MVA

Geometry Processing and Geometric Deep Learning

Today

- Practical Information
- Introduction to the course
- **Actual content:**
  - Surfaces and Shape Analysis
  - Surface features, Discrete representations, Discrete Laplace-Beltrami operator, applications in shape comparison and shape analysis

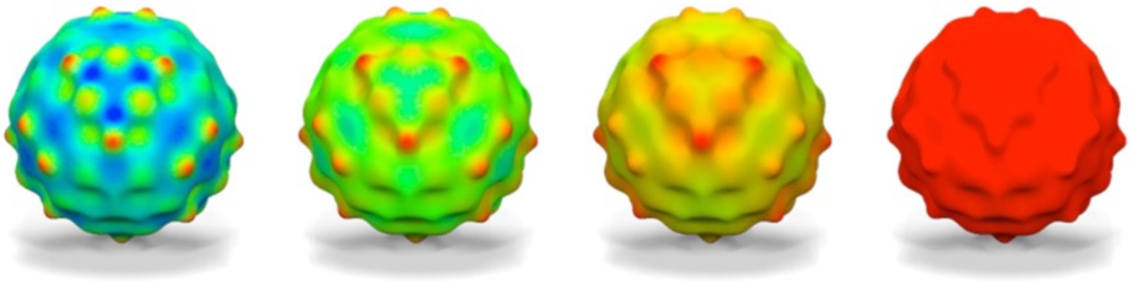

Practical Information – Team

Lectures **TDs**

Maks Ovsjanikov

maks@lix.polytechnique.fr

Etienne Corman

etienne.corman@cnrs.fr

Julie Digne

julie.digne@cnrs.fr

Emery Pierson

emery.pierson@polytechnique.edu

http://www.lix.polytechnique.fr/~maks/ https://members.loria.fr/ECorman/ https://perso.liris.cnrs.fr/julie.digne/

Practical Information

Course website: https://jdigne.github.io/mva_geom/

- 6 Lectures and Practical Sessions (TD)

  Lectures: Wednesdays 13:00 – 15:20

  TD's: Wednesdays 15:40 – 17:40

Lecture slides before each lecture

Please register: https://forms.gle/PNmMpJyfwjZEyT1M8

Register!

Practical Information

Final Exam:

Paper Presentations:

Online, early December?

Evaluation: (tentative)

3 graded TD's: 20% (10% each, best 2)
3 Quizzes: 20% (10% each , best 2)
Final presentation: 60%

Graded TD's: 2nd, 4th and 6th

We will accept submissions up to 1 week after the TD.

Graded quizzes: based on the material of 1st, 3rd and 5th lectures.

15 minutes At the beginning of 2nd, 4th and 6th TDs.

Register!

Practical Information

Final presentation:

- **Research paper** presentations:
  - Main goal: **read and understand a recent paper**.
  - OK to work in a team, **but** at most 2 people
  - Every topic: at most 2 teams (**pick early!**)

Presentation should highlight your **detailed understanding**.

We will ask questions about both the paper and possibly related course content.

Introduction to the course

1. What is Geometry processing and Geometric deep learning?
2. Why is it useful?
3. What are its main challenges?
4. What will we learn?





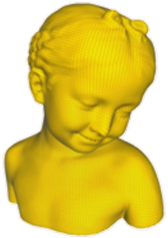
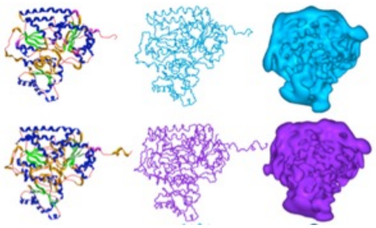

Evolution of Multimedia

New types of data are constantly being acquired, digitized and manipulated.

