Rasmus R. Paulsen

AI Driven Surface Analysis

Take your telephone or computer – and go here!

PollEv.com/rasmuspaulse538

Just skip the registration

What do we see in this photo?

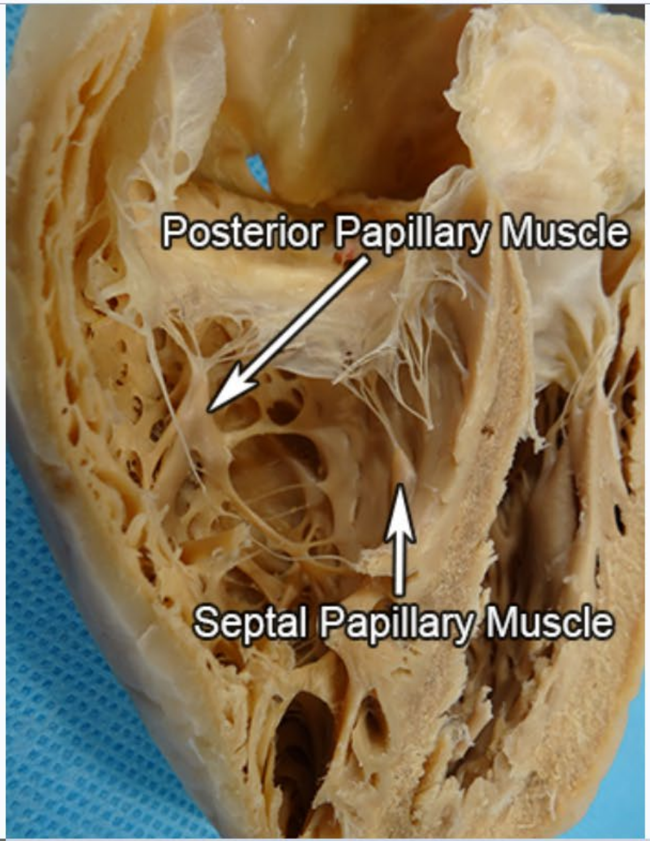

- Left atrial appendage and its muscles?
- The lower part of the stomach?
- The left ventricle with trabeculation?
- The aortic arch?
- The left kidney?

What do we see in this photo?

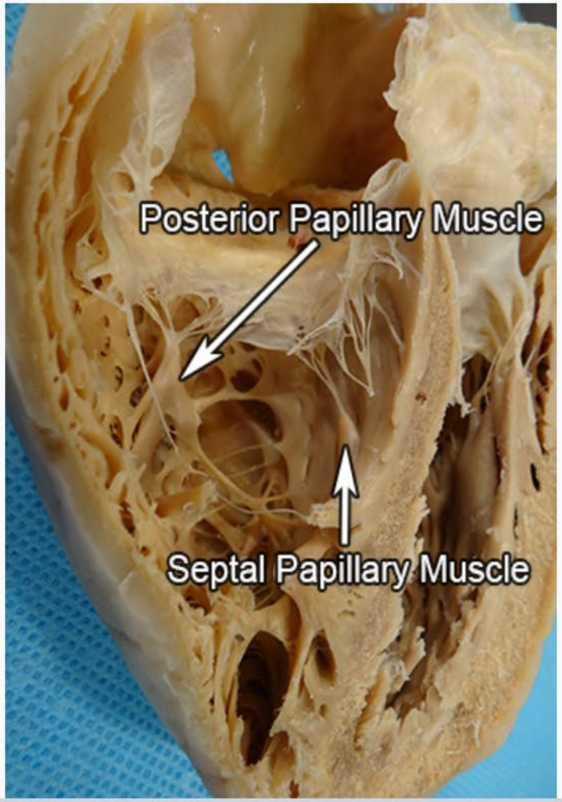

Left atrial appendage and its muscles?	20%
The lower part of the stomach?	6%
The left ventricle with trabeculation?	60%
The aortic arch?	9%
The left kidney?	6%

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

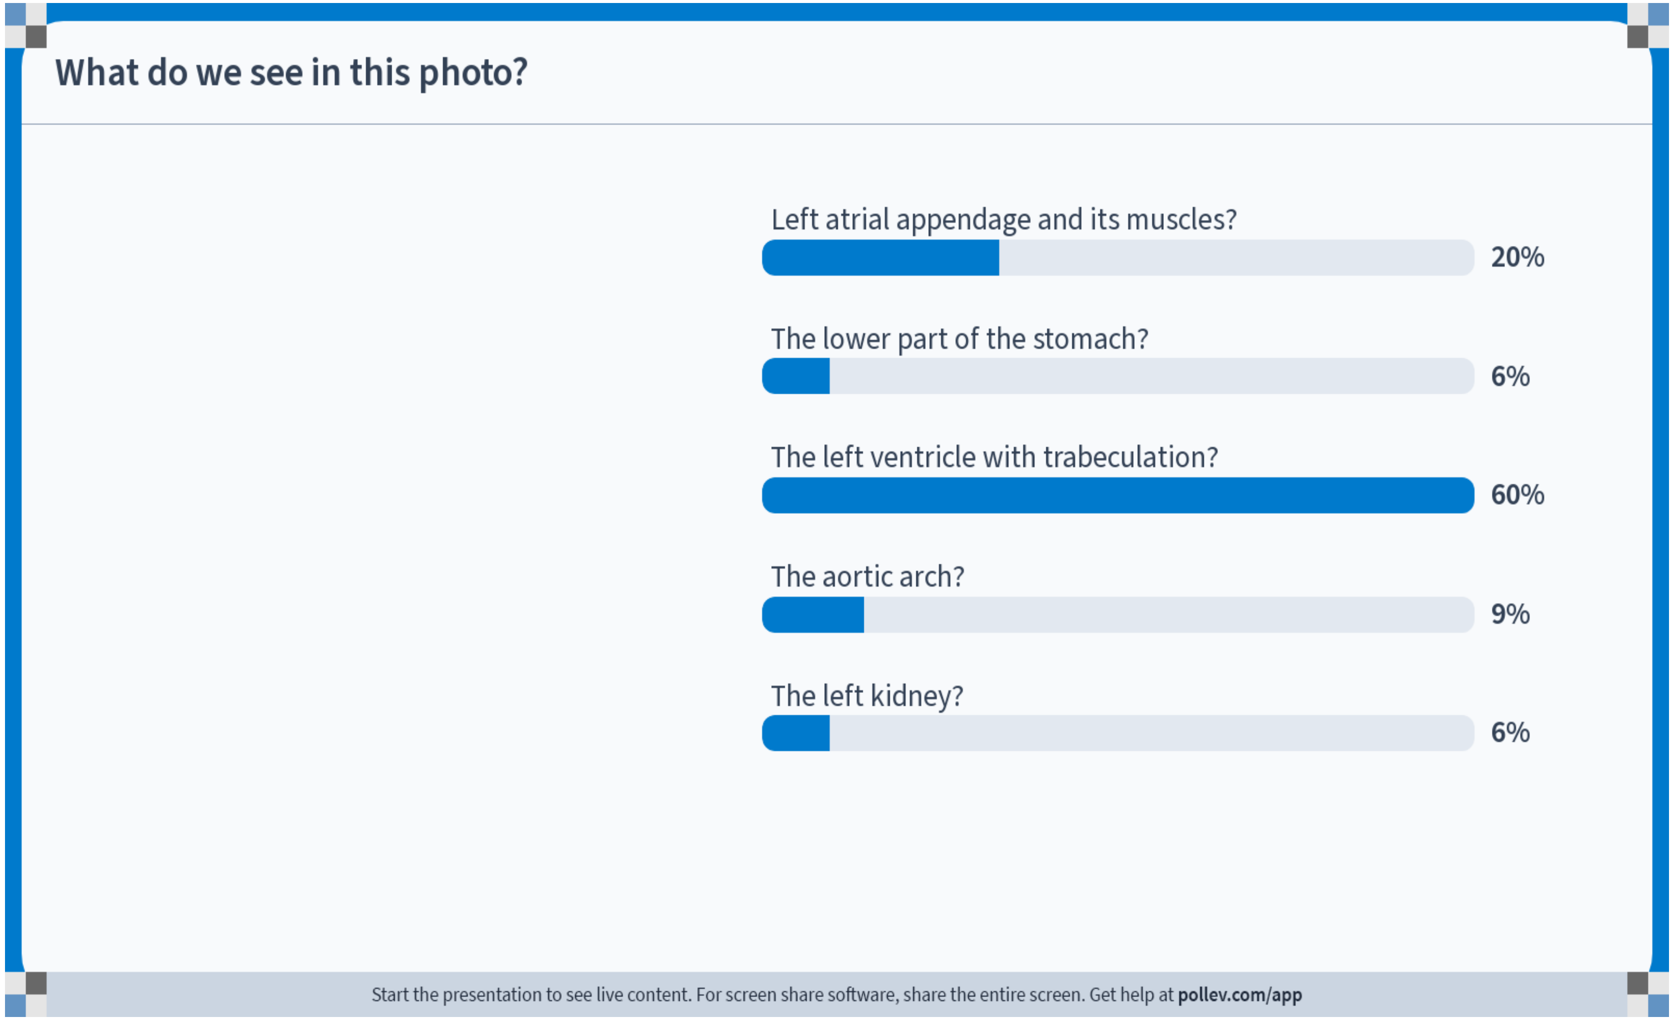
What do we see in this photo?
Left atrial appendage and its muscles?
20%
The lower part of the stomach?
6%
The left ventricle with trabeculation?
60%
The aortic arch?
9%
The left kidney?
6%
Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

Complex geometries – left ventricular blood pool

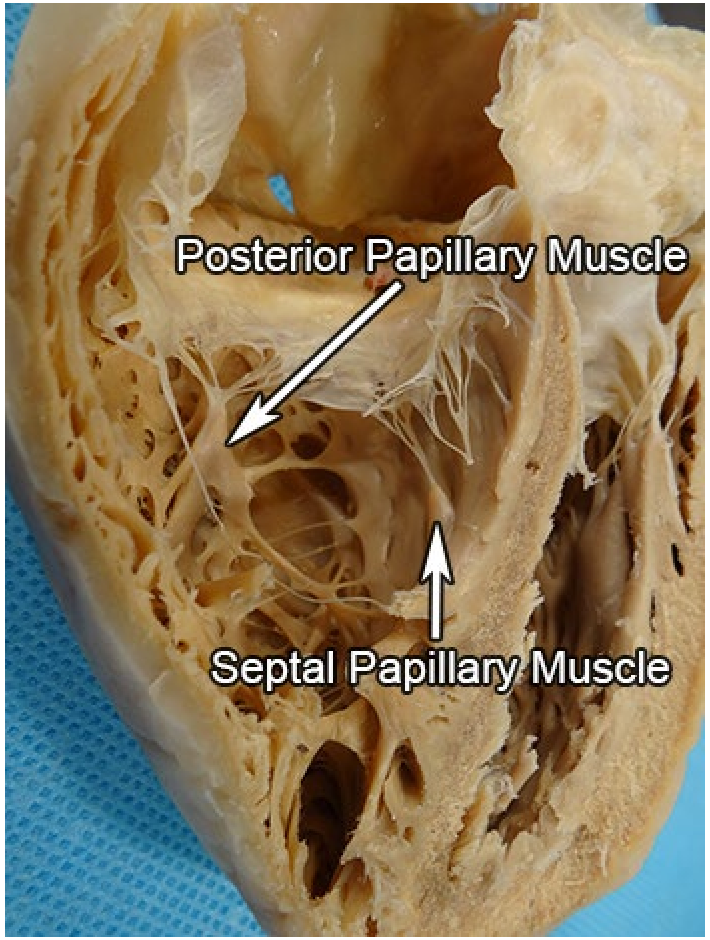

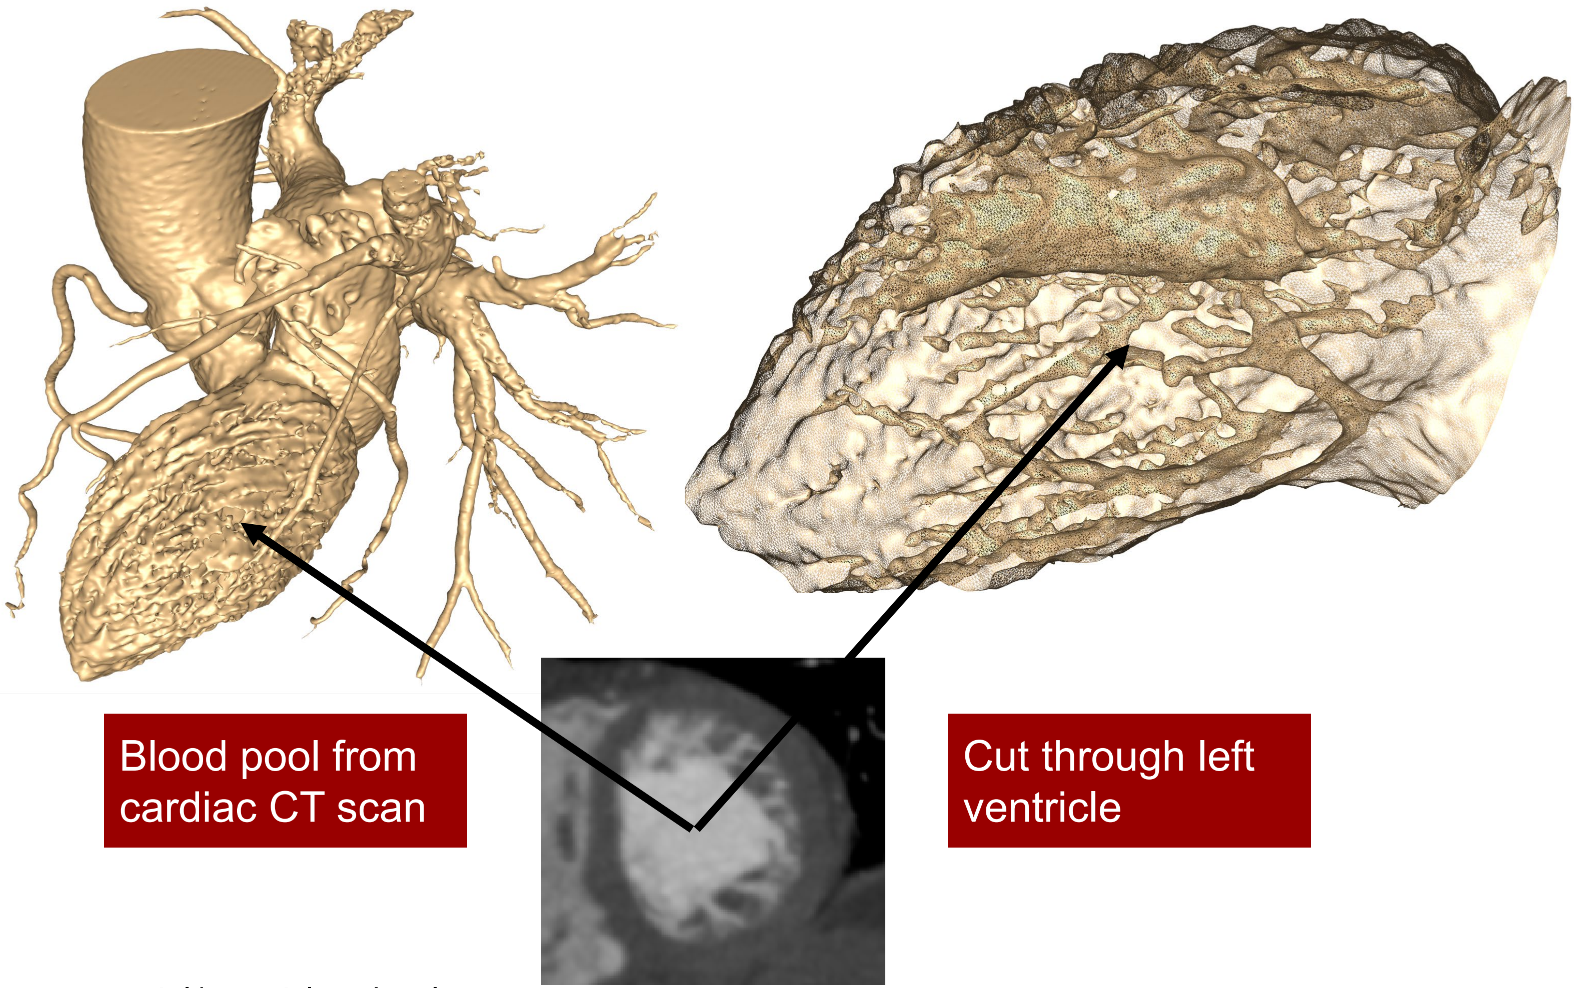

http://www.vhlab.umn.edu/atlas/comparative-anatomy-tutorial/ventricles.shtml

Statistics on complex shapes

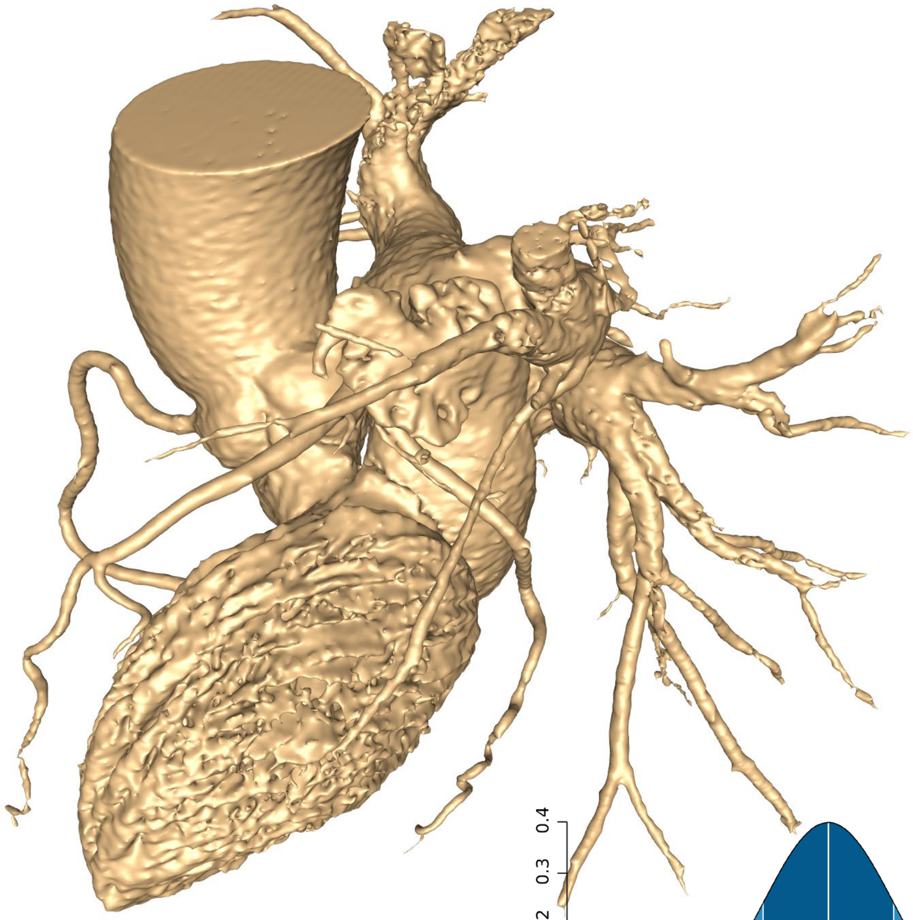

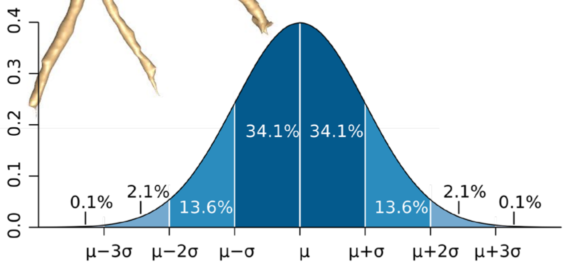

- Research questions:
  - How to parameterize complex geometries
  - How do we make meaningful statistical distributions of these shapes?
  - How do we test if a given patient is closer to one distribution or another?
  - How to compute risk scores using shapes?

