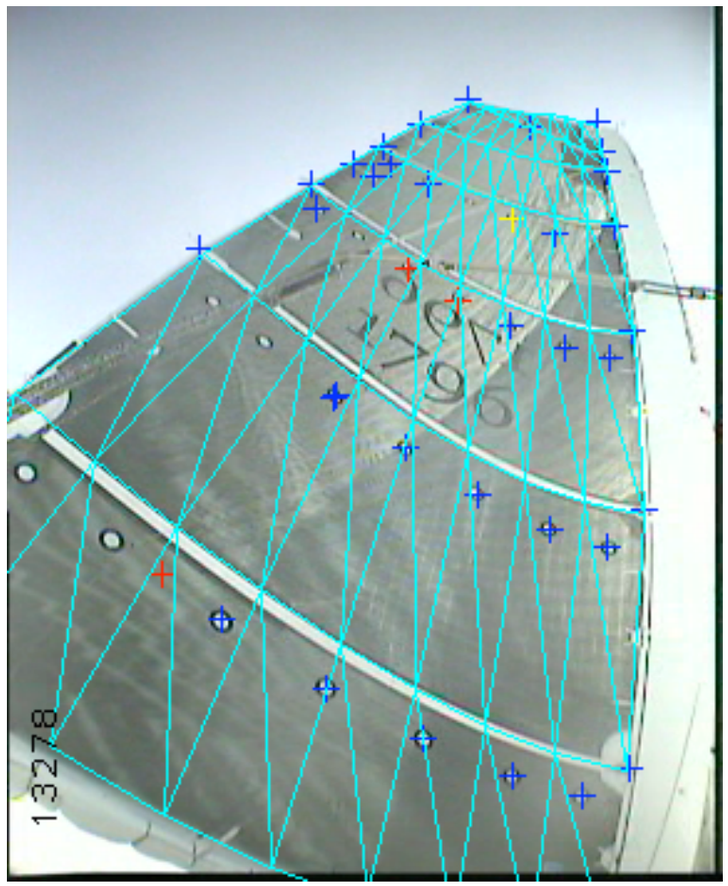
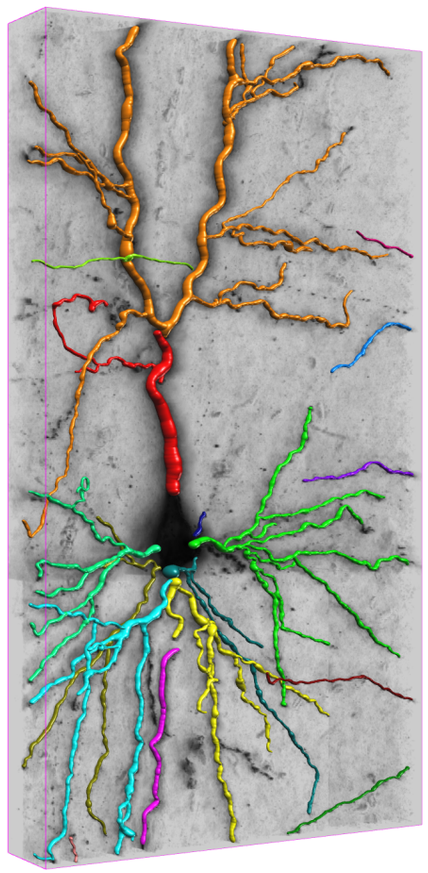
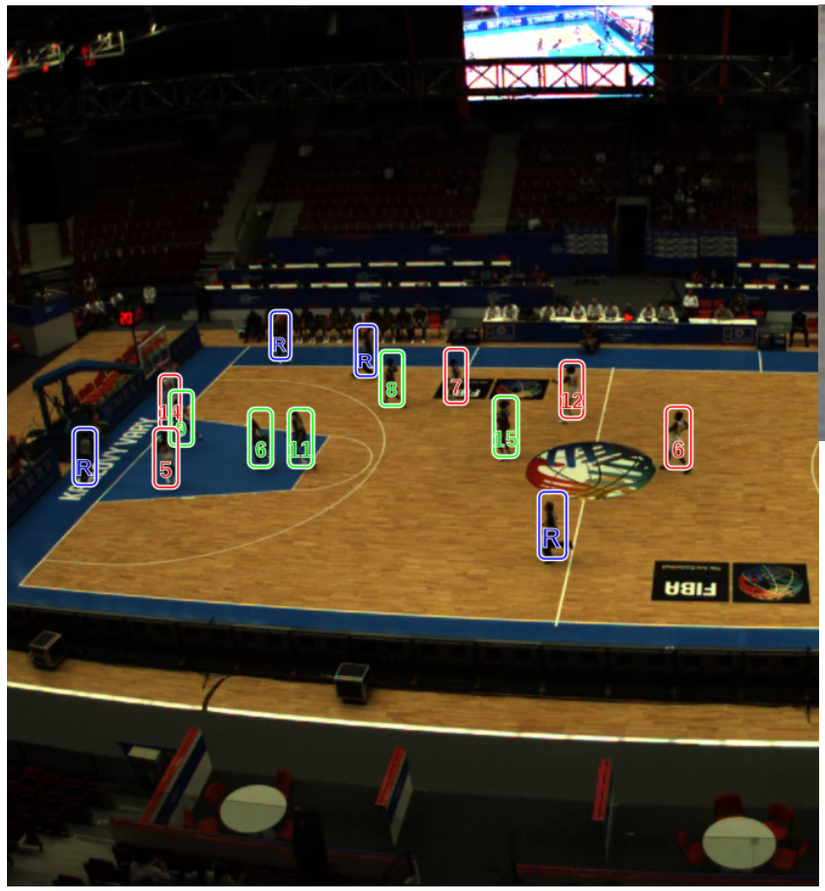

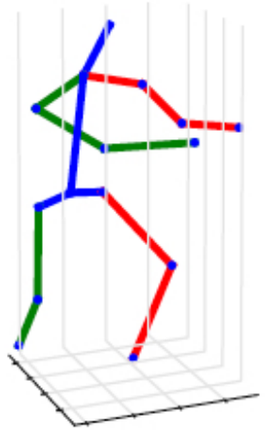

EPFL Computer Vision Lab
Lausanne
Switzerland

EPFL CVLab

The research team:

- 1 Professor
- 1 Senior scientist
- 6 Post docs
- 16 PhD candidates

Research Topics:

- Shape and motion from video
- People tracking
- Neural Structures from Microscopy

What do You See?

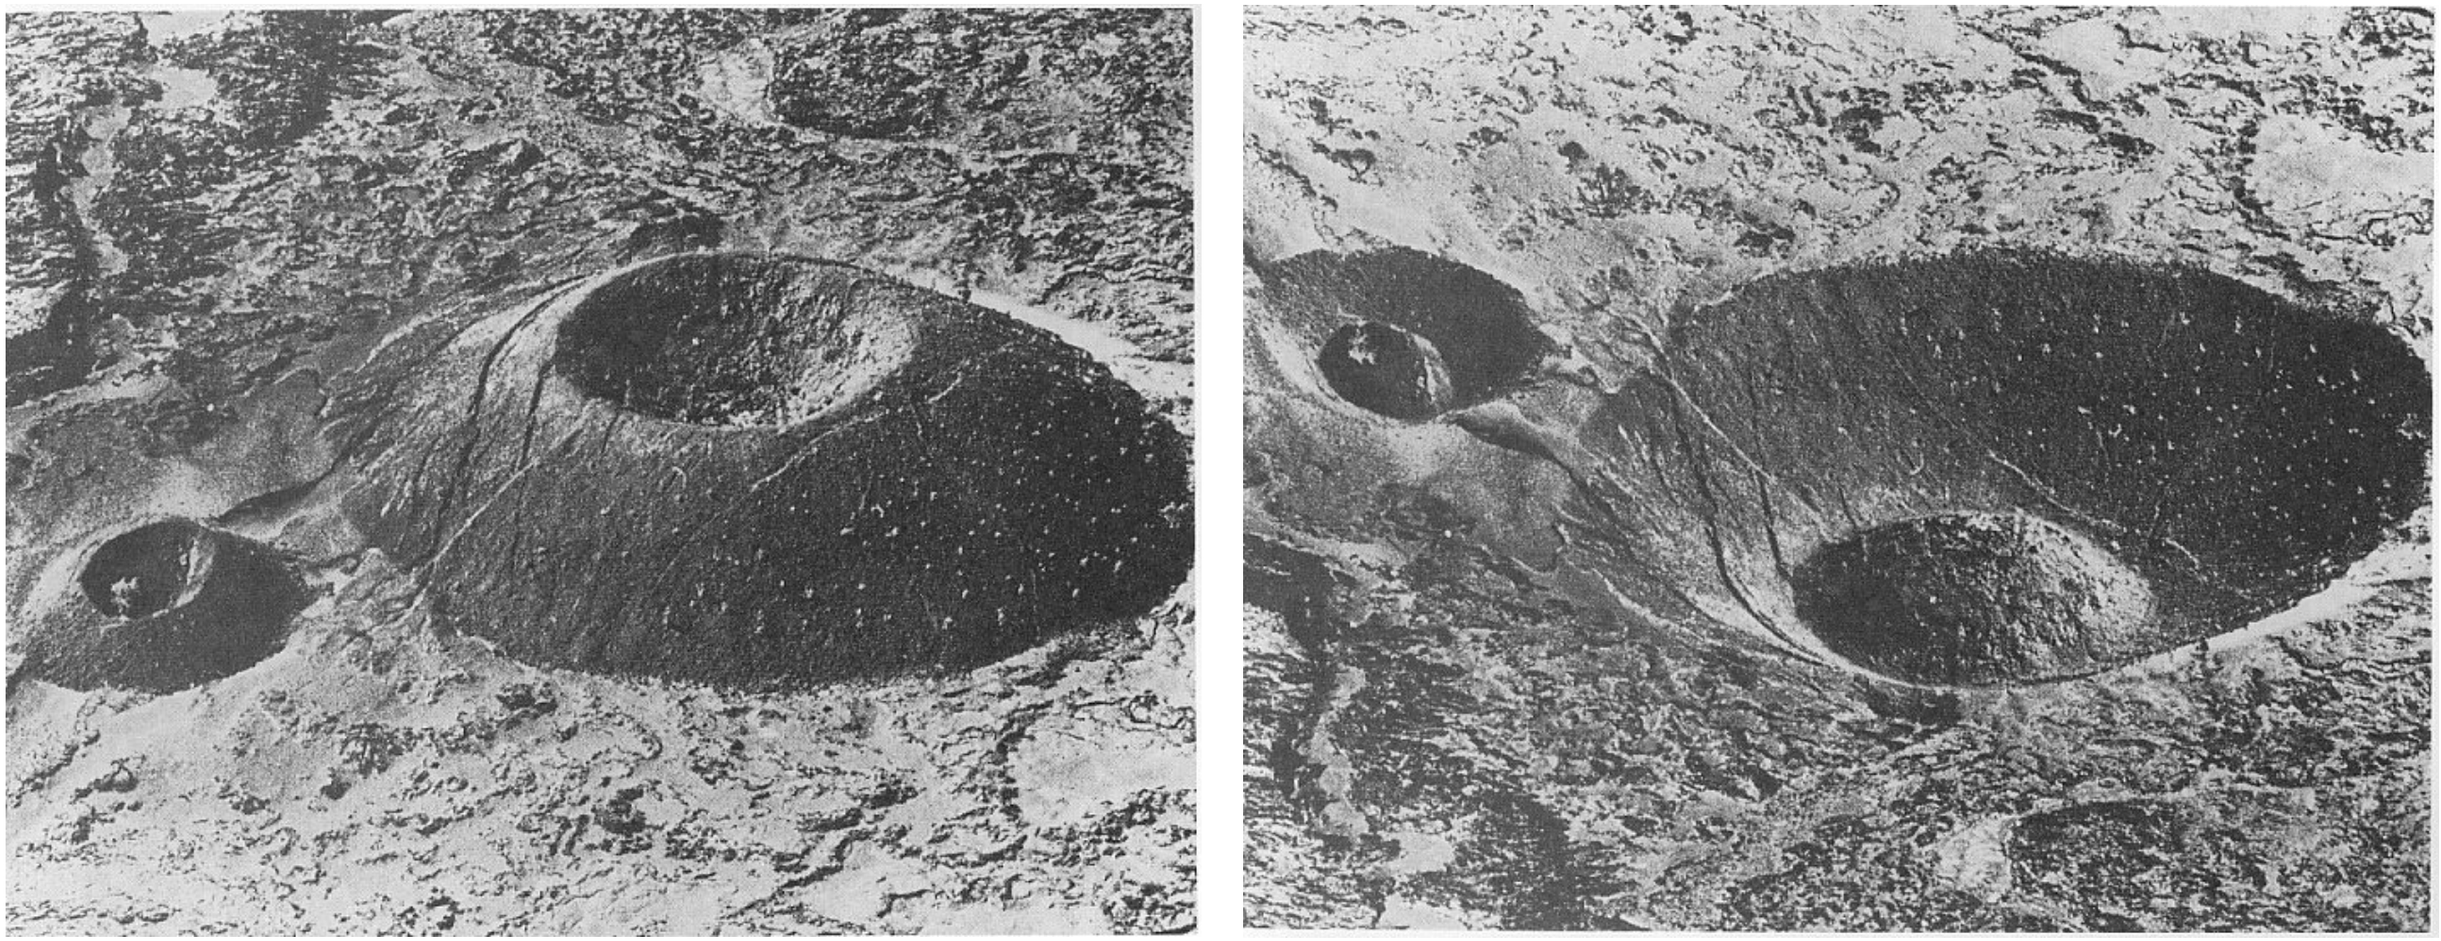

Scene understanding requires:

- Scene models
- Fitting these models to noisy data

—> Machine Learning, Signal Processing, Numerical Optimization, Geometric modeling ….

