

Geometry Processing

7 Data-driven Approach

Ludwig-Maximilians-Universität München

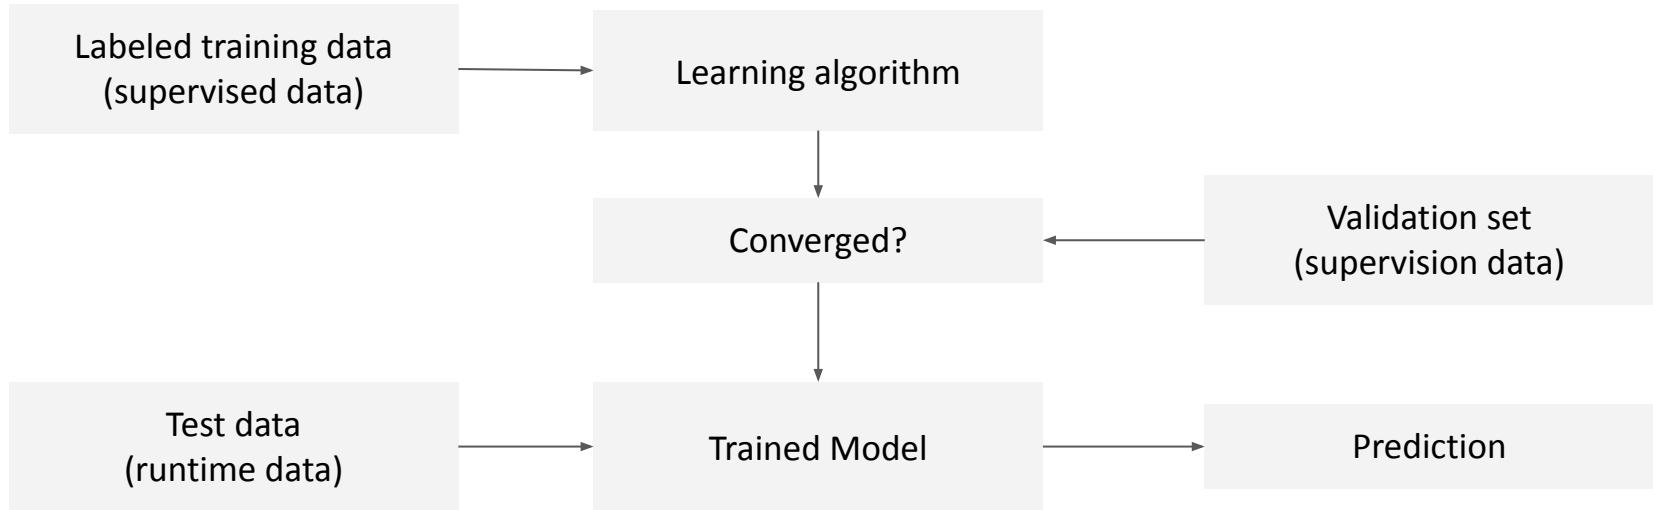
Session 7: Data-driven Approach

- Statistical Learning Schemes
- Representations for Learning in 3D
- Trends and Challenges with 3D Data
 - Dealing with Defects and Flaws Inputs
 - Ground Truth User Expectations
- Summary

Statistical Learning Schemes

- Supervised learning
- Unsupervised learning
- Semi-supervised learning
- Generative learning
- (Inverse) Deep Reinforcement learning
- Active Learning
- ...