Heart structures: Myocardium and left ventricle

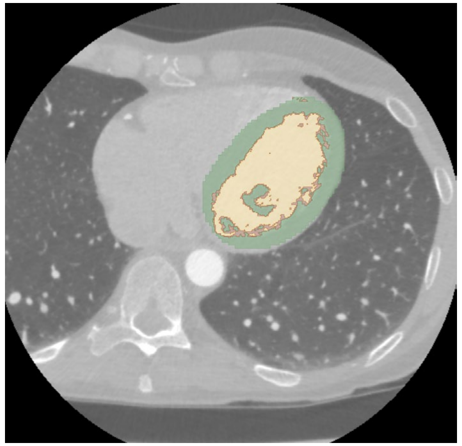

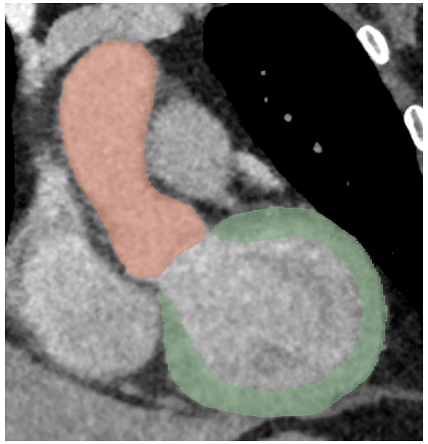

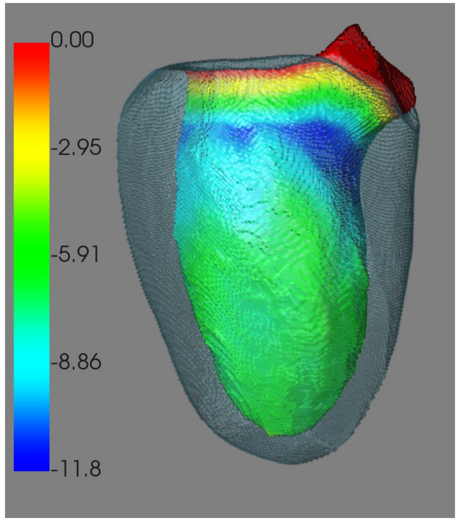

- The shape and appearance of the heart muscle (myocardium) is a known predictor for cardiac death
- Not trivial to define the borders between
  - Heart muscle
  - Left ventricular blood volume
  - Trabeculation

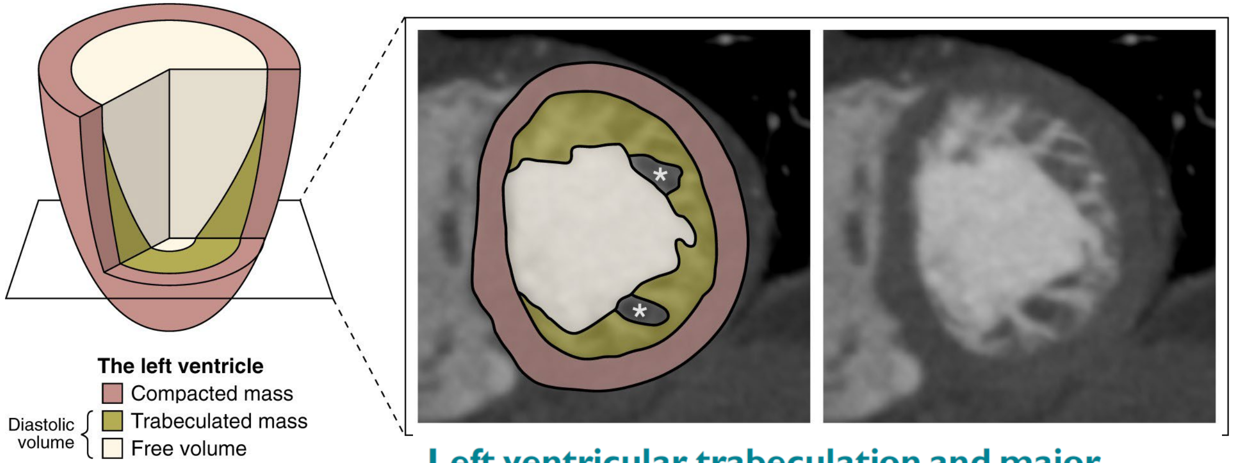

Left ventricular trabeculation and major adverse cardiovascular events: the Copenhagen General Population Study

Per E. Sigvardsen [1,2], Andreas Fuchs [1], Jørgen T. Kühl [1], Shoaib Afzal [2,3], Lars Køber [1,2], Børge G. Nordestgaard [2,3], and Klaus F. Kofoed [1,2,4]*

What is the most similar shape to this left atrial appendage (LAA)

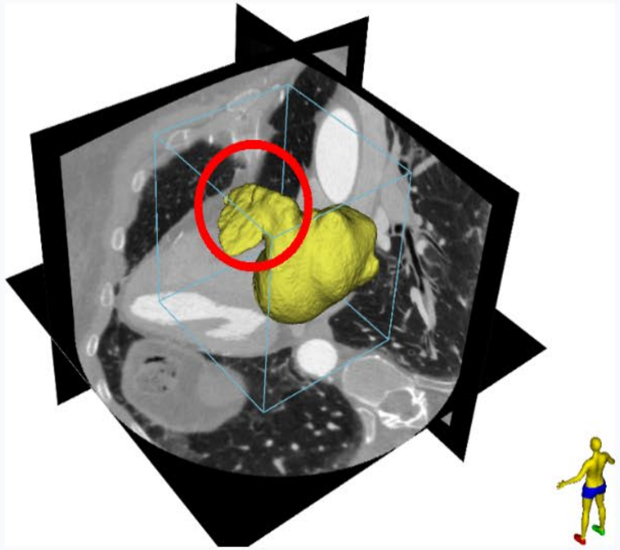

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

What is the most similar shape to this left atrial appendage (LAA)

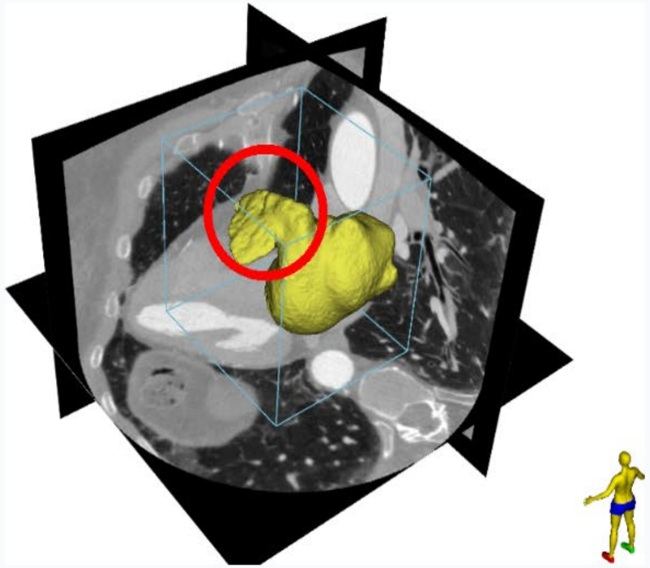

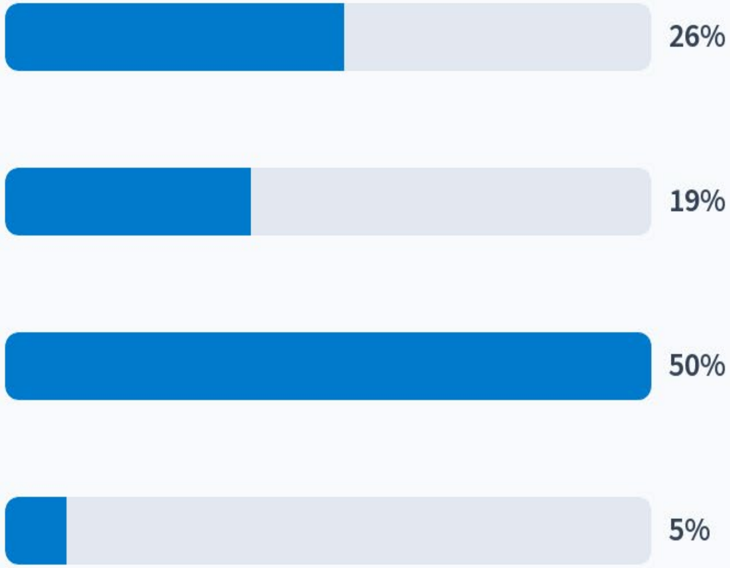

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

What is the most similar shape to this left atrial appendage (LAA)

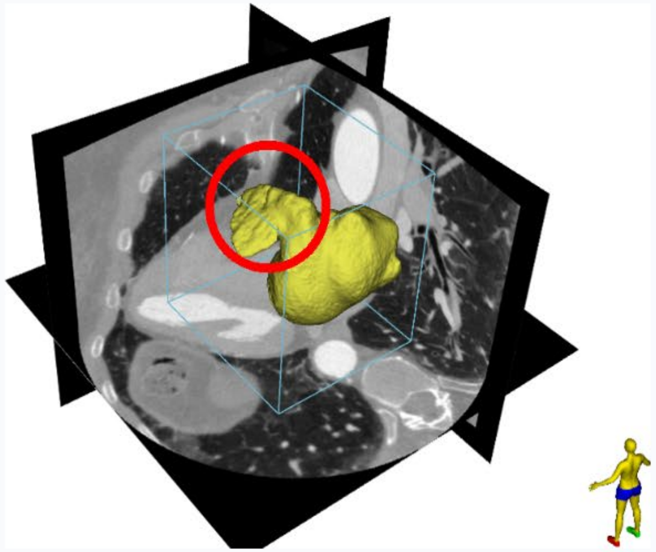

26%

19%

50%

5%

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

The shape of the left atrial appendage and stroke risk

Higher stroke risk

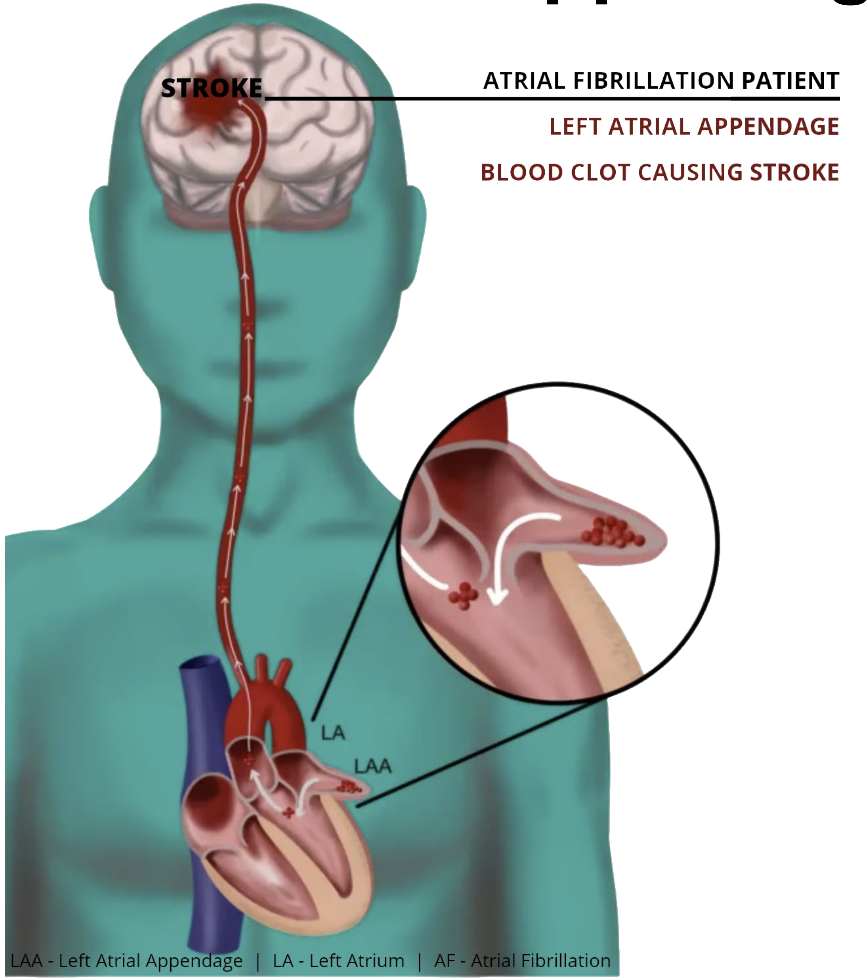

Lower stroke risk

More than 90% of thrombus accumulation occurs in the left atrial appendage (LAA) (for atrial fibrillation related strokes)

Stroke prevention

- It is possible to reduce the stroke risk
  – medicine (anticoagulants) or surgery (left atrial appendage closure)

- **Is it possible to identify patients at risk?**
- **Is it possible to optimise the surgical intervention?**

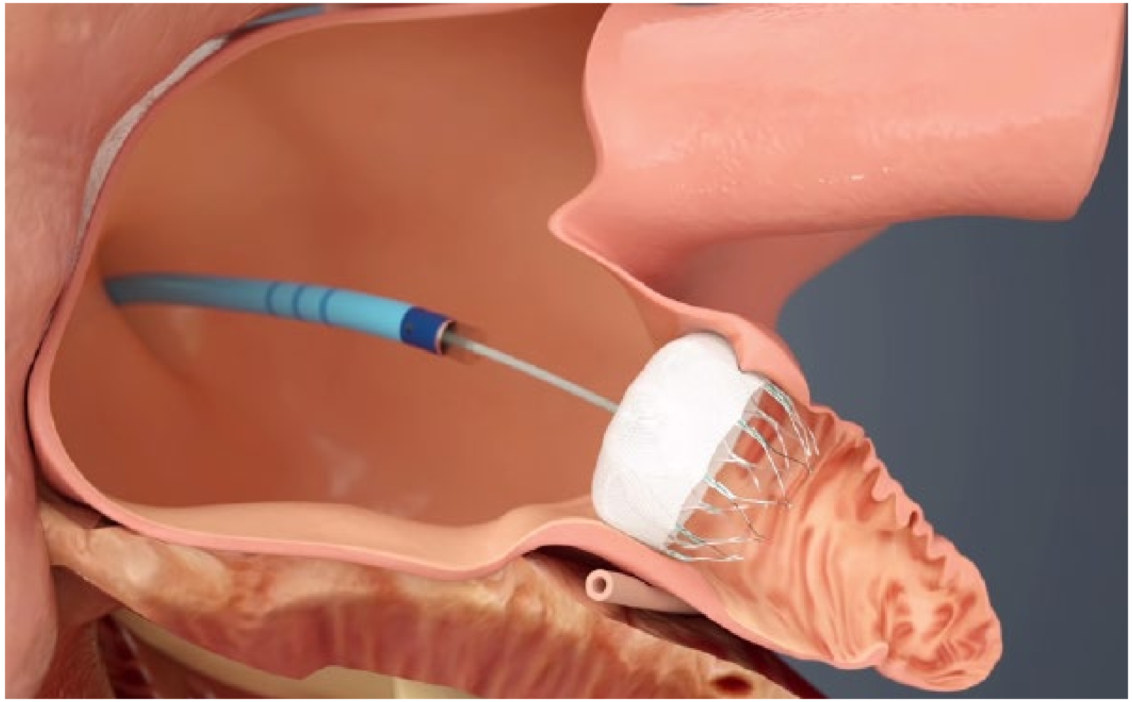

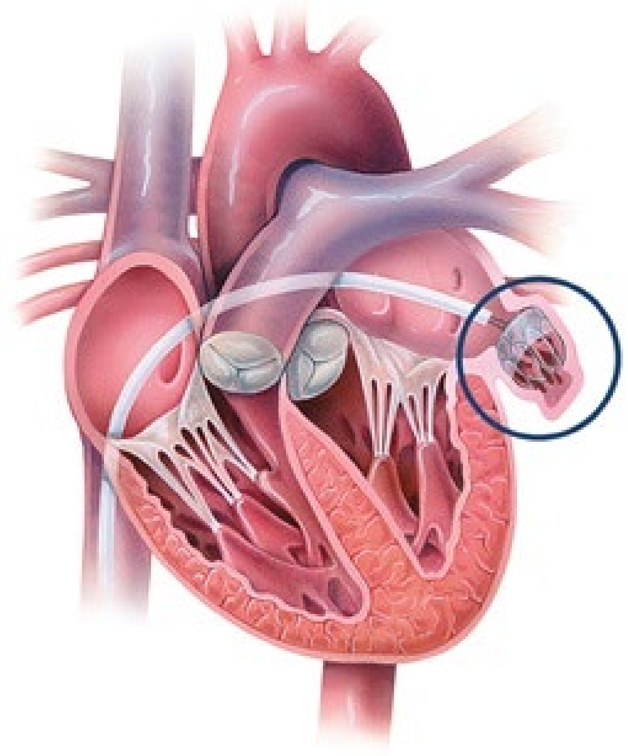

Statistics on complex biological shapes

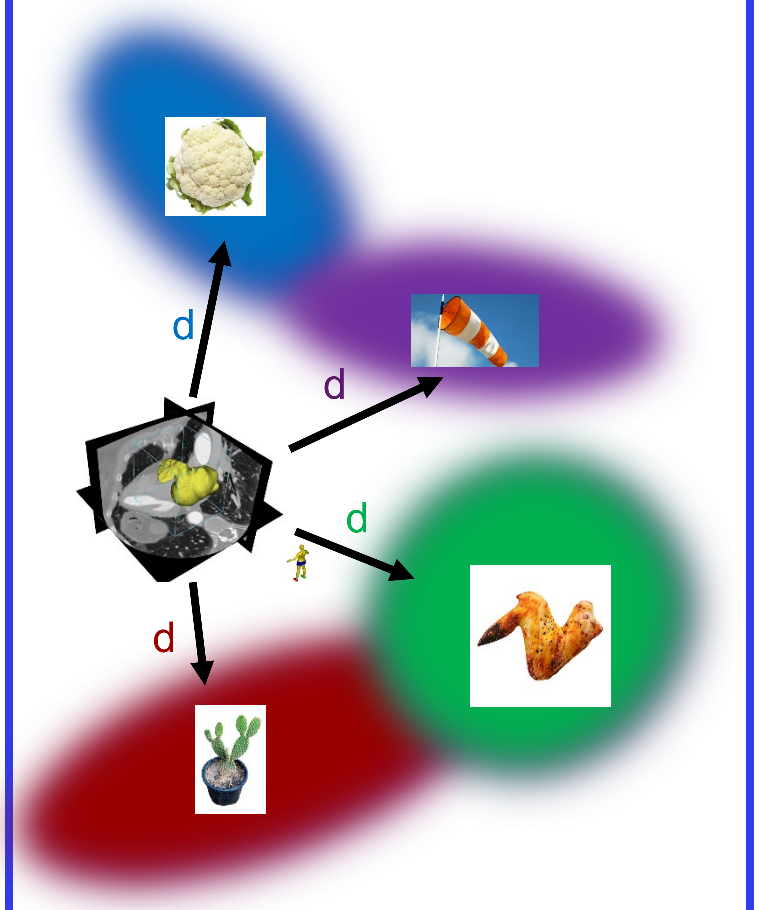

Shape space – latent space

- **Research questions**
  - How to parameterise complex 3D shapes to be able to do machine learning?
  - How to map complex 3D shapes to low-dimensional spaces (latent spaces)
  - How to compute meaningful distances in latent spaces
  - Supervised and unsupervised clustering and classification of complex 3D shapes
  - Prediction based on 3D shapes:
    - Risk scores
      - Risk of stroke based on your LAA shape
    - Device selection and deployment strategies
    - Procedural outcome prediction

How does a CT scan look like?

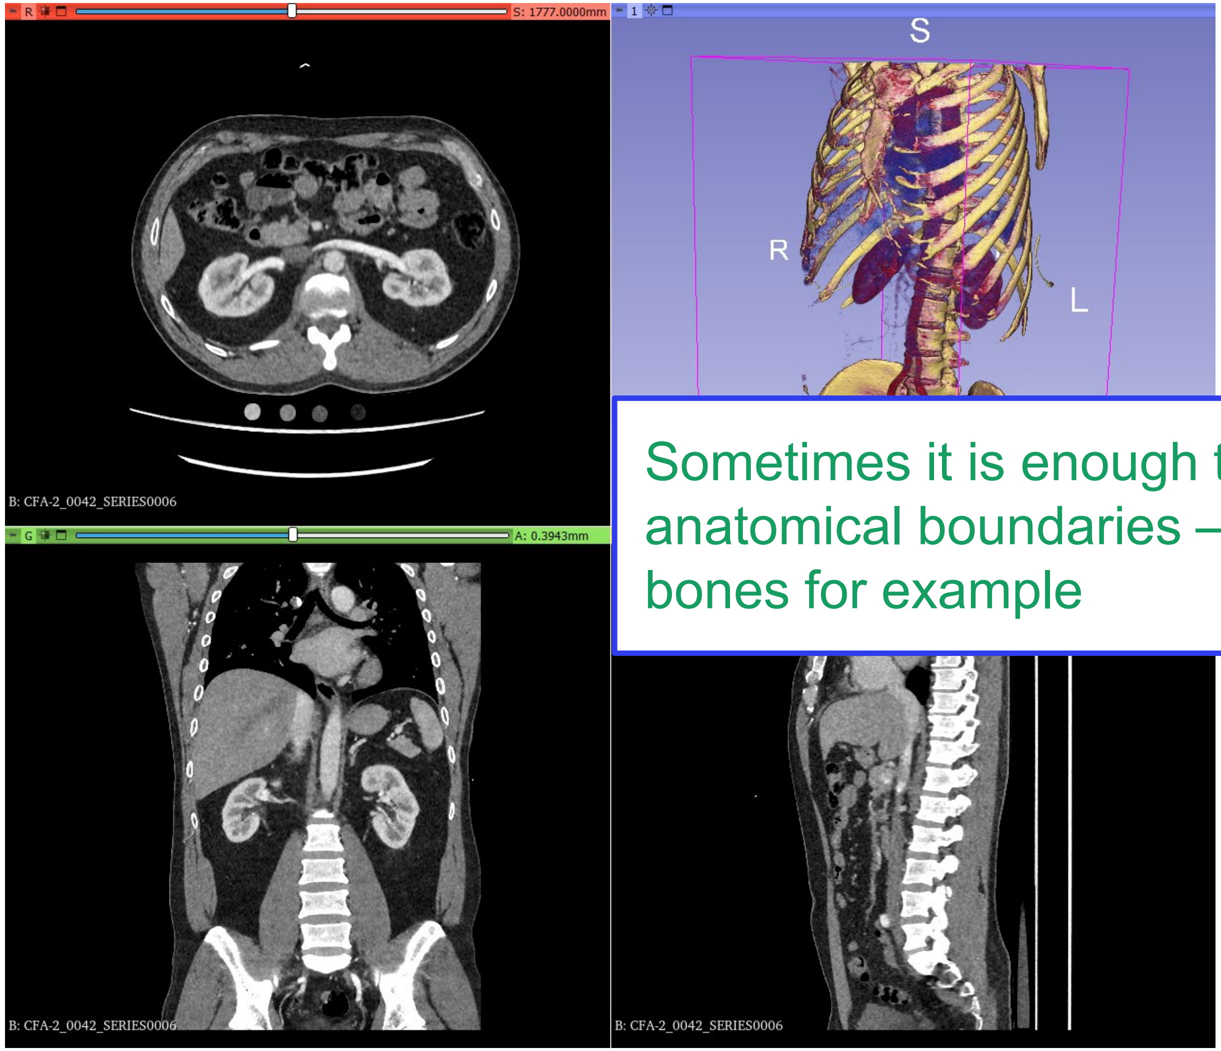

- A 3D volume consisting of small cubes (voxels)
- The value in each voxel reflects the amount of X-ray radiation that is absorbed
  - Bone: A lot of absorption (bright voxels)
  - Soft-tissue: Medium absorption (grey voxels)
  - ption (dark voxels)

ed CT-scan
  - A liquid is injected just before the CT scan
  - The liquid makes blood light up on the CT scan
  - Blood pools, arteries and veins become clearly visible

Sometimes it is enough to model the anatomical boundaries – blood pools or bones for example

What do I want?