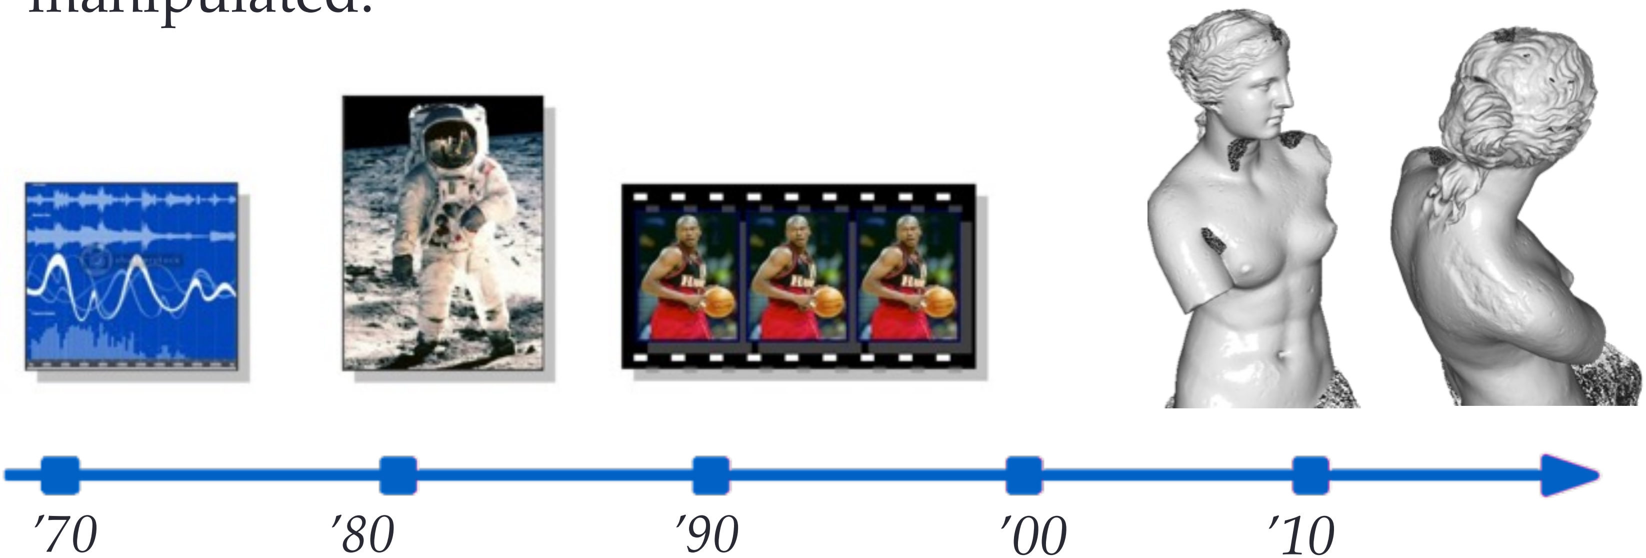

Growing demand for acquisition, processing and analysis of 3D geometric data.

Motivation: acquisition of 3D data

The first efforts in 3D acquisition focused on capturing *individual objects*.

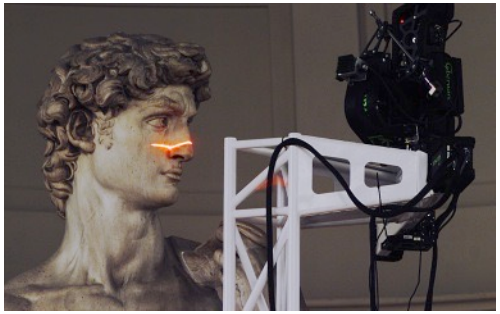

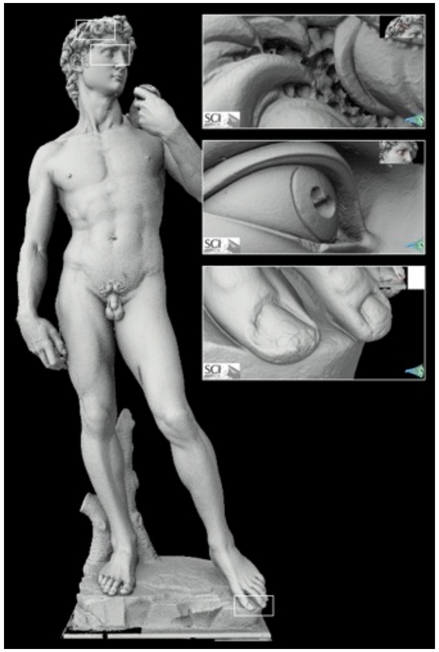

Digital Michelangelo Project (1998-1999): approximate cost 2M USD.

Acquisition of 3D data

Scans of Hannover (ca. 2007): approximate cost 200,000 USD.

Acquisition of 3D data

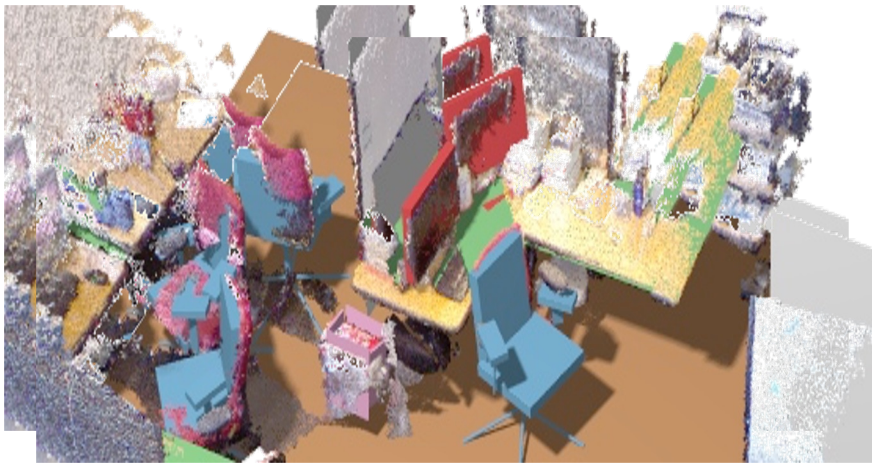

2010 Microsoft Kinect (100$)
3D scanner – gadget for Xbox

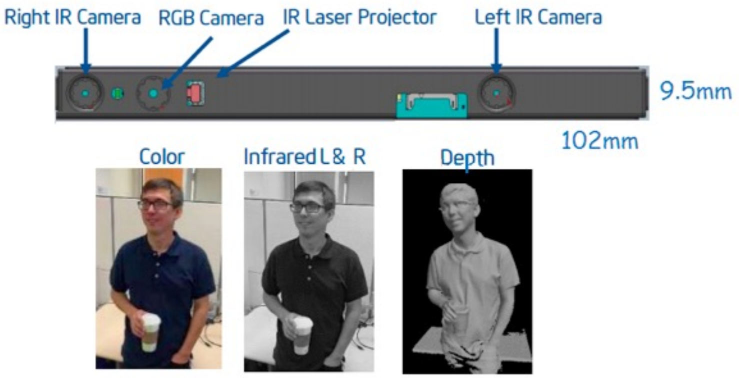

2014 Intel RealSense
integrated 3D **scanner**

Geometry is not isolated

Large **collections** of 3D shapes are becoming available.

ABC: A Big CAD Model Dataset For Geometric Deep Learning, CVPR 2019

TURBO SQUID

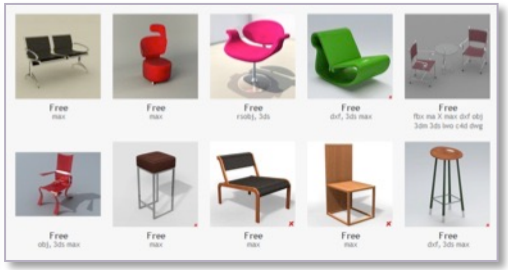

Trimble.

