Deep 3D Surface Meshes

P. Fua

EPFL Computer Vision Lab
Lausanne
Switzerland

Deforming 3D Surfaces

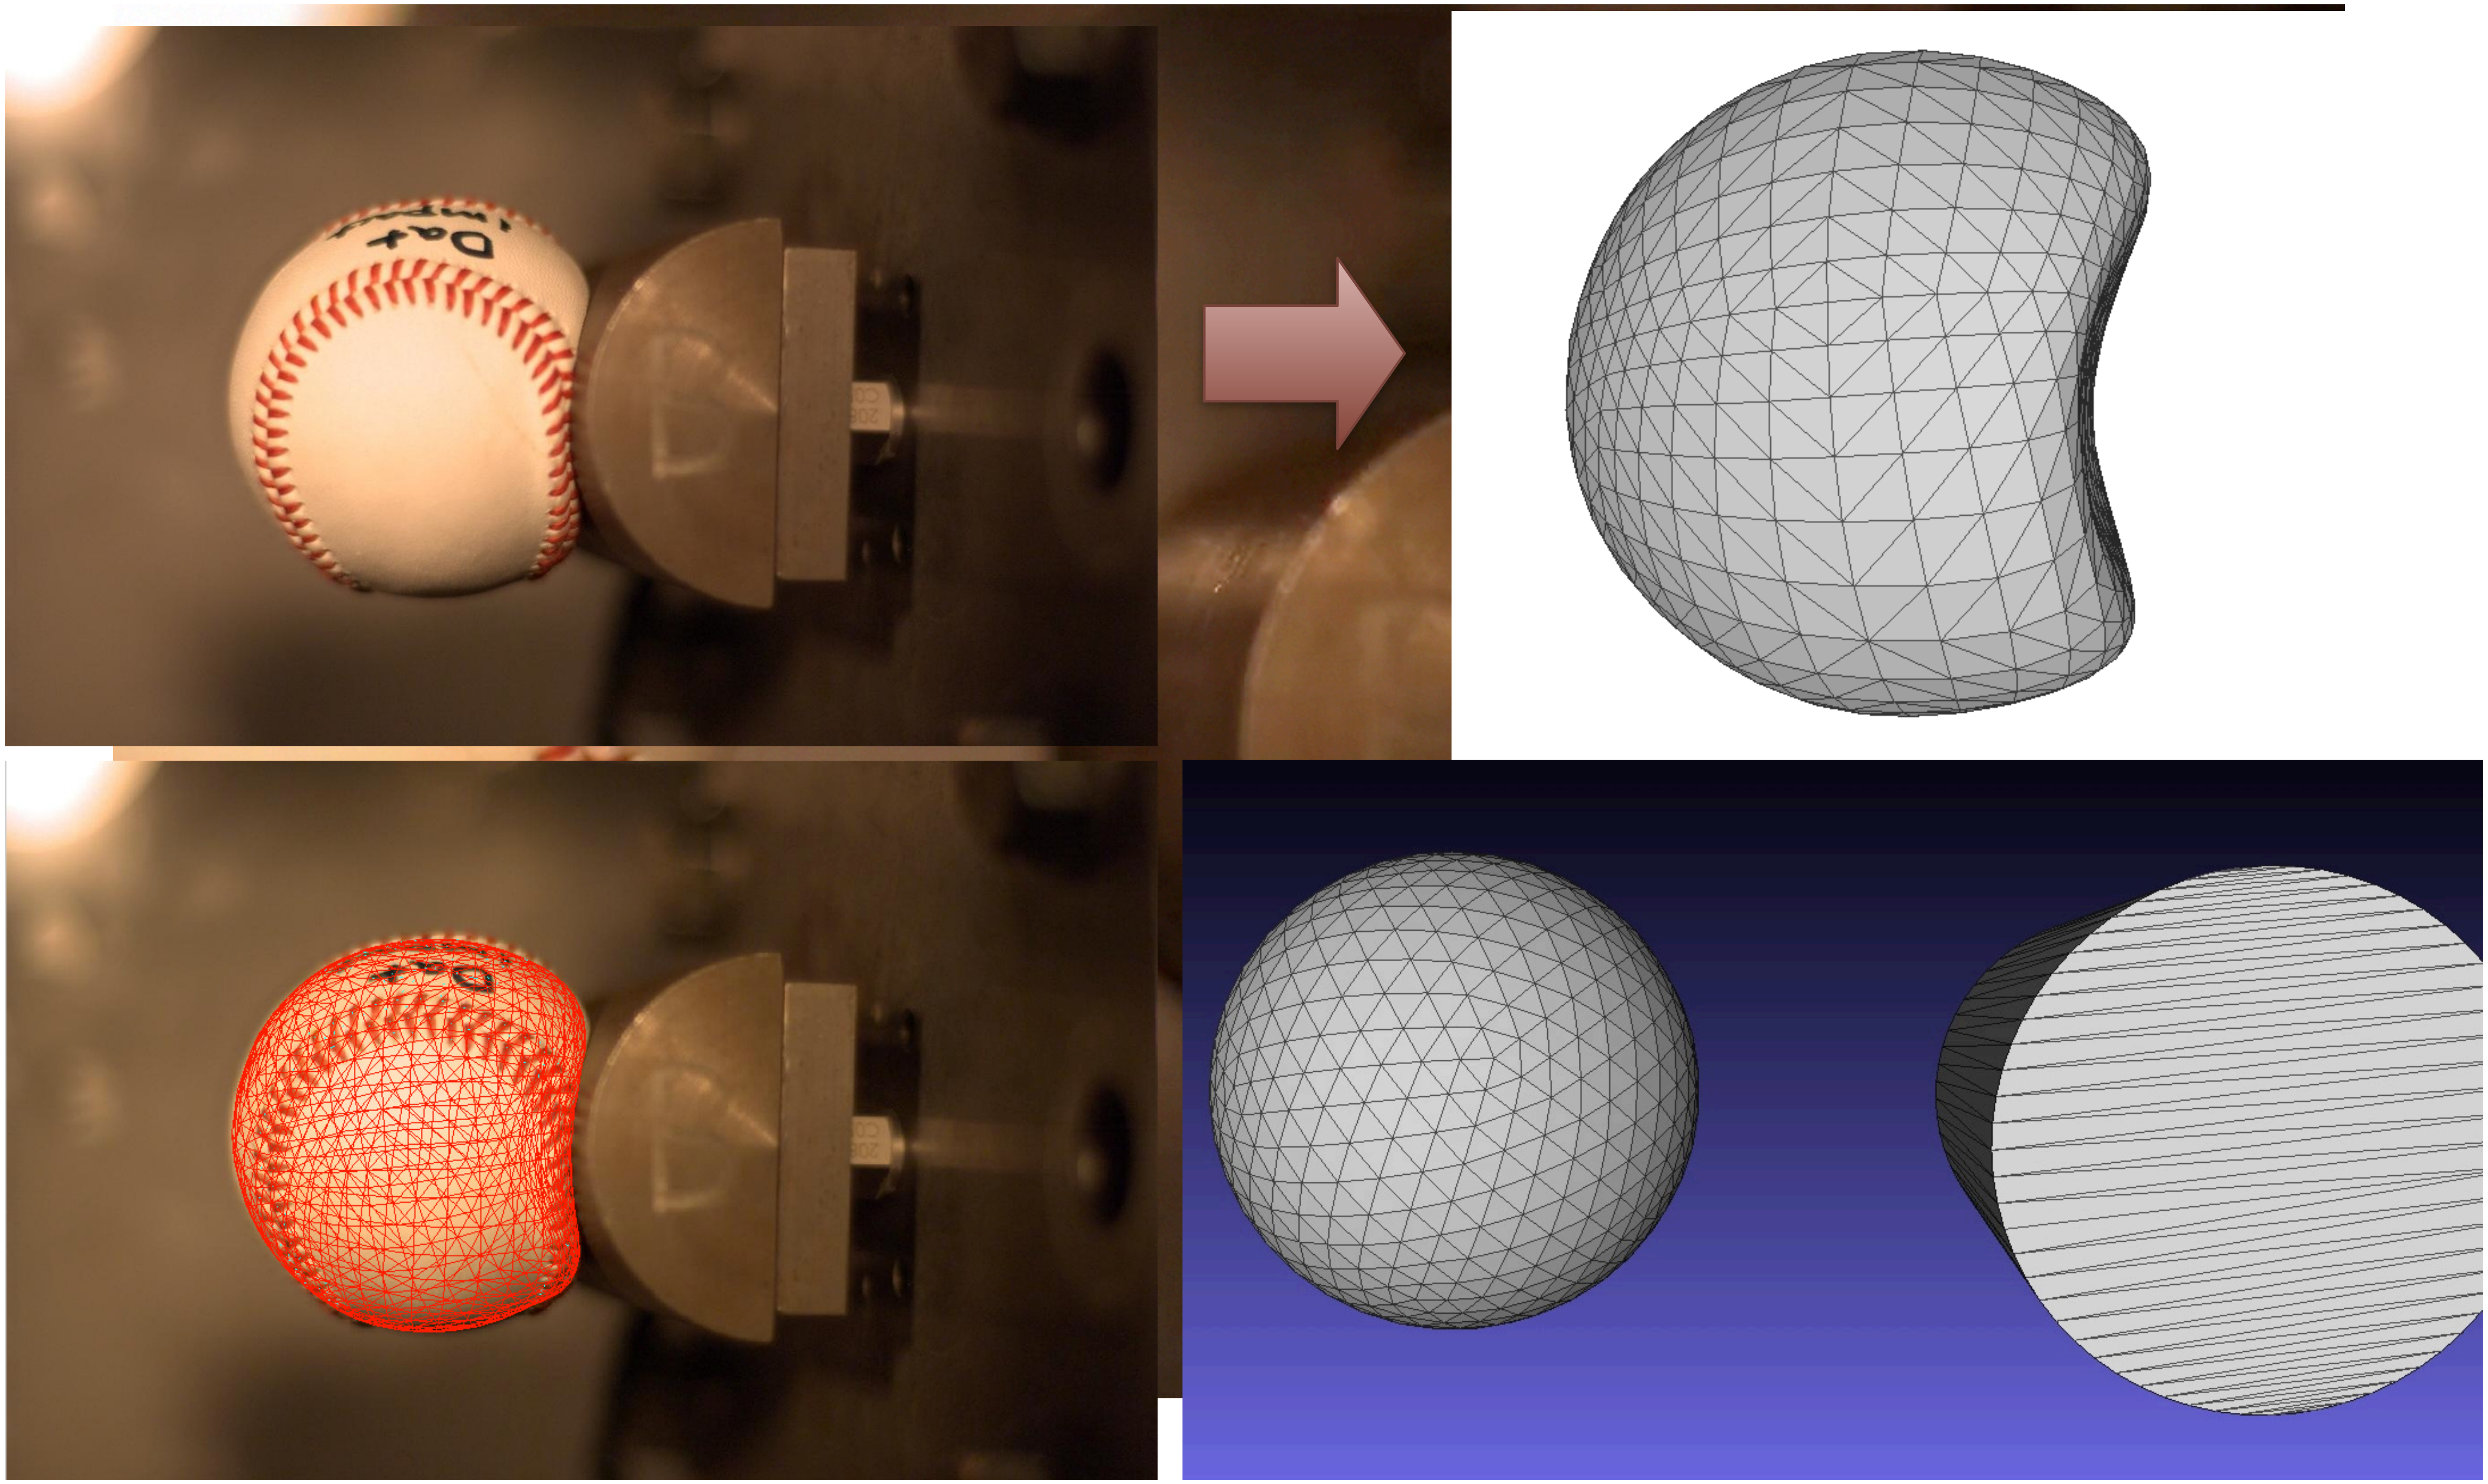

Ngo, PAMI'16

Augmented Reality

Real Time Implementation on an iPad.