- What would I like (this is not a standard list of requirements)
  - **Information preserving**: does not remove or filter geometric information
  - **Compact:** Uses a minimum of parameters
  - **Consistent:** There should not be (too many) ways the same surface can be represented by the parameterization
  - **Rotationally invariant:** The parameterisation is invariant to rotations.
  - **Can represent all topologies:** Works with non-manifolds, open surfaces and holes
  - **Can be used in ML frameworks:** It should be possible to feed the representation into a deep learning framework

What about surface meshes?

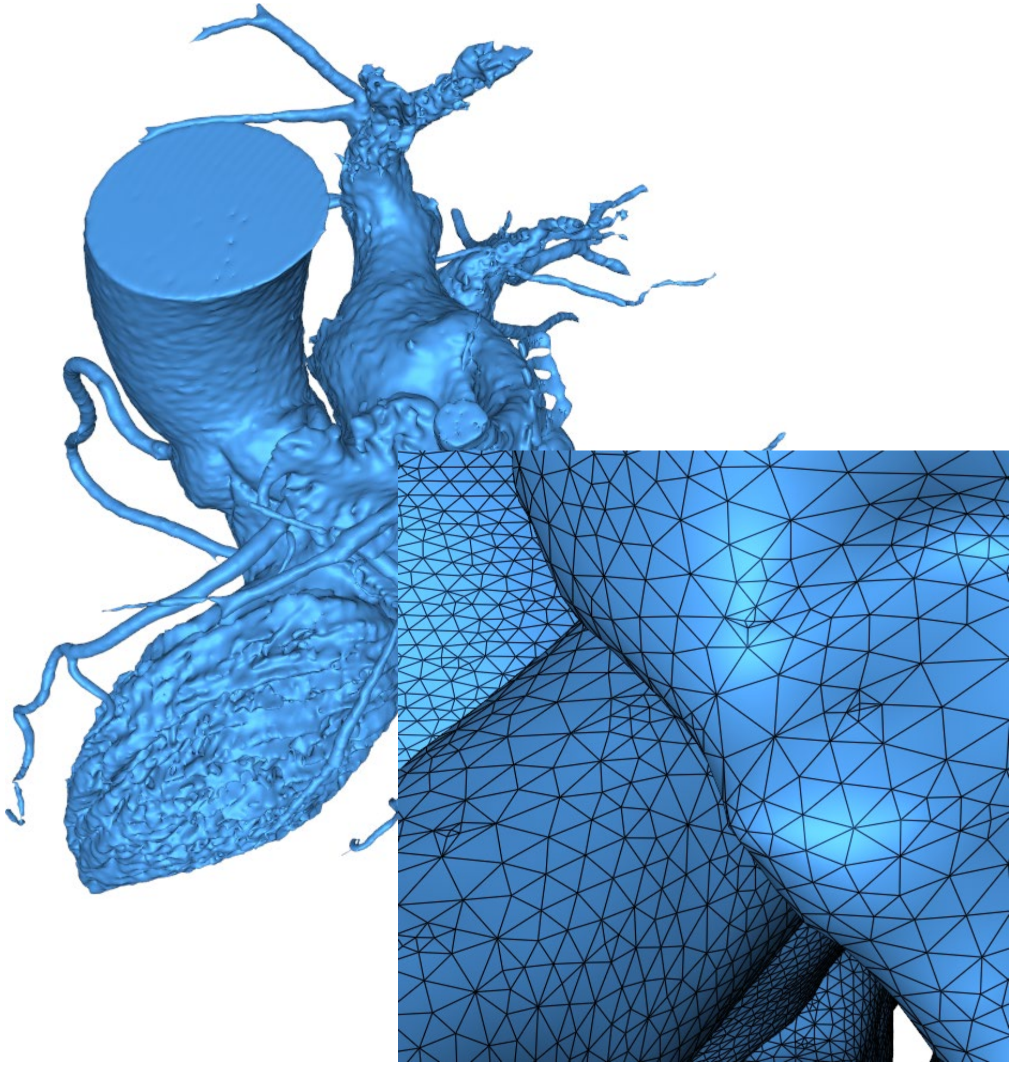

- Meshes:
  - **Information preserving**: does not remove or filter geometric information
  - **Compact:** Uses a minimum of parameters
  - **Consistent:** There should not be (too many) ways the same surface can be represented by the parameterization
  - **Rotationally invariant:** The parameterisation is invariant to rotations.
  - **Can represent all topologies**: Works with non-manifolds, open surfaces and holes
  - **Can be used in ML frameworks:** It should be possible to feed the representation into a deep learning framework

Deep learning directly on 3D meshes

SparseMeshCNN with Self-Attention for Segmentation of Large Meshes

Bjørn Hansen*[1], Mathias Lowes*[1], Thomas Ørkild[1], Anders Dahl[1], Vedrana Dahl[1], Ole de Backer[2], Oscar Camara[3], Rasmus Paulsen[1], Christian Ingwersen†[1,4], and Kristine Sørensen†[1]

[1]Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark
[2]The Heart Center, Rigshospitalet, University of Copenhagen, Copenhagen, Denmark
[3]BCN MedTech, Universitat Pompeu Fabra, Barcelona, Spain
[4]Trackman A/S, Vedbæk, Denmark

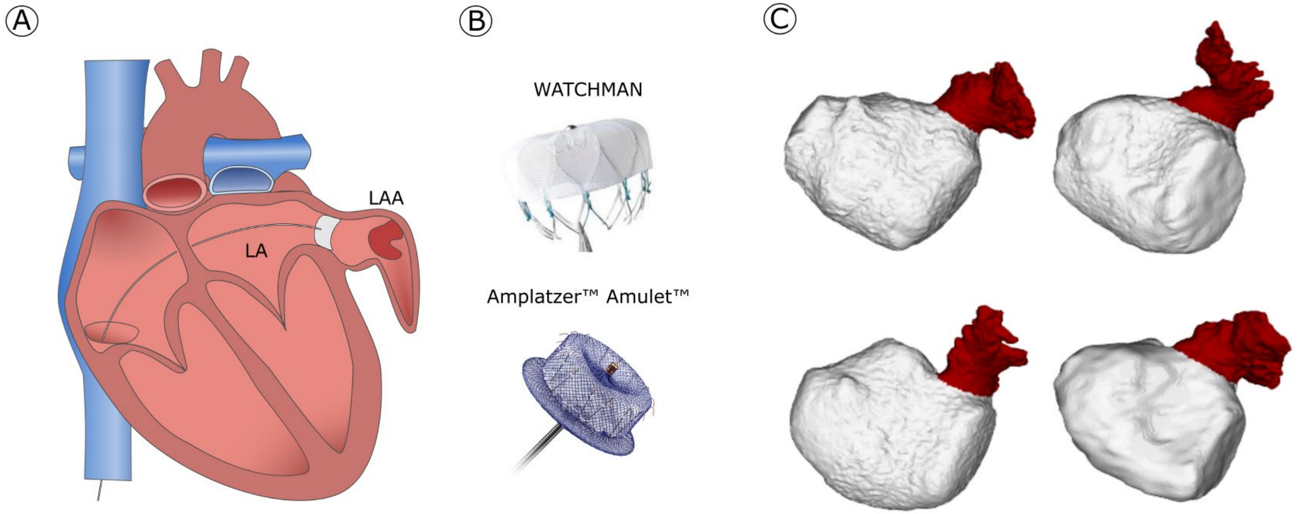

Prediction of intersection between the left atrium and the left atrial appendage in the human heart. For simulation of surgical device insertion.

Implicit shape descriptions

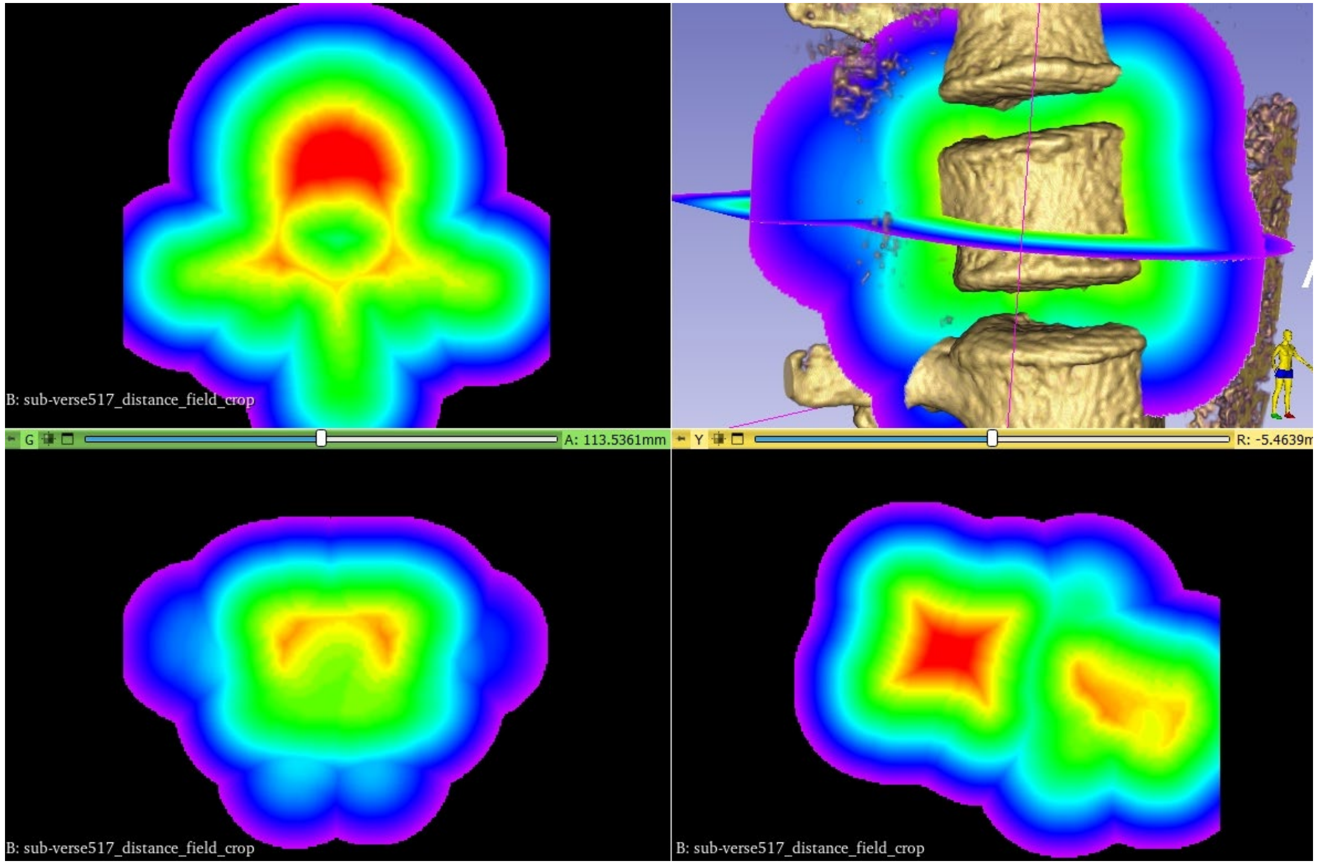

- Implicit shape description
- Carries information about the shape in the entire field
- In the simplest version it is just a 3D voxel grid
  – A distance field



Intersection between image analysis and computer graphics



Neural Representation of Open Surfaces

Christiansen, T. V., Bærentzen, J. A., Paulsen, R. R. & Hannemose, M. R., 2023, (Accepted/In press) In: Computer Graphics Forum. 13 p., e14916.

Signed distance fields?

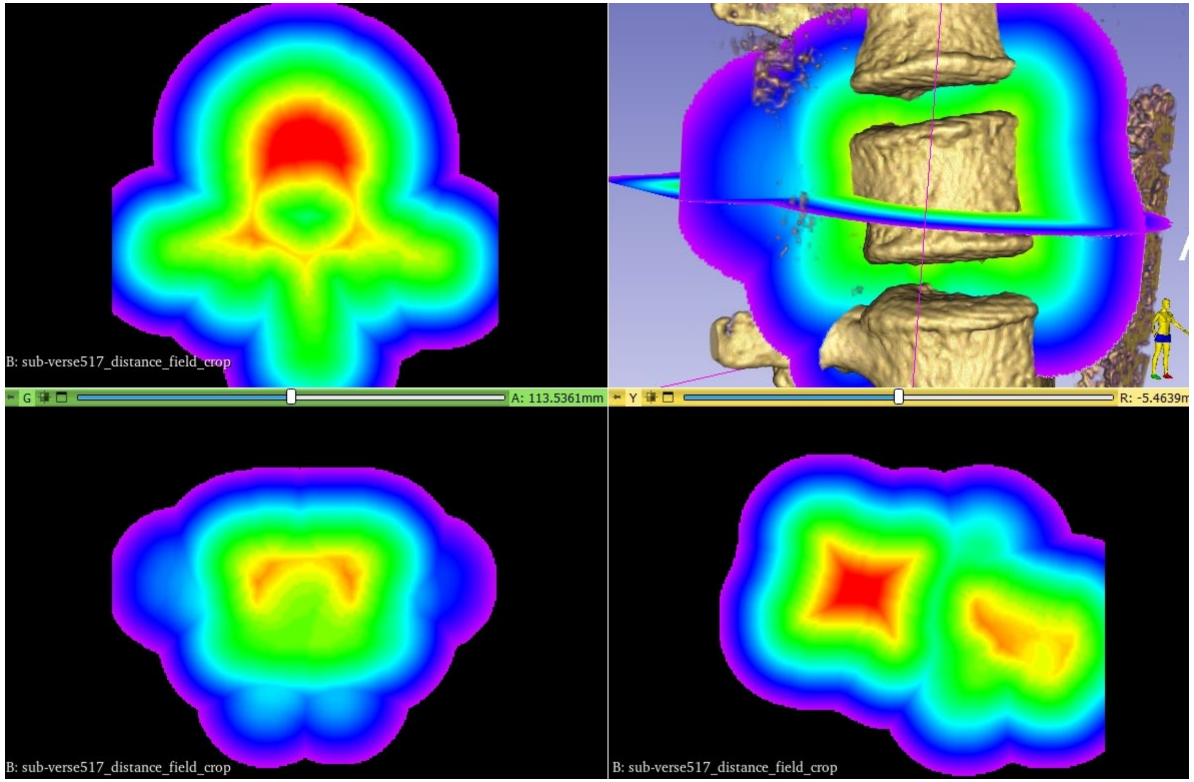

- Signed distance fields:
  - **Information preserving:** does not remove or filter geometric information
  - **Compact:** Uses a minimum of parameters
  - **Consistent:** There should not be (too many) ways the same surface can be represented by the parameterization
  - **Rotationally invariant:** The parameterisation is invariant to rotations.
  - **Can represent all topologies:** Works with non-manifolds, open surfaces and holes
  - **Can be used in ML frameworks:** It should be possible to feed the representation into a deep learning framework

Unsigned distance fields?

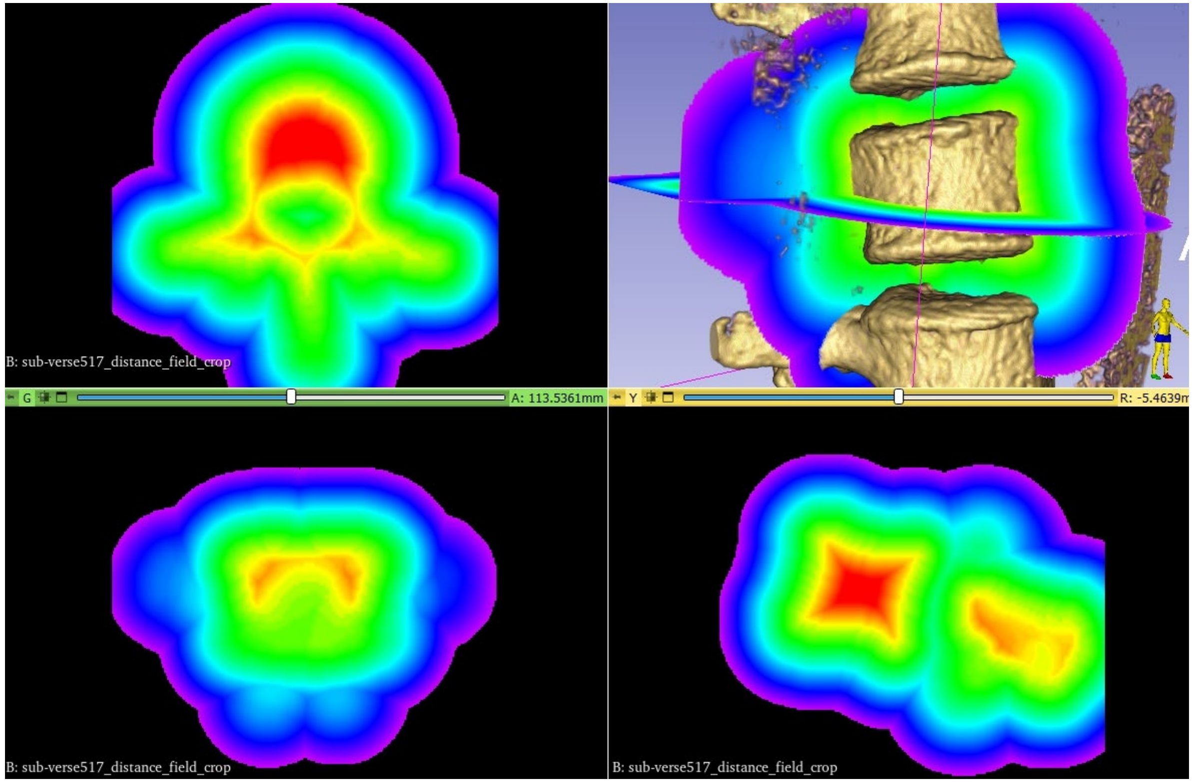

- Unsigned distance fields:
  - Information preserving: does not remove or filter geometric information
  - Compact: Uses a minimum of parameters
  - Consistent: There should not be (too many) ways the same surface can be represented by the parameterization
  - Rotationally invariant: The parameterisation is invariant to rotations.
  - Can represent all topologies: Works with non-manifolds, open surfaces and holes
  - Can be used in ML frameworks: It should be possible to feed the representation into a deep learning framework

Implicit Neural Distance representations

MICCAI 2021

Implicit Neural Distance Representation for Unsupervised and Supervised Classification of Complex Anatomies

Kristine Aavild Juhl[1(✉)], Xabier Morales[2], Ole de Backer[3], Oscar Camara[2],
and Rasmus Reinhold Paulsen[1]

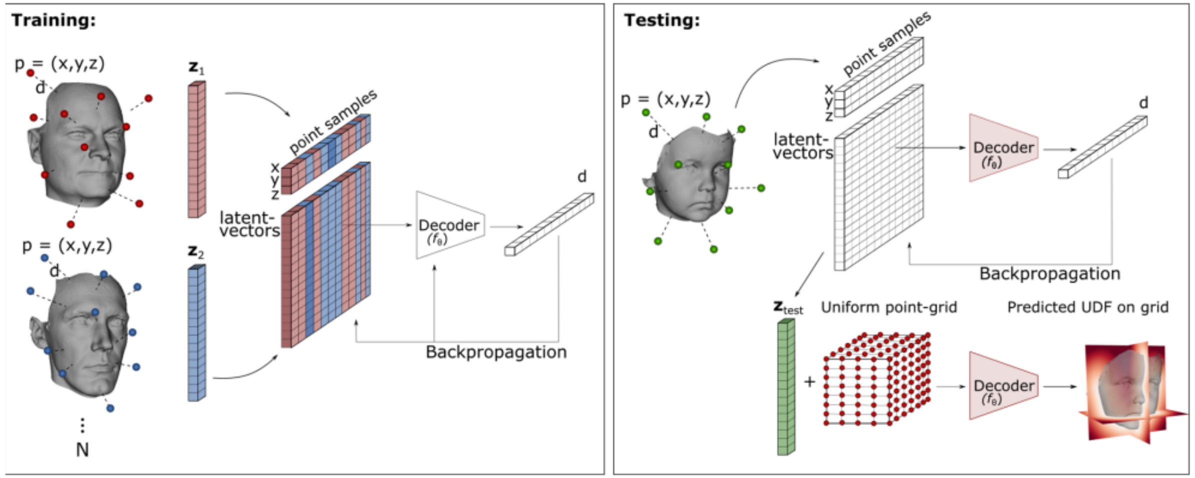

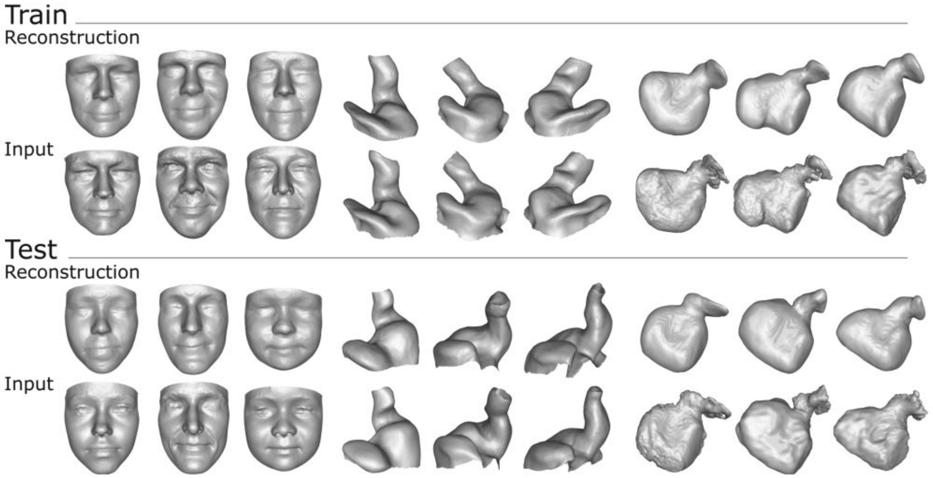

Distance fields

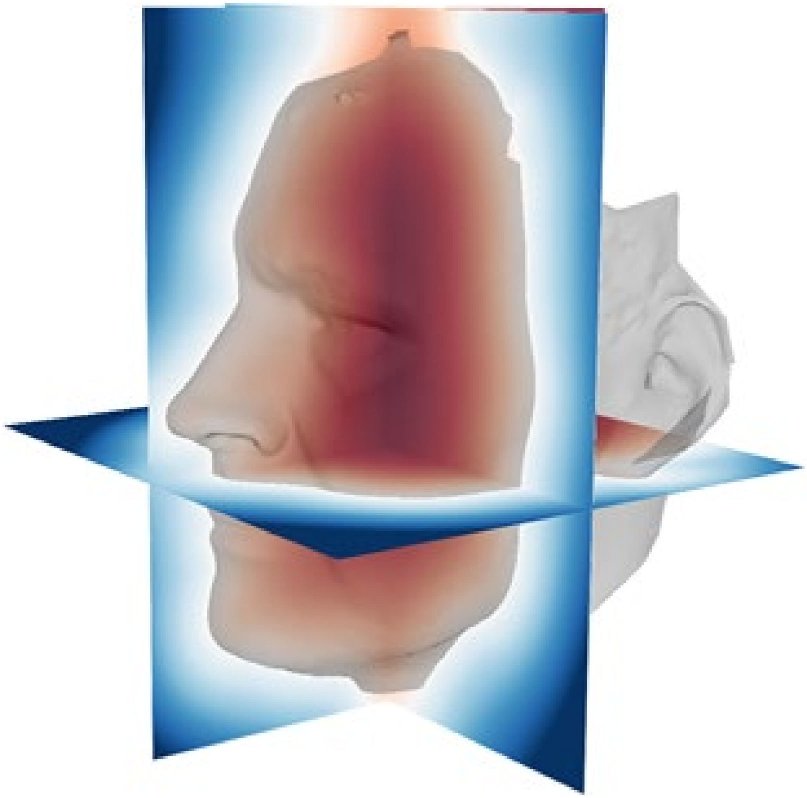

Signed distance fields

+ Easy surface extraction at zero-level isosurface
+ Differentiable at all points
- Surface must be closed