Example: Supervised Approach



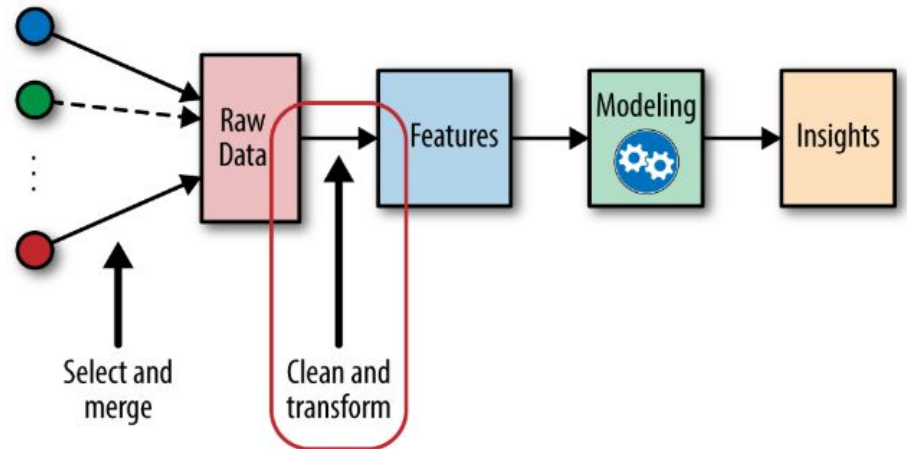
Traditional: Feature Engineering

Manually extract features

- Laplacian matrix: well understood intrinsic surface representation
 - Uniform/Cotan/MVC weights

Feature Selection

- Average geodesic distance
- Gaussian curvature
- Conformal factor
- Shape contexts
- Shape diameter function
- ...



Hierarchical: Deep Neural Networks

Input: a sort of representation of 3D shapes


Output: whatever we want, such as vertex informations (normals, uvs, etc.)

Procedure in three major steps:

- Design network architecture: Determines *hypothesis space*
- Design loss function: Determines *loss (distance) function* between *hypothesis*
- Design optimization strategy: Determines the path of searching hypothesis

Issue: Overfitting (Memorization) v.s. Generalization

- Different perspective: Overfit and memorize all data
- See [understanding generalization in deep learning](#)



Example: Convolutional NN (CNN) on Grid-based Images

AlexNet: CNN to the mainstream


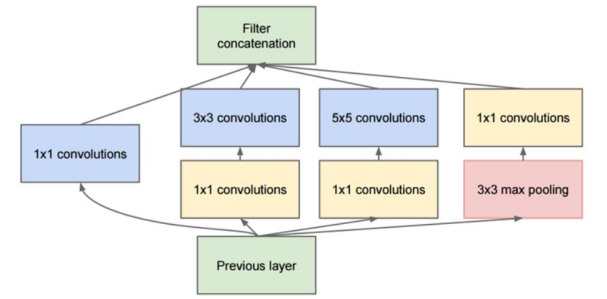

VGG: Deep and simple

Inception: 1x1 feature pooling

ResNet: Skip connection and residual block (powerful, for network performance)

CapsuleNet: Dynamic routing to dealing with rotation invariance

...