3D Surface Representations

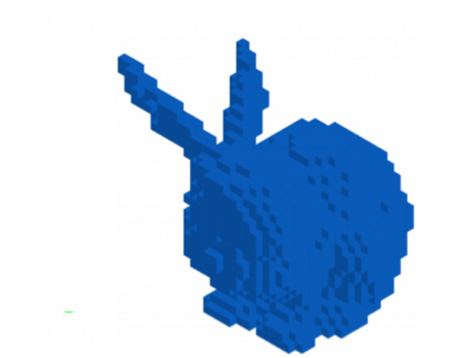

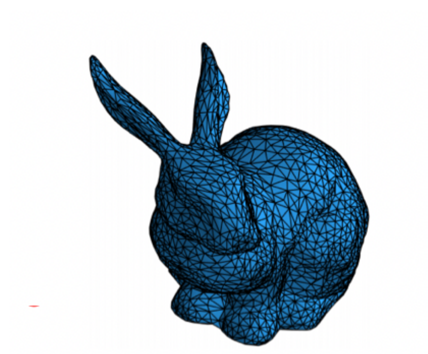

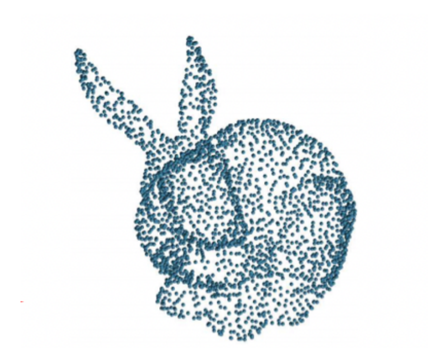

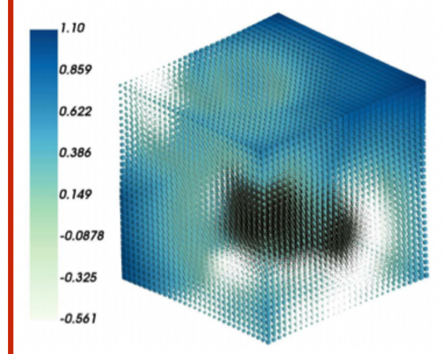

	Voxels	Explicit surface mesh	Point sets	Continuous implicit fields
High frequency details?	--	++	+	++
Arbitrary topology?	+	-	+	++
Regularity?	+	+	-	++

There are 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

Signed Distance Fields (SDF)

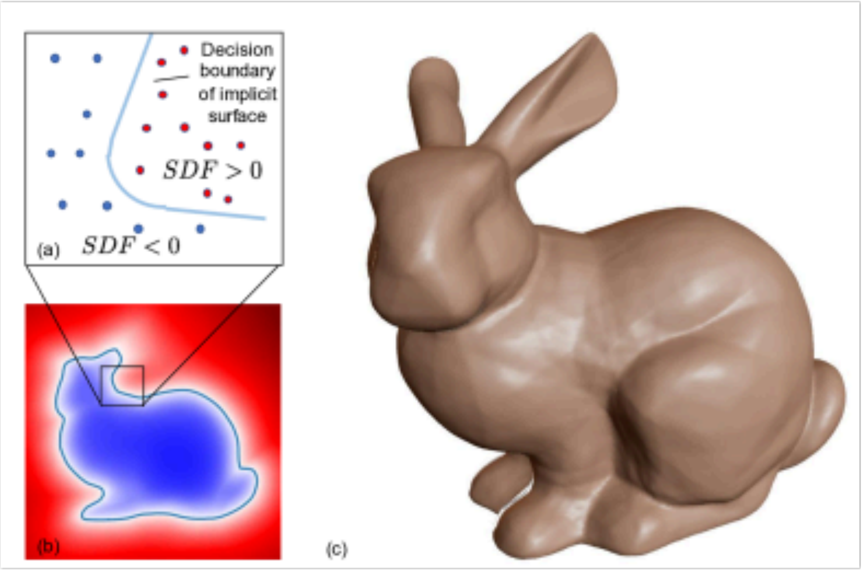

- Represent a 3D surface S by the zero crossings of a signed distance function

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

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

- 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

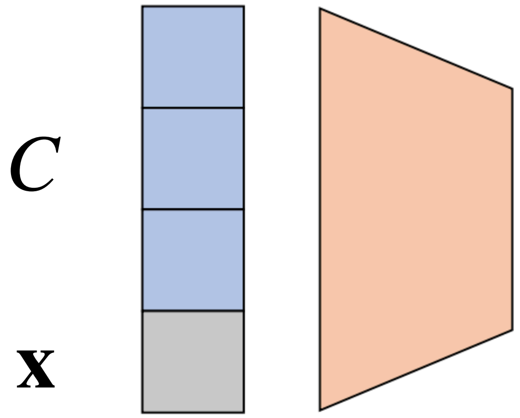

$s = f_\theta(\mathbf{x} \mid C)$

Coded Shape DeepSDF

[Park et al., CVPR'19]

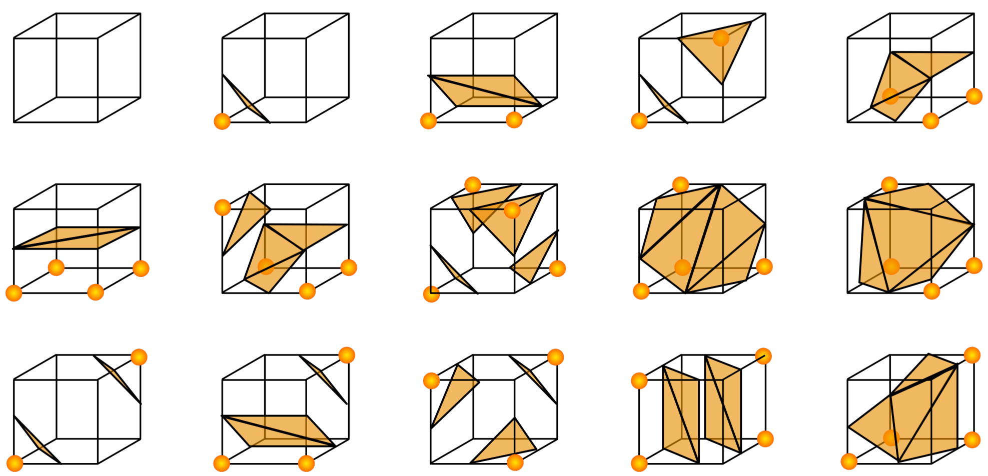

[Liao et al., CVPR'18]

- But one bottleneck remains: If an explicit surface representation is required, one has to run a marching-cube style algorithm, which is **not differentiable** and often **slow**.
- We have shown that its possible to differentiate nevertheless so that it can be incorporated into an end-to-end trainable framework.