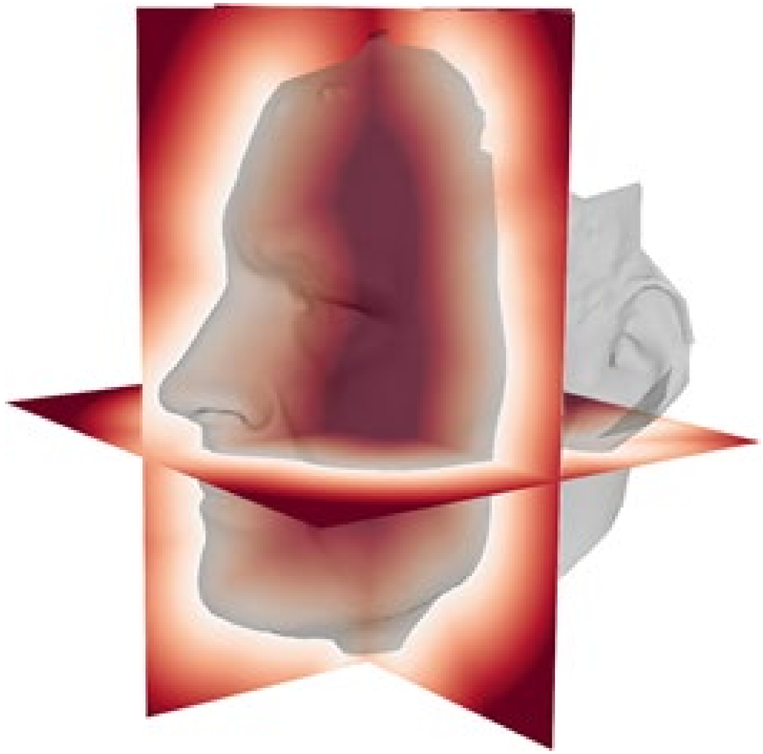

Unsigned distance fields

+ Can represent arbitrary topologies
- Undifferentiable near surface
- More advanced methods needed for surface extraction

Distance fields and deep learning

Single shape representation

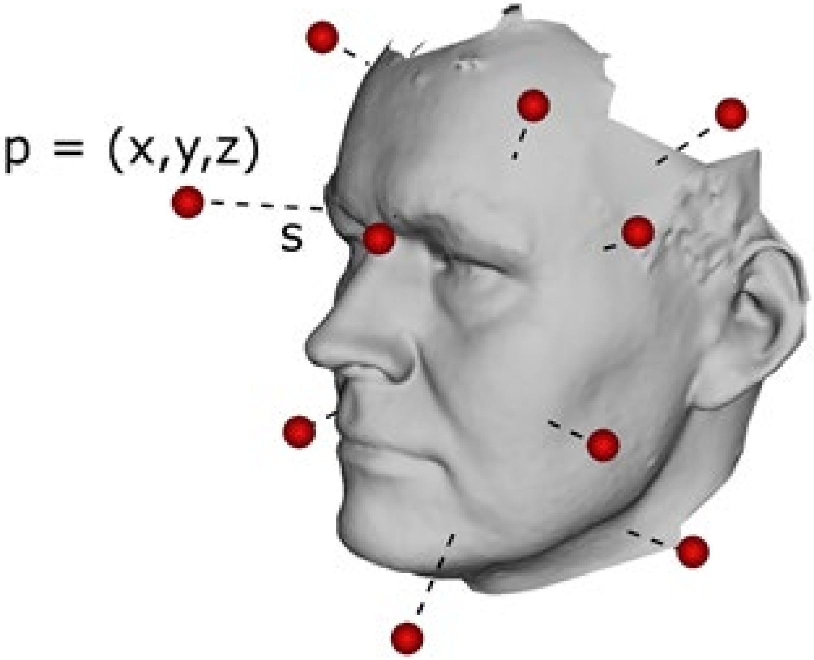

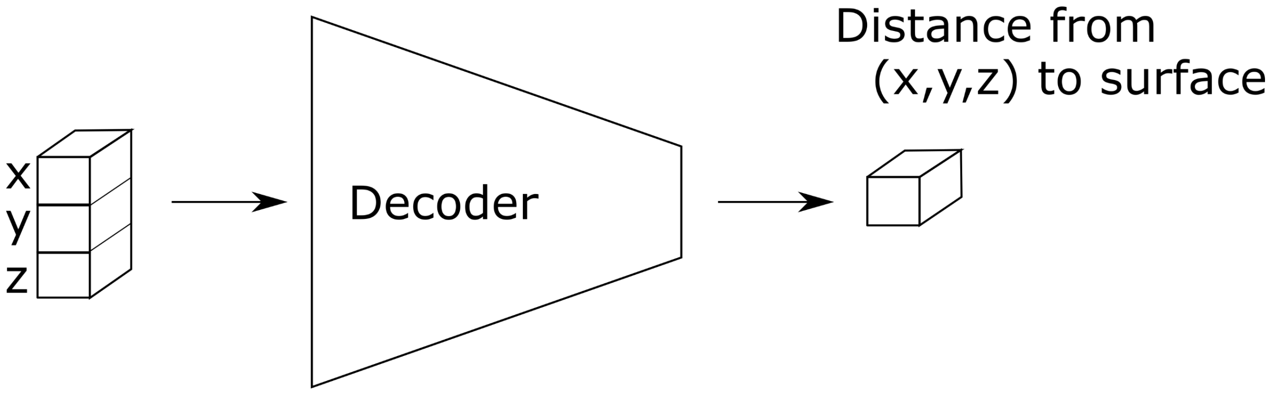

$$f(\boldsymbol{p}) \mapsto s, \qquad \boldsymbol{p} \in \mathbb{R}^3, s \in \mathbb{R}$$

Predicts a signed distance to a surface given a position

No need for a voxel grid

Park et.al., "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation ", CVPR2019

Deep signed distance field?

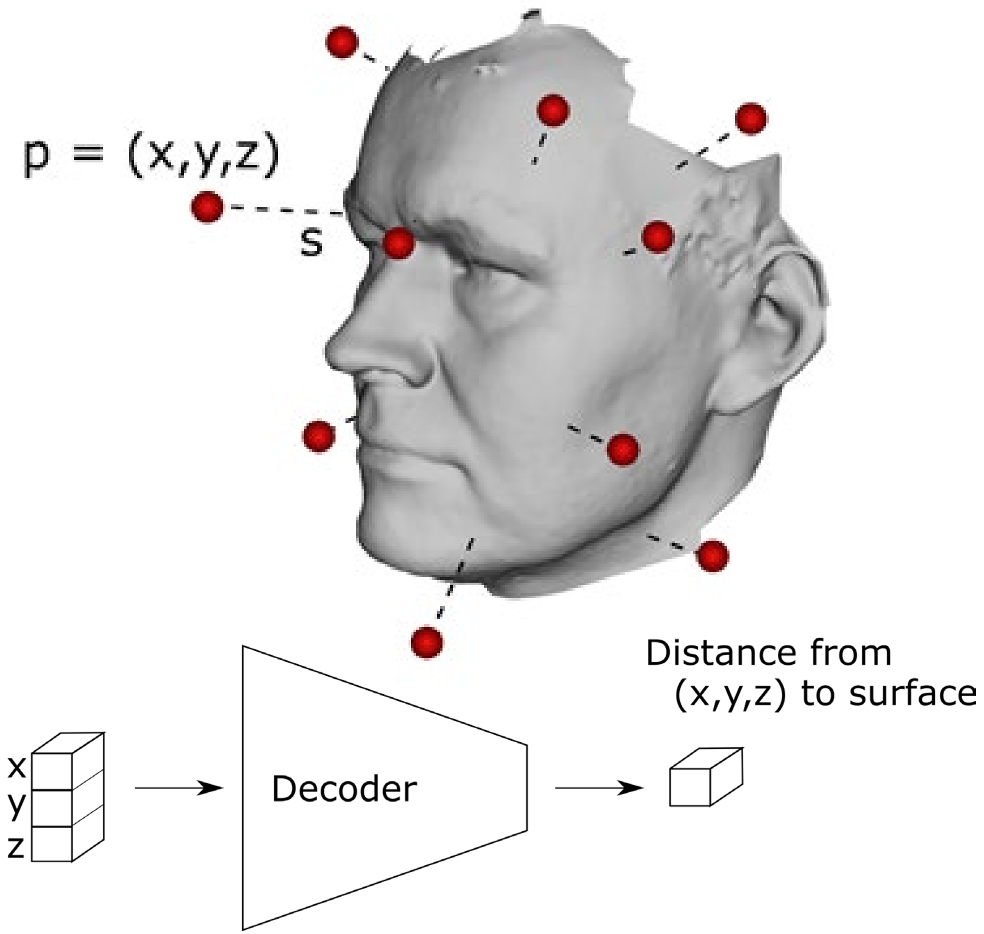

- Unsigned distance fields:
  - **Information preserving:** does not remove or filter geometric information
  - **Compact:** Uses a minimum of parameters
  - **Consistent:** There should not be (too many) ways the same surface can be represented by the parameterization
  - **Rotationally invariant:** The parameterisation is invariant to rotations.
  - **Can represent all topologies:** Works with non-manifolds, open surfaces and holes
  - **Can be used in ML frameworks:** It should be possible to feed the representation into a deep learning framework

Distance fields and deep learning

Single shape representation

What about more shapes?

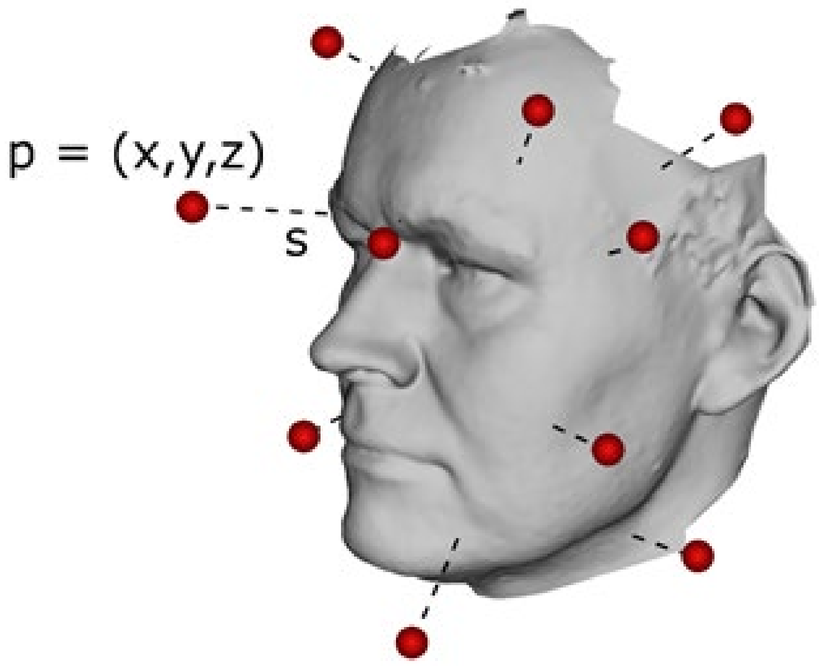

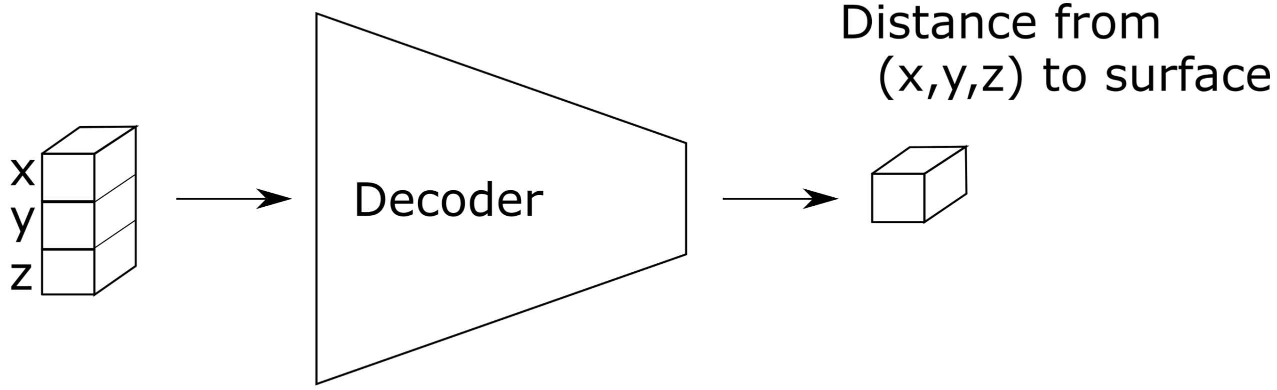

$$f(\boldsymbol{p}) \mapsto s, \qquad \boldsymbol{p} \in \mathbb{R}^3, s \in \mathbb{R}$$

Predicts a signed distance to a surface given a position

No need for a voxel grid

Park et.al., “DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation ”, CVPR2019

Distance fields and deep learning

Multi shape representation

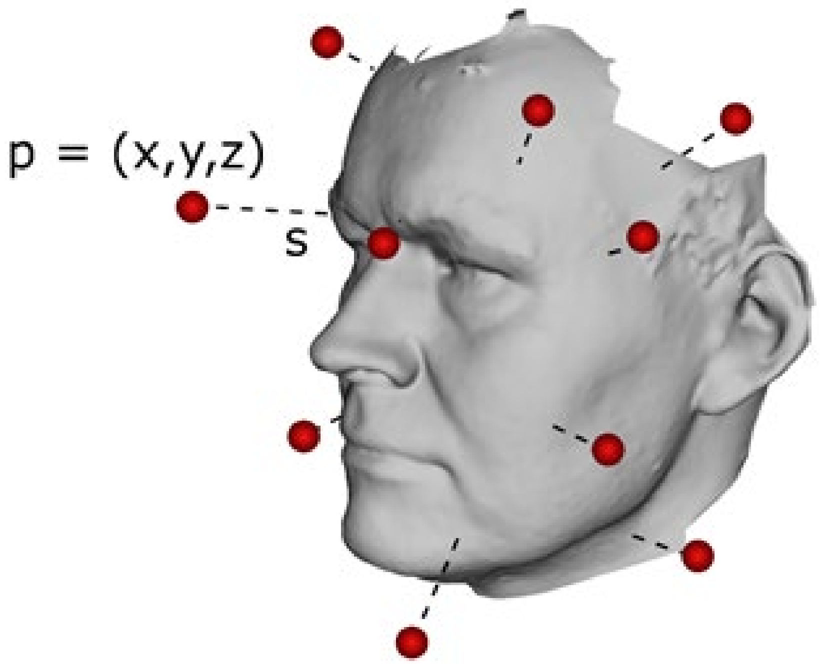

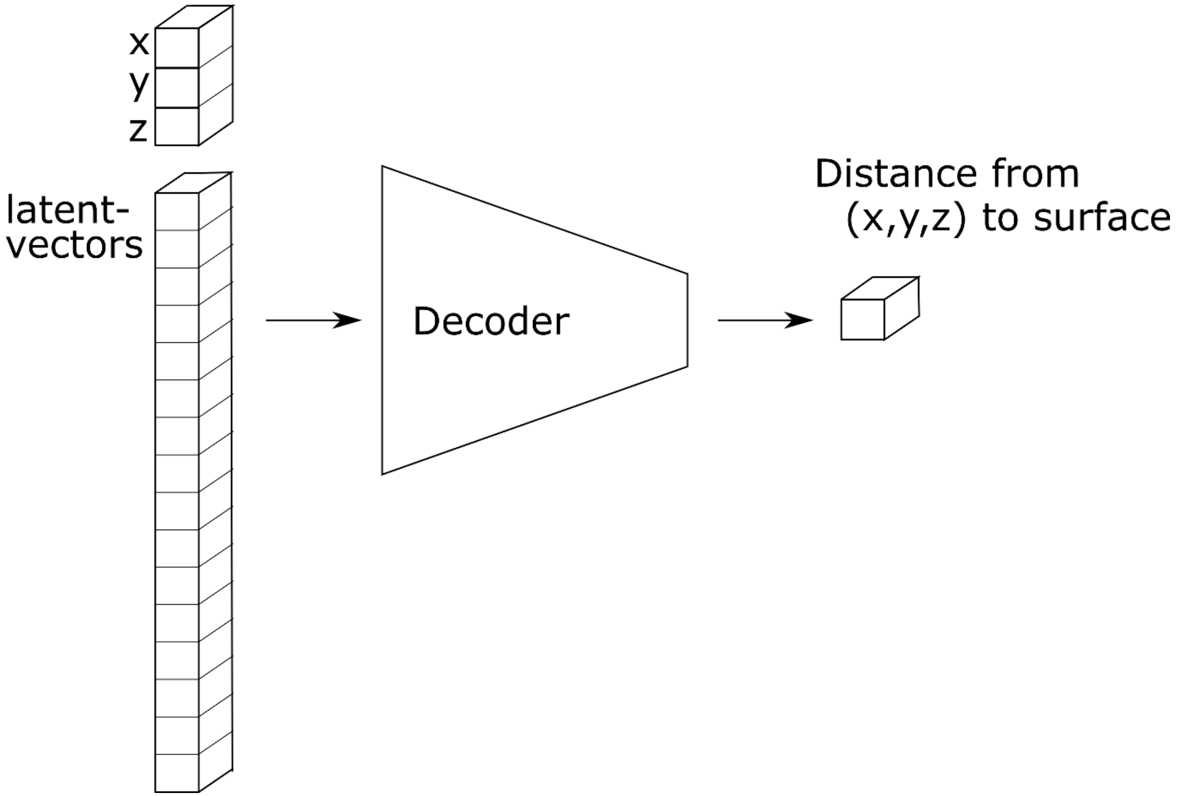

$$f(\boldsymbol{p}) \mapsto s, \qquad \boldsymbol{p} \in \mathbb{R}^3, s \in \mathbb{R}$$

Park et.al., "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation ", CVPR2019