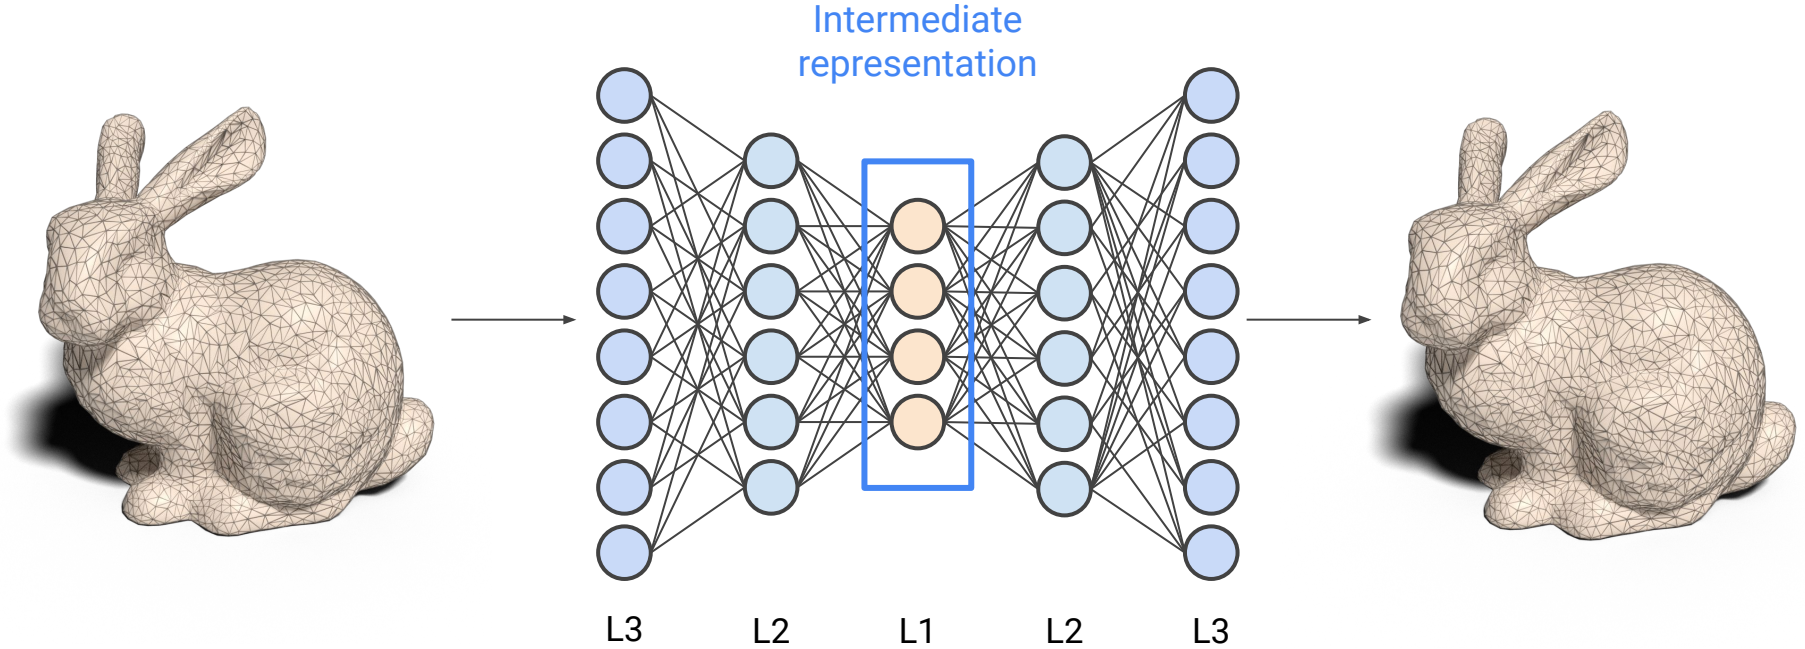


Session 7: Data-driven Approach

- Statistical Learning Schemes
- Representations for Learning in 3D
- Trends and Challenges with 3D Data
 - Dealing with Defects and Flaws Inputs
 - Ground Truth User Expectations
- Summary


Representation Learning

Train an autoencoder that learns an intermediate representation




Representation Learning (2)