[Remelli et al., NeurIPs'20]

From Genus 0 to Genus 1

Application: Single View Reconstruction

3D Model from Image

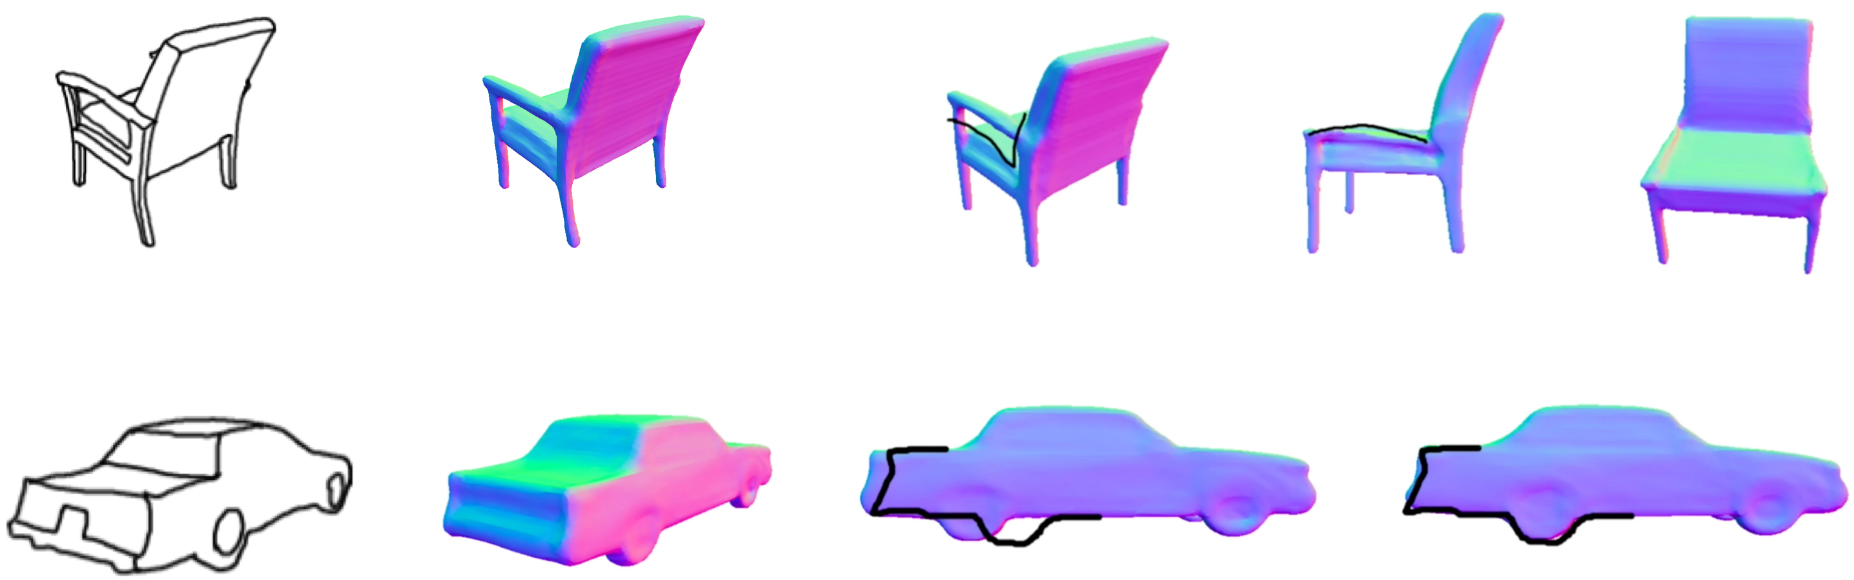

Editable 3D Models from Sketches

Application: Shape Optimization

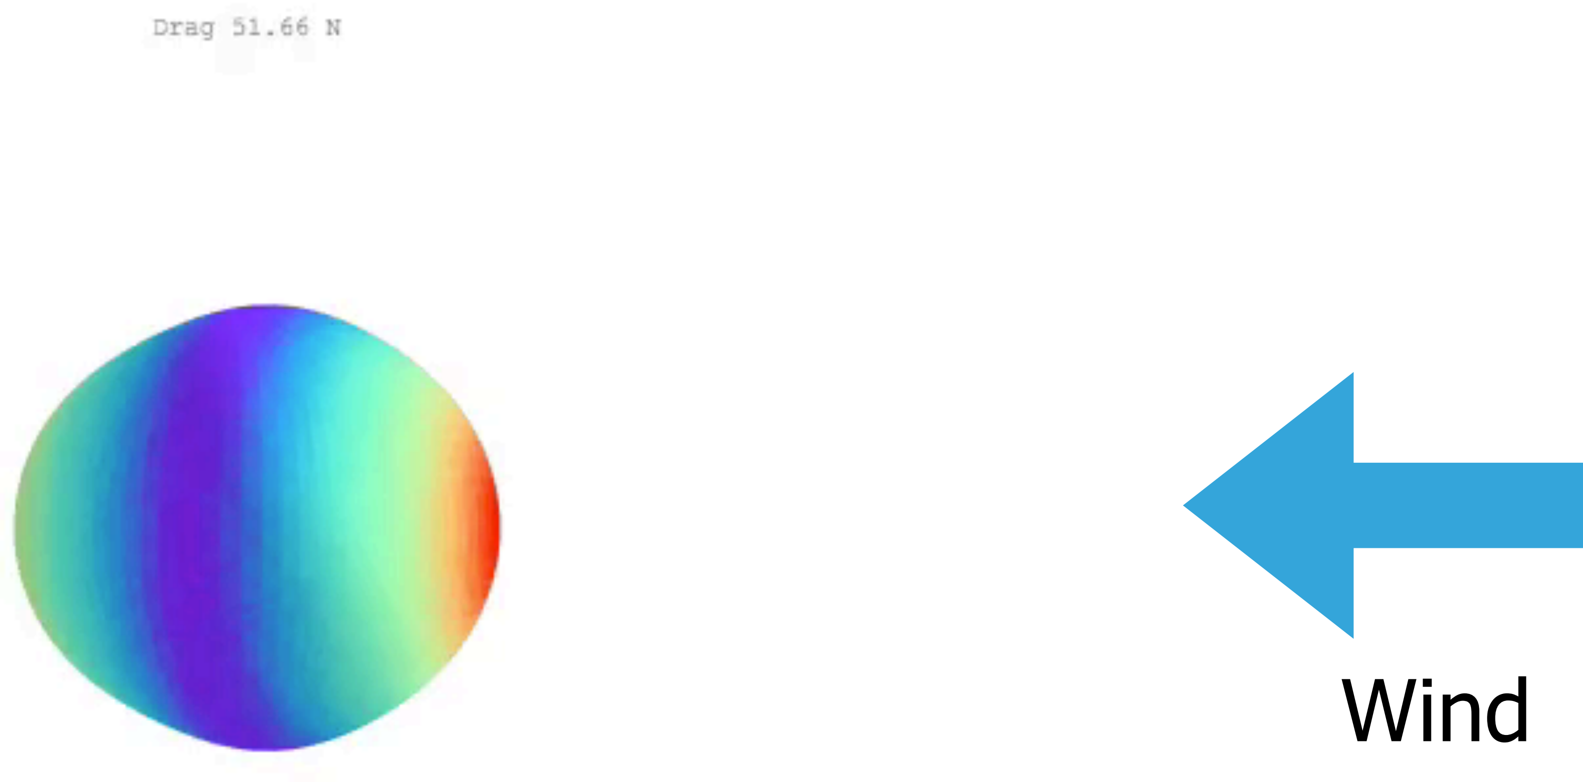

Minimizing drag while enclosing a sphere.