Aerial Reconstruction

http://www.pix4d.com/

6D Pose Estimation in Space

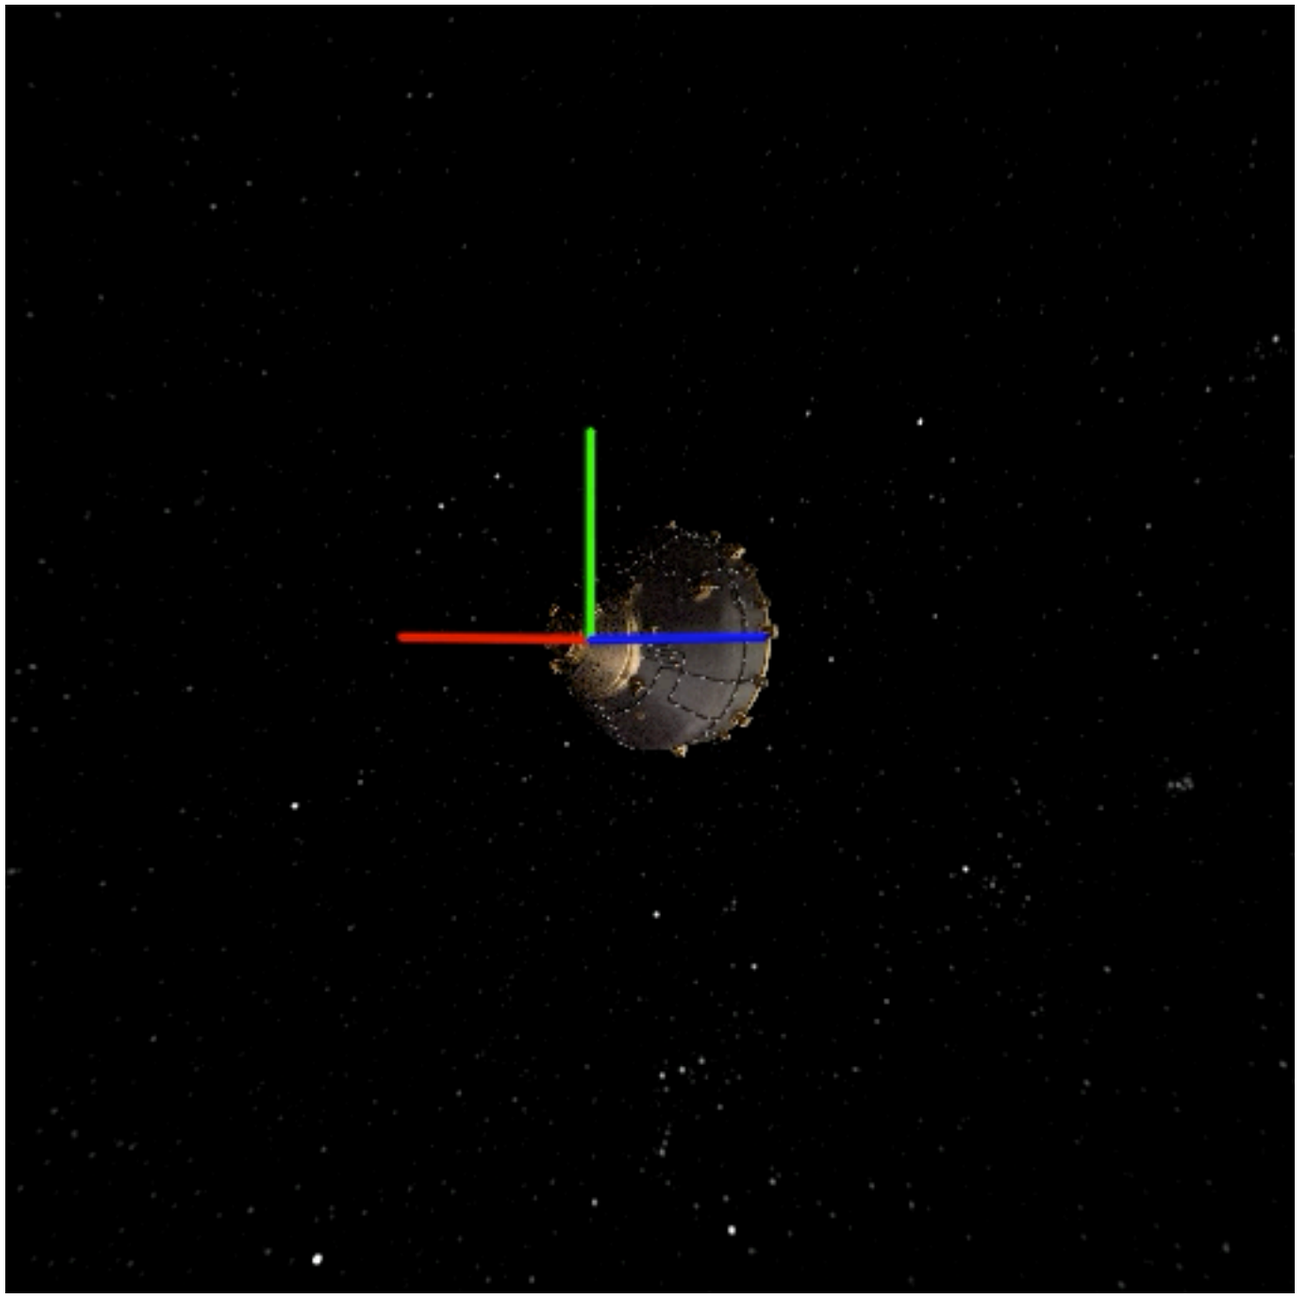

clearspace.today

Deformable Reconstruction and AR

[Dat et al., TPAMI'15]

Mapping the Brain

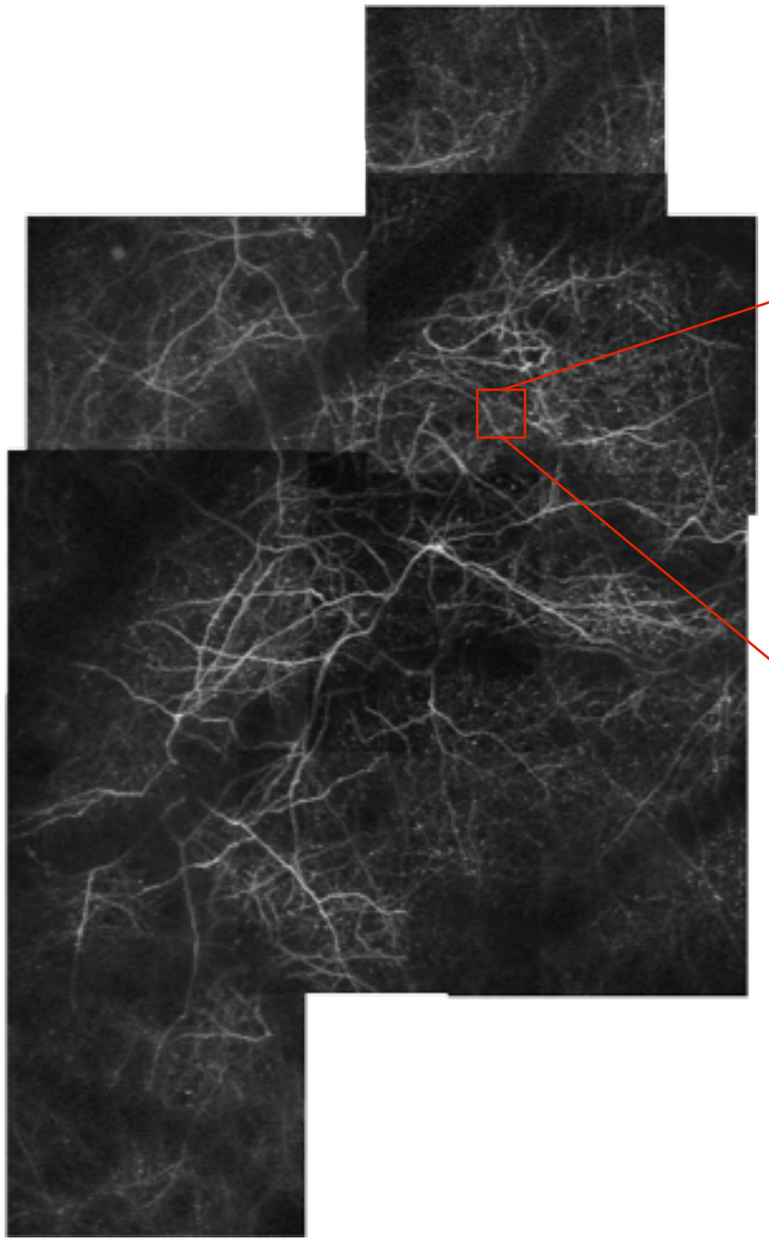

Fluorescent neurons in vivo in the adult mouse brain.

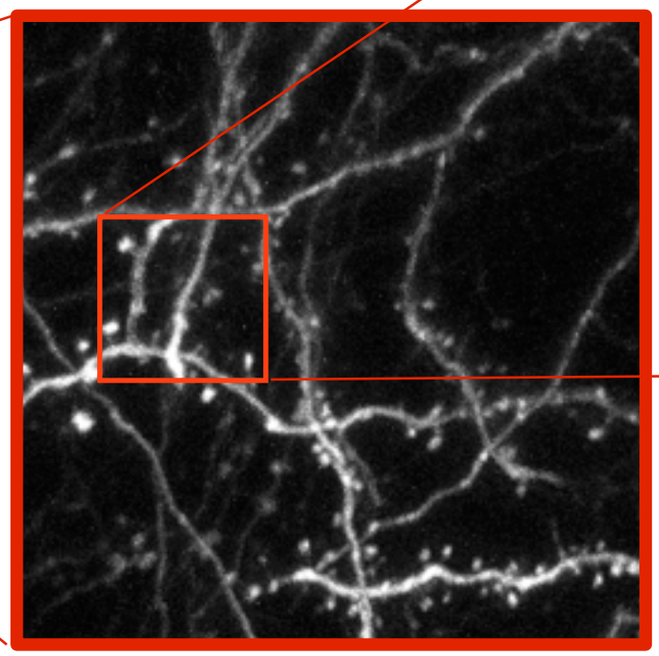

Imaged through a cranial window using a 2-photon microscope.

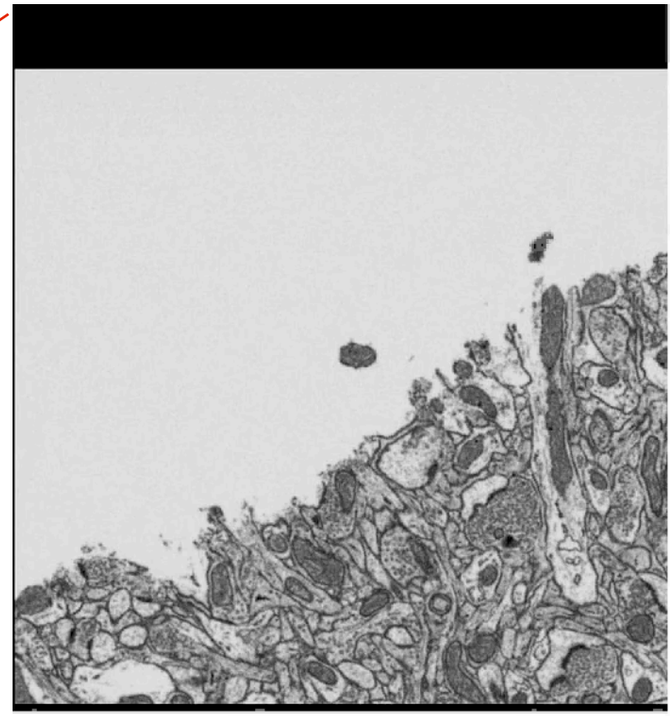

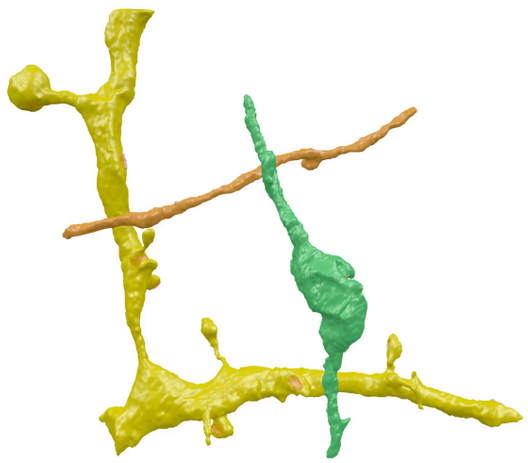

FIB stack and reconstructed neurites.