Cutting output layers to reduce the dimensionality of outputs



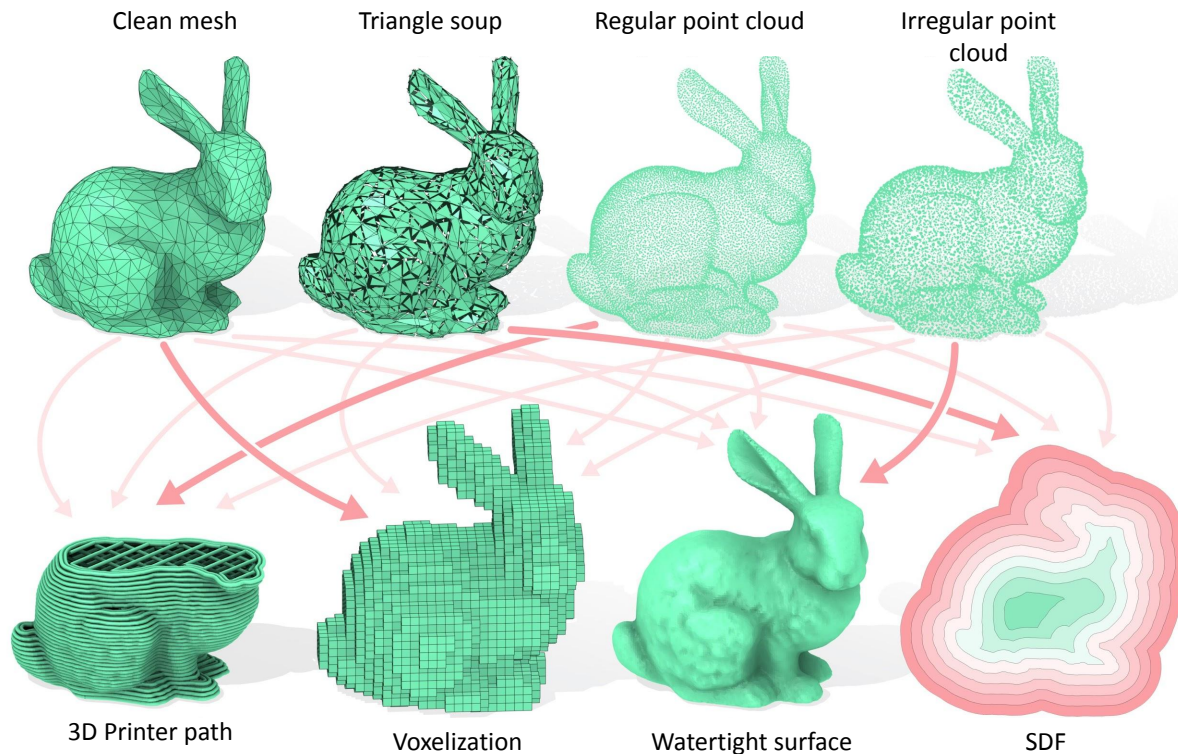
Representation Learning (3)

Changing output layers to produce different representations



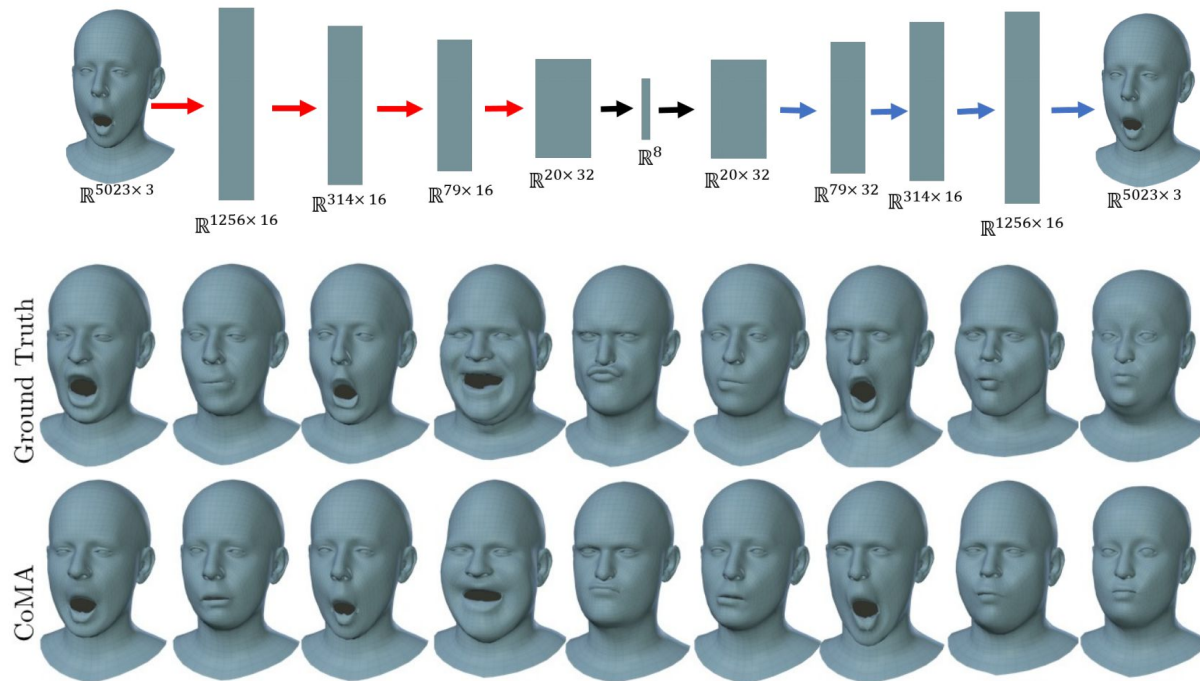
Representation Learning (4)