Auto-decoders

Auto-encoder

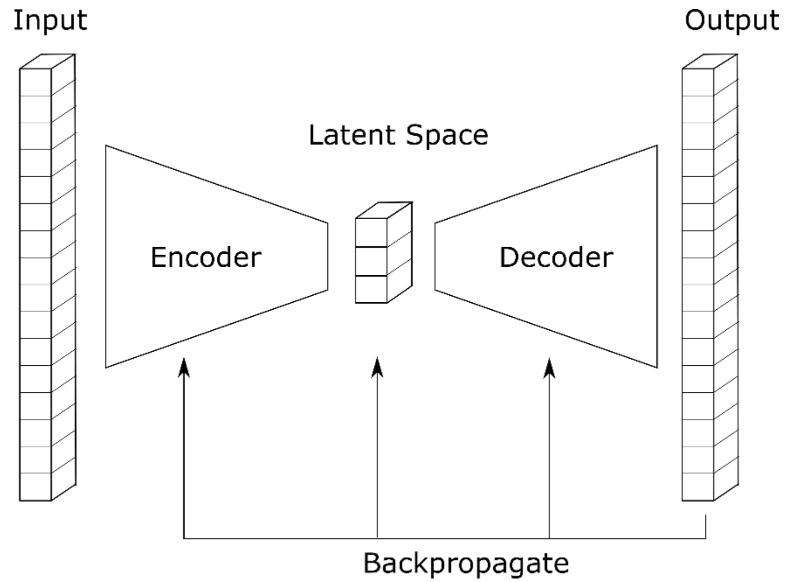

Auto-decoder

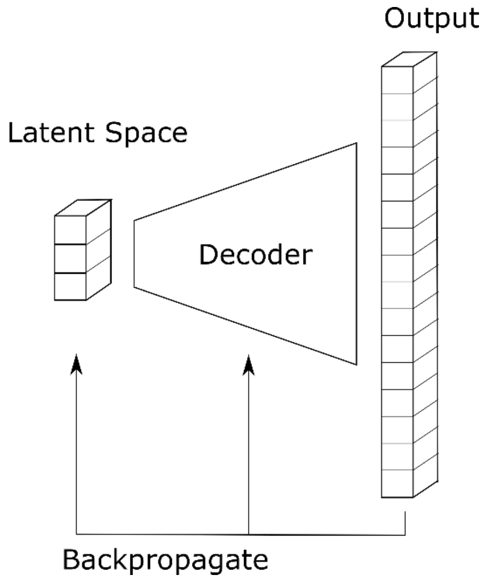

Distance fields and deep learning

Multi shape representation - training

$$X_i = \{(\boldsymbol{p}_j, s_j): \quad s_j = DF^i(\boldsymbol{p}_j)\}$$

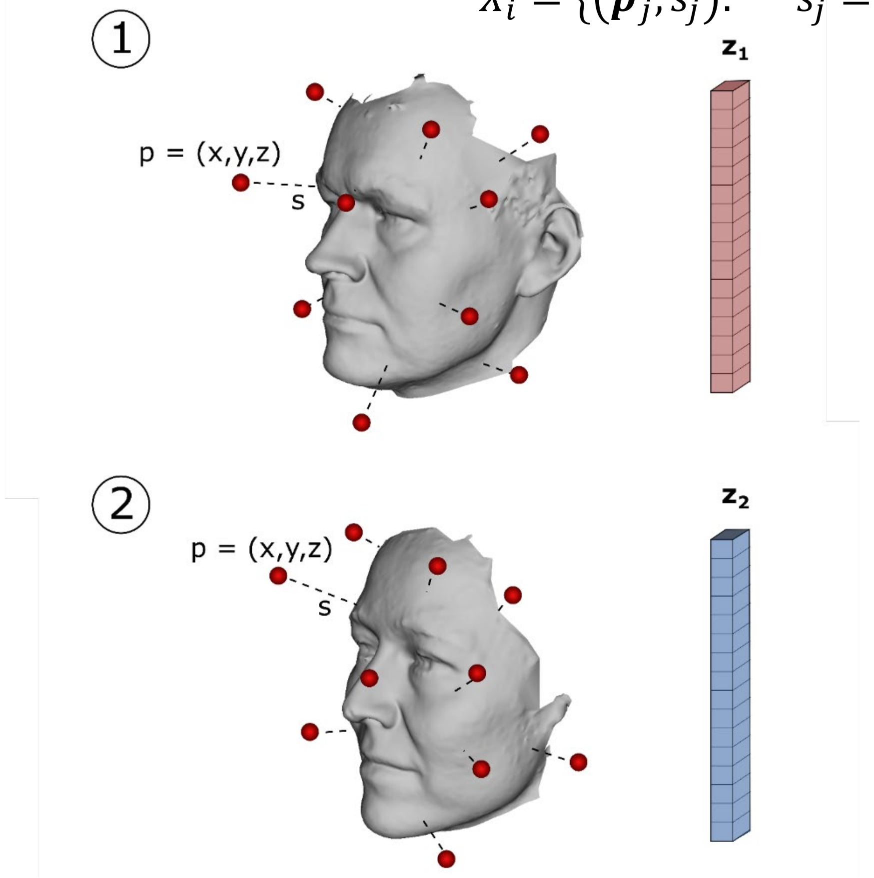

$$\underset{\theta,\{\boldsymbol{z}_i\}_{i=1}^{N}}{\arg\min} \sum_{i=1}^{N} \left(\sum_{j=1}^{K} \mathcal{L}(f_\theta(\boldsymbol{z}_i, \boldsymbol{p}_j), s_j) + \frac{1}{\sigma^2} \|\boldsymbol{z}_i\|_2^2 \right)$$

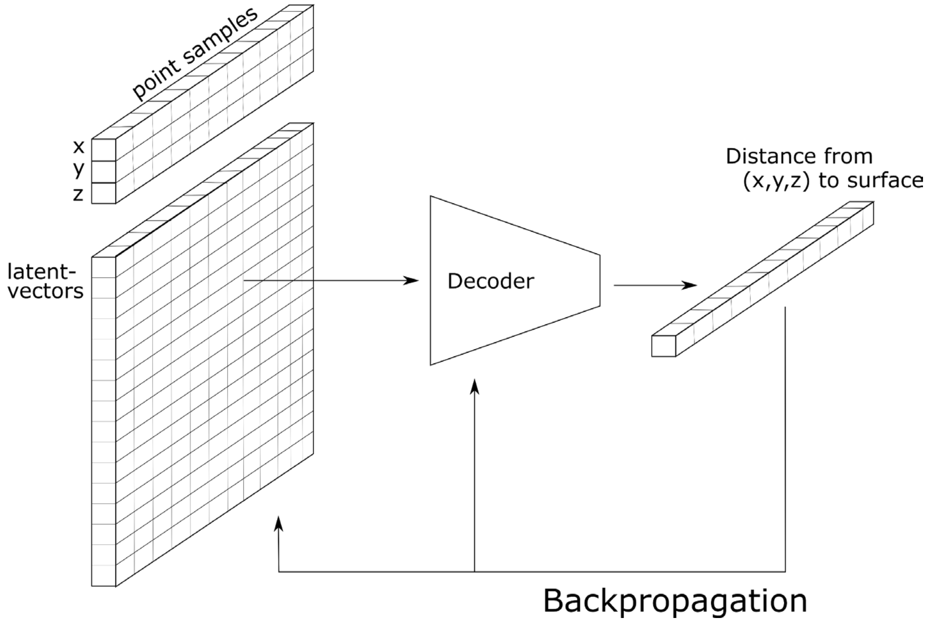

Distance fields and deep learning

Multi shape representation - training

- Training:

$$X_i = \{(\boldsymbol{p}_j, s_j): \quad s_j = DF^i(\boldsymbol{p}_j)\}$$

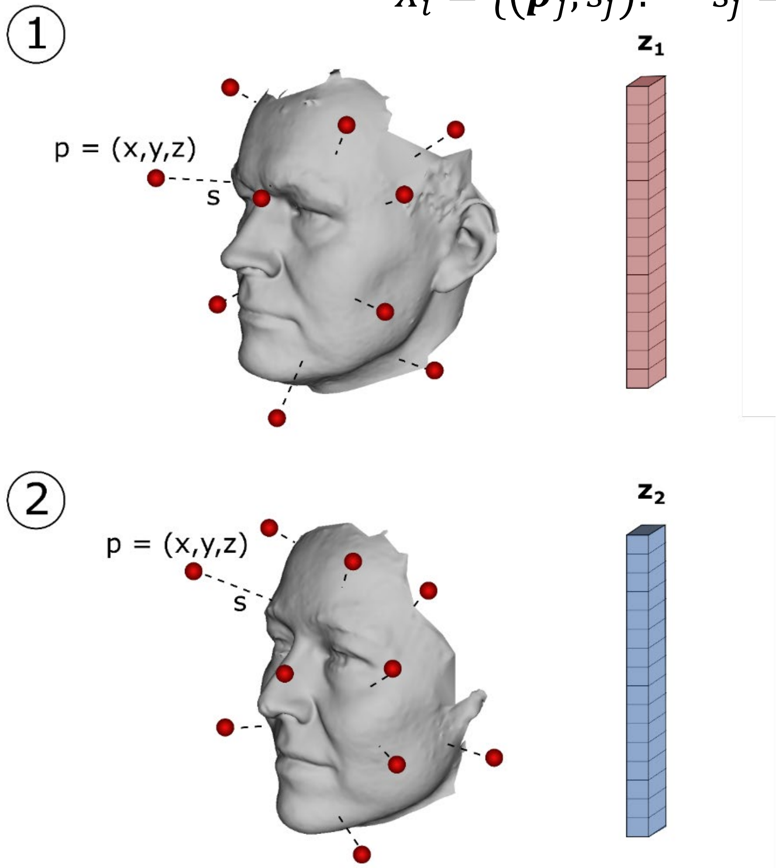

Clamped L1-distance

$$\underset{\theta,\{\boldsymbol{z}_i\}_{i=1}^{N}}{\arg\min} \sum_{i=1}^{N} \left(\sum_{j=1}^{K} \mathcal{L}(f_\theta(\boldsymbol{z}_i, \boldsymbol{p}_j), s_j) + \frac{1}{\sigma^2} ||\boldsymbol{z}_i||_2^2 \right)$$

Regularization

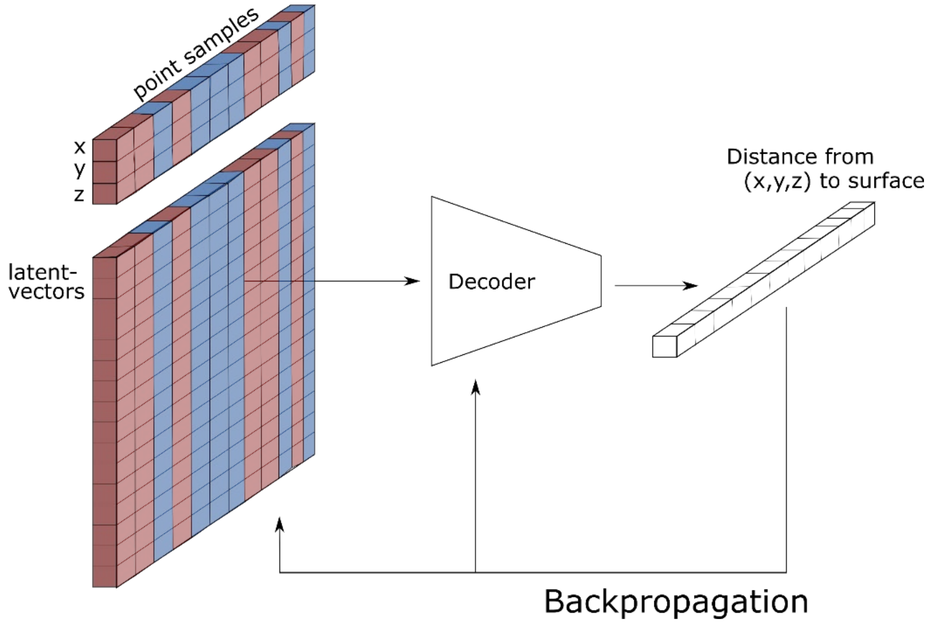

Distance fields and deep learning

Multi shape representation – testing with unseen shapes

- Testing with unseen examples:

$$X_{test} = \{(\boldsymbol{p}_j, s_j): \quad s_j = DF^{test}(\boldsymbol{p}_j)\}$$

$$\hat{\boldsymbol{z}} = \arg\min_{\boldsymbol{z}} \sum_{(\boldsymbol{x}_j, s_j) \in X}^{N} \mathcal{L}(f_\theta(\boldsymbol{z}_i, \boldsymbol{p}_j), s_j) + \frac{1}{\sigma^2} ||\boldsymbol{z}_i||_2^2$$

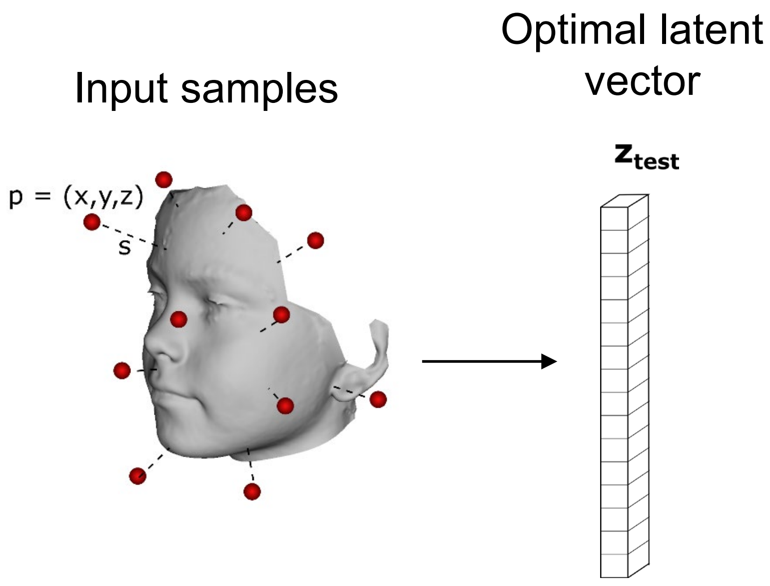

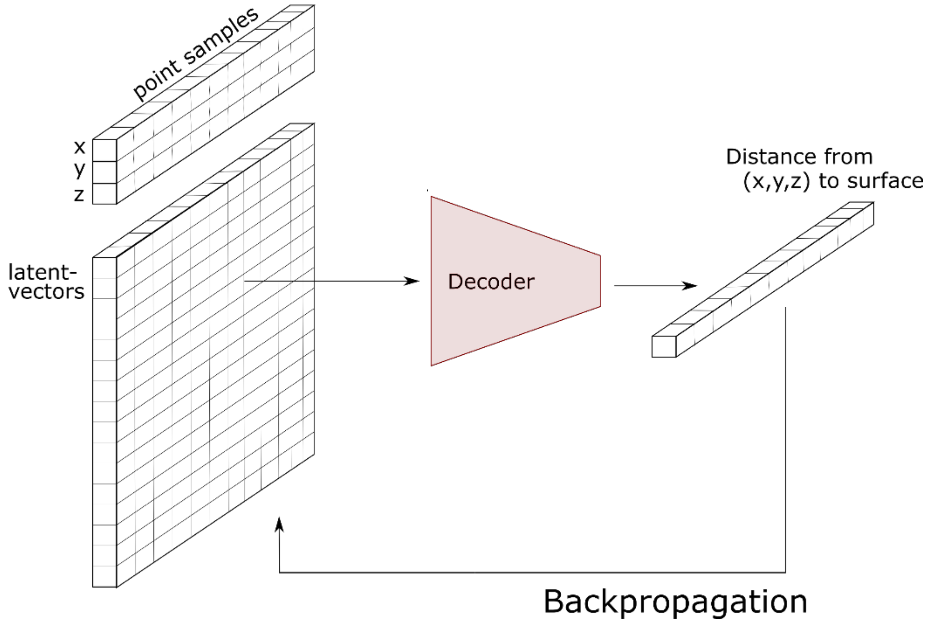

Distance fields and deep learning

Multi shape representation – testing with unseen shapes

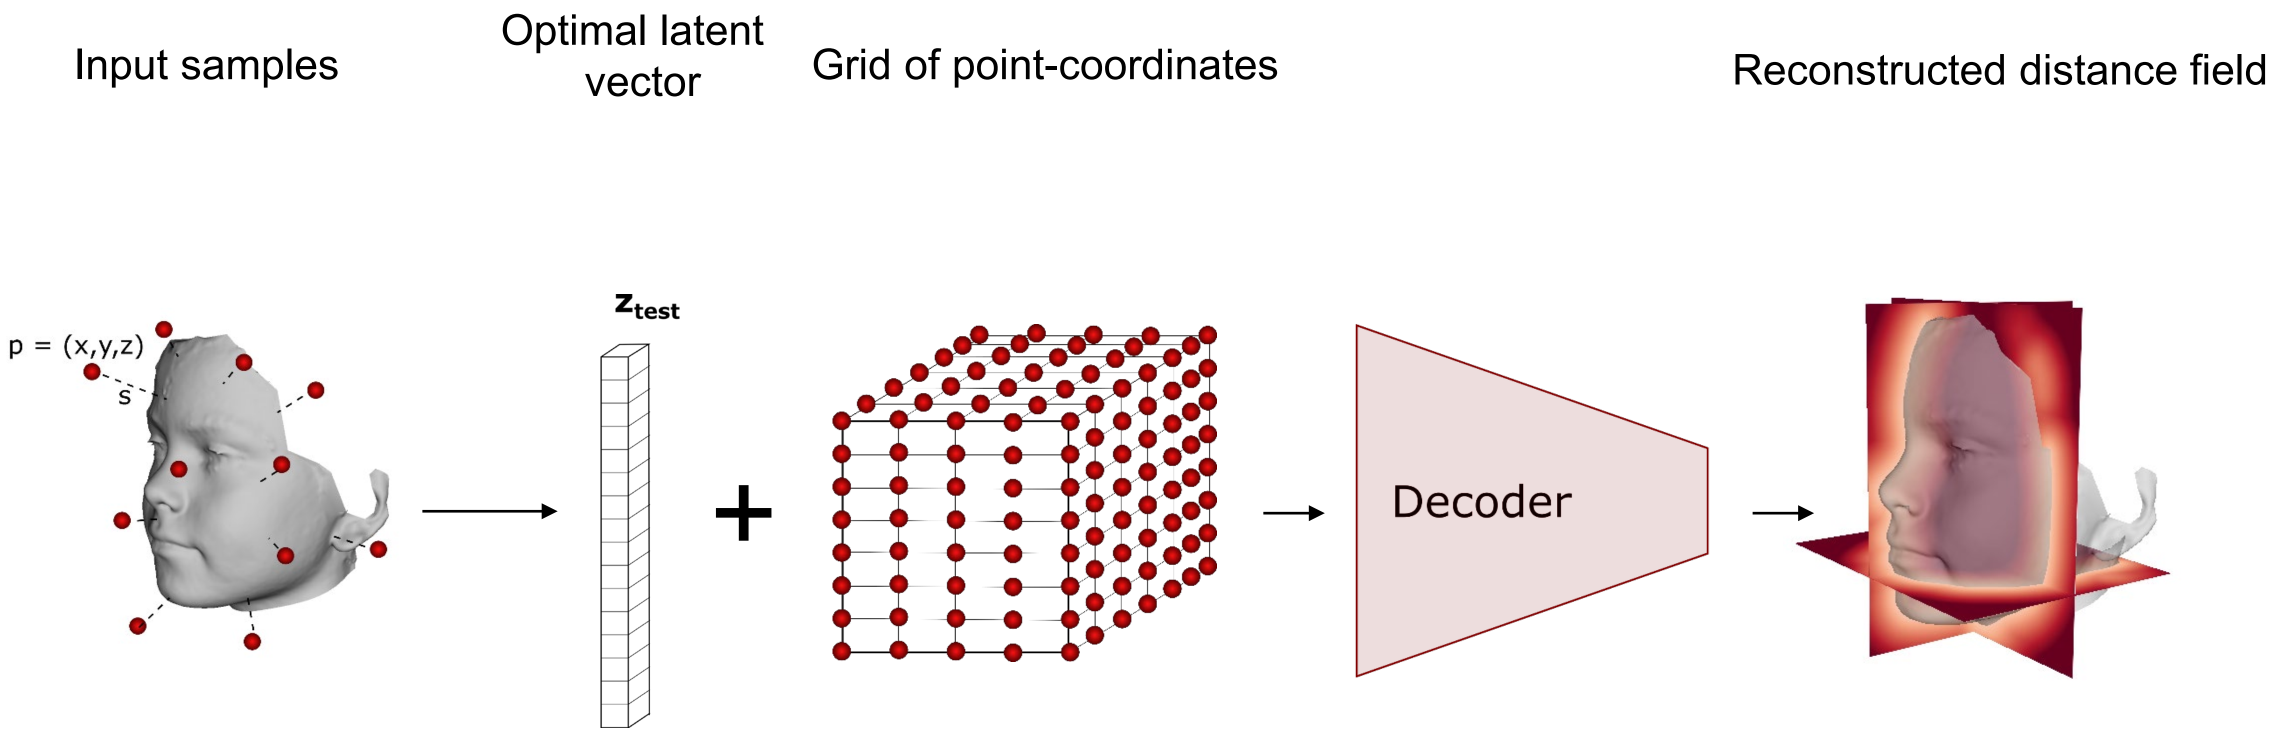

Distance fields and deep learning

Multi shape representation

- Decoder architecture:
  - Deep Feed Forward network with 8 layers
  - Latent vector and coordinates are reintroduces at the 4th layer
  - Latent space size: 64/128