Courtesy of G. Knott

Neural Structures in Light Microscopy

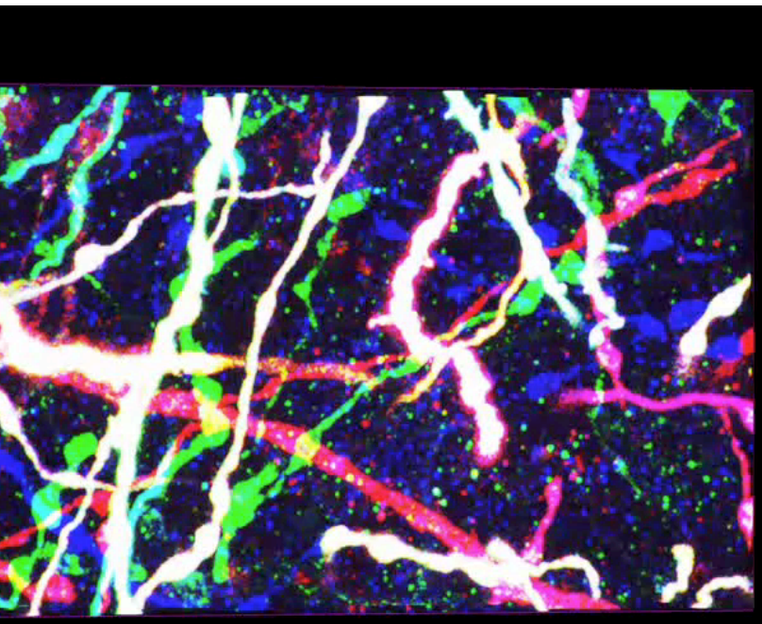

Brainbow Stack

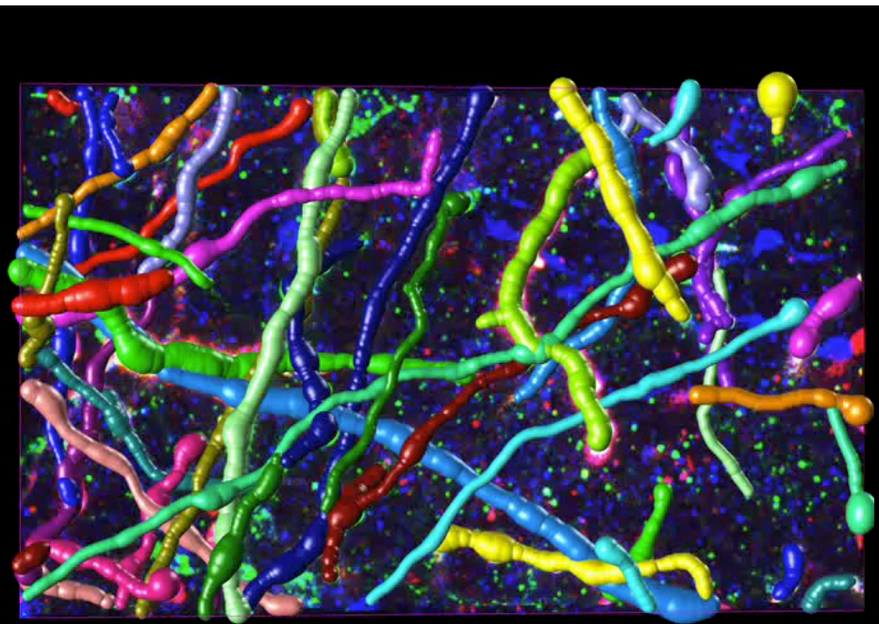

Ground Truth

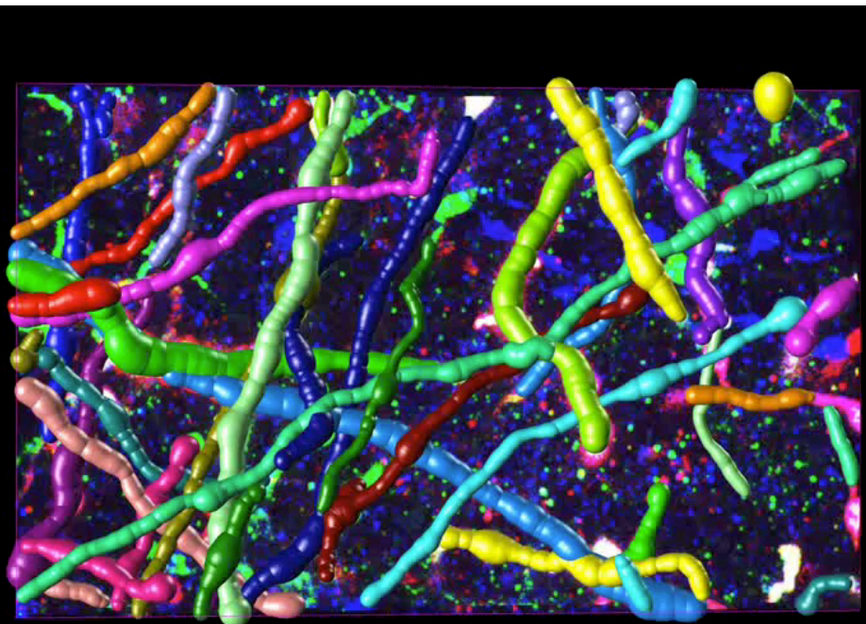

QMIP reconstruction

Courtesy of J. Lichtman