Mixing input representations and producing other representations



3D Representations

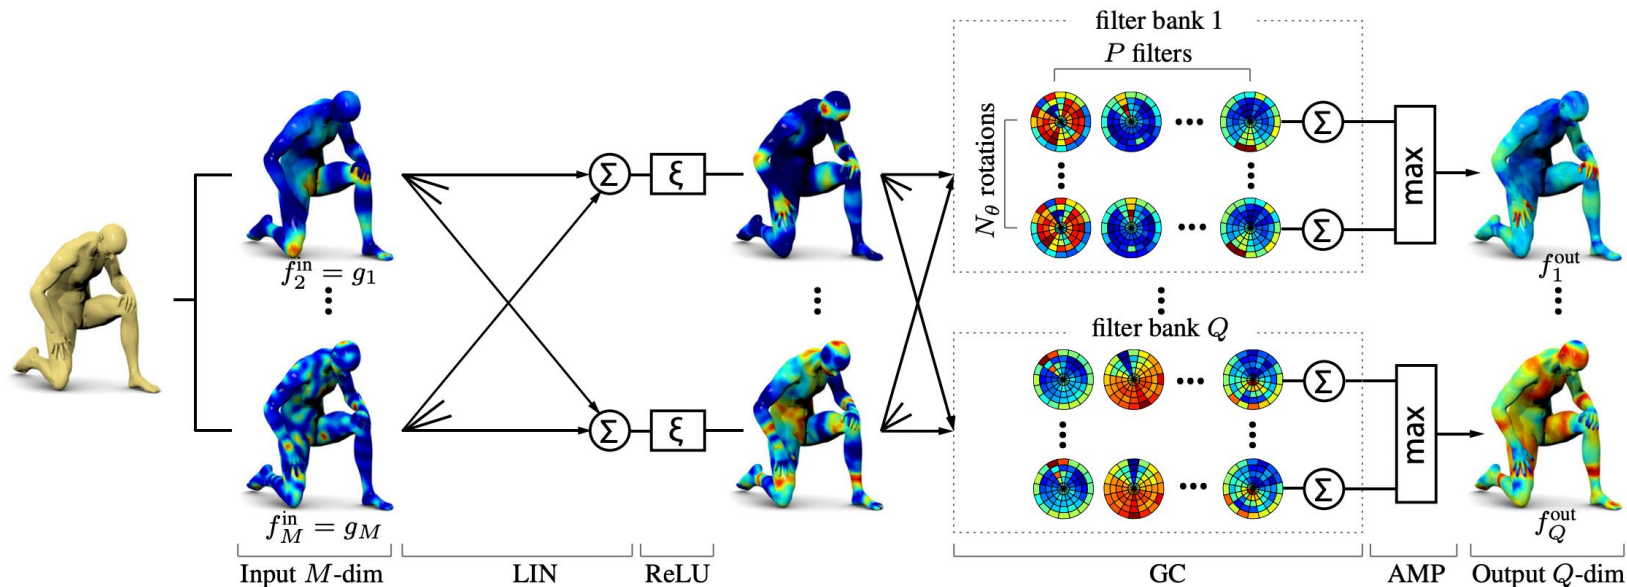
- Surface-based (*major focus in this course*)
 - e.g. mesh autoencoder for deformation