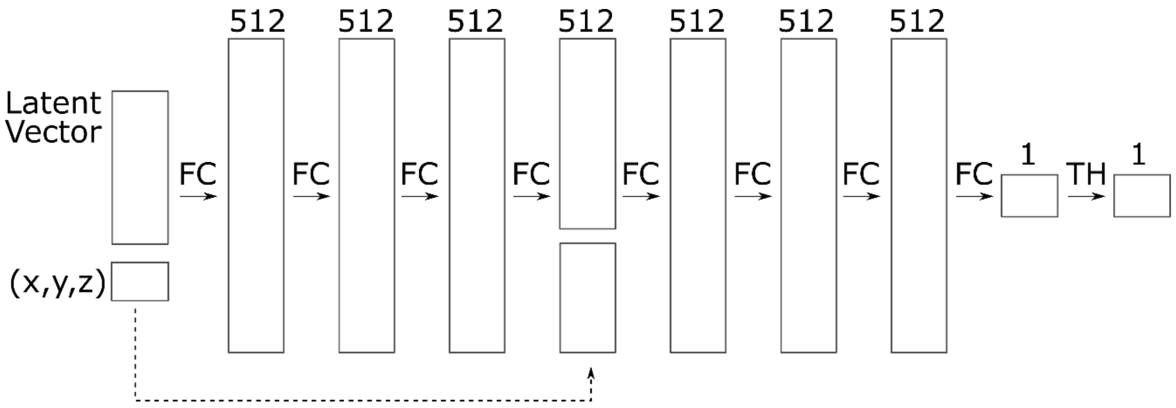

Experiments with neural unsigned distance functions

Input samples

Optimal latent vector

Grid of point-coordinates

Reconstructed distance field

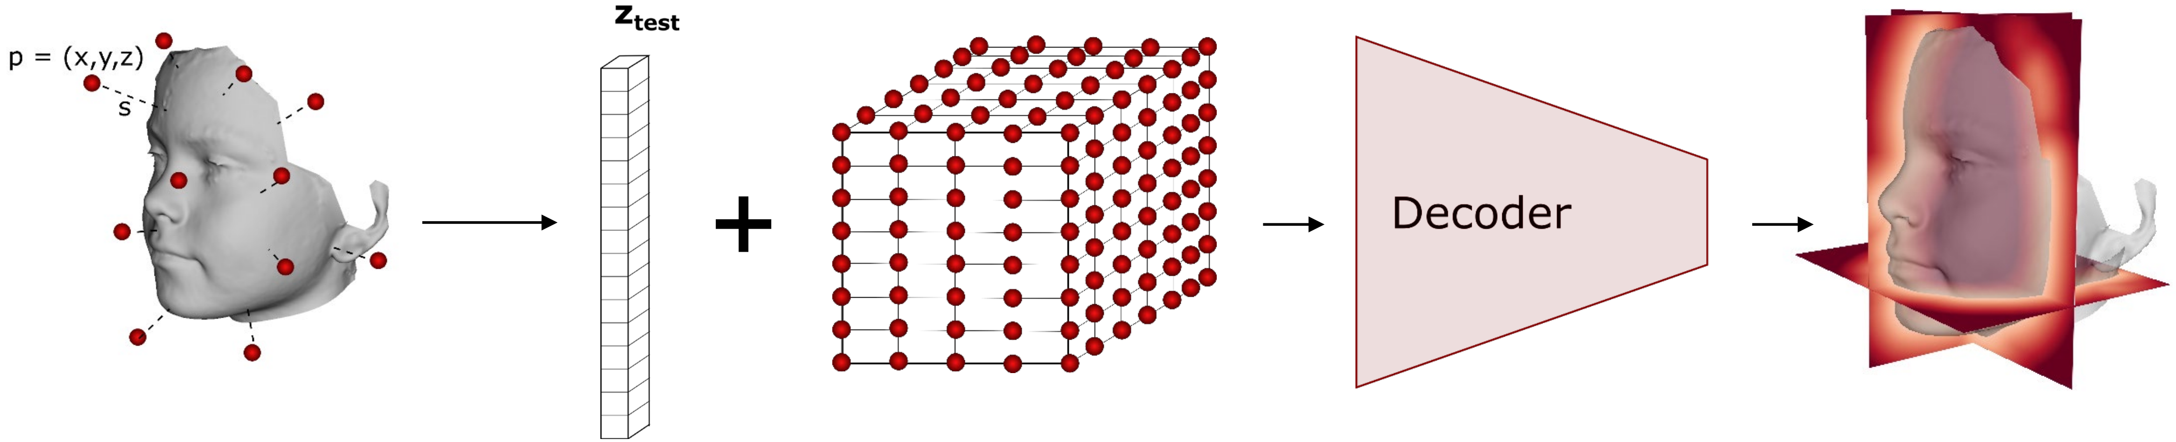

Data

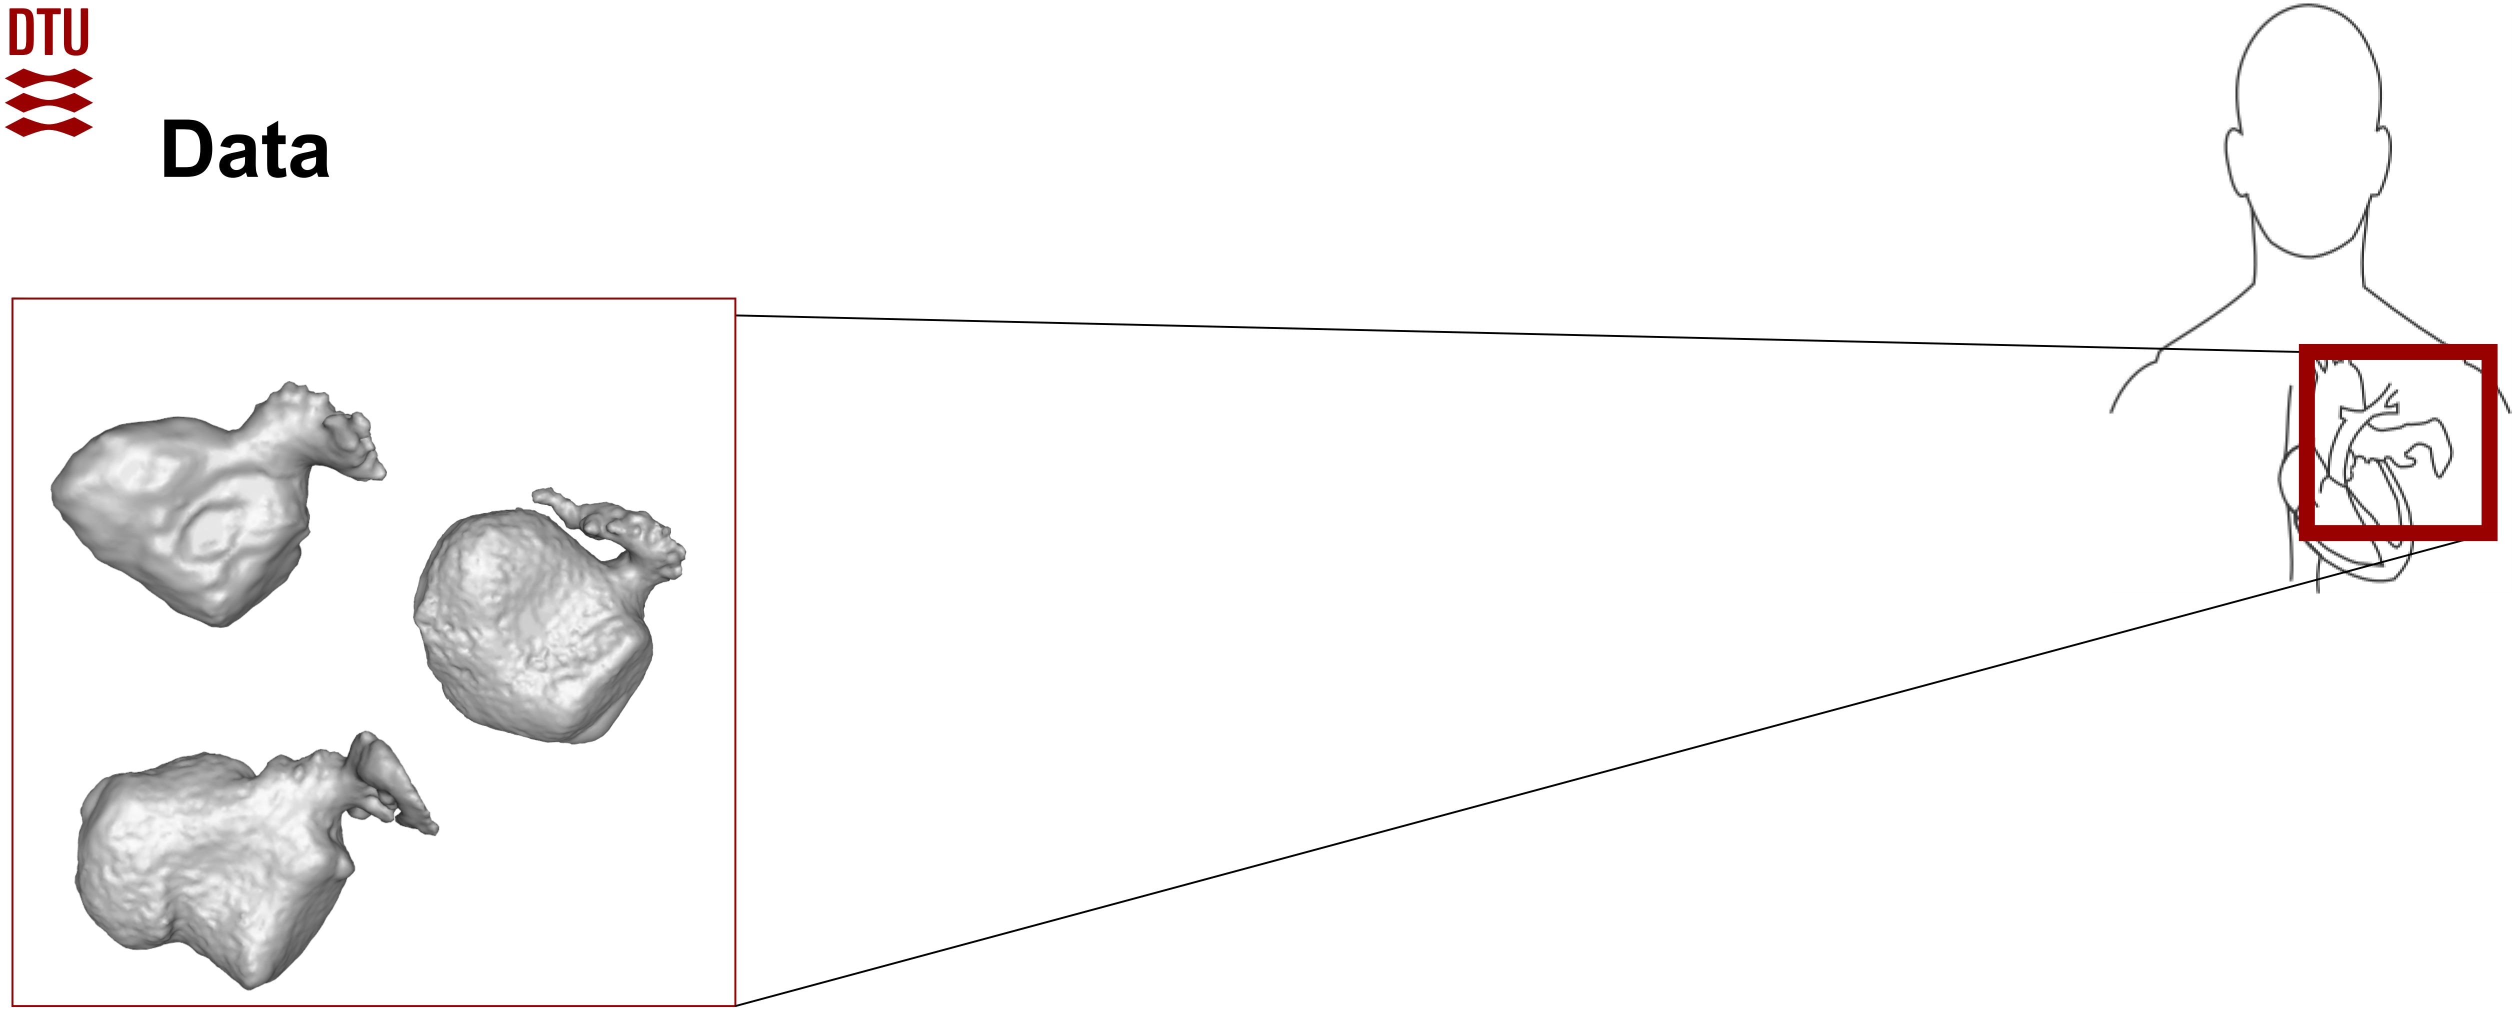

LA

- 106 surfaces
- Topology: spheres

EARS

- 571 ears (259/312 left/right)
- Topology: Tubes

ESOF

- 394 faces
  (192/202 male/female, 0-84 years)
- Topology: Planes

Data

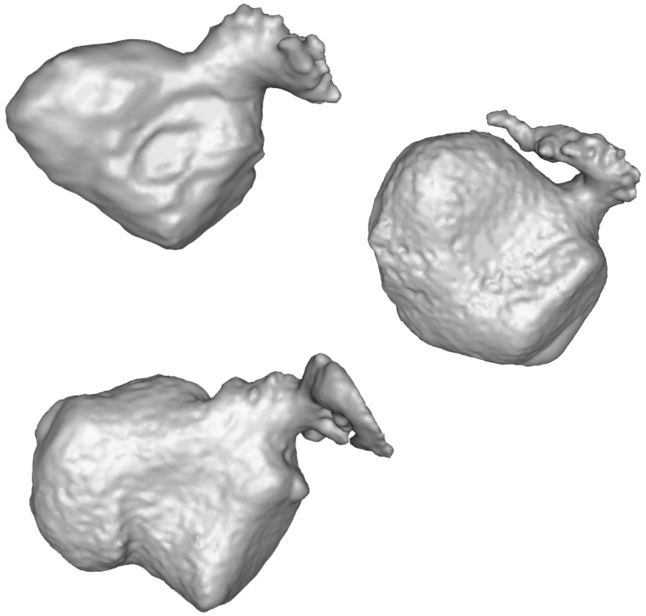

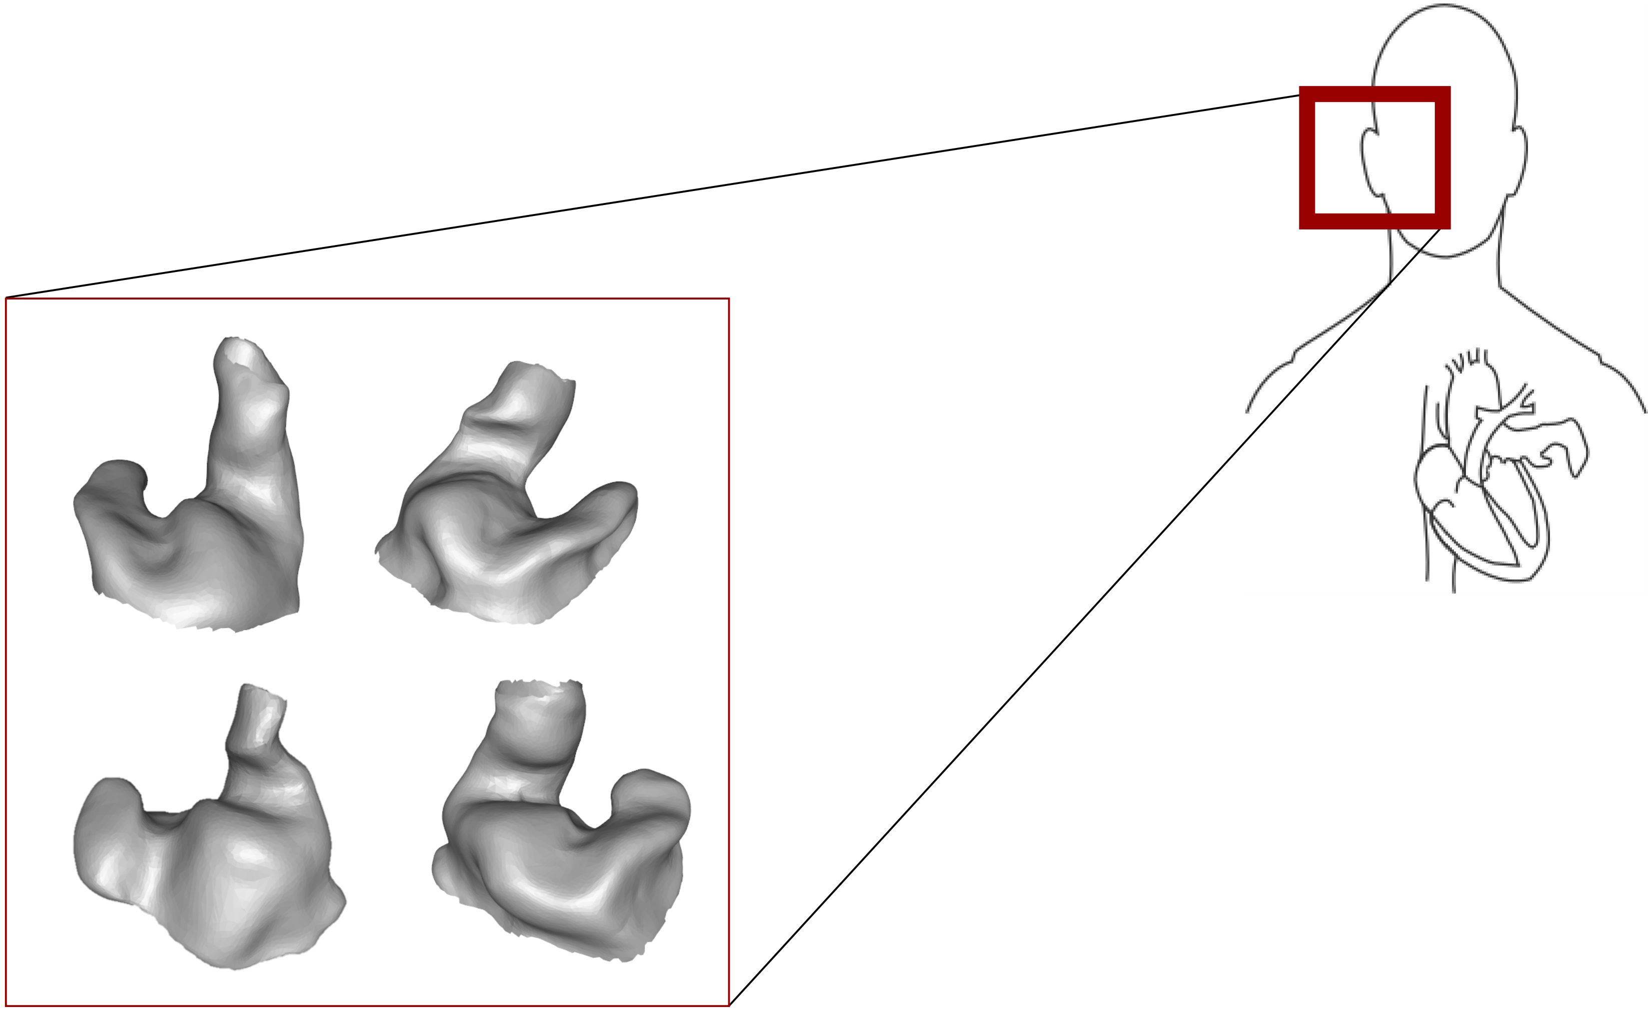

LA
- 106 surfaces
- Topology: spheres

EARS
- 571 ears (259/312 left/right)
- Topology: Tubes

ESOF
- 394 faces
  (192/202 male/female, 0-84 years)
- Topology: Planes

Data

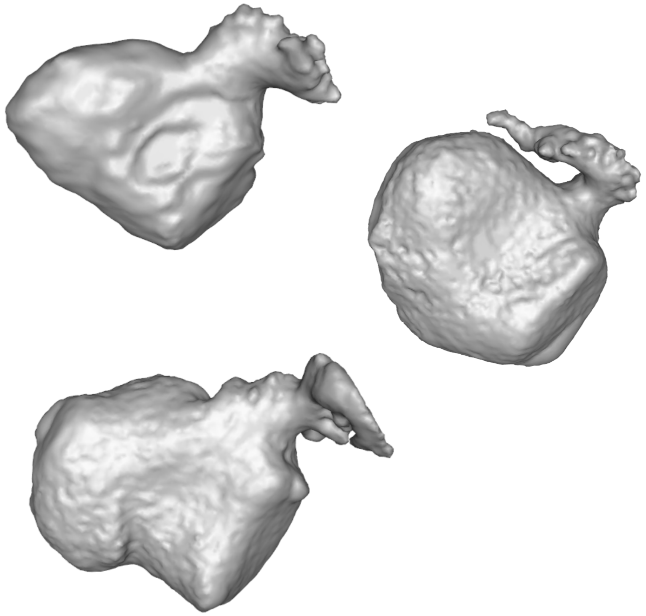

LA
- 106 surfaces
- Topology: spheres

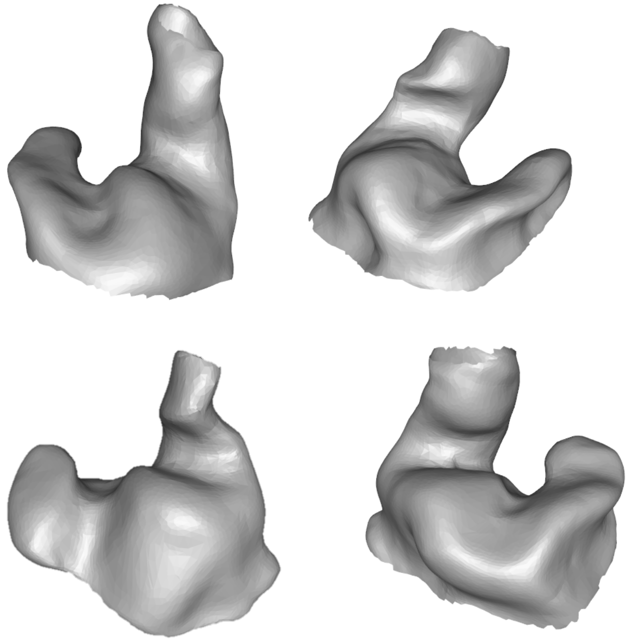

EARS
- 571 ears (259/312 left/right)
- Topology: Tubes

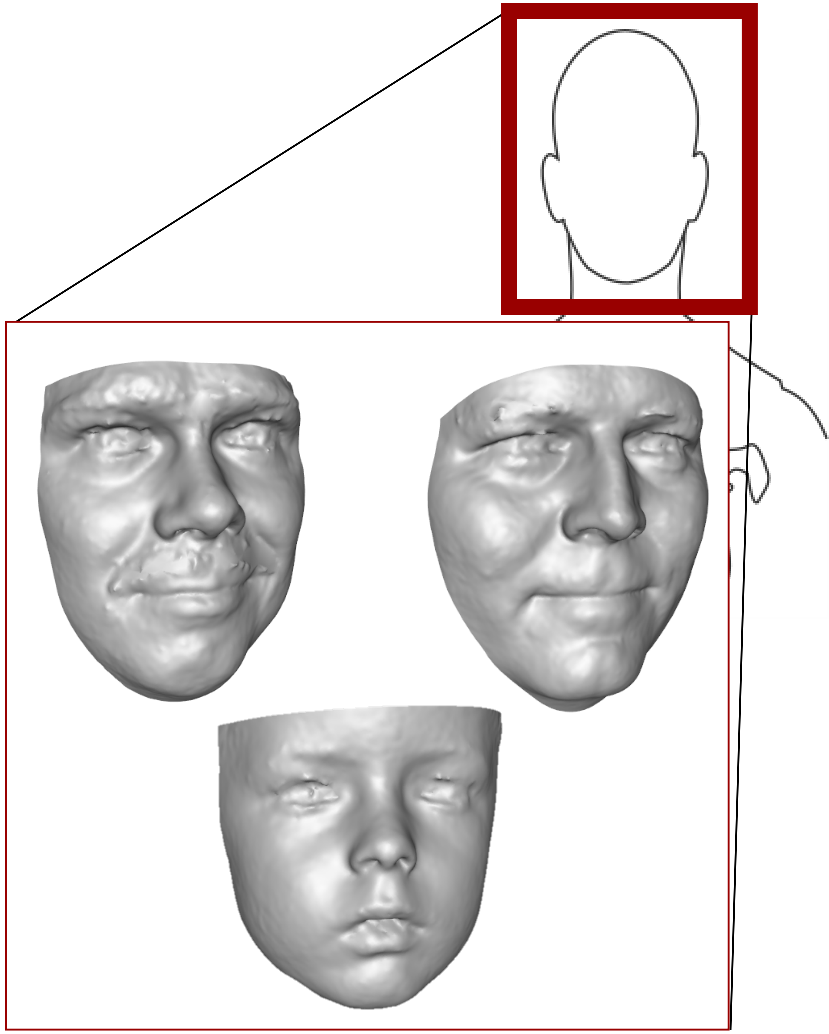

ESOF
- 394 faces
  (192/202 male/female, 0-84 years)
- Topology: Planes

Experiment 1: Reconstructing complex anatomies

