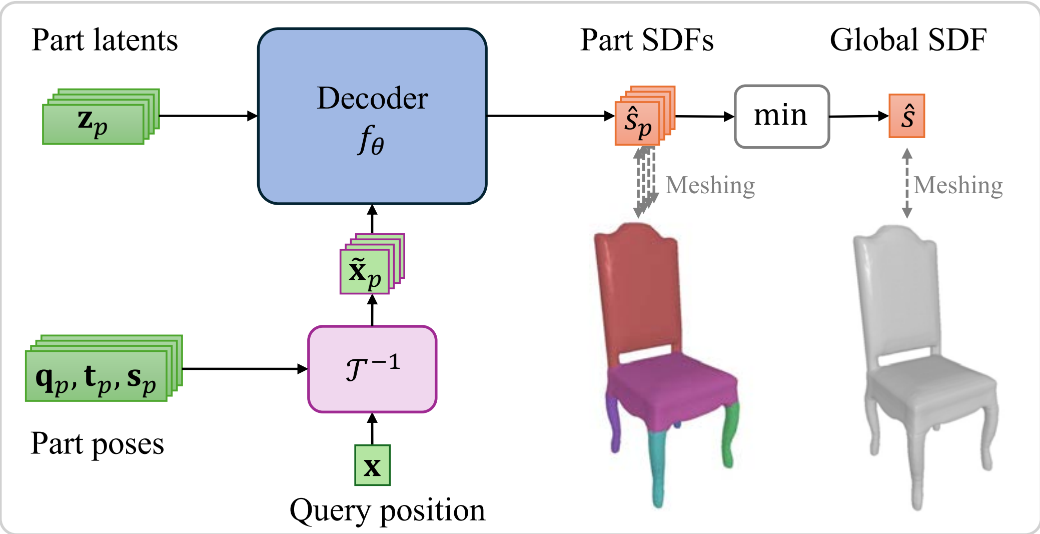

Composite shape decoder

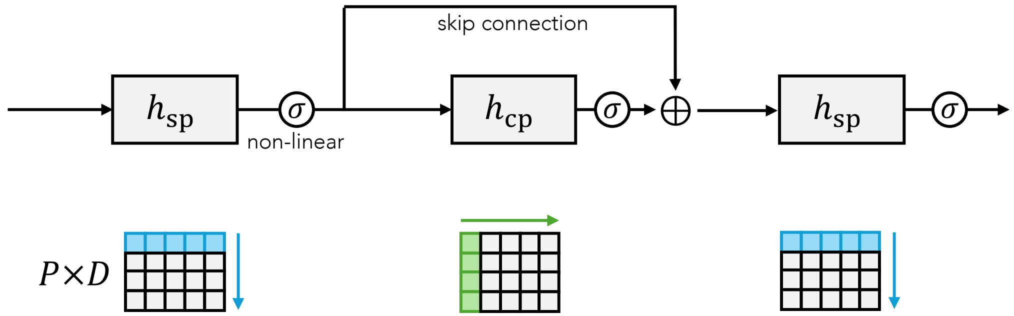

Alternates convolutions
within parts and across parts

- Uniform architecture
- Part consistency enforced
- Manual editing possible

EPFL CVLab

Conclusion

- Deep Signed and Unsigned Distance Functions can be used to implement 3D surface meshes that can change their topology while preserving end-to-end differentiability.
- The formalism can be extended to composite shapes to represent complex objects.

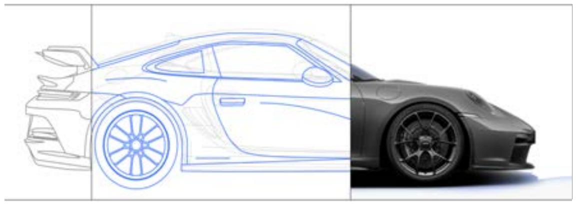
Automotive


Space


Electrification

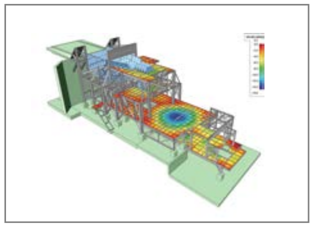
Architecture

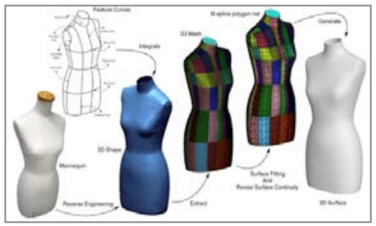
Textile

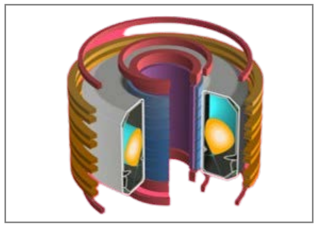
Nuclear fusion

—> Major Impact in Computer Assisted Engineering