Shape priors improve lung segmentation in ARDS-afflicted subjects

Wrik Bhadra[1], Muhammad F. A. Chaudhary, PhD[1], Sarah E. Gerard, PhD[1]
[1]Department of Biomedical Engineering, The University of Iowa

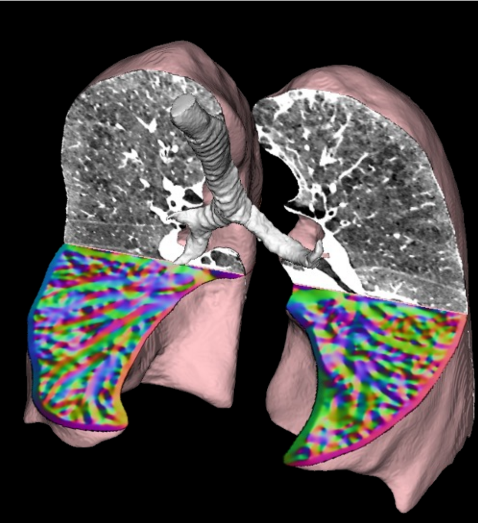

Pulmonary Imaging & Computer Vision Lab

Background

Acute respiratory distress syndrome (ARDS) is a life-threatening condition despite advancements in medical care. Computed tomography (CT) is used to diagnose ARDS and study lung heterogeneity. Automated segmentation of lungs in CT images is necessary for efficient quantitative-CT analysis. However, **existing state-of-the-art lung segmentation algorithms fail in severe cases of ARDS**, where it is difficult to discriminate consolidated lung from surrounding soft tissue.

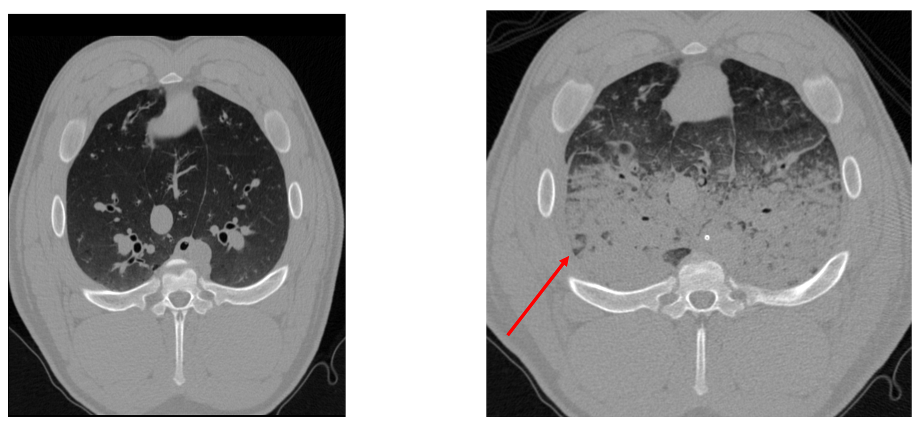

CT scans of two porcine subjects. Healthy lungs (left) appear darker on CT compared to injured lungs which have regions that appear brighter (right). The intensity of the injured region is similar to the surrounding soft tissue, making it difficult to identify the lung boundary.

Objective

Incorporate learnable anatomical shape priors in a deep learning framework to improve automatic lung segmentation in CT images of multiple species with ARDS.

Methods

First (A), a variational autoencoder (VAE) model is trained using the ground truth (GT) segmentations to encode lung shapes in a compressed latent space (z). Next (B), a CNN is trained that uses the learned shape priors to generate predicted (pred) lung segmentations. For comparison, high resolution and low resolution UNet models are trained without shape priors.