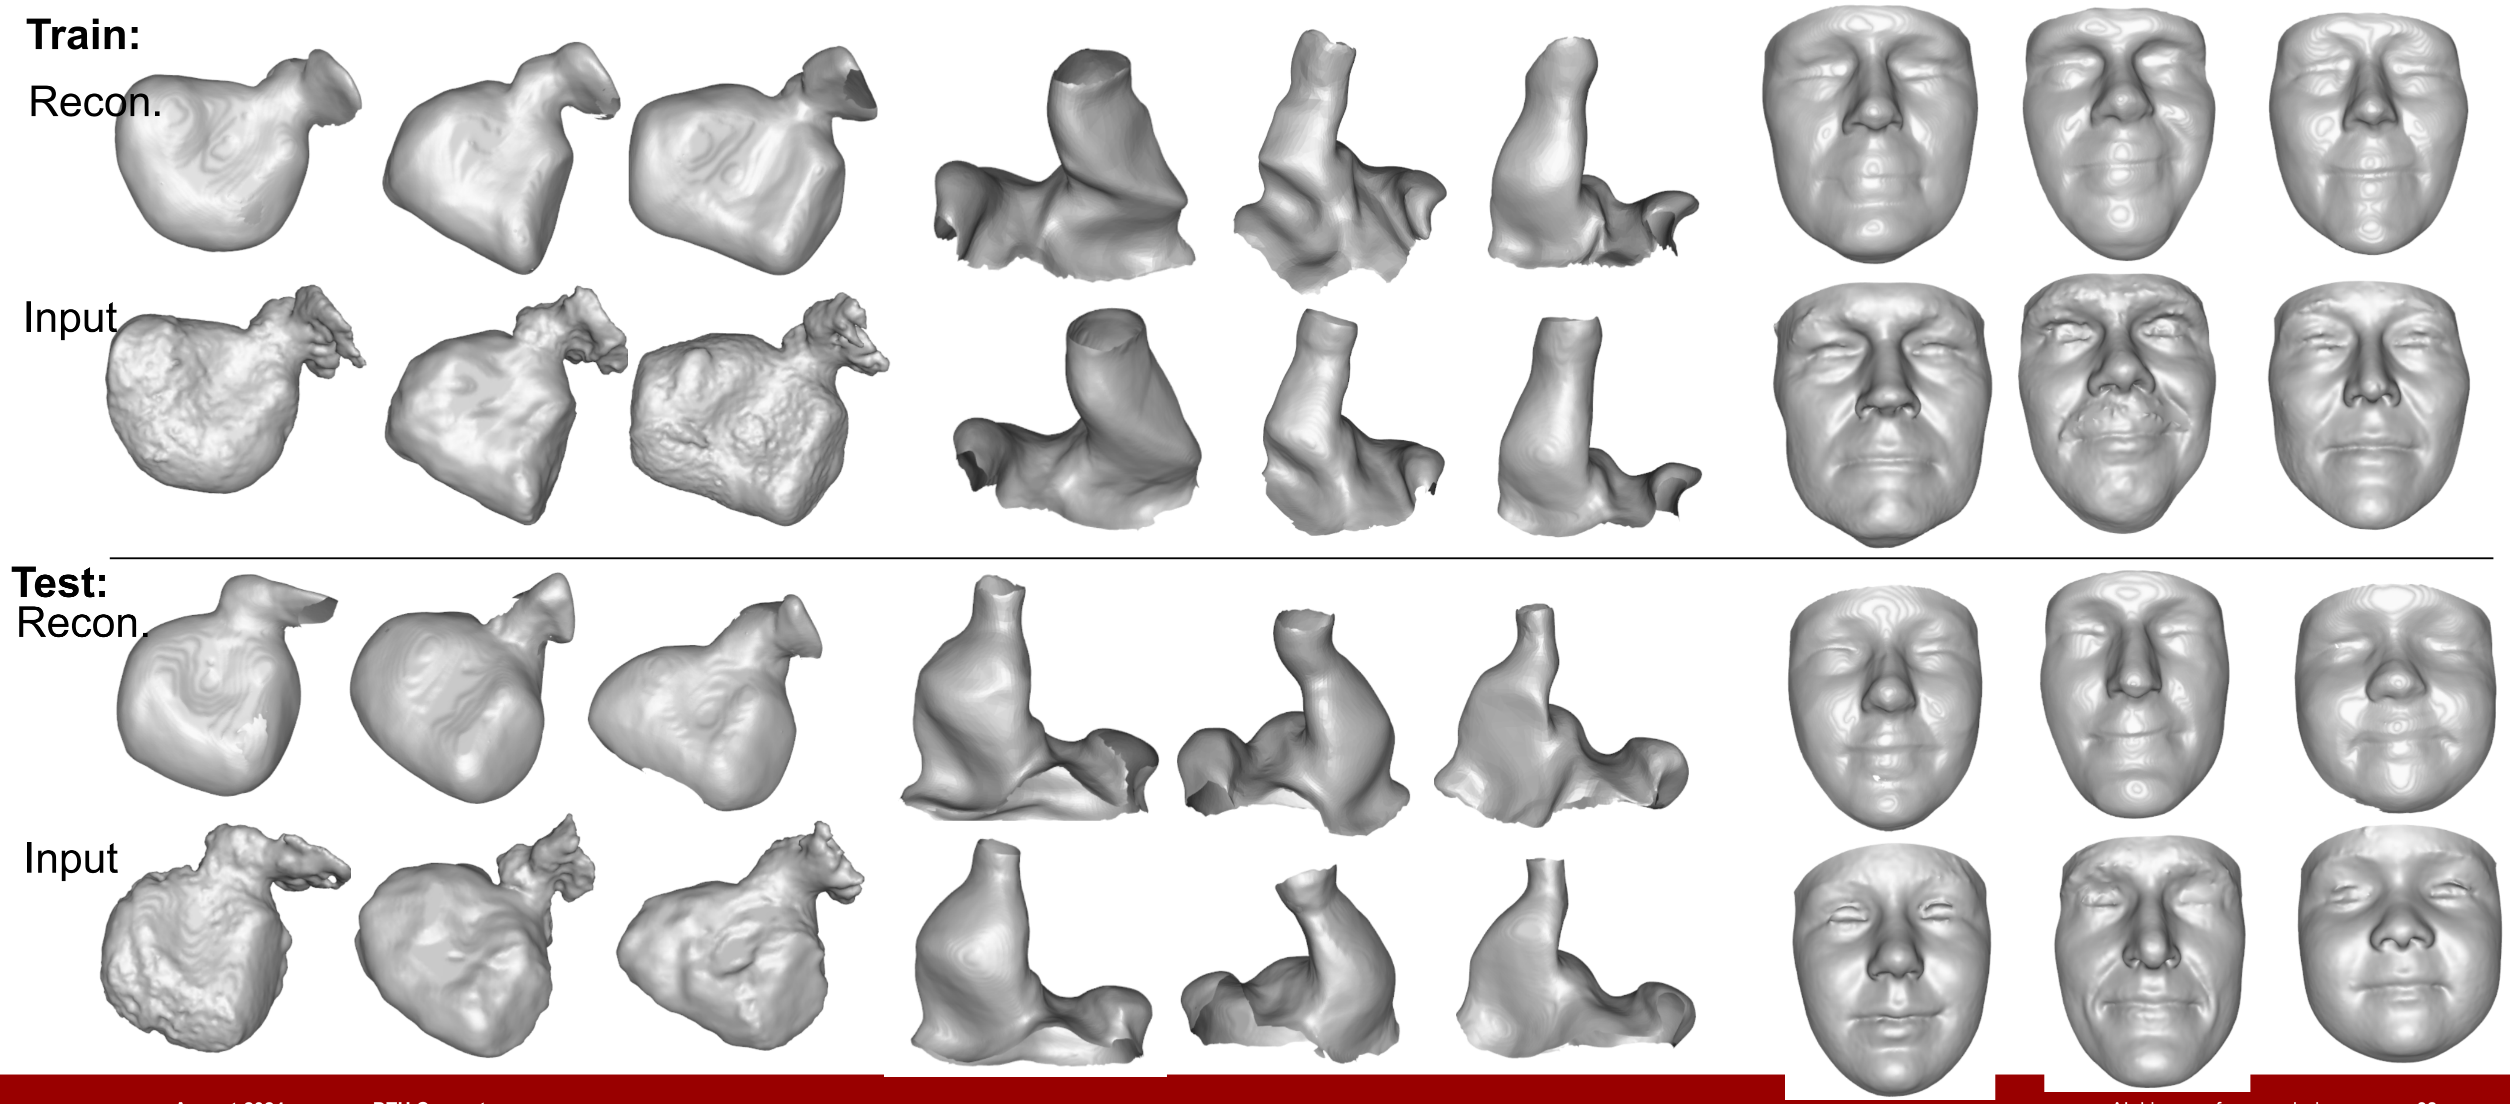

Experiment 1: Unsupervised clustering

- 128 dim. latent space

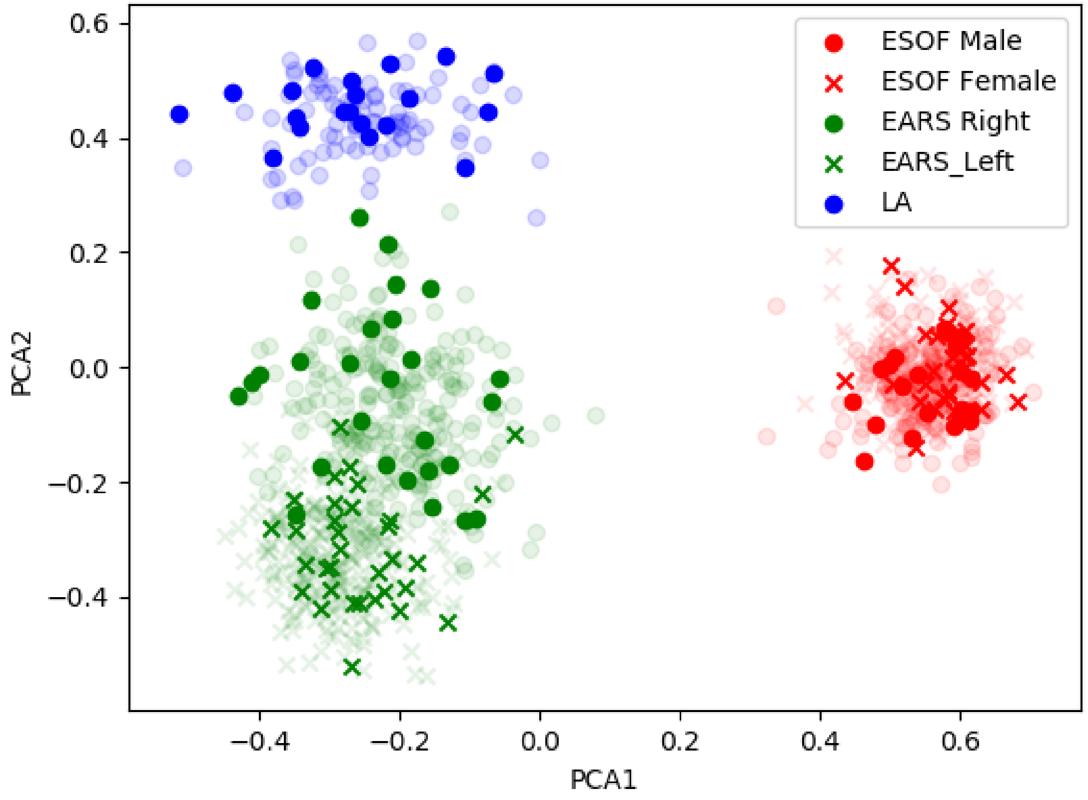

Experiment 1: Unsupervised clustering

- K-means with 3 clusters
  - Dataset accuracy: **99.23%**

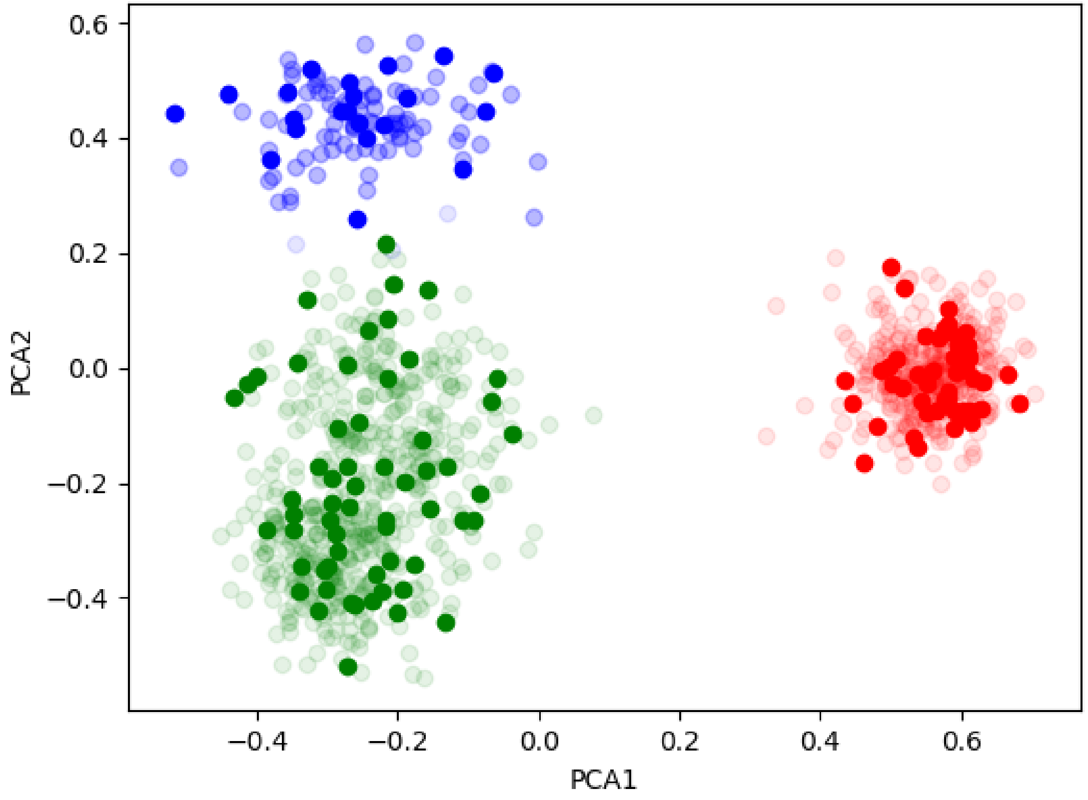

Experiment 1: Unsupervised clustering

- K-means with <u>4 clusters</u>
  - Dataset accuracy: **100%**
  - EARS left/right accuracy: **100%**

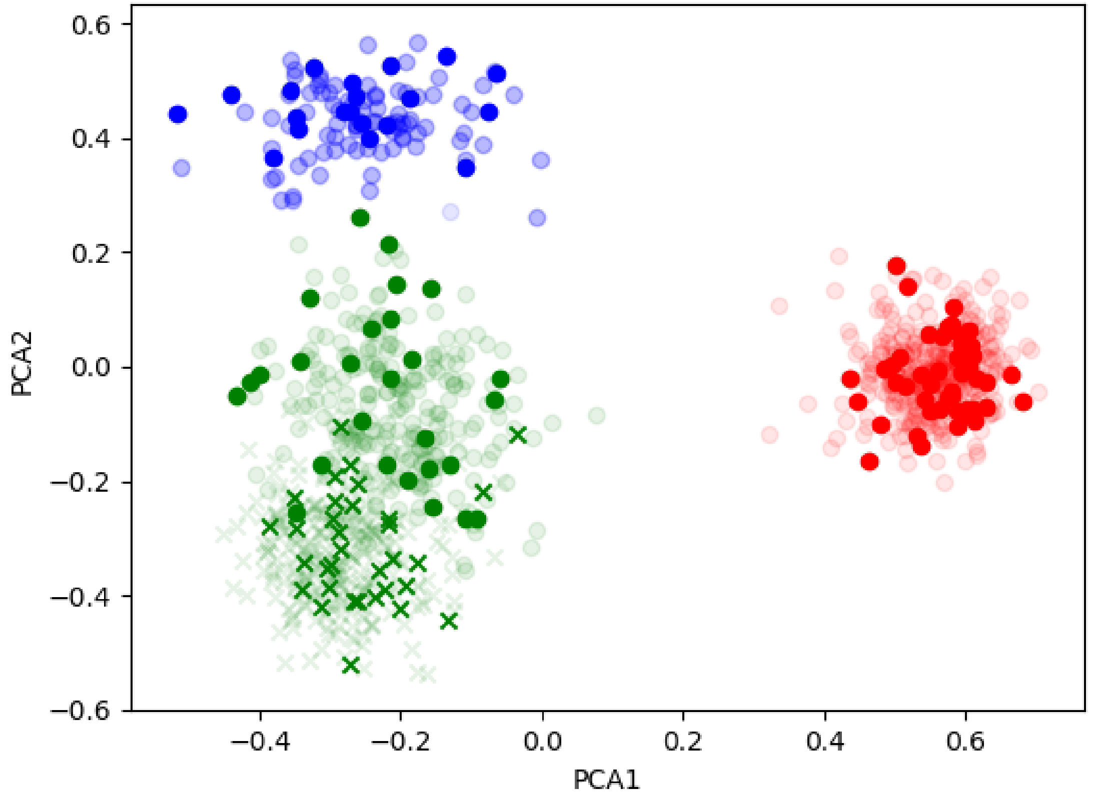

Summary – neural (un)signed distance functions

- Neural unsigned distance functions can represent complex anatomies with arbitrary topologies.
- The self-optimized latent space holds important global shape information and can be used to classify complex anatomies.

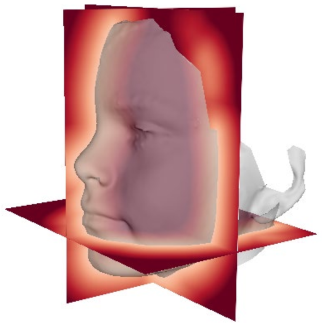

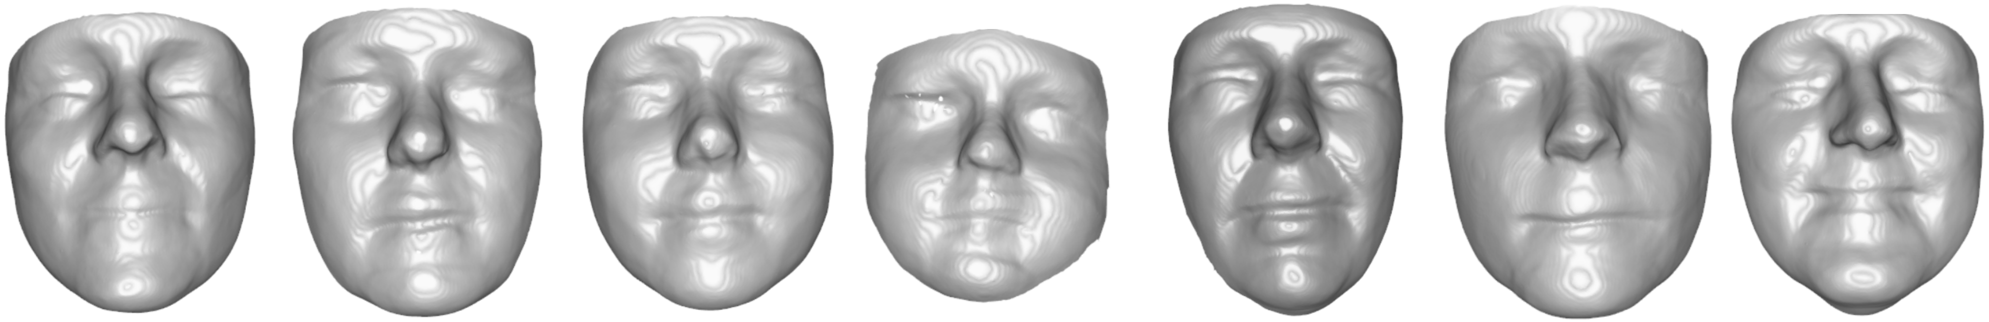

Cardiac movement

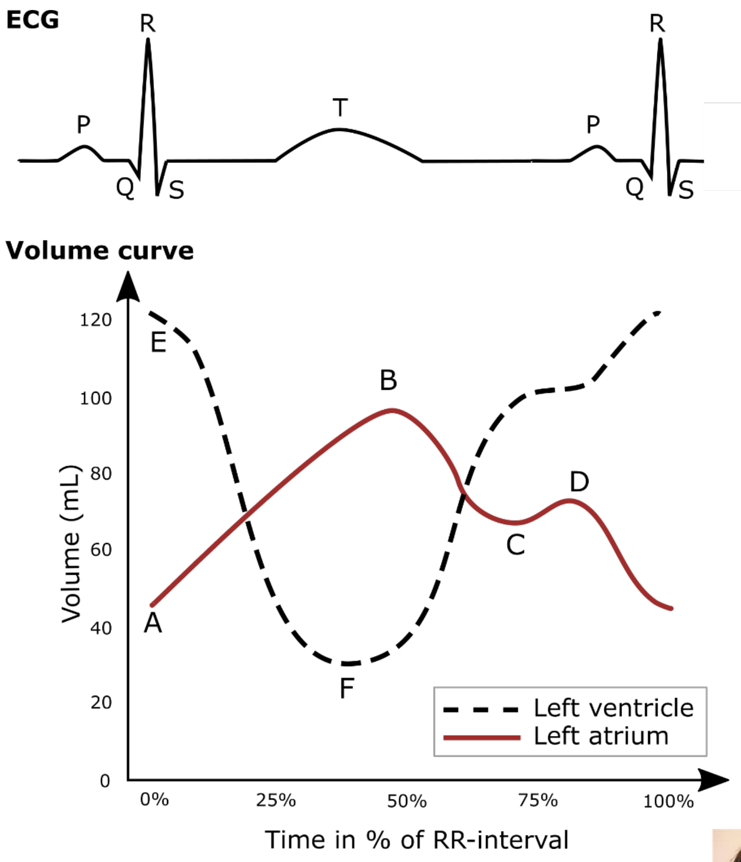

P	Atrial contraction (depolarization)
QRS	Ventricular contraction (depolarization)
T	Ventricular relaxation (depolarization)