Voxel to Mesh

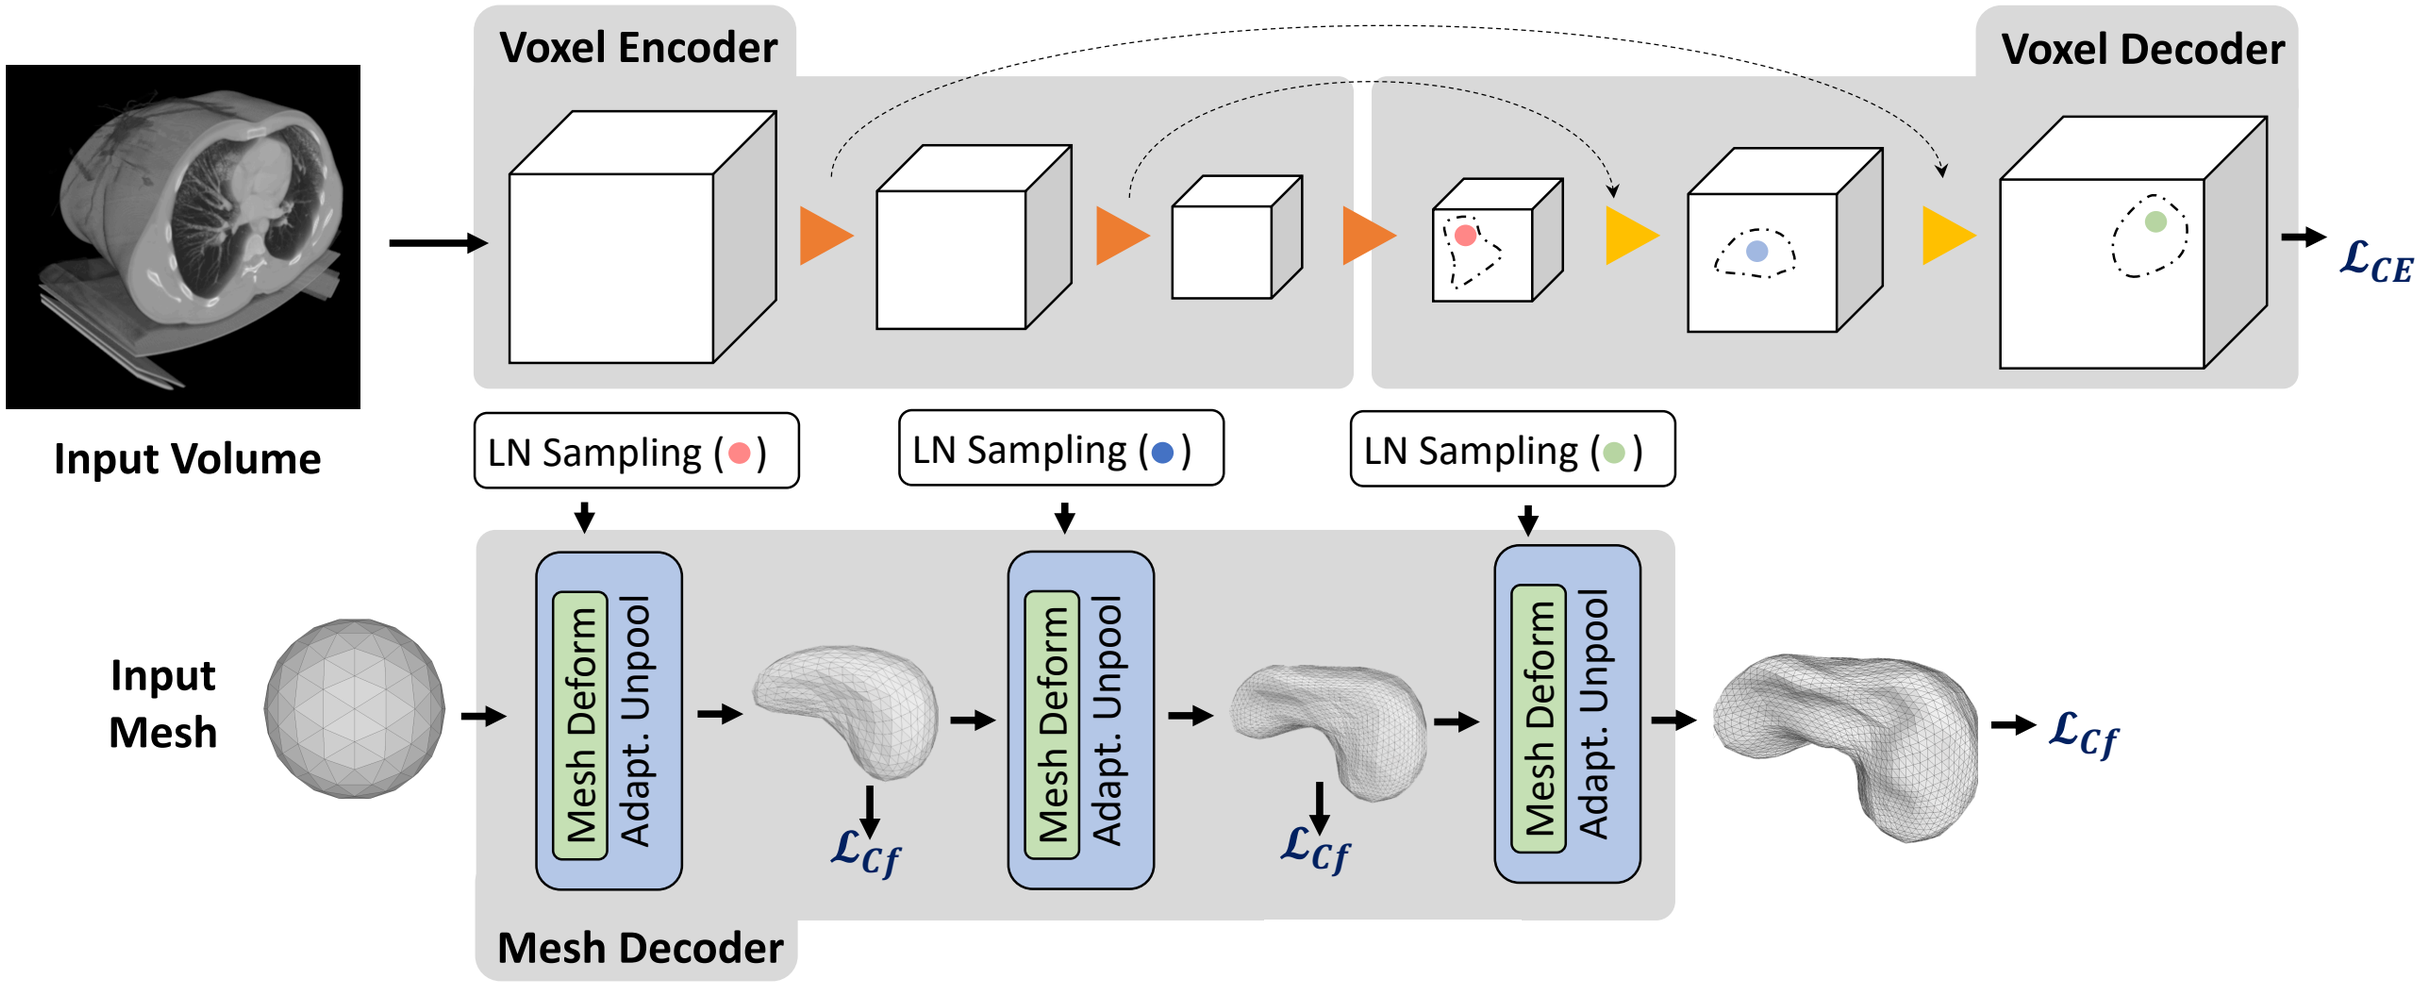

- Input: Image stack and spherical mesh
- Output: Deformed mesh

Deep Active Shape Models

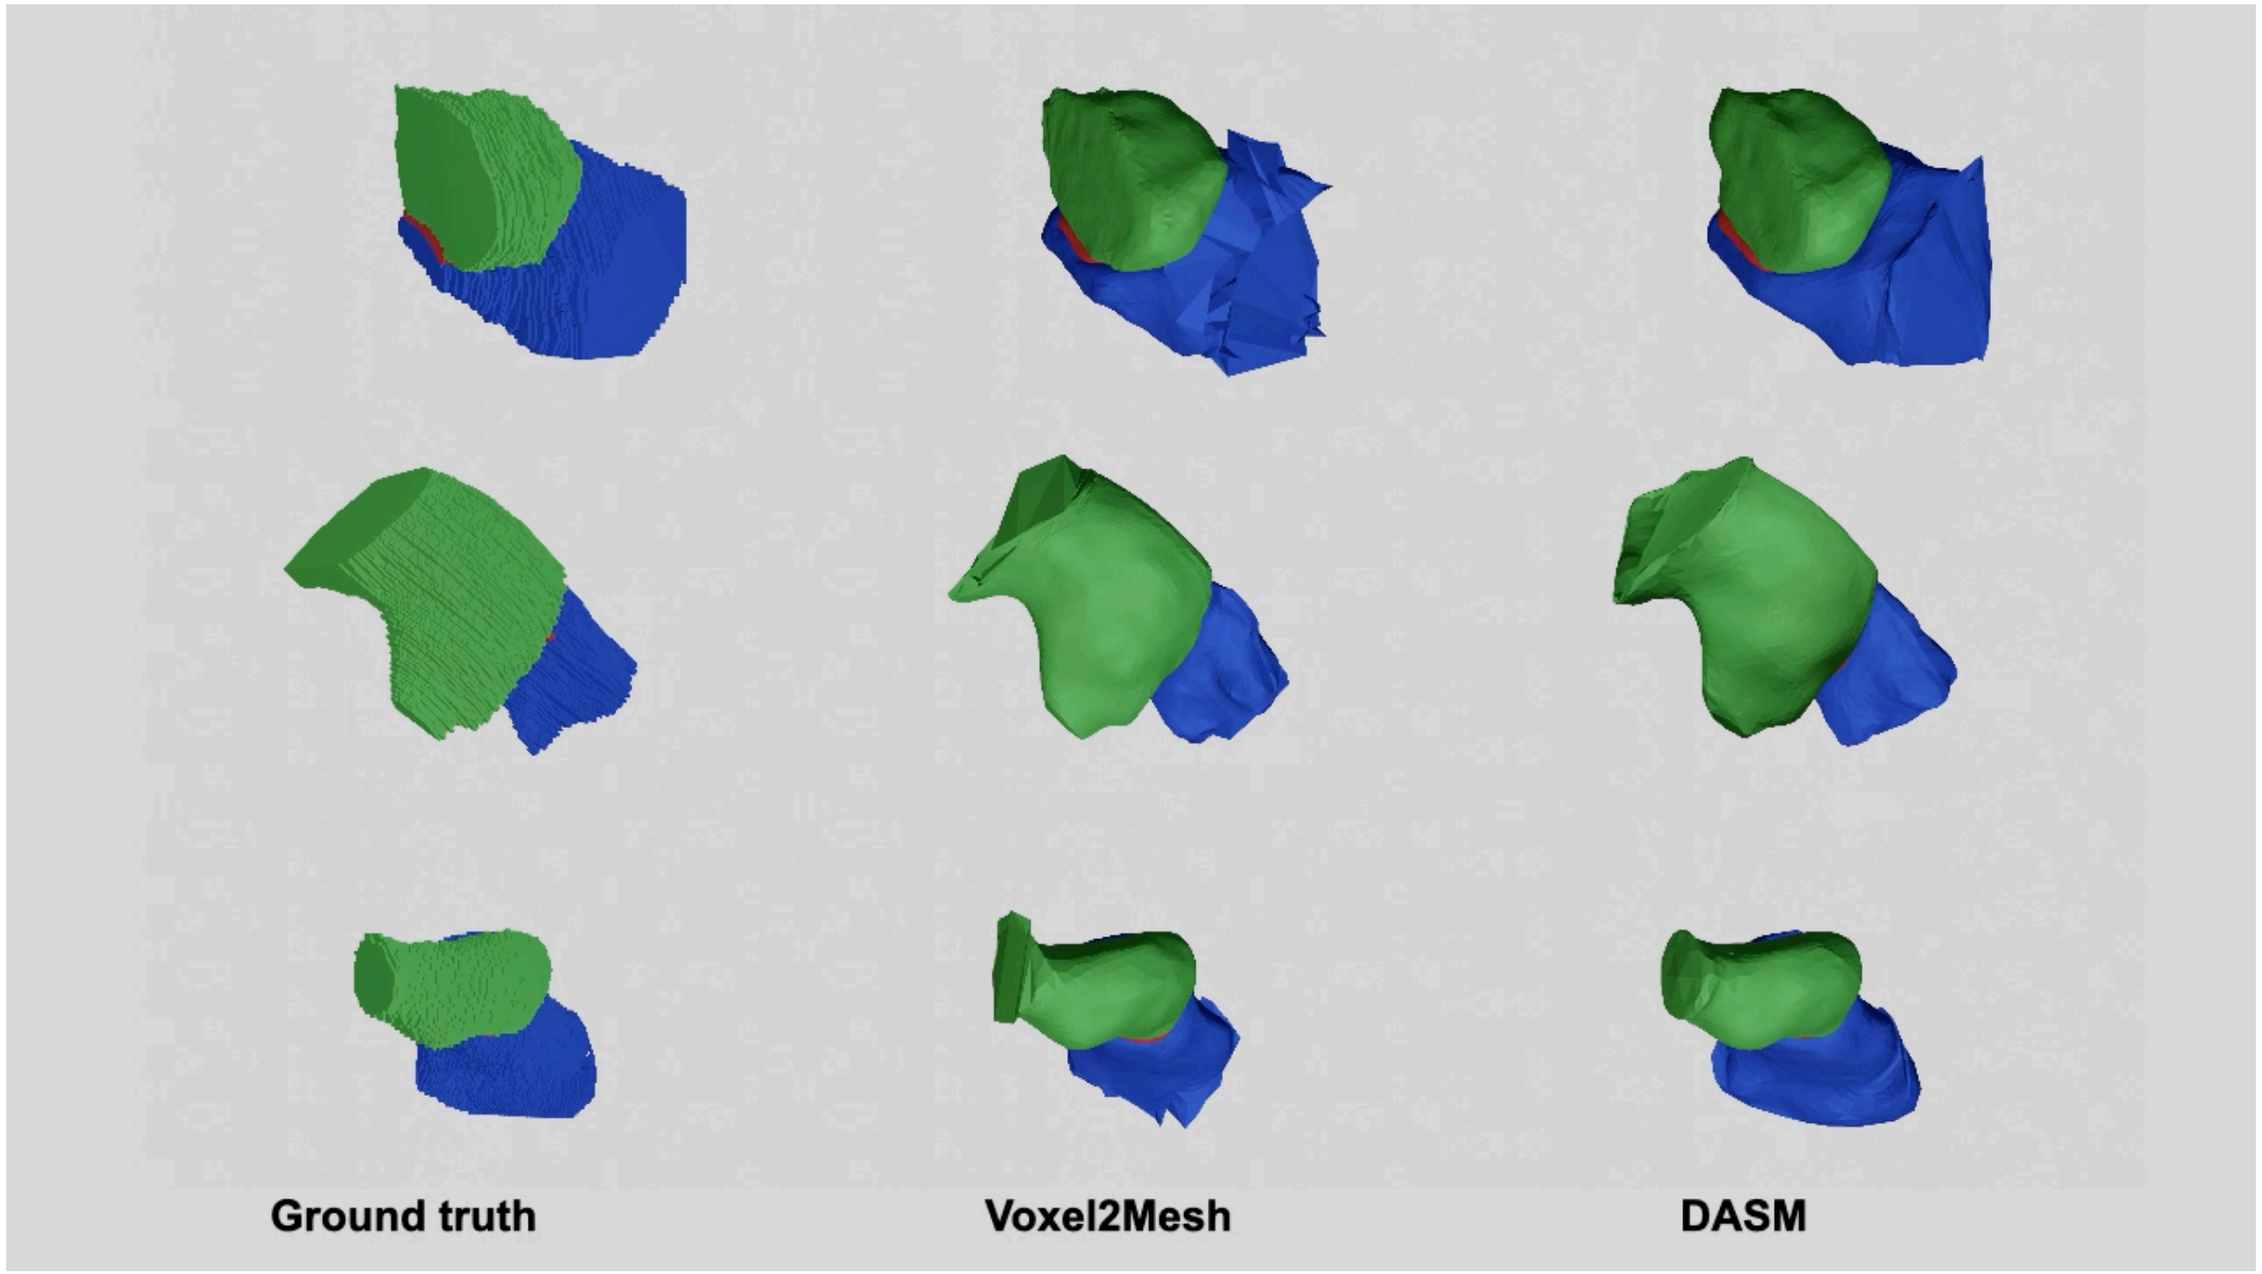

Active Shape Models layers embedded in the network

PhD students needed!