Interactive Design

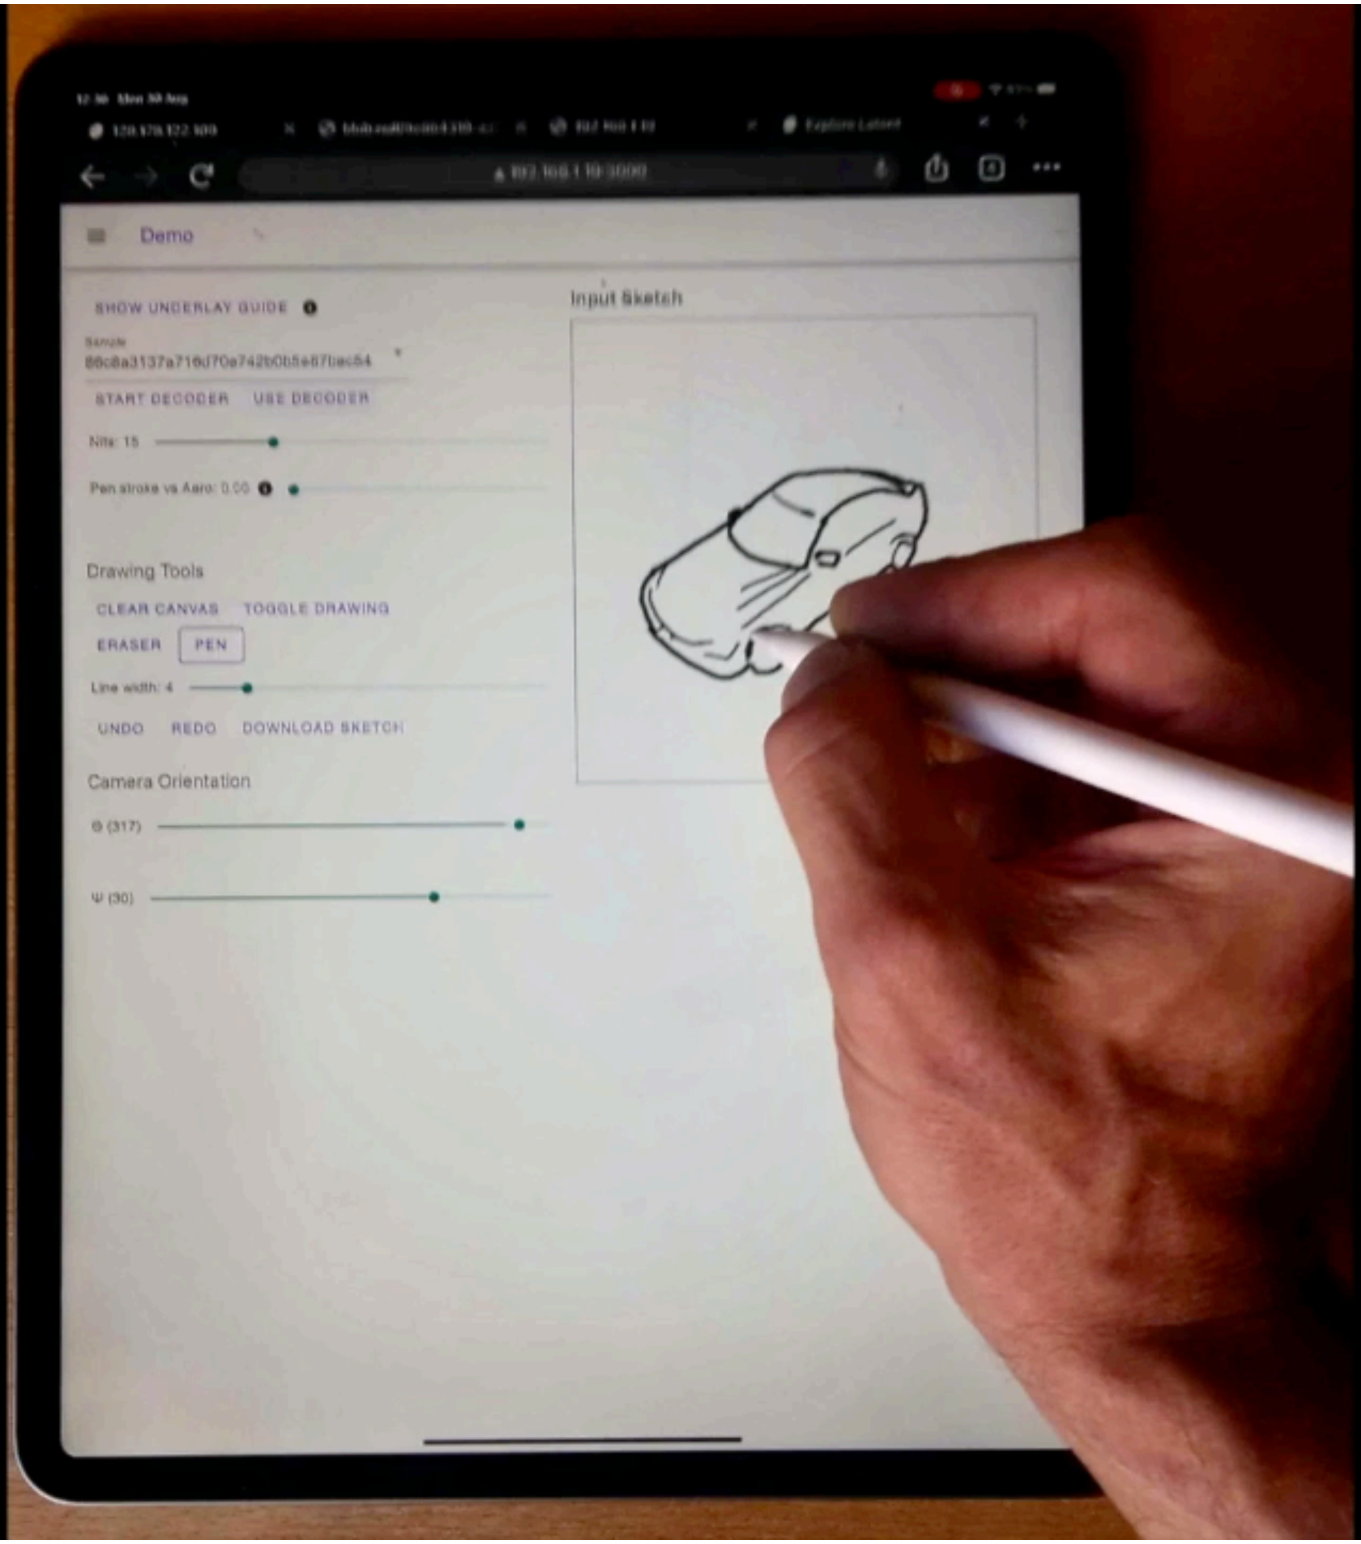

Bicycle Shell

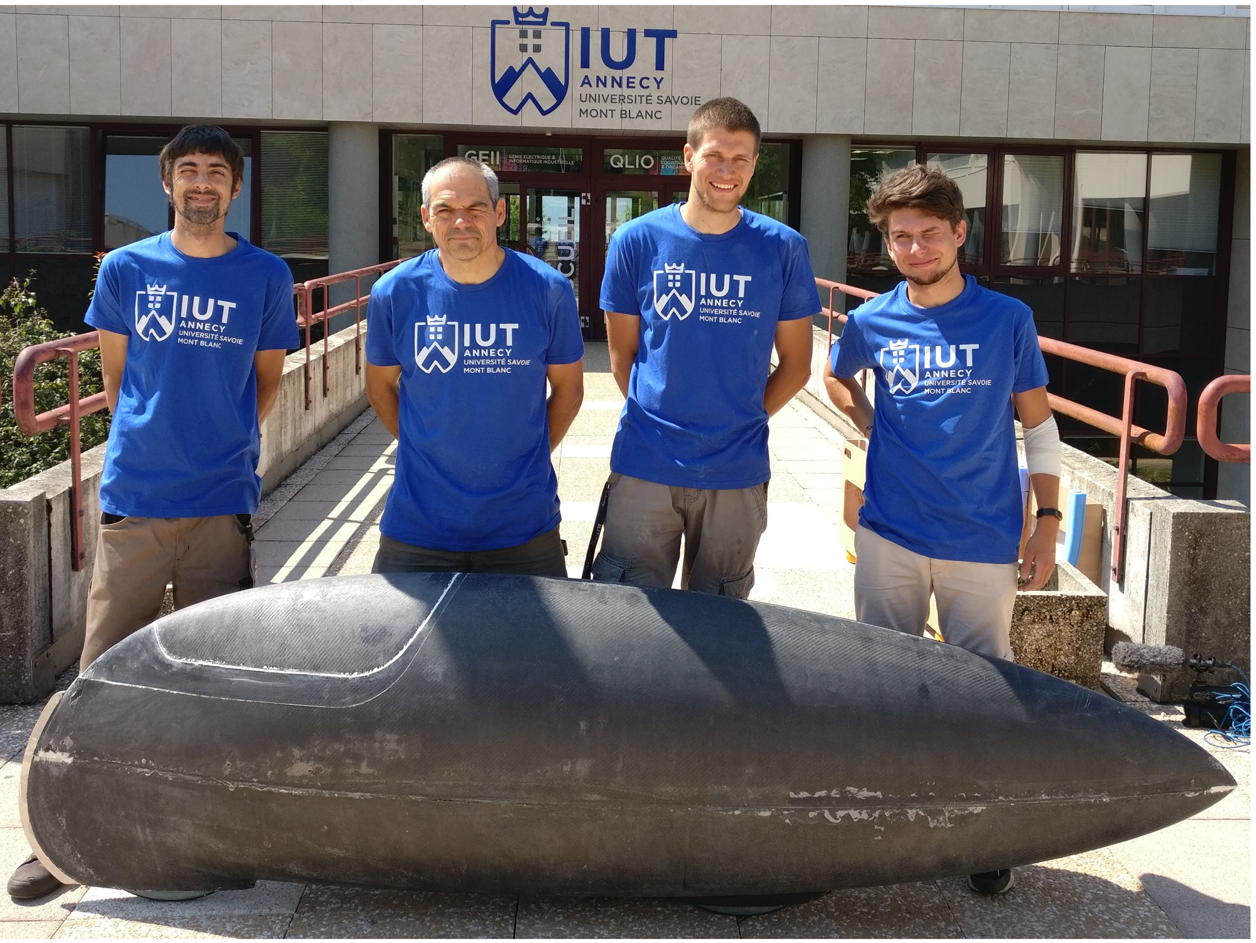

Altair 6, IUT Annecy, 2018

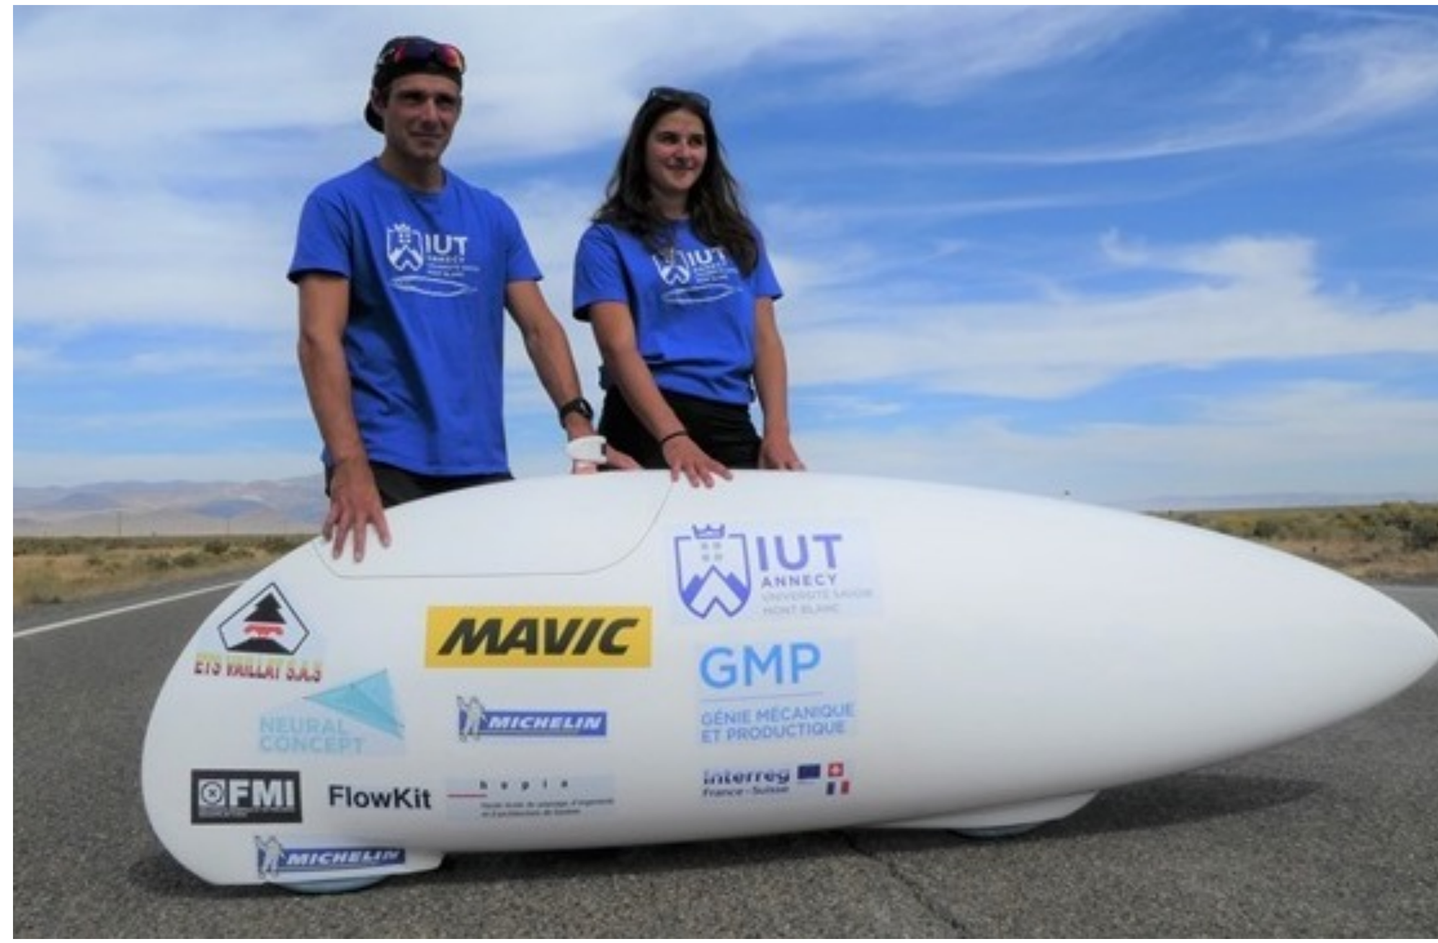

World Human Powered Speed Challenge
Battle Mountain Nevada, 2019

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

Dynamic Soaring

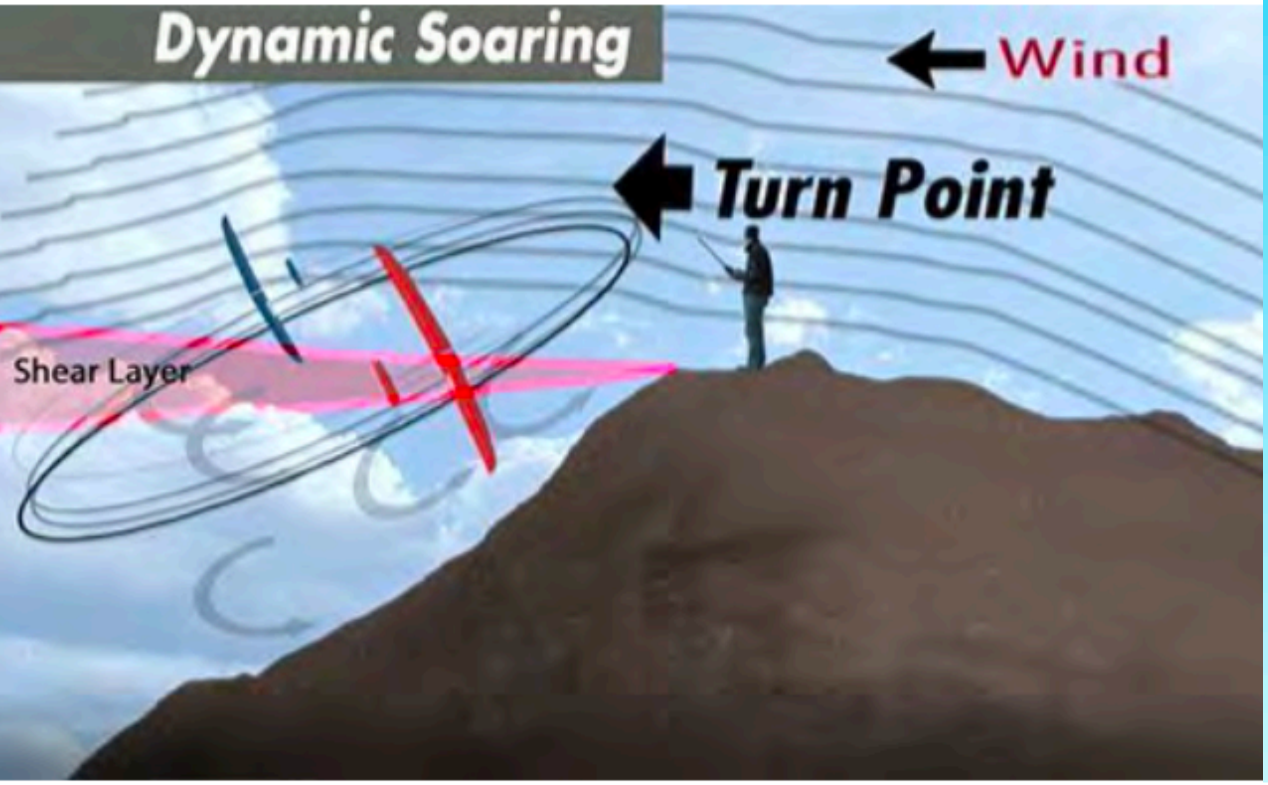

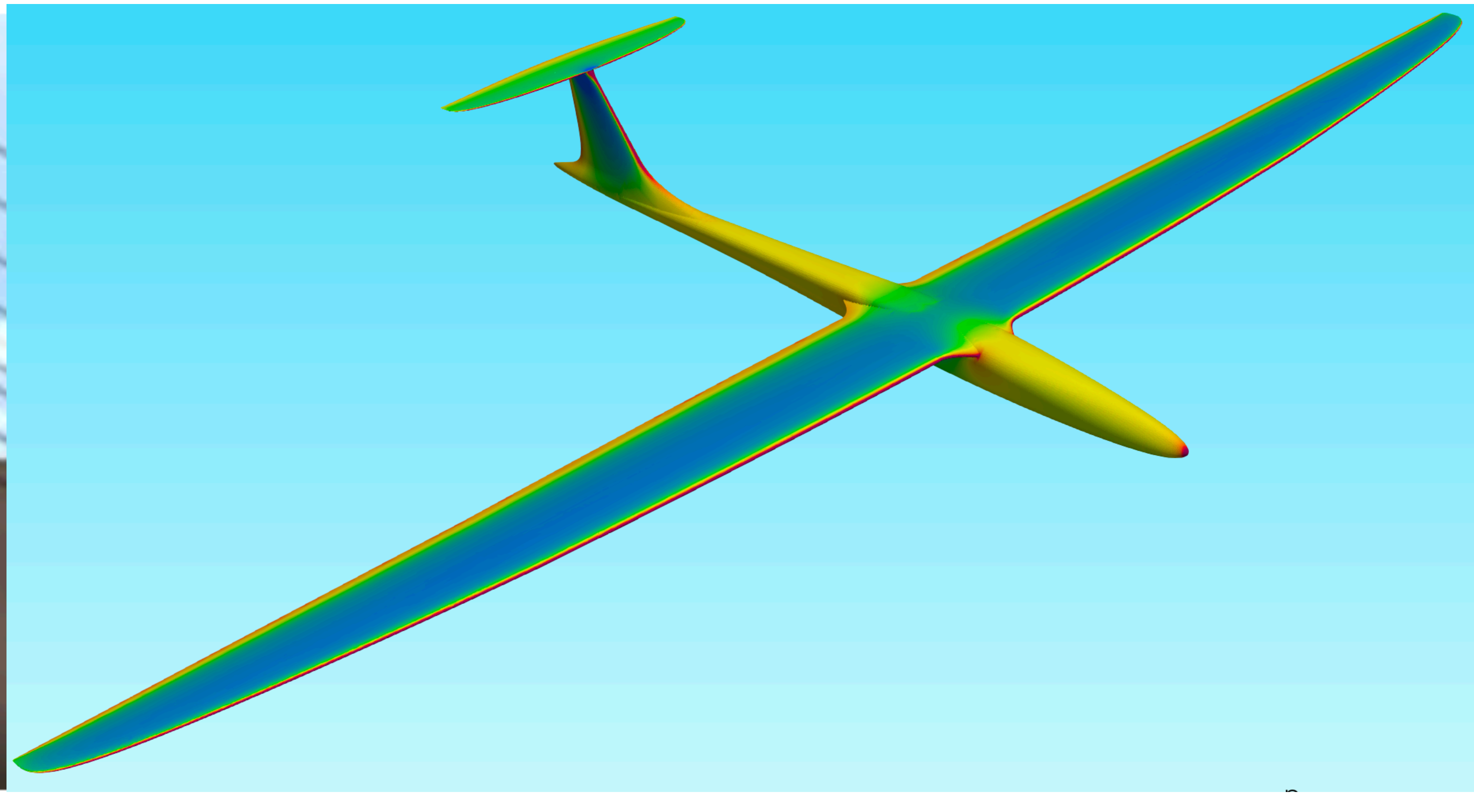

- We plan to design for ease of control.
- We will use dynamic soaring to prove the concept.

Real World Example

Aerodynamics:

- They are complex,
- and different from those seen in other fields.

Structural issues:

- All the components have to be strong enough.

Power control:

- Rotation speed but be properly controlled.
- Torque limits cannot be exceeded.

—> Ideally all these elements should be taken into consideration at the same time but this currently beyond the state-of-the-art.

Conclusion

- Combining explicit and implicit representations early makes it possible to exploit the strength of both representations.

- Deep Signed Distance Functions can be used to implement 3D surface meshes that can change their topology while preserving end-to-end differentiability.

—> This opens the door for new applications in fields as diverse as Computer Assisted Design and Medical Imaging.

Many Thanks To

- Timur Bagautdinov (NeuralConcept)
- Pierre Baque (NeuralConcept)
- Benoît Guillard (EPFL)
- Graham Knott (EPFL)
- Artem Lukoianov (NeuralConcept)
- Edoardo Remelli (EPFL)
- Stephan Richter (Intel)
- Udaranga Wickramasinghe (EPFL)
- Pierre Yvernay (EPFL)