3D warehouse

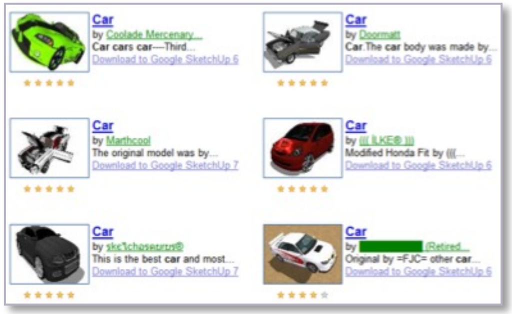

Millions of 3D shapes

Why Geometric Modeling **Now?**

3D Scanning capabilities in recent devices

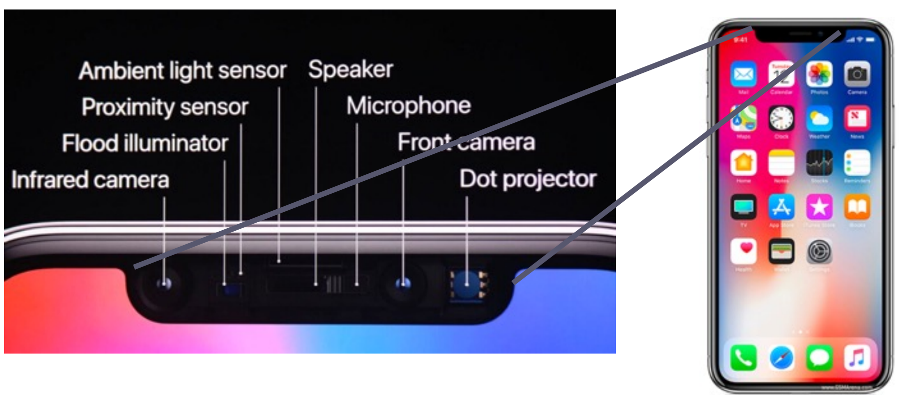

Apple iPhone X, 2017

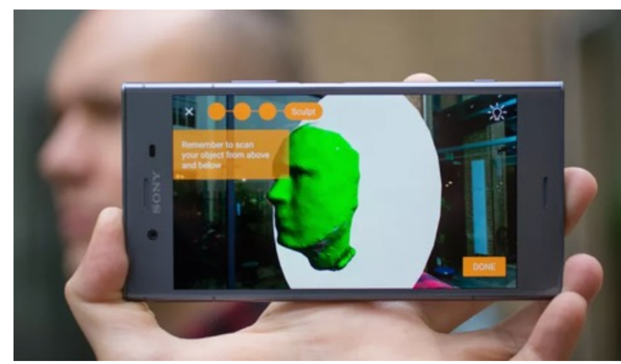

Sony Xperia XZ1, 2017

Why Geometric Modeling **Now?**

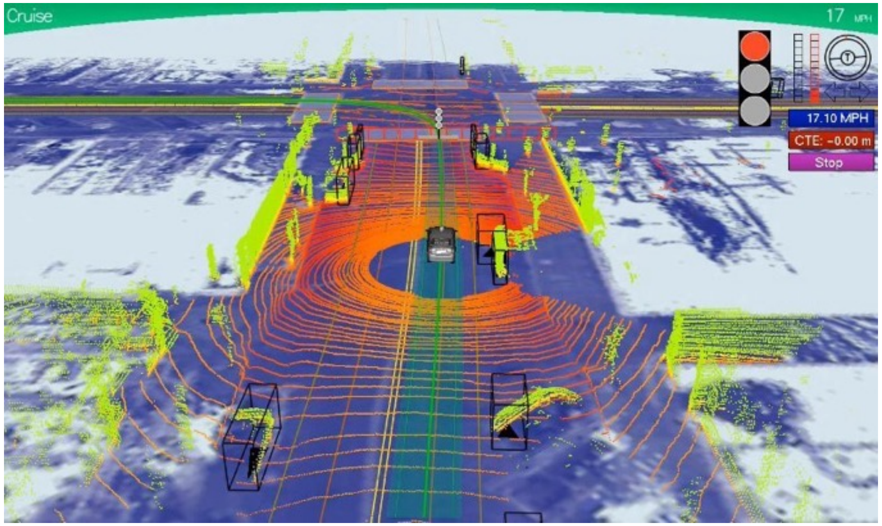

LIDAR sensors on self-driving cars

3D shape analysis tasks

Fundamental tasks: *classification, segmentation, correspondence, reconstruction, alignment,* etc. on 3D data.