3D Deformable Surface

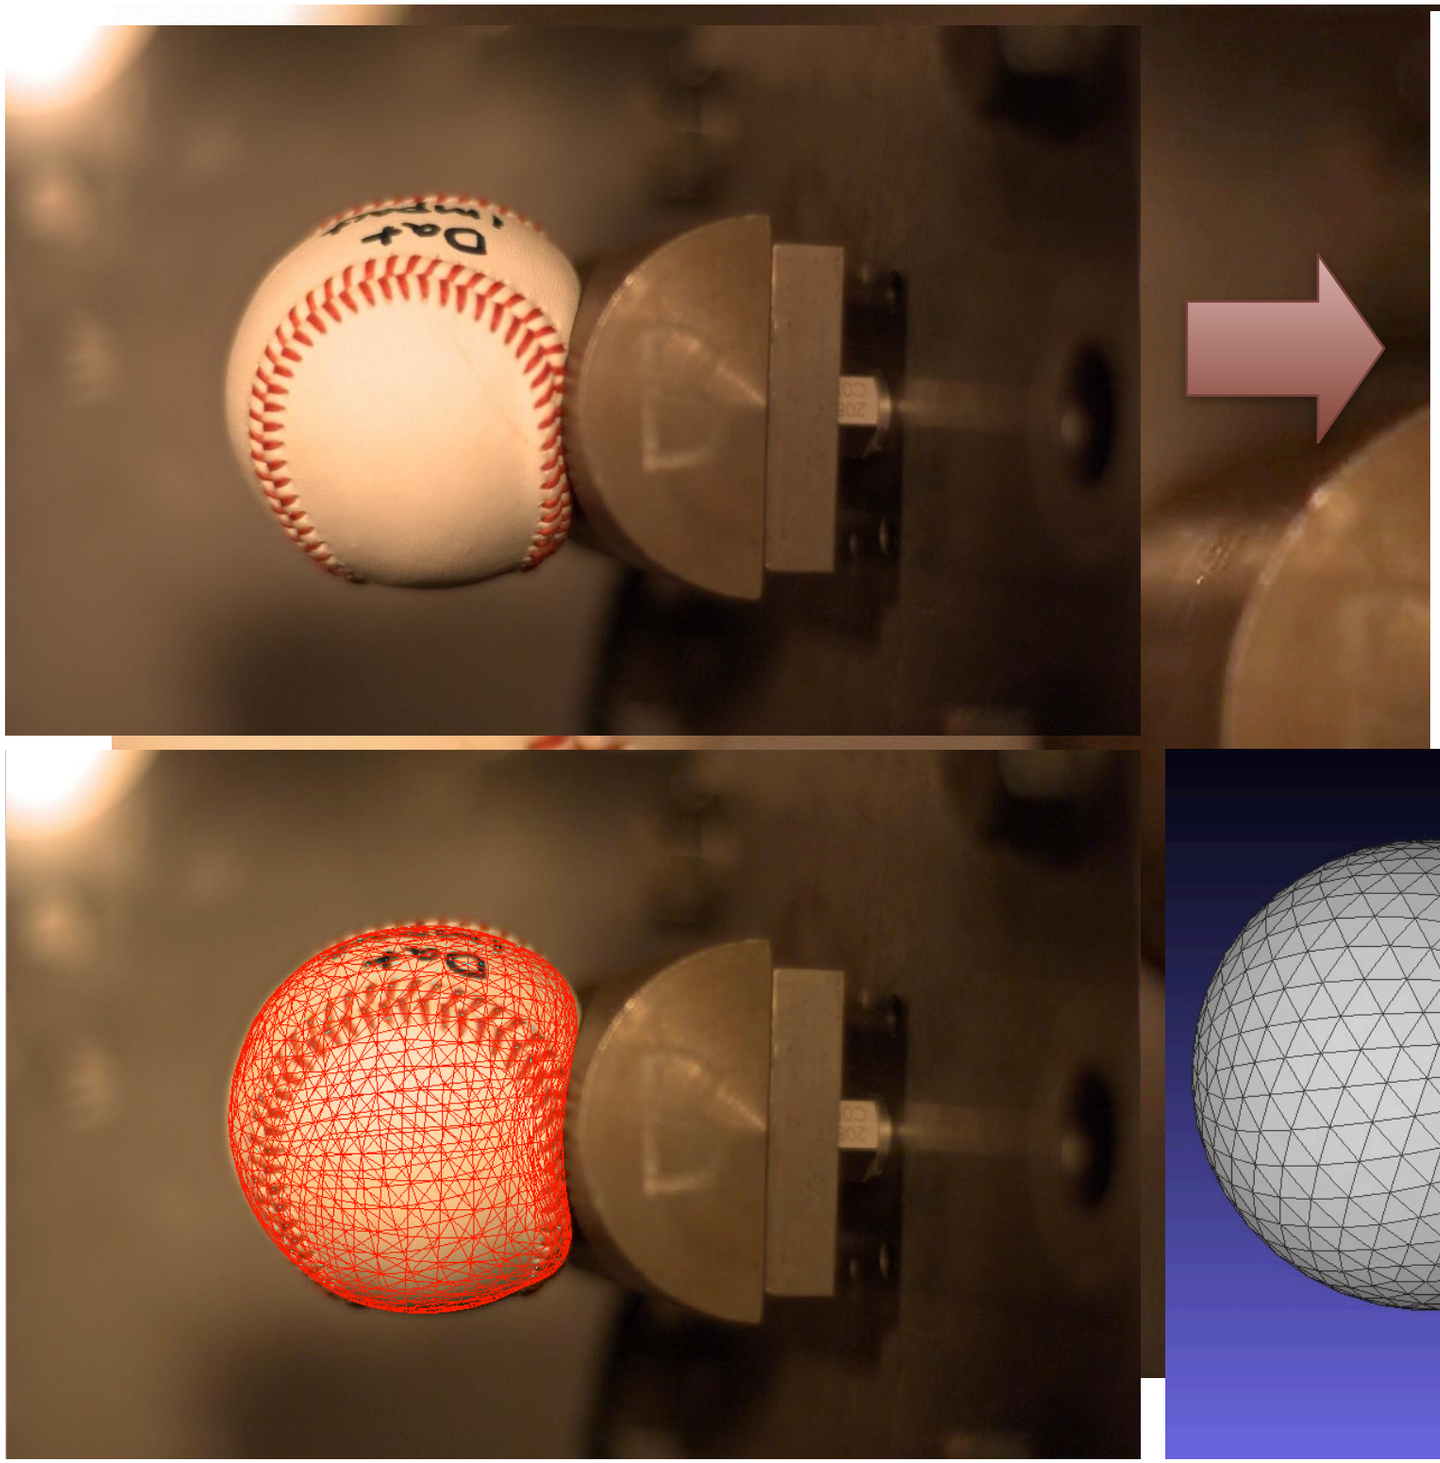

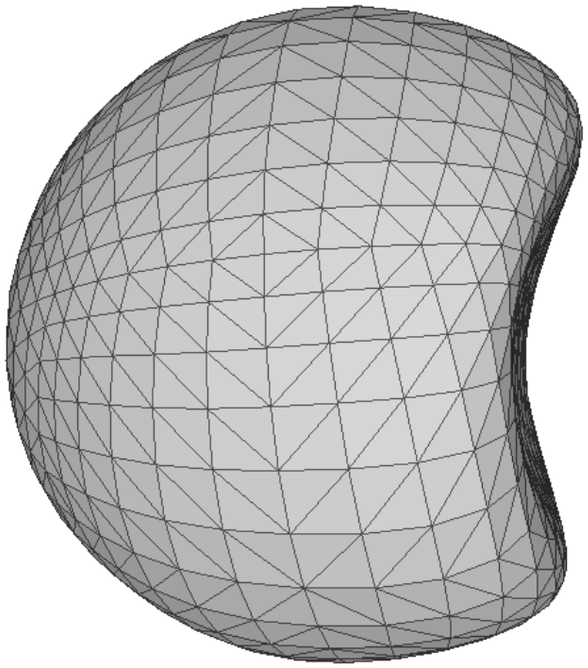

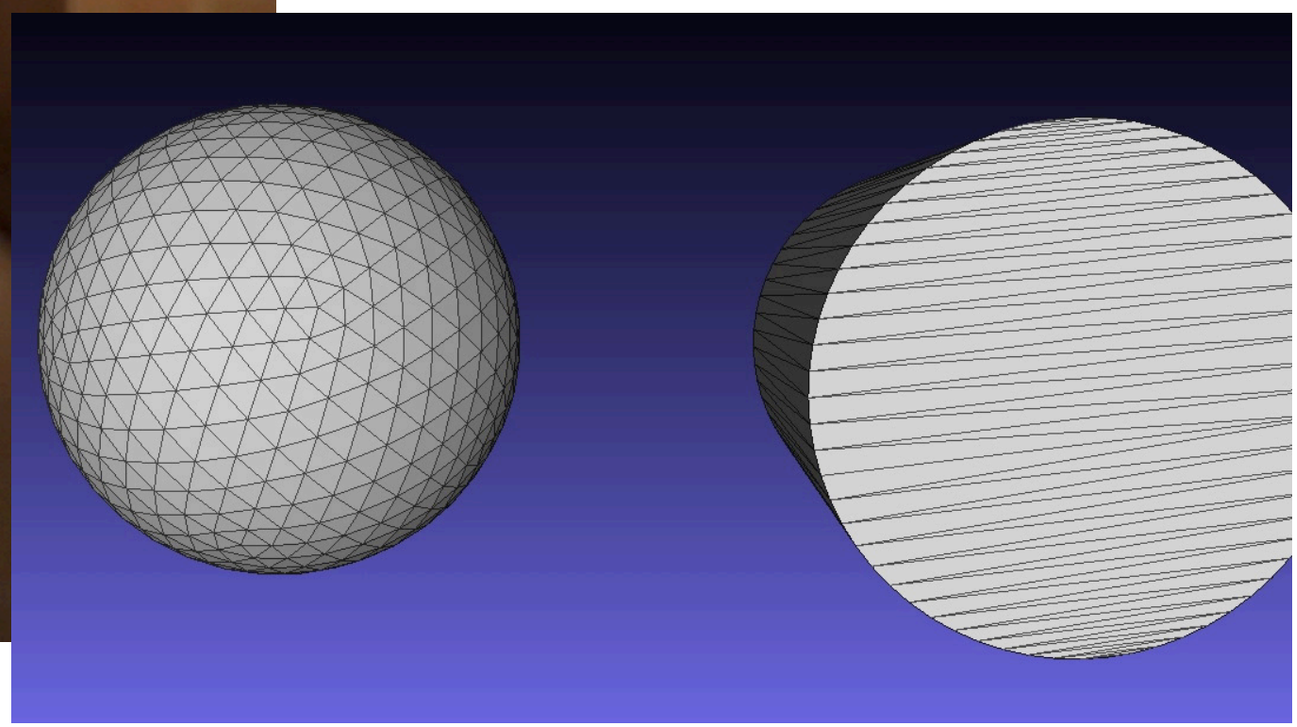

Deep Surrogate Method

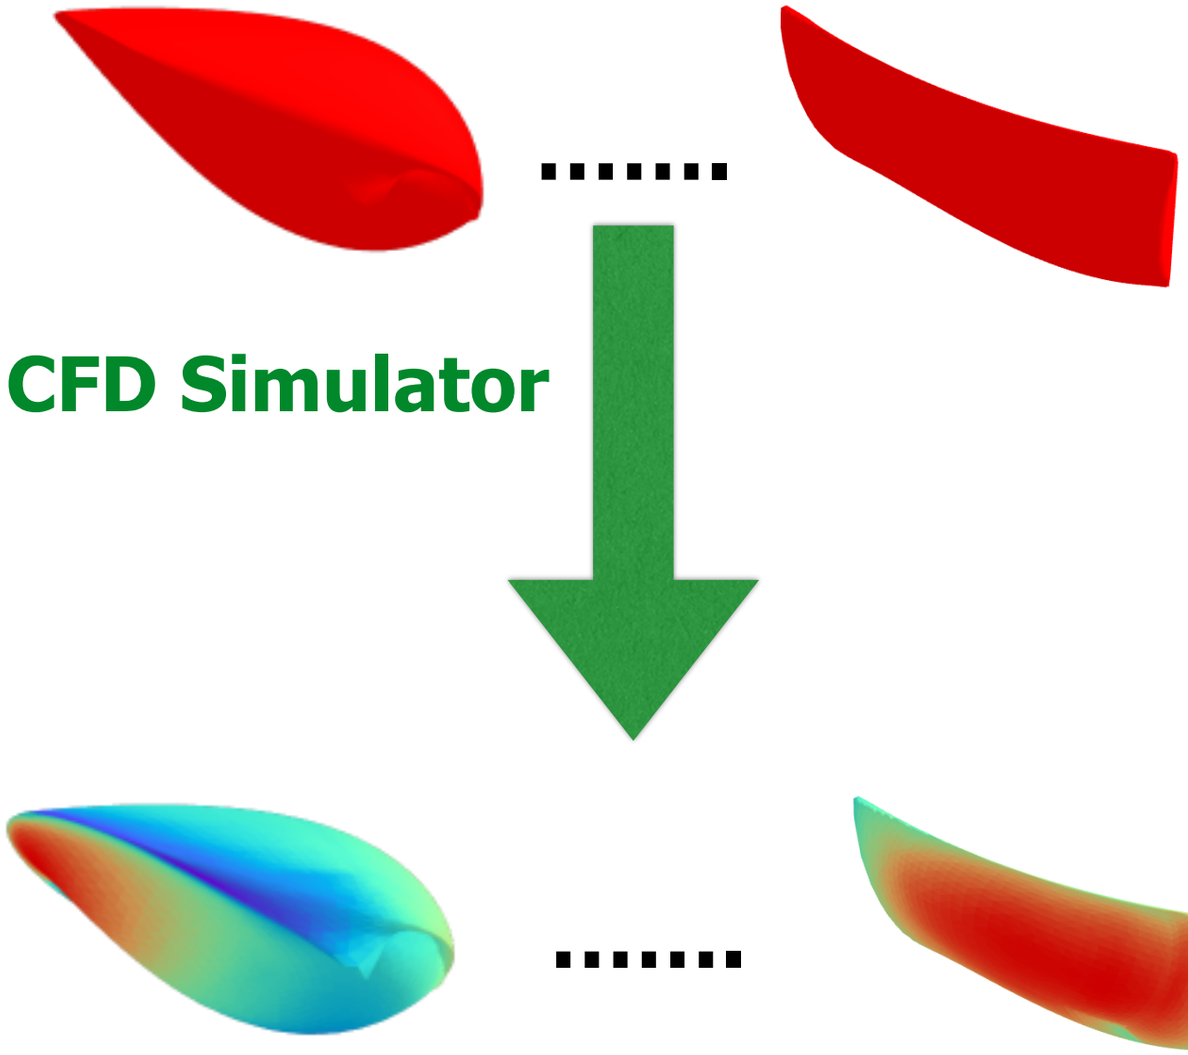

- Drag
- Pressure Coefficients
- Boundary Layer Velocities
- …

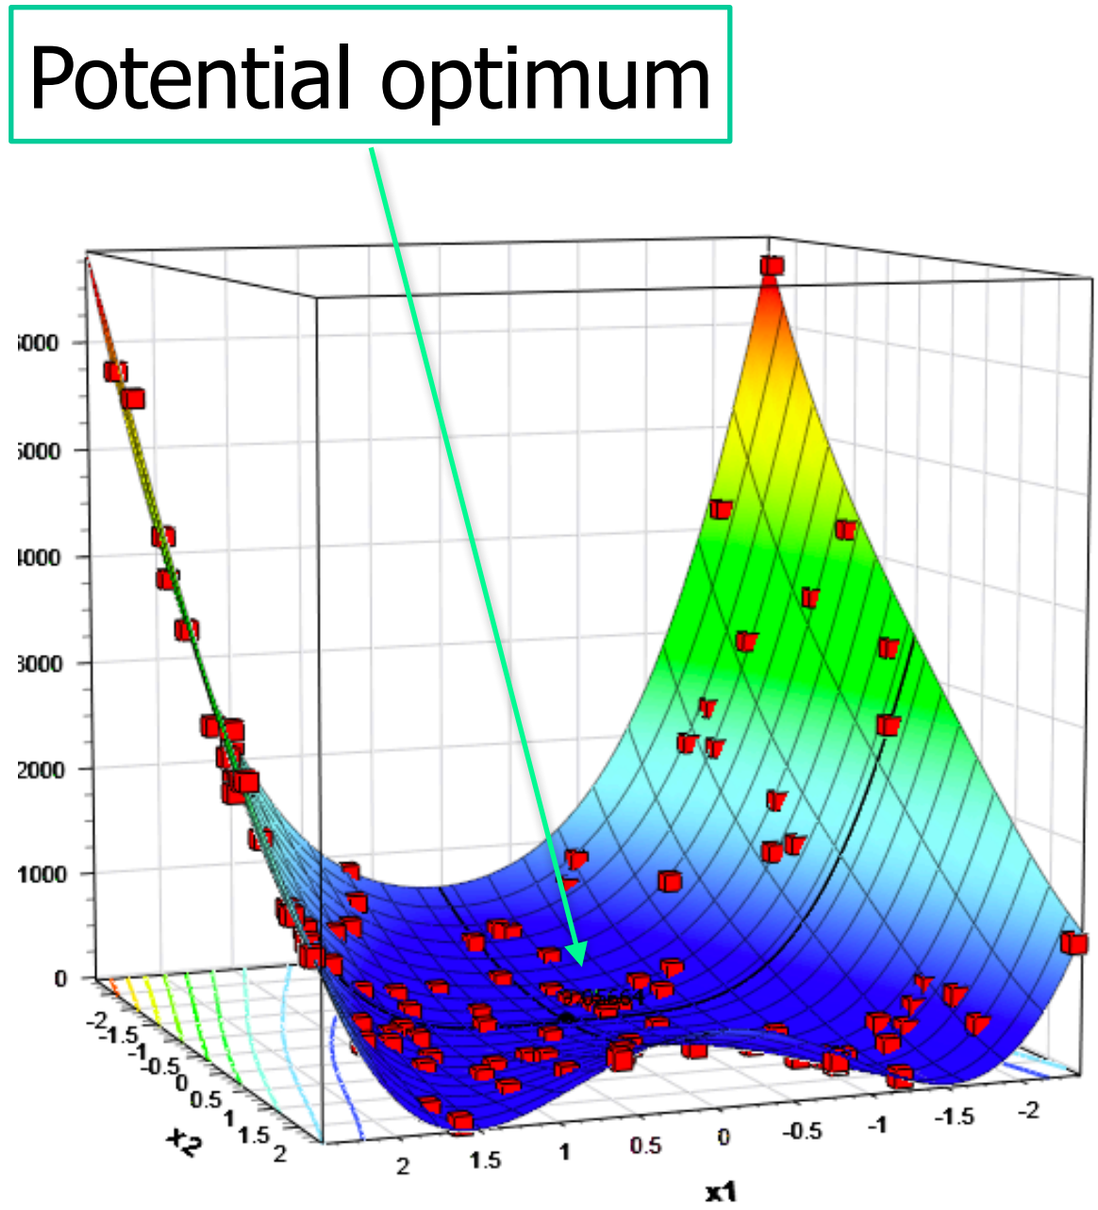

The response surface can be approximated by a geodesic CNN trained to do so.

[Baque et al., ICML'18]