Left ventricular function	
Left ventricular stroke volume	LVSV = E - F
Left ventricular ejection fraction	LVEF = (E - F) / F
Left atrial filling	
Cyclic change	CC = B - A
Fractional change	FC = (B - A) / B
Left atrial passive emptying	
Left atrial reserve volume	LARV = B - C
Left atrial active emptying	
Left atrial stroke volume	LASV = D - A
Left atrial ejection fraction	LAEF = (D - A) / D

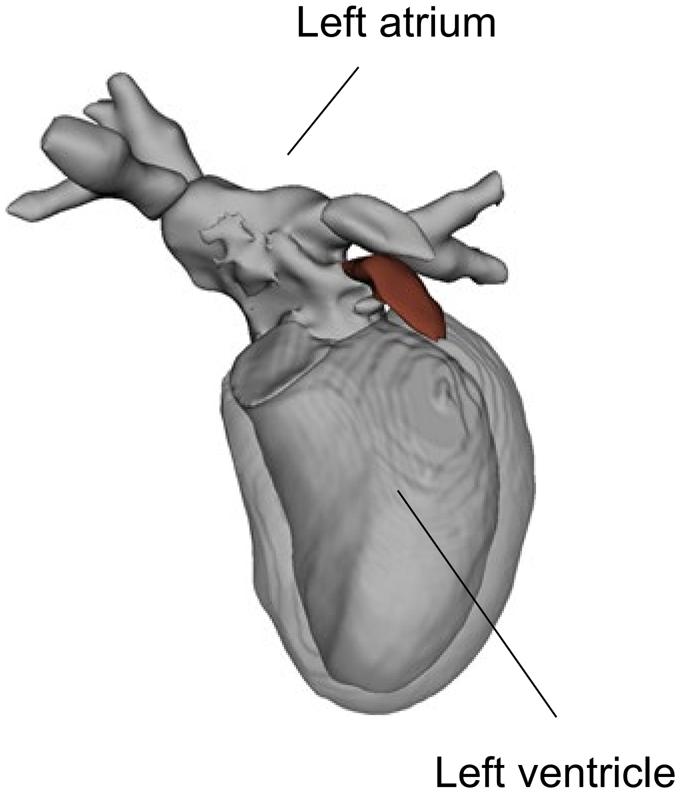

K. Sørensen et al. Spatio-temporal neural distance fields for conditional generative modeling of the heart. MICCAI 2024

Temporal CT scans

	LAA100	CFA
Image type	CCTA	4D CT
Number of images	108	679
Acquisition period	2010-2013	2016-2021
Image dimensions	$512 \times 512 \times$ (512-640)	$512 \times 512 \times$ (50-160)
Image resolution [mm]	$0.429 \times 0.429 \times 0.250$	(0.171-0.724) $\times$ (0.171-0.724) $\times$ (1.0-2.0)
Annotation method	Semi-manual (See below)	Automatic (See Paper C)
Additional information	Fully anonymized	Clinical demography (Age, gender SBP)

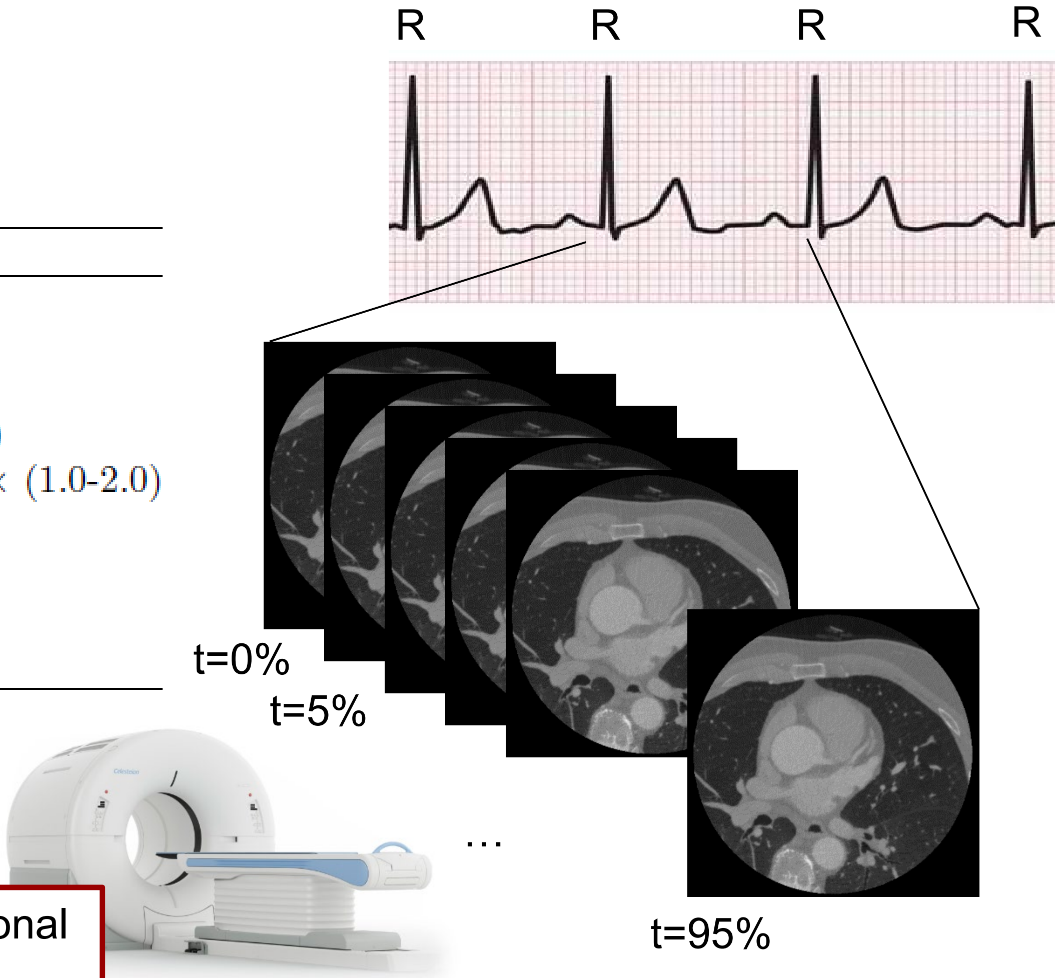

K. Sørensen et al. Spatio-temporal neural distance fields for conditional generative modeling of the heart. MICCAI 2024

Neural Distance fields

- Signed distance field (SDF)

Neural network

$$f_\theta(\boldsymbol{p}) \mapsto d, \qquad \boldsymbol{p} \in \mathbb{R}^3, d \in \mathbb{R}$$

- Surface as the zero-level isosurface

$$S = \{\boldsymbol{p} \in \mathbb{R}^3 | f_\theta(\boldsymbol{p}) = 0\}$$

- How do we approximate the SDF for a **set of N shapes**?

$$\underset{\theta,\{\boldsymbol{z}_n\}_{n=1}^{N}}{\arg\min} \sum_{n=1}^{N} \sum_{j=1}^{K} \mathcal{L}\left(f_\theta\left(\boldsymbol{z}_n \otimes \boldsymbol{p}_{n,k}\right), d_{n,k}\right) + \frac{1}{\sigma^2} \left|\left|\boldsymbol{z}_n\right|\right|_2^2$$

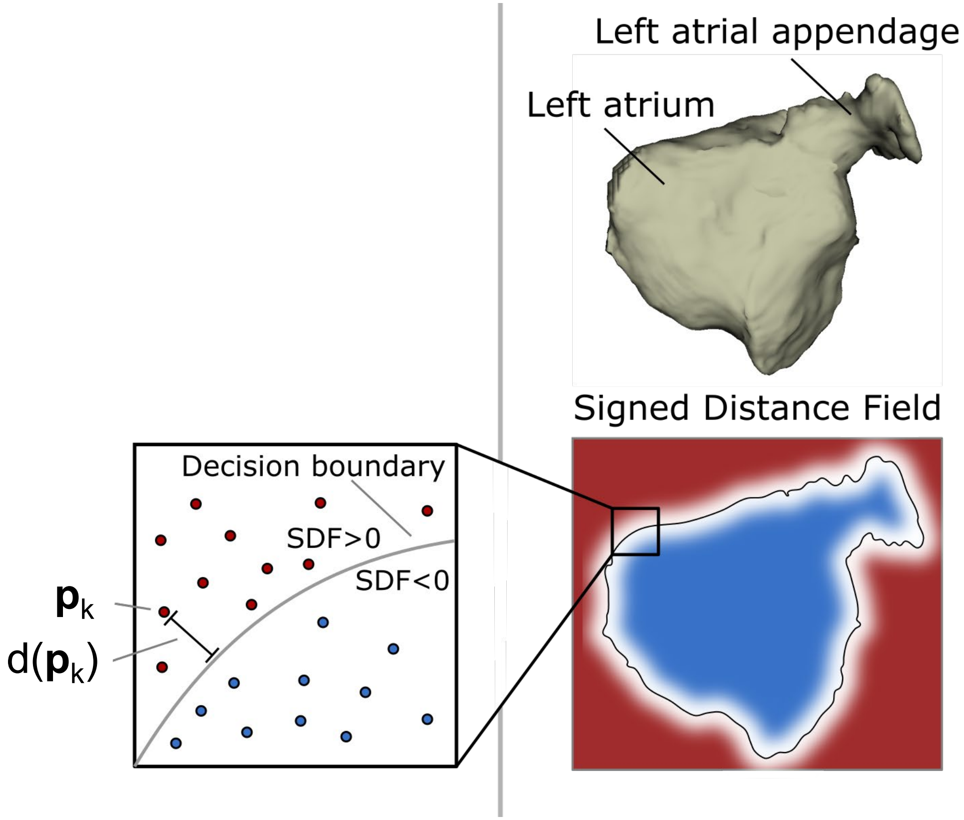

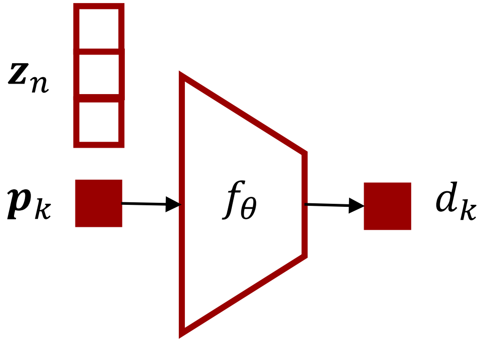

K. Sørensen et al. Spatio-temporal neural distance fields for conditional generative modeling of the heart. MICCAI 2024

Temporal Neural Distance fields

- Signed distance field (SDF)

Neural network

$$f_\theta(\boldsymbol{p}) \mapsto d, \qquad \boldsymbol{p} \in \mathbb{R}^3, d \in \mathbb{R}$$

- Surface as the zero-level isosurface

$$S = \{\boldsymbol{p} \in \mathbb{R}^3 | f_\theta(\boldsymbol{p}) = 0\}$$

- How do we approximate the SDF for a set of N **temporal sequences**?

$$\underset{\theta,\{\boldsymbol{z}_n\}_{n=1}^{N}}{\arg\min} \sum_{n=1}^{N} \sum_{j=1}^{K} \sum_{t=0}^{T} \mathcal{L}\big(f_\theta\big(\boldsymbol{z}_n \otimes (\boldsymbol{p}_{n,k}, t_{n,k})\big), d_{n,k,t}\big) + \frac{1}{\sigma^2} ||\boldsymbol{z}_n||_2^2$$

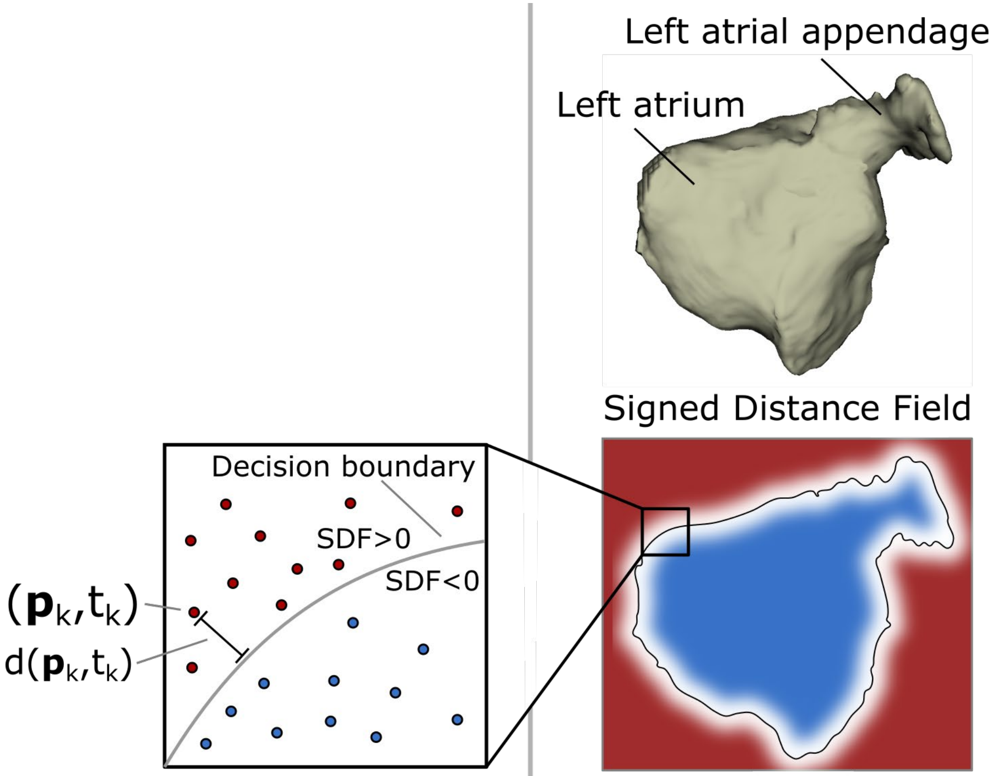

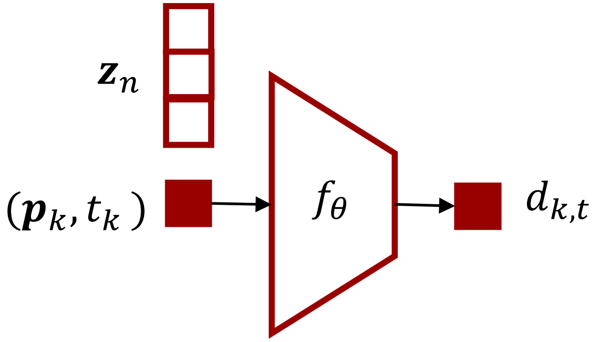

K. Sørensen et al. Spatio-temporal neural distance fields for conditional generative modeling of the heart. MICCAI 2024

Temporal Sequence completion

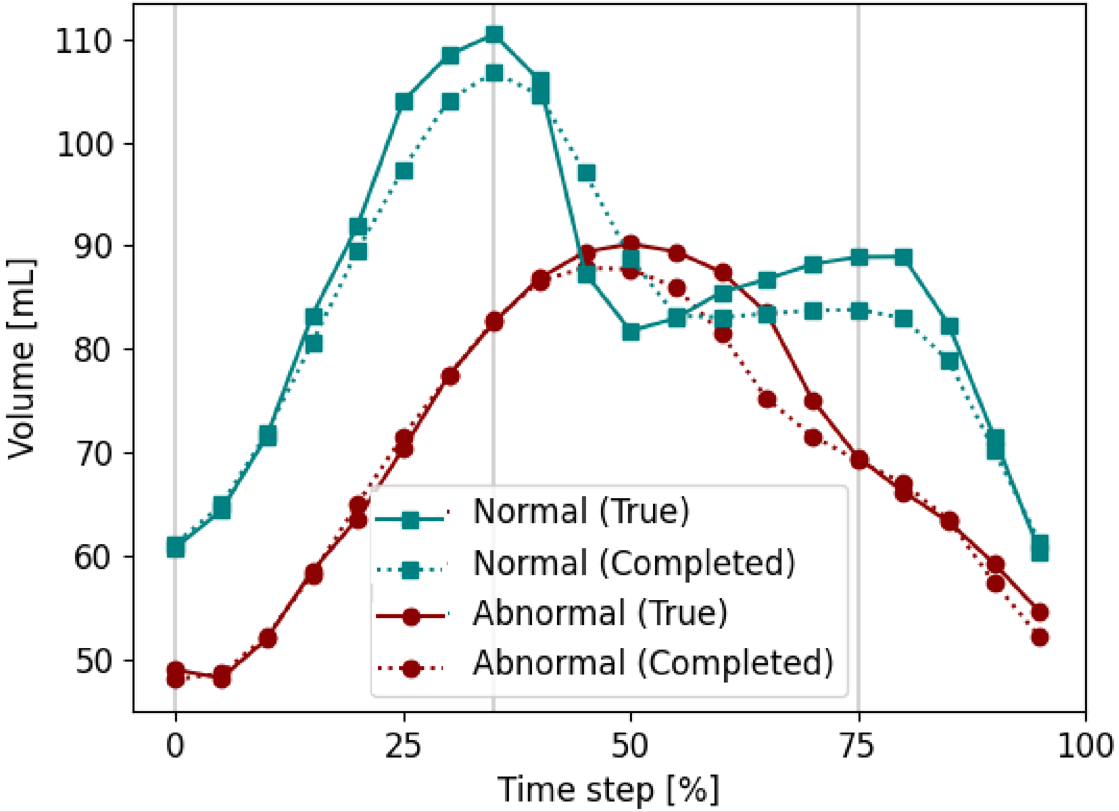

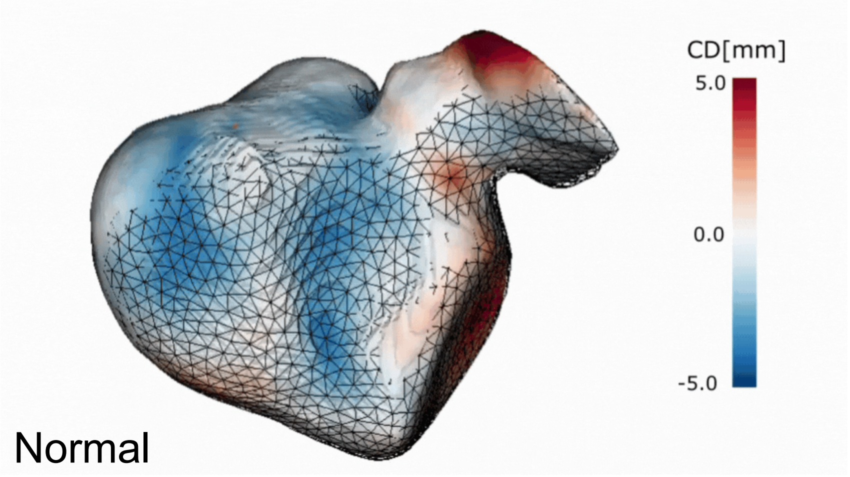

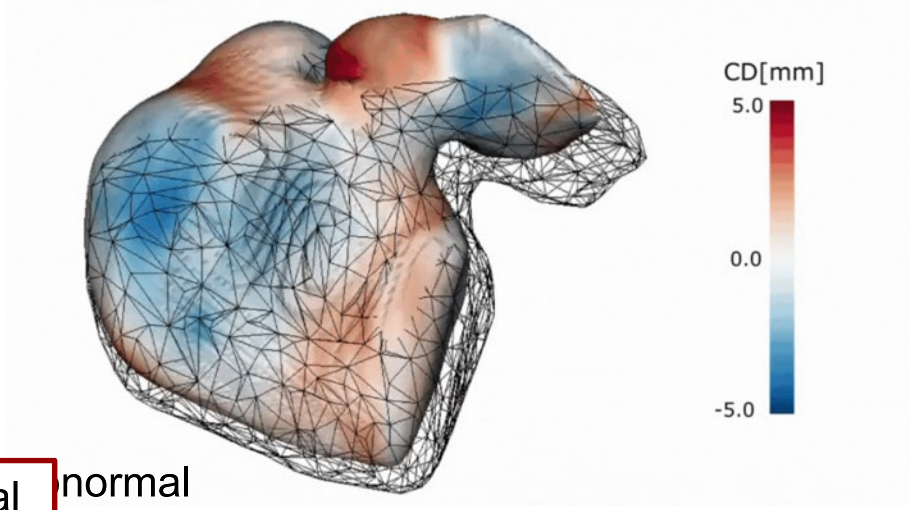

K. Sørensen et al. Spatio-temporal neural distance fields for conditional generative modeling of the heart. MICCAI 2024

Summary

- It is important to choose a representation that represents your data as well as possible
  - Geometry / topology / level of details
- You should be aware of the limitations of the representation
  - Rotational invariance
- What type of operations can you do with the representations
  - Statistical measures
  - Convolutions
  - Pooling

- Neural representations like neural distance fields is a major research focus
  - Come with both benefits and weaknesses