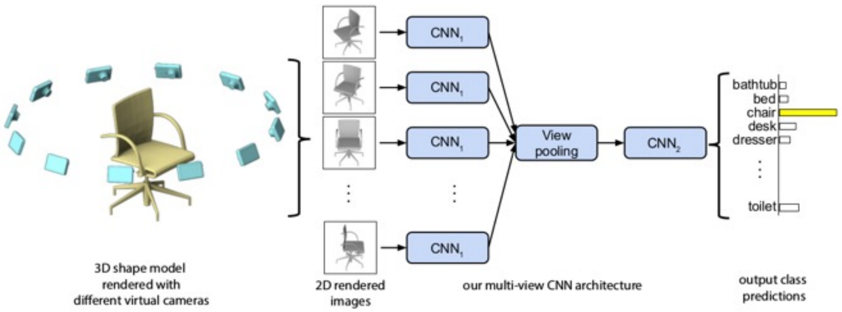

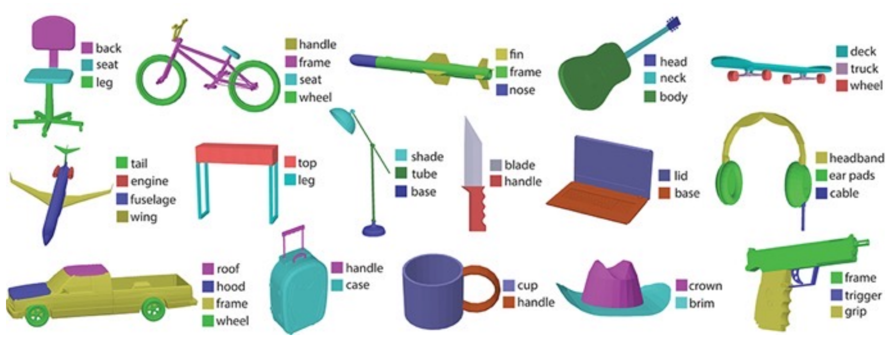

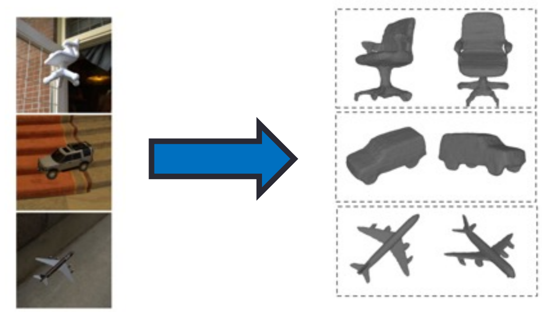

Learning on images

Standard 2D Computer Vision Deep Learning "ingredients":