Shape Optimization

Changing the topology

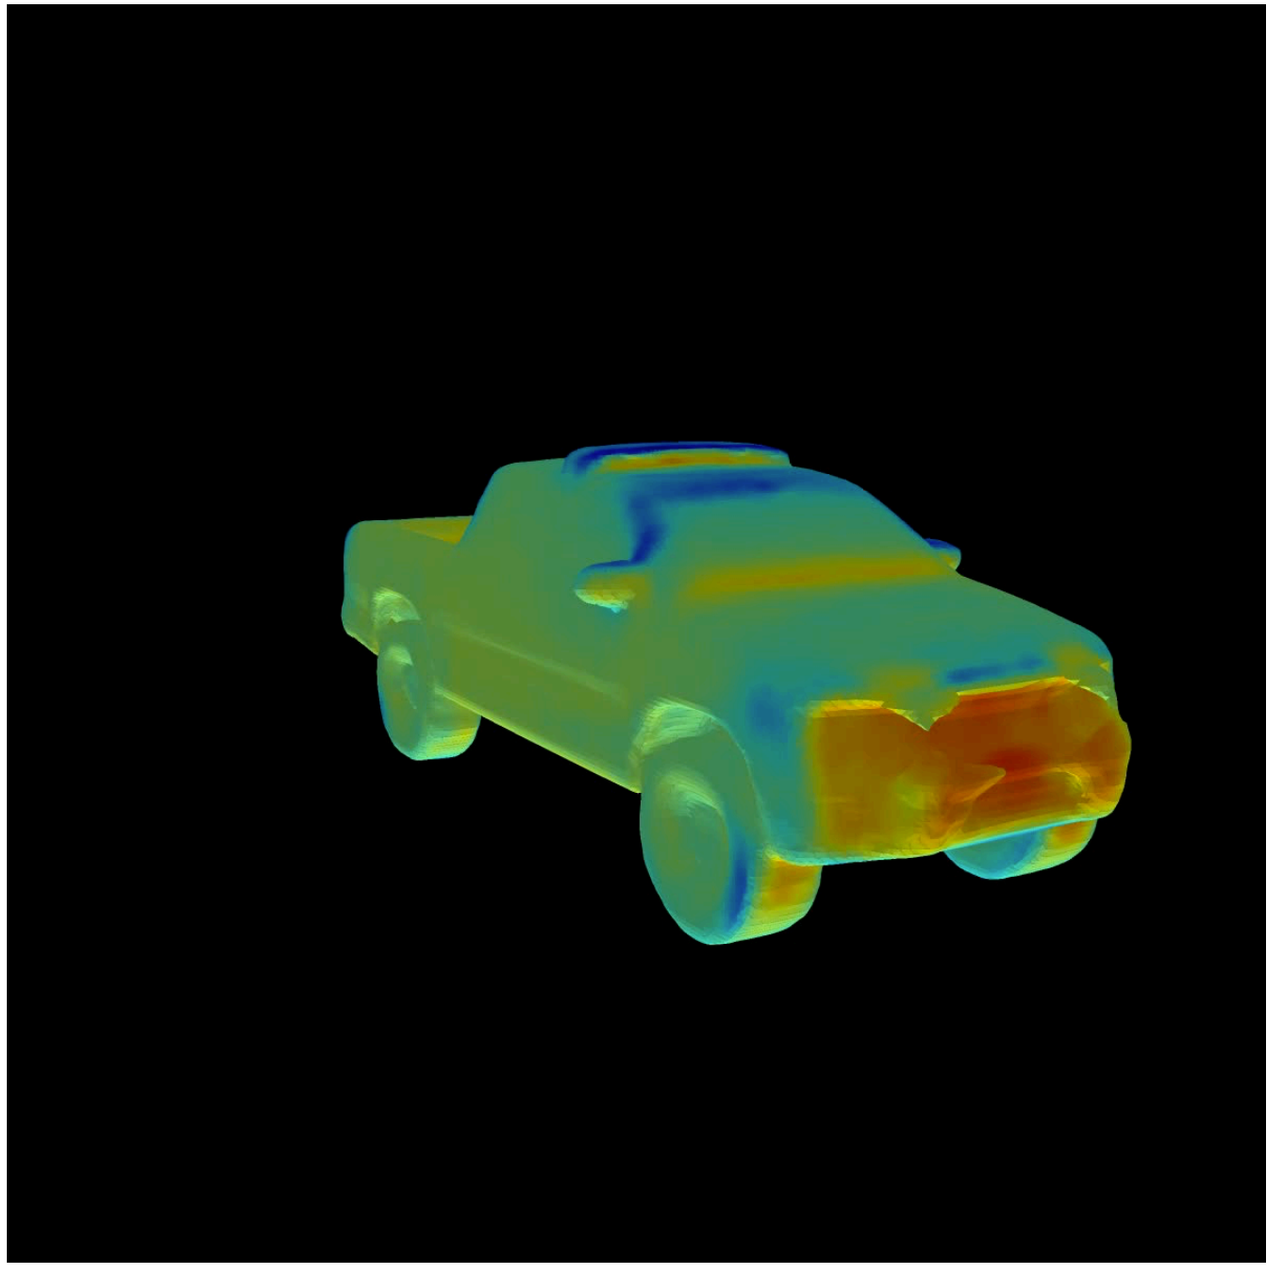

Minimizing the drag

Bicycle Shell

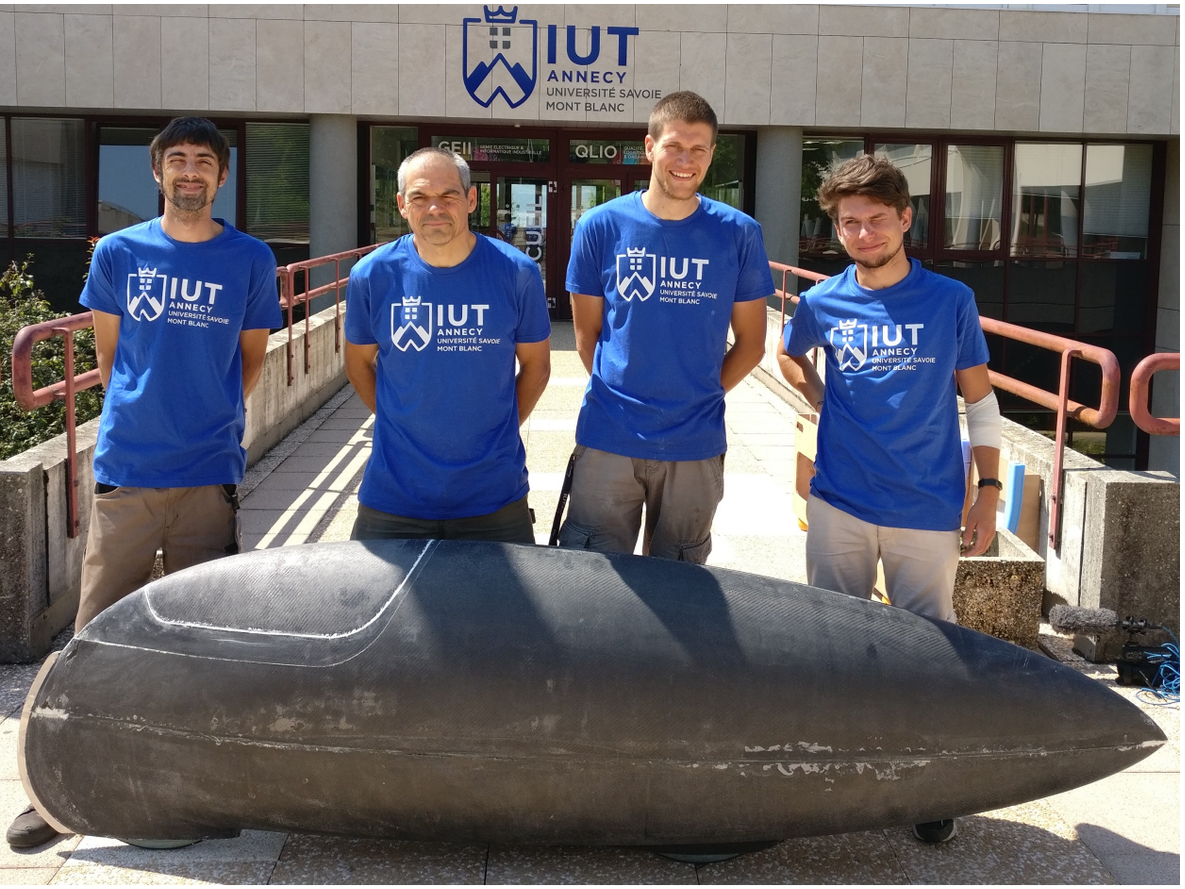

Altair 6, IUT Annecy, 2018

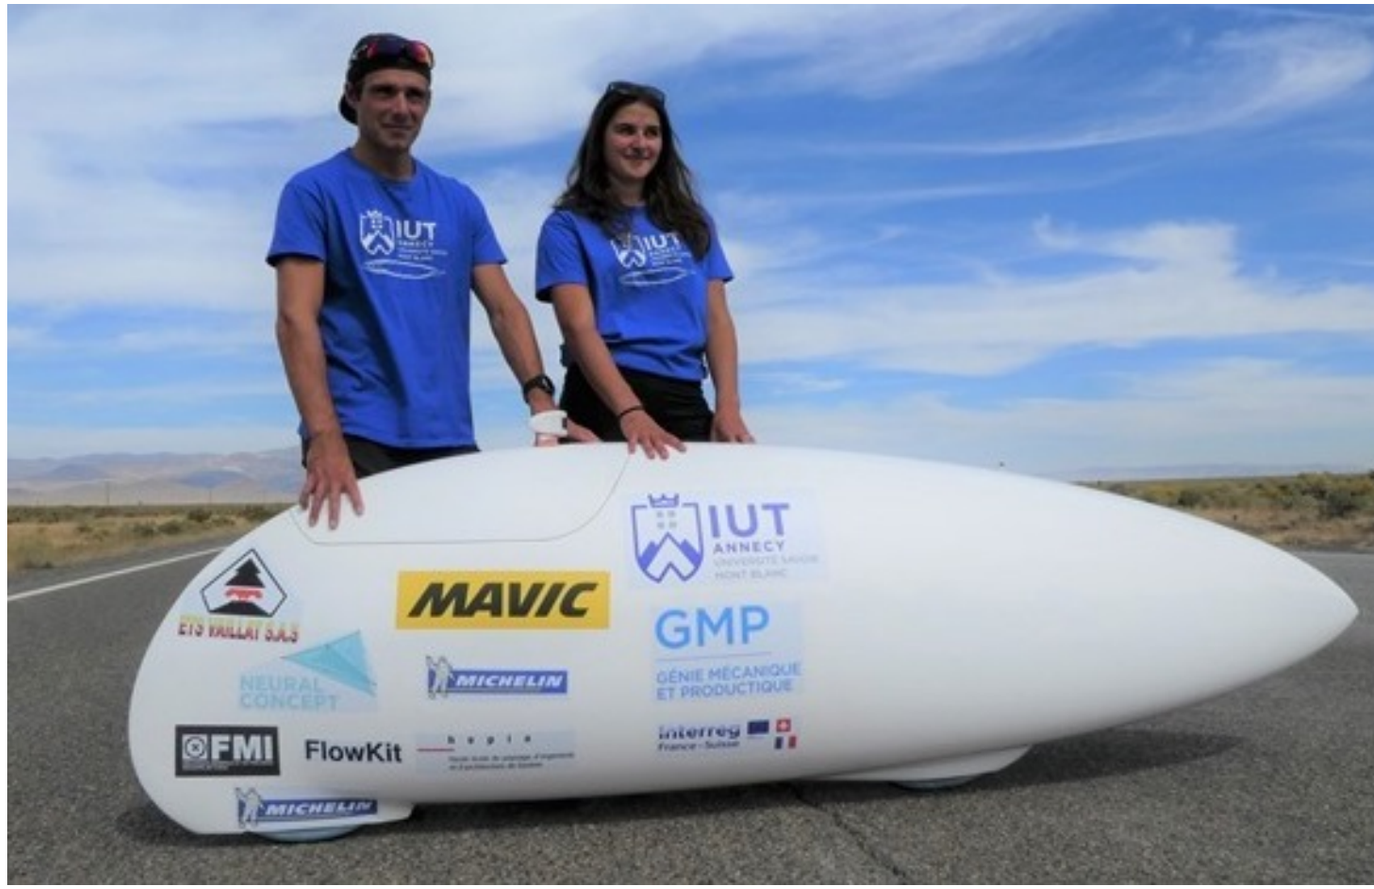

World Human Powered Speed Challenge
Battle Mountain Nevada, 2019

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

From Satellite Systems to Nuclear Fusion

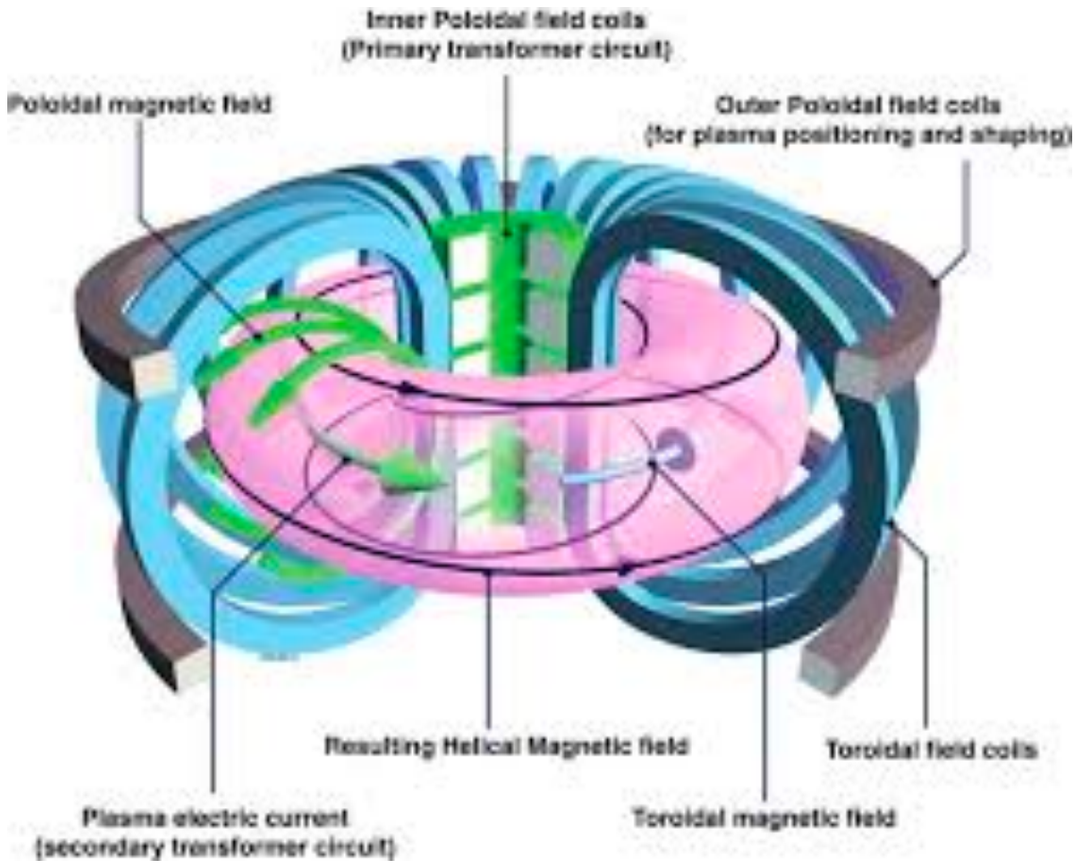

PhD students needed!

People Tracking

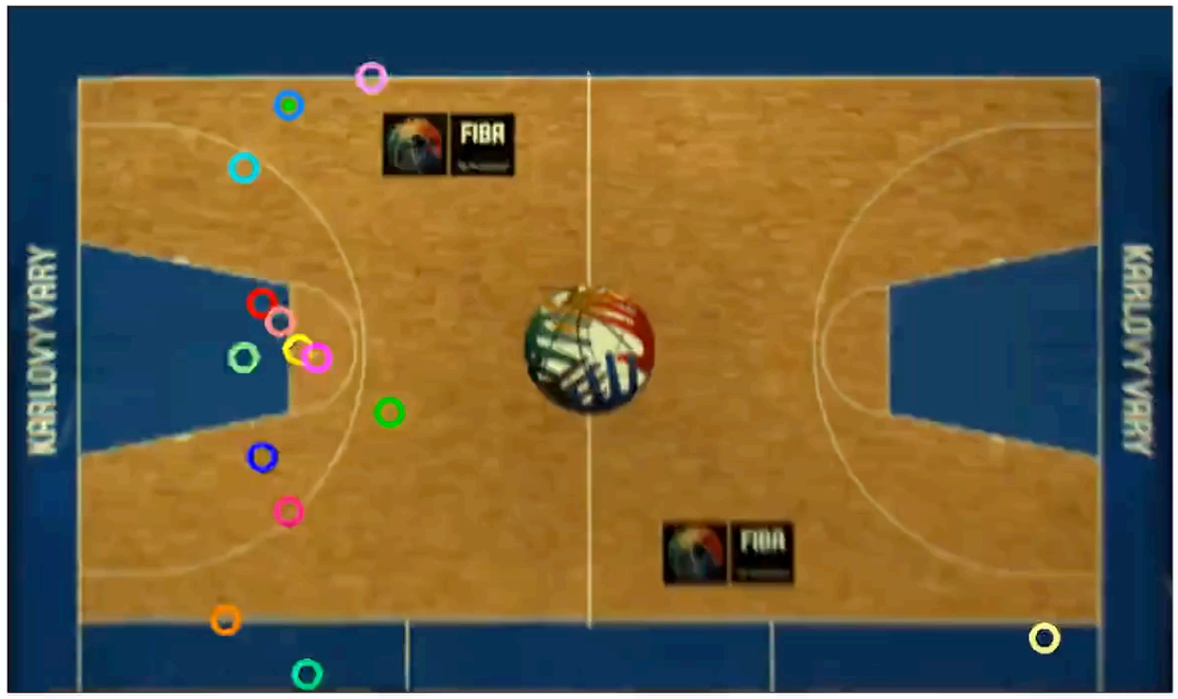

Two-step process:

1. Detection in each time-frame independently using a generative model
2. Linking detections across frames using an LP solver

-> System is very robust to occlusions and occasional detection failures

From White Board to NBA and Soccer

- 2005: First ICCV paper published
- 2014: PlayfulVision founded
- 2015: PlayfulVision acquired by SecondSpectrum
- 2016: NBA Official Optical Tracking Provider
- 2017: System deployed in NBA arenas
- 2019: Premier League Optical Tracking Provider

Denser Crowds

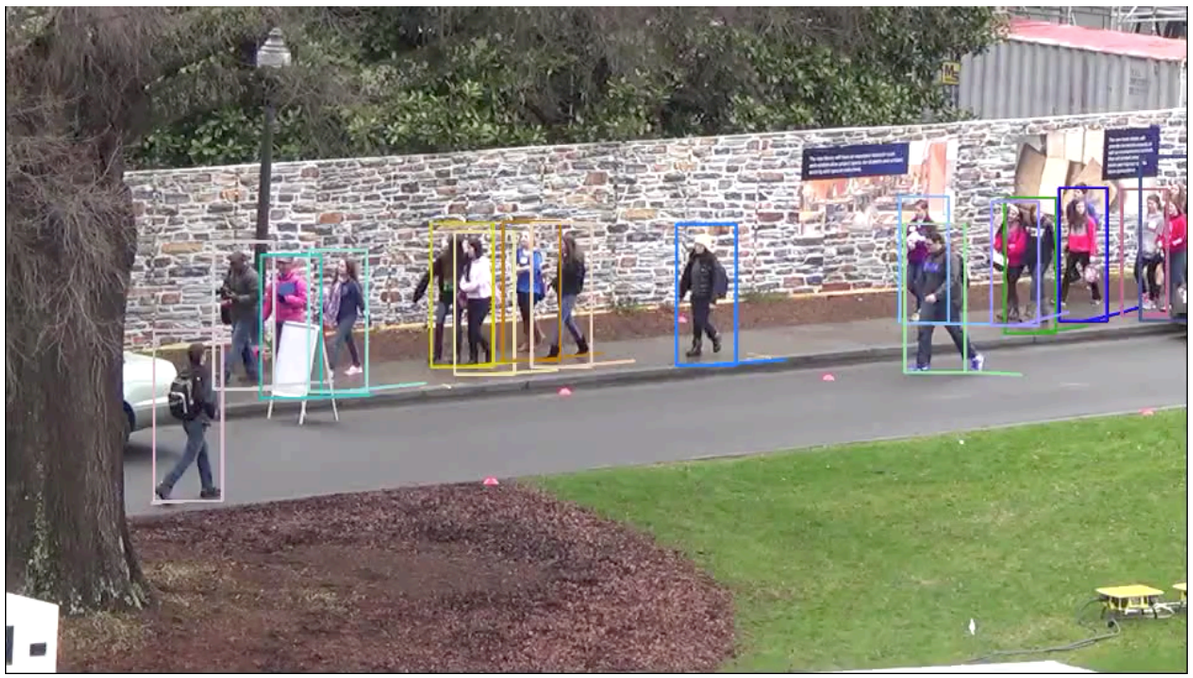

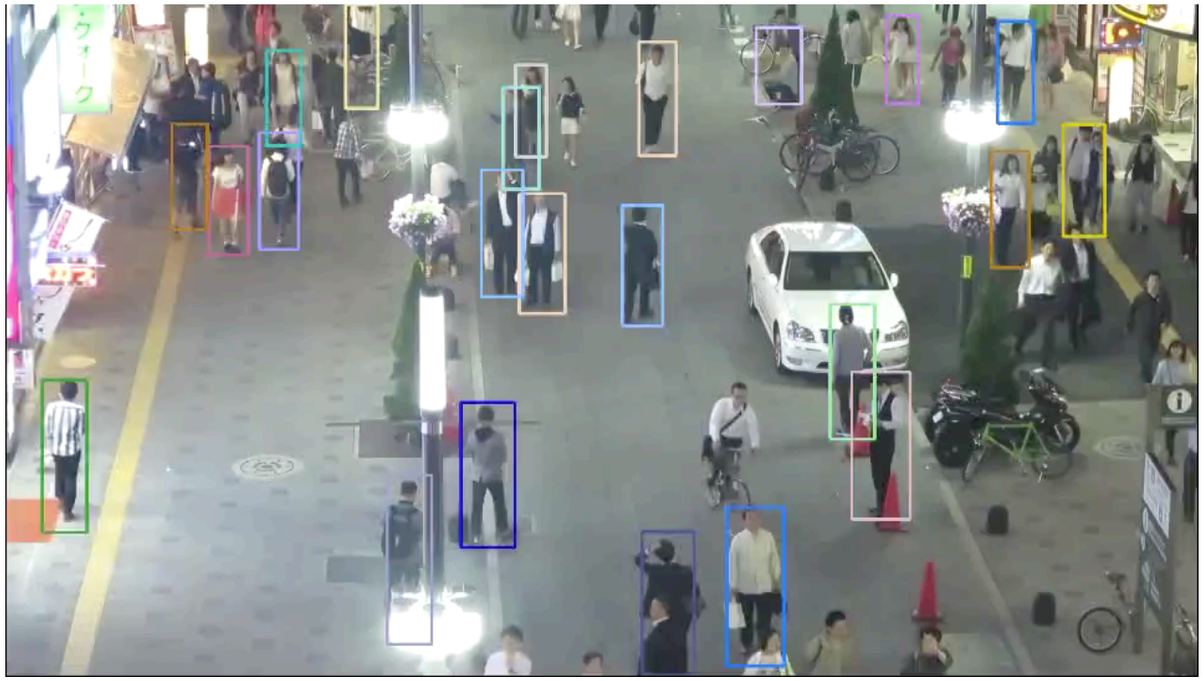

- Bounding boxes and last 3s of tracking
- Long tails denote identity preservation

… and even Denser

Venice

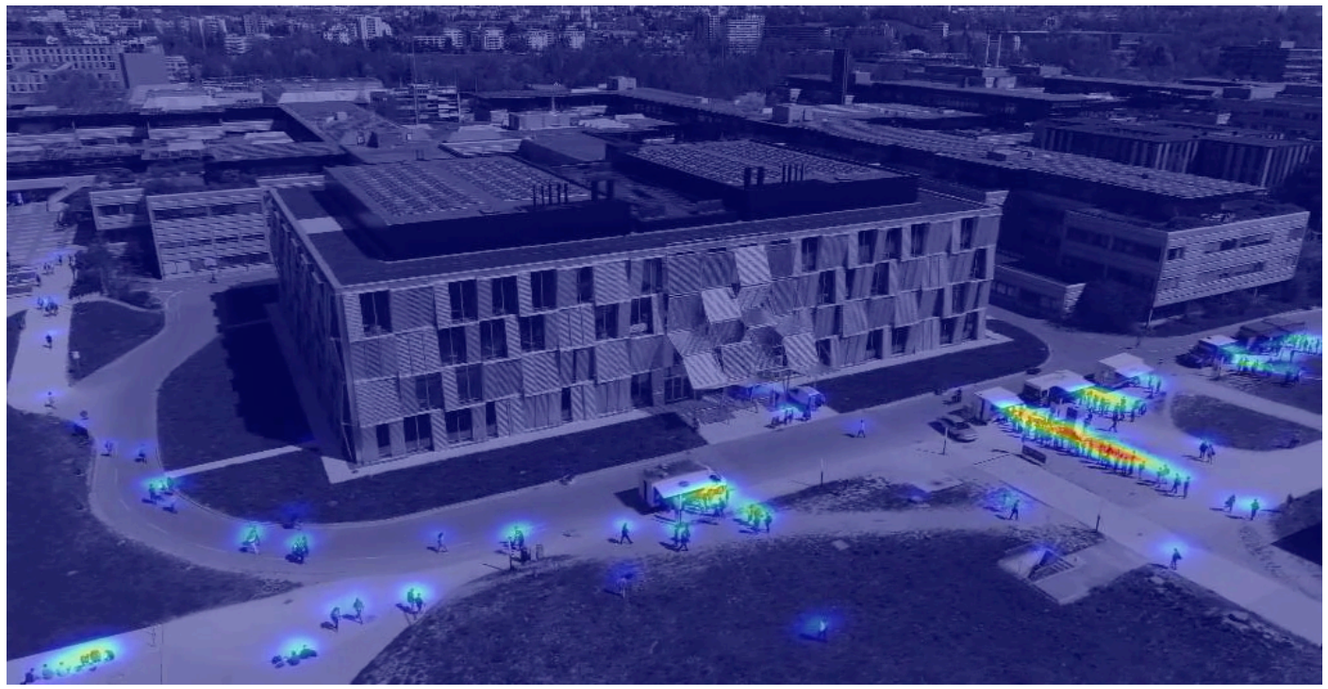

Lausanne

- Regress from image to people density
- Impose geometric constraints
- Impose temporal consistency on the flow

Monocular 3D Pose Recovery

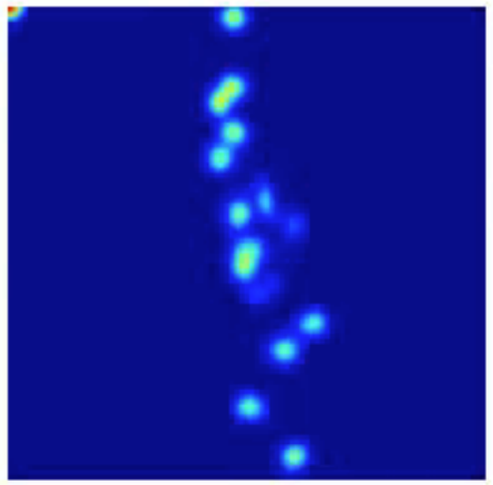

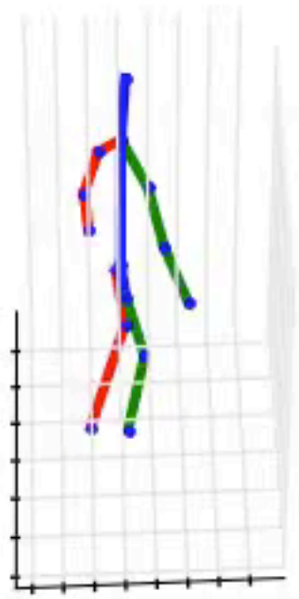

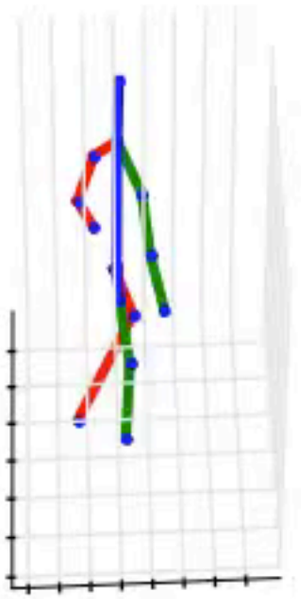

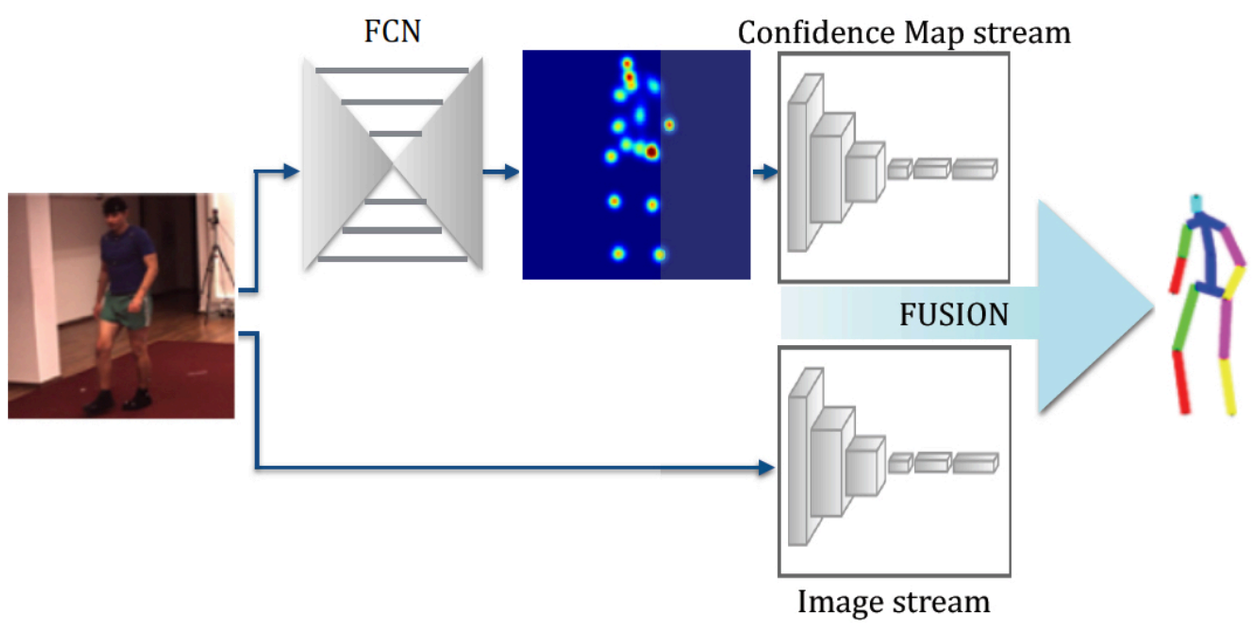

Qualitative results on skiing

Training on the ski dataset with weak multi-view supervision improves accuracy.