[Ranjan et al. 2018]

3D Representations

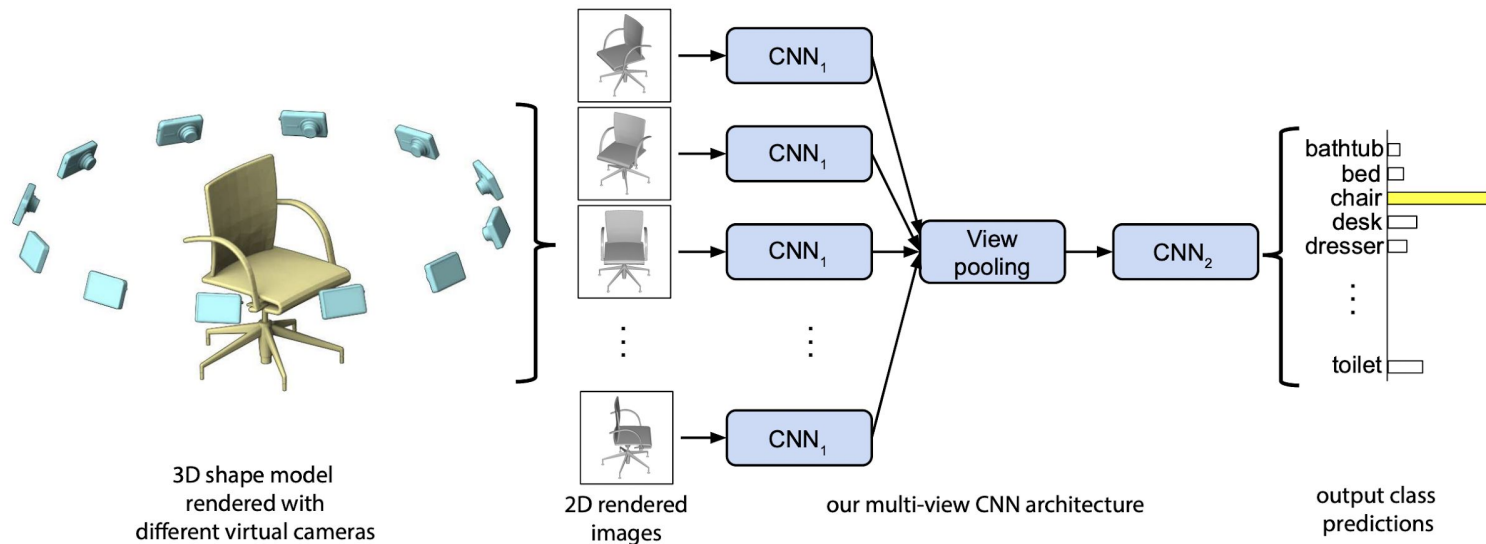
- Surface-based (*major focus in this course*)
 - e.g. mesh autoencoder for deformation
 - How to encode manifold representation and feed into a NN?



[Masci et al. 2015]

3D Representations



- Image-based



[Su et al. 2015]

3D Representations

- Point-based



[Wang et al. 2019]

3D Representations: Pros and Cons

- Image-based
 - Pros: good performance, easy to transfer knowledge
 - Cons: rendering is slow and memory-heavy, not geometric
- Surface-based
 - Pros: parameterize+image networks(intrinsic representation)
 - Cons: suffers from parameterization artefacts (local vs global distortion), require good quality mesh
- Point-based
 - Pros: native processing, directly applicable to scans
 - Cons: memory hungry, missing connectivity
- Volumetric and Implicit (SDF or occupancy): different stories