1. A lot (!) of labeled training data
2. Convolutional Neural Networks (CNNs)

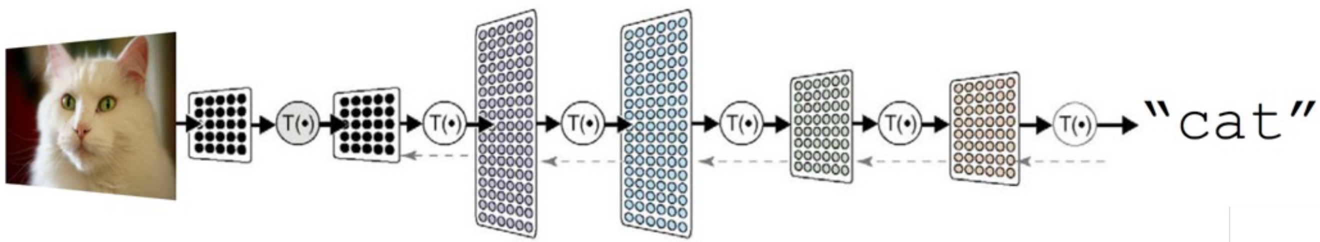

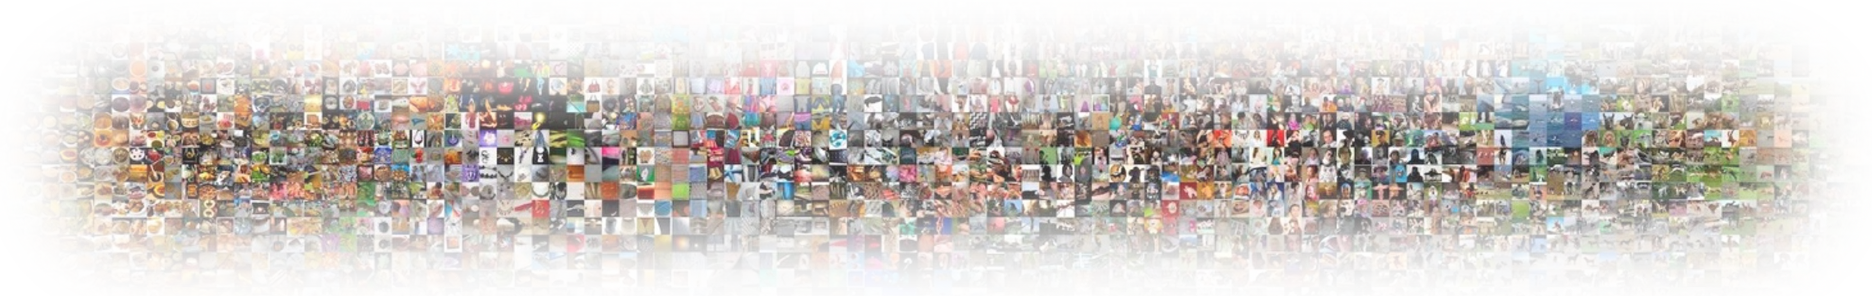

Deep Learning for 3D shapes

- Conv-Nets in 2D

Fundamental operation: convolution

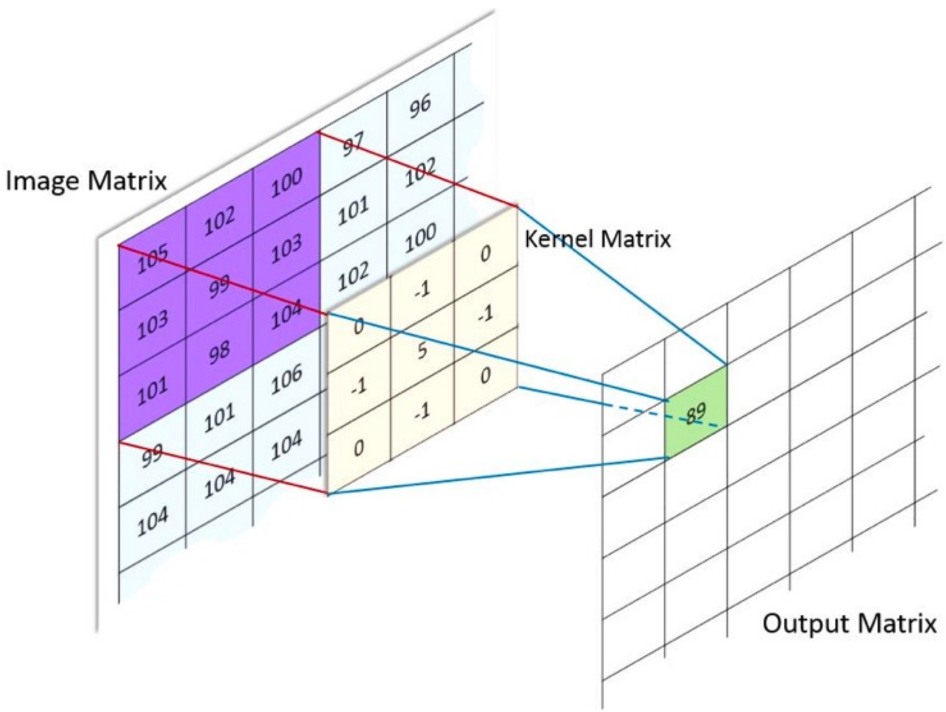

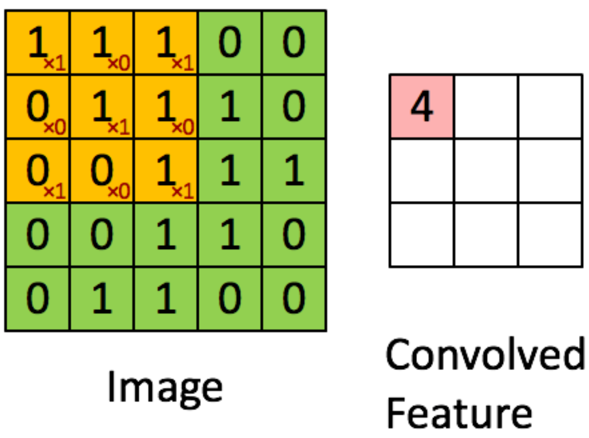

Deep Learning for 3D shapes

- Main Challenge

3D shapes (typically) do not have a canonical (grid-like) representation!

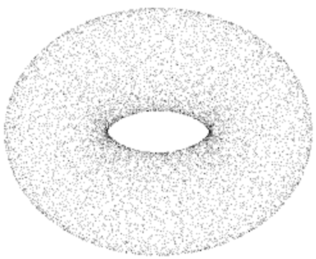

3D point cloud: an *unorganized collection of 3D coordinates*

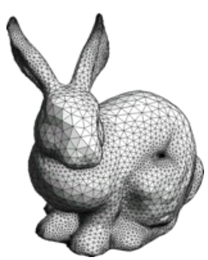

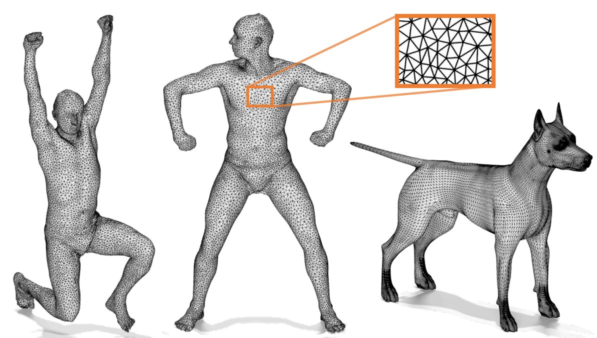

3D mesh: a *collection of points* and *triangles* connecting them.