Cropped input

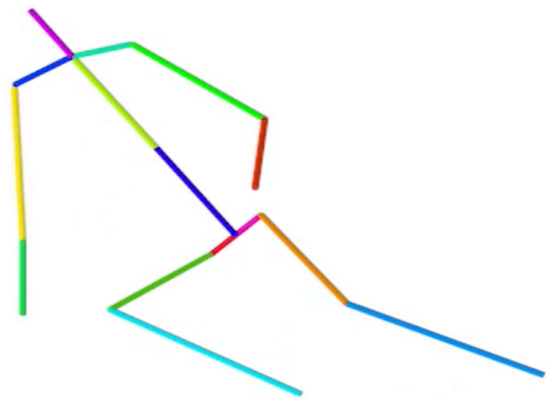

Our method*
(trained on new ski dataset)

Baseline*
(trained on H3.6M)

*smoothed temporally with a Gaussian window of std=1

Multiple People

Multi-person 3D pose estimation with NSD

Input

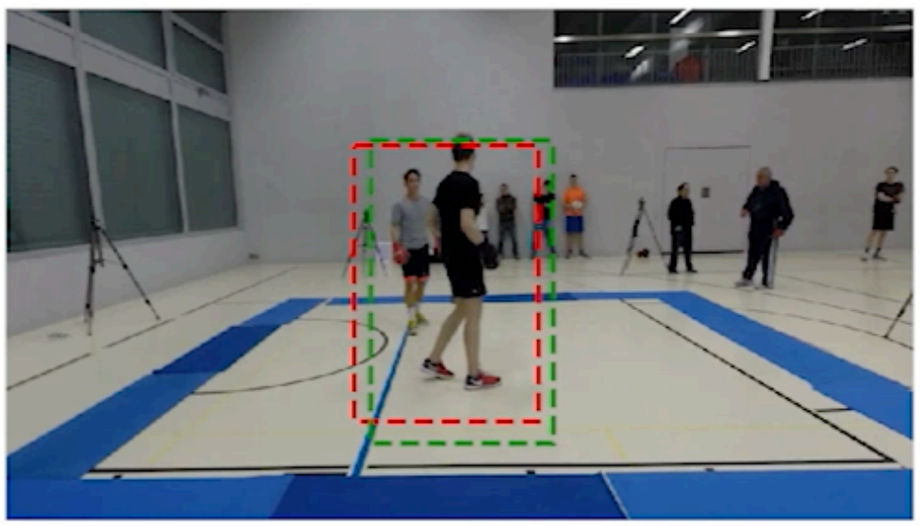

Bounding box

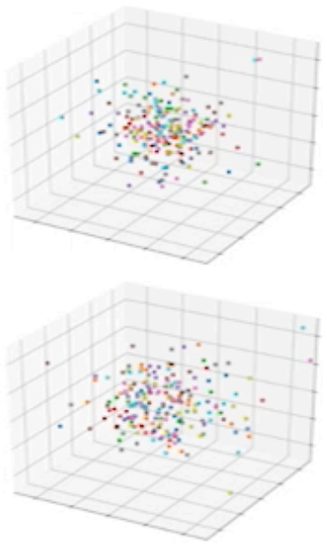

3D latent spaces

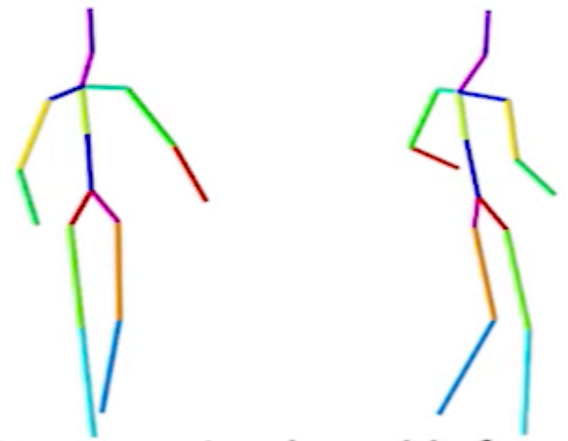

GT 3D poses (ordered left to right)

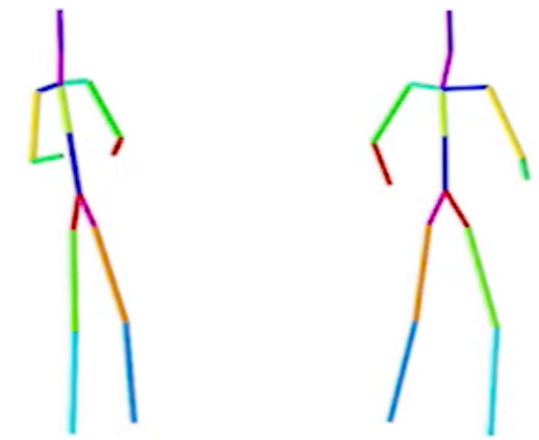

Our 3D pose prediction
(ordered left to right)

Temporal Consistency and Physics

Modeling Clothed People

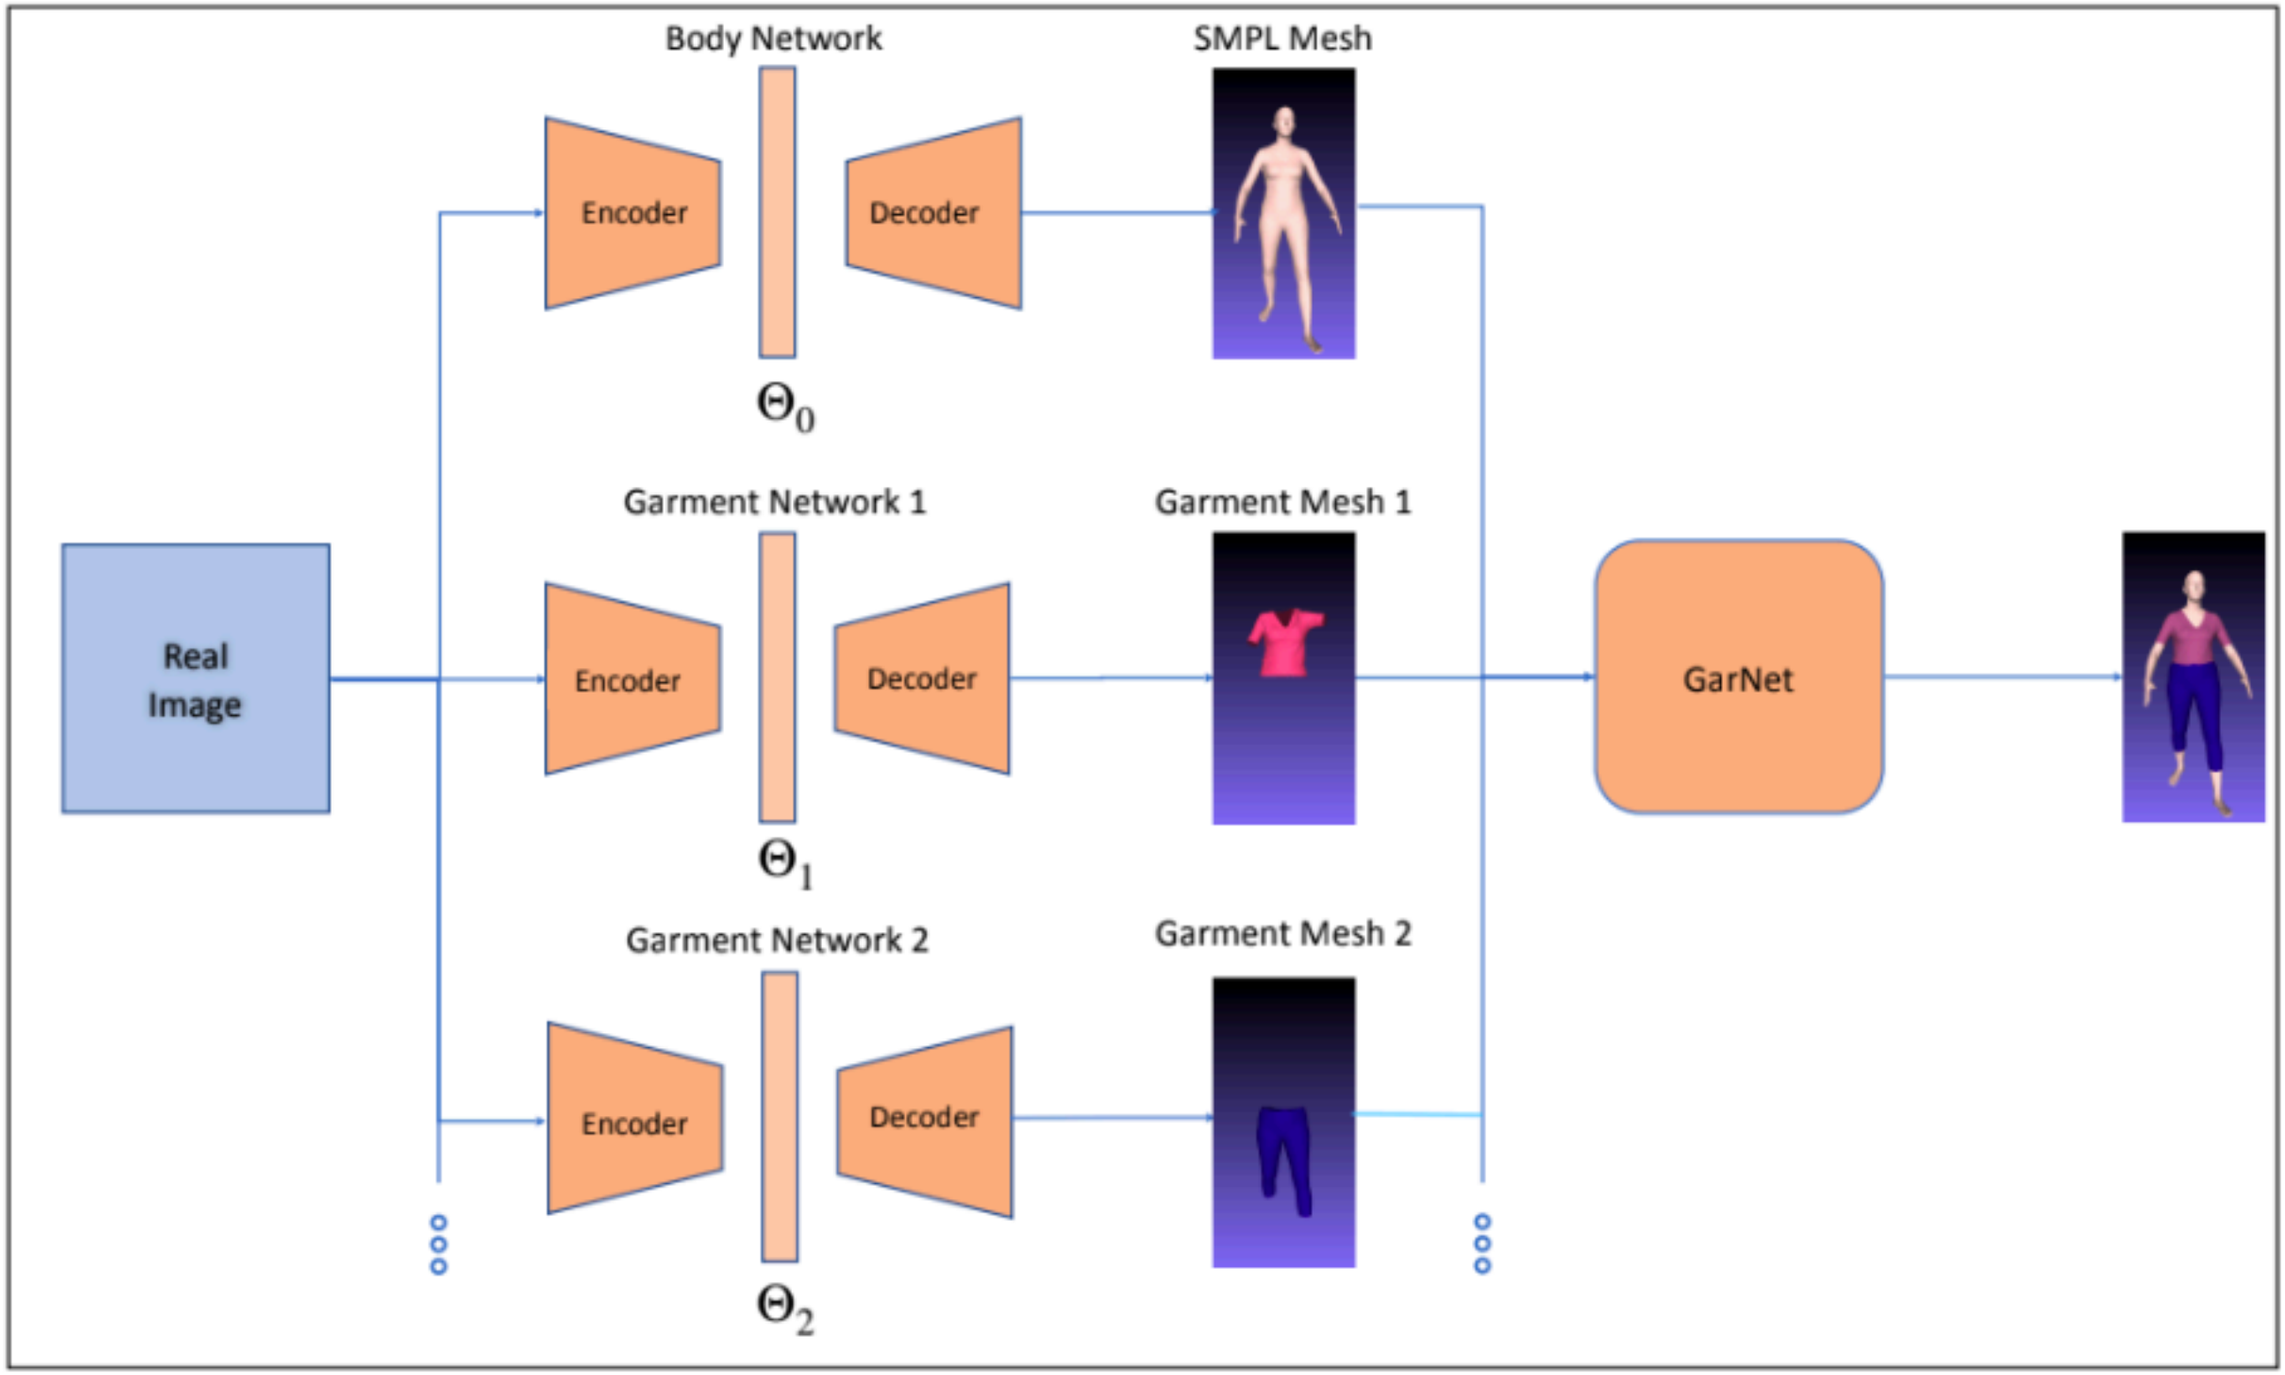

PhD students needed!

Future Work

Deep Learning now is at the heart of most of our algorithms but this not the end of the story:

- Geometry and physical constraints still matter
- The required training databases are not always available
- Guarantees and explanations are necessary
- For many tasks, we are still far from human level