Session 7: Data-driven Approach

- Statistical Learning Schemes
- Representations for Learning in 3D
- Trends and Challenges with 3D Data
 - Dealing with Defects and Flaws Inputs
 - Ground Truth User Expectations
- Summary

Challenges When Deal with Representations

1. Main Question: How to feed 3D data into a neural network?
2. Neighborhood information (one-ring in previous sessions, maybe more?)
 - Who are the neighbouring elements
 - How are the elements ordered
 - ...
3. Extrinsic v.s. intrinsic representation (Differential form on surface embedding, or Euclidean embedding)

Challenges: Dealing with “Bad” Inputs

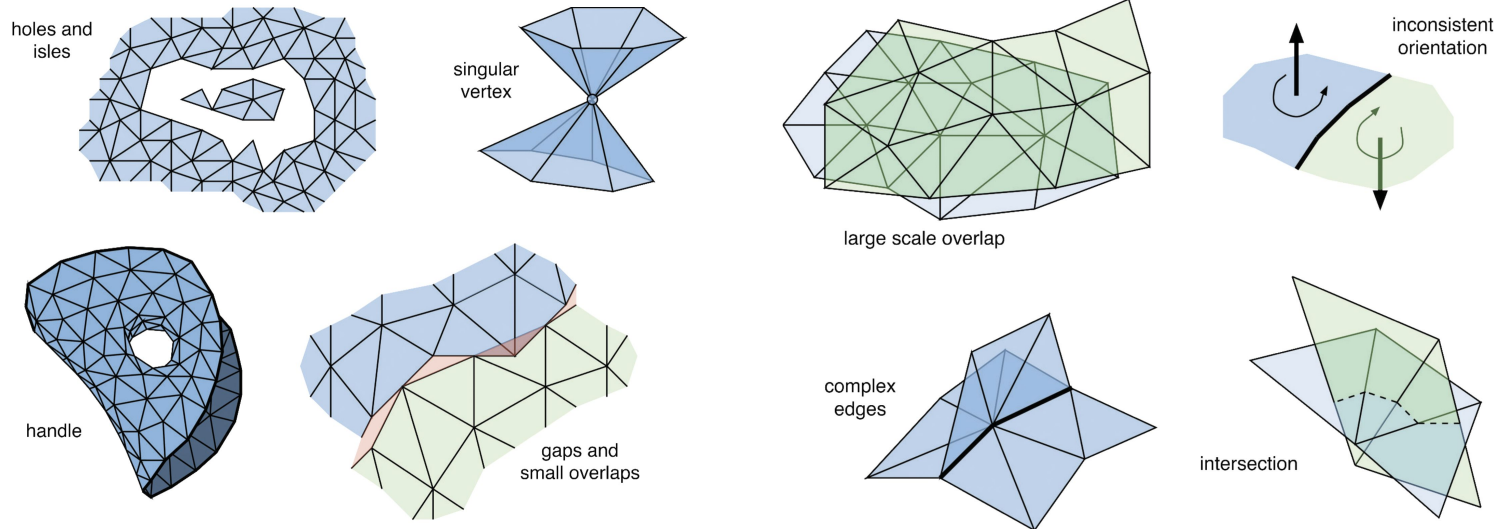
General goals are clear but very tricky to find an answer:

- Prevent input with artifacts
- Prevent producing outputs that contains artifacts

Flaw Inputs with Artifacts

Artifacts does not well fit traditional theory

- Laplacian equation does not work with non-manifolds
- Quadrics are not invertible in mesh simplification
- ...



[Botsch et al. 2006]

Upstream and Downstreams in The Processing Pipeline

Upstream producer determines characteristics and defects of outputs

The origin of defects in mesh

- Nature: (physical) real-world data, e.g. statuary (noise, holes, chamfered feature, topological noise)
- Approach: algorithm itself does not guarantee or implementation specific
- ...

Downstream consumer determines requirements on their inputs

- Visualizations: rendering v.s. printing
- Modeling: surface properties and further animations
- ...

Challenges: Repairing Artifacts

The process of dealing with bad inputs is often *tedious* and had to be done manually

Traditional wisdoms

- *Artifacts repairing is expected to be eliminated if all algorithms does not produce bad inputs*
- Unfortunately, algorithms does not guarantee to produce high quality mesh

Example:

- Noisy point cloud \Rightarrow Denoising and reconstruction
- Mesh with holes \Rightarrow Filling holes
- ...

Neural networks (may) intrinsically removes the flaws from inputs

Does artifacts really important for data-driven processing pipeline?

Challenges: Data Augmentation

- Infinite inputs in image-based representations
 - Render images from different scene, camera, illumination settings
- Transformed (deformed) meshes as inputs
 - Is it a chicken egg problem? NN learns the algorithm instead of ground truth

Q: What is ground truth, where and how to obtain it?

Challenges: Ground Truth User Expectations

User expectations are application dependent: Where to obtain ground truth labels?

- What is the target user for the models? Low-fidelity Gaming? Filming? Industrial design?
- What exactly contributes to "artifacts"?
- When do "people" (regular users or experts) satisfy with the model for "further processing" or "final use"?
- How to properly evaluate user expectations? e.g. equal loudness contour and head-related transfer function for audio measurement and evaluation
- ...

3D Datasets

ABC Dataset [Koch et al. 2019] : A collection of one million Computer-Aided Design (CAD) models for research of geometric deep learning methods and applications <https://deep-geometry.github.io/abc-dataset/>



More:

<https://pytorch-geometric.readthedocs.io/en/latest/modules/datasets.html>

Session 7: Data-driven Approach

- Statistical Learning Schemes
- Representations for Learning in 3D
- Trends and Challenges with 3D Data
 - Dealing with Defects and Flaws Inputs
 - Ground Truth User Expectations
- Summary

Summary

- Selecting and learning 3D representations remains open problem
- Evaluating inputs and user expectations remains open problem

TLDR: Large and rich research opportunities!

Further Reading Suggestions

[**Su et al. 2015**] Su, Hang, et al. "Multi-view convolutional neural networks for 3d shape recognition." Proceedings of the IEEE international conference on computer vision. 2015.

[**Maturana et al. 2015**] Maturana, Daniel et al. "Voxnet: A 3d convolutional neural network for real-time object recognition." 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2015.

[**Wang et al. 2016**] Wang PS, et al. Mesh denoising via cascaded normal regression. ACM Trans. Graph.. 2016 Nov 11;35(6):232-1.

[**Ranjan et al. 2018**] Ranjan, Anurag, et al. "Generating 3D faces using convolutional mesh autoencoders." Proceedings of the European Conference on Computer Vision (ECCV). 2018.

[**Wang et al 2018**] Wang, Peng-Shuai, et al. "Adaptive o-cnn: a patch-based deep representation of 3d shapes." ACM Transactions on Graphics (TOG) 37.6 (2018): 1-11.

[**Mohri et al. 2018**] Mohri, Mehryar, Afshin Rostamizadeh, and Ameet Talwalkar. Foundations of machine learning. MIT press, 2018.

[**Koch et al. 2019**] Koch, Sebastian, et al. ABC Dataset A Big CAD Model Dataset For Geometric Deep Learning. CVPR 2019.
<https://deep-geometry.github.io/abc-dataset/>

[**Wang et al. 2019**] Wang, Yue, et al. "Dynamic graph cnn for learning on point clouds." Acm Transactions On Graphics (TOG) 38.5 (2019).

[**Hanocka et al. 2019**] Hanocka, Rana, et al. "MeshCNN: a network with an edge." ACM Transactions on Graphics (TOG) 38.4 (2019): 1-12.

... and many more :)

Open Positions

- Work as a tutor in [Computer Graphics 1](#)
 - Teaching is a further step of learning

- An *Einzelpraktikum* or a *Thesis* in this area
 - Feel free to contact me :)