Approaches for 3D Deep-Learning

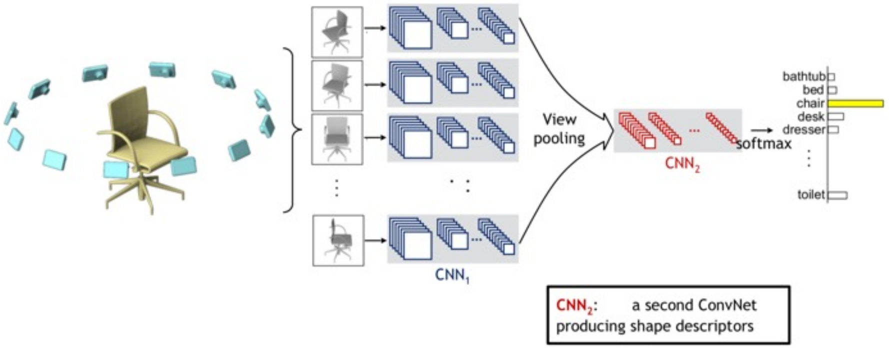

Multi-view based

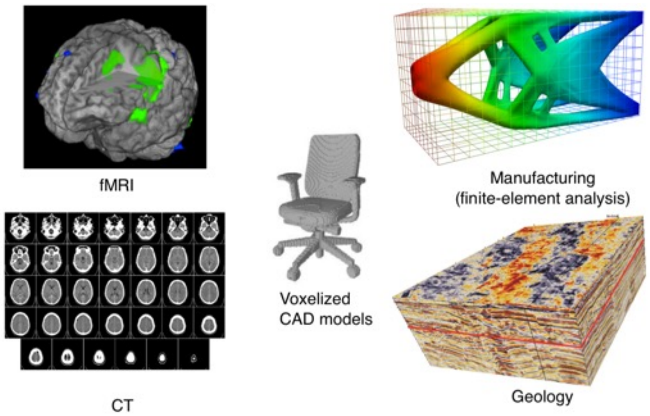

Volumetric

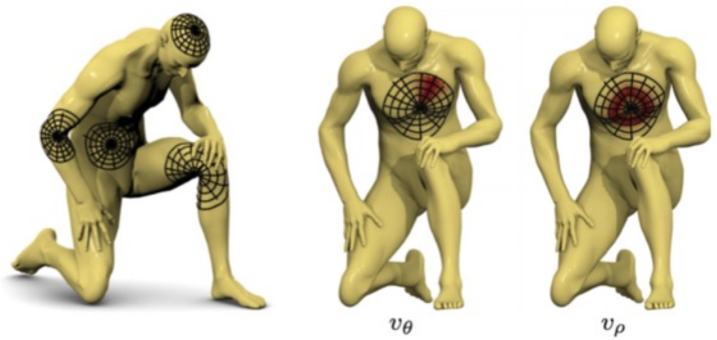

Intrinsic (surface-based)

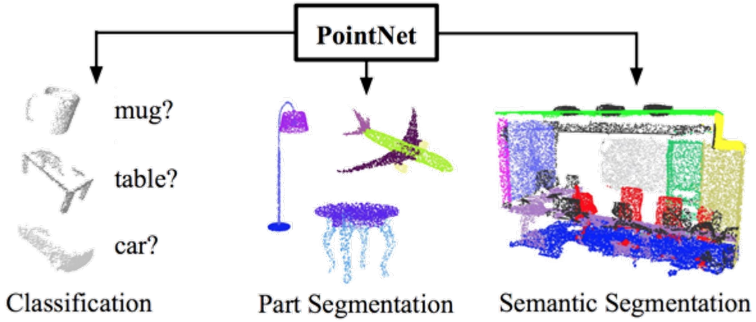

Point-based

3D Shape Analysis and Learning

Main questions:

- How to *represent* the 3D shapes to enable learning?
- How to design *robust* and *principled* data analysis approaches?

General insight:

- (Often) the more *mathematically founded* methods are the better they tend to perform.

What are we going to learn?

Lecture 1. Calculus on surfaces:

Functions, derivatives (gradients), integration, **Laplacian**, Spectral quantities, Diffusion, Descriptors.

meshes

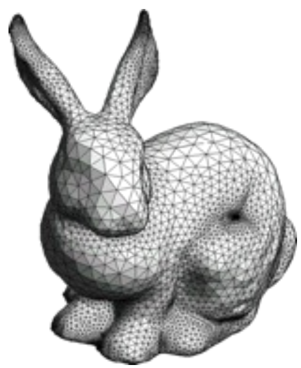

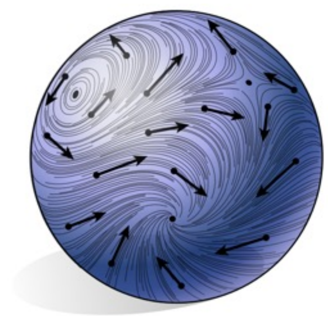

gradients

Diffusion

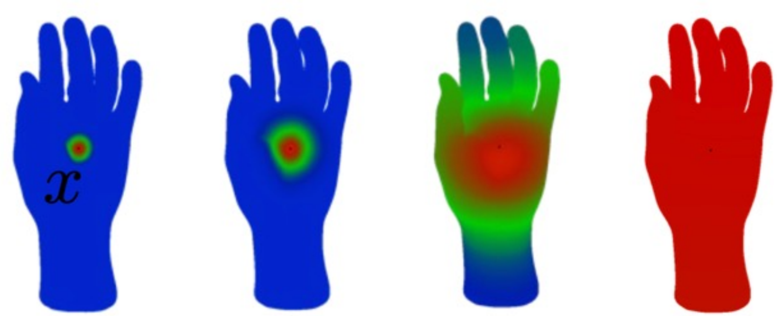

$t = 0.001$ $t = 0.02$ $t = 3$ $t = 40$

What are we going to learn?

Lecture 2. Optimization of geometric energies.
Surface parameterization. Mappings between surfaces, deformation. Basic surface topology, 3D learning via 2D.

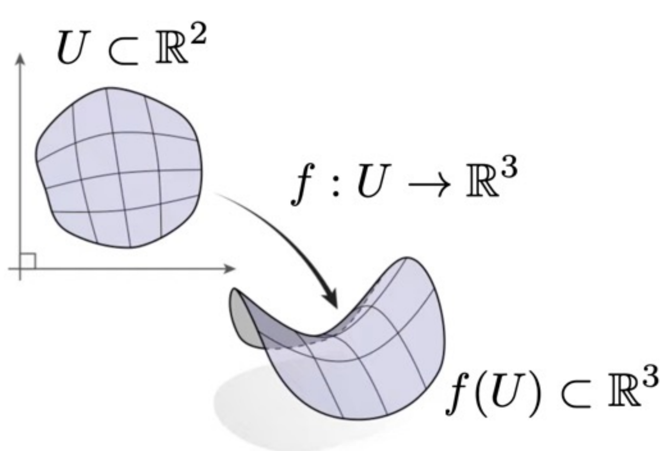

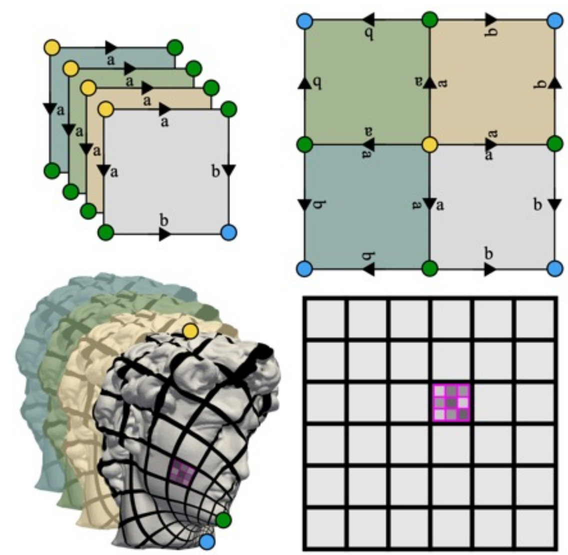

What are we going to learn?

Lecture 3. Analysis and machine learning on point clouds. Point-based architectures. Information propagation on point clouds. Learnable kernels. Normal estimation & denoising.

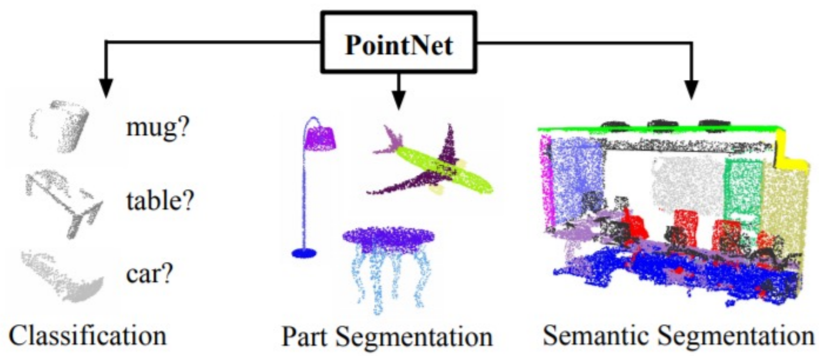

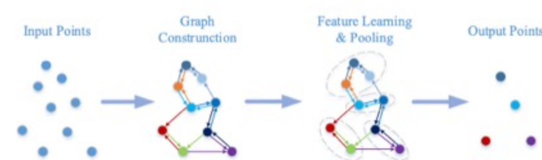

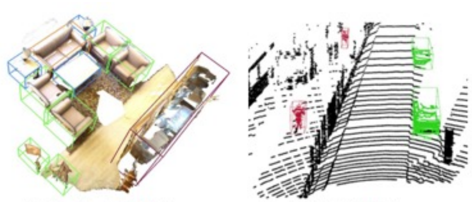

(a) ScanNetV2 [11] dataset (b) KITTI [14] dataset

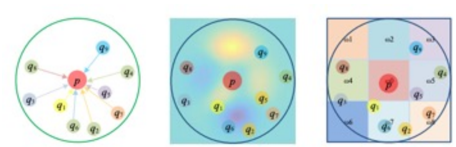

(a) 3D neighboring points (b) 3D continuous convolution (c) 3D discrete convolution

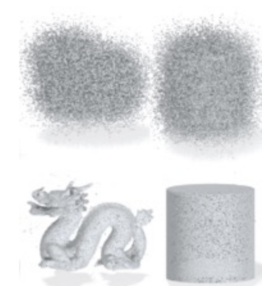

What are we going to learn?

Lecture 4. Deep learning on curved surfaces.
Extrinsic vs. intrinsic convolution, Geodesic CNNs and their variants. Effective diffusion-based learning methods.

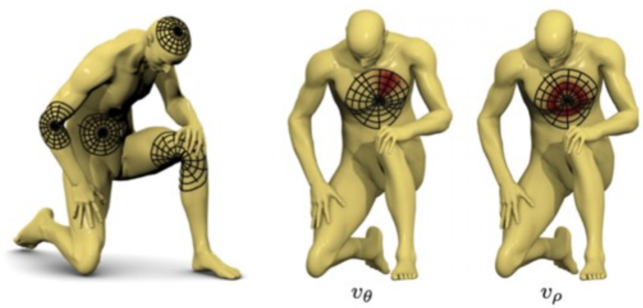

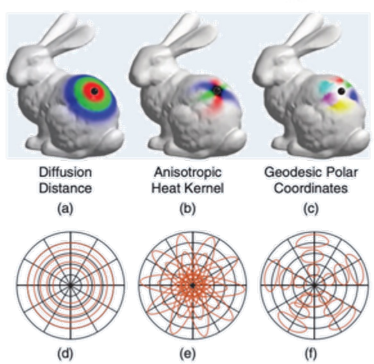

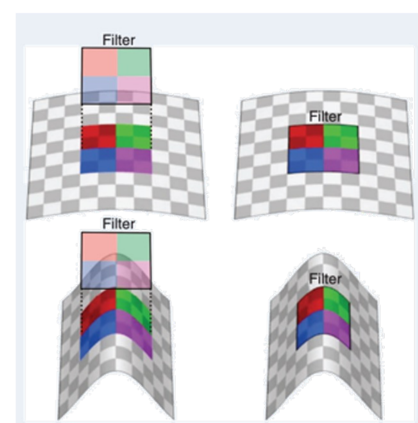

What are we going to learn?

Lecture 5. Neural field for surface representation.
Neural radiance field and neural fields regularization.
DeepSDF, Occupancy network

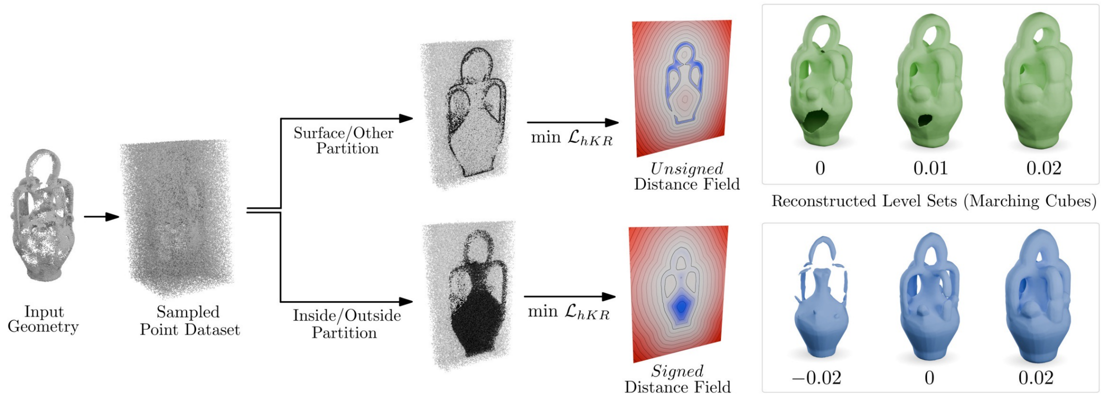

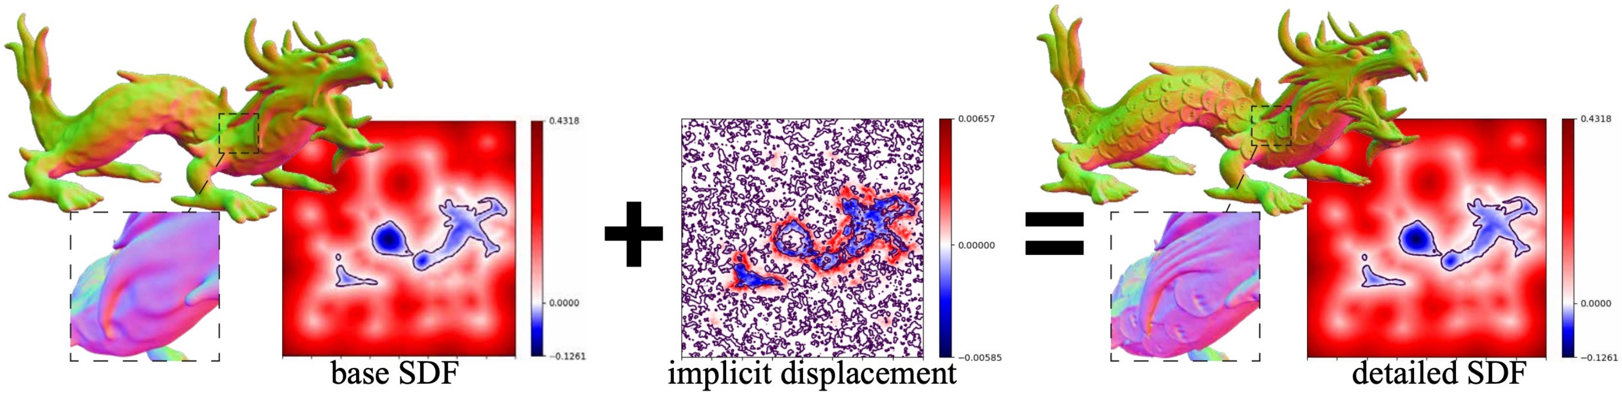

What are we going to learn?

Lecture 6. Generative modelling.
How to generate the surface structure? Geometric texture synthesis. Inpainting. Mesh generation. Differential meshing.

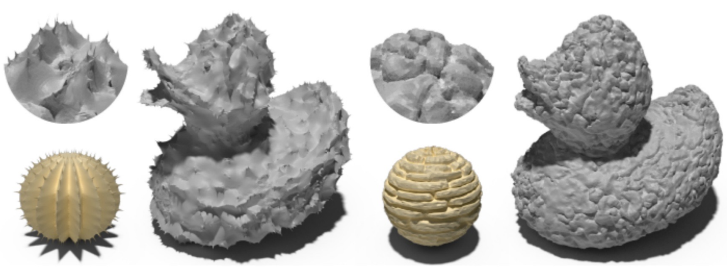

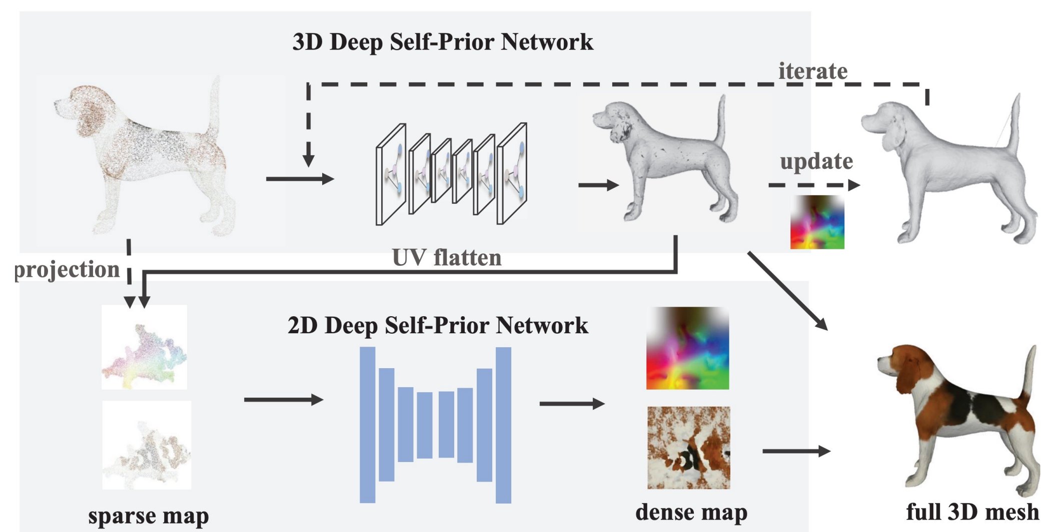

Introduction

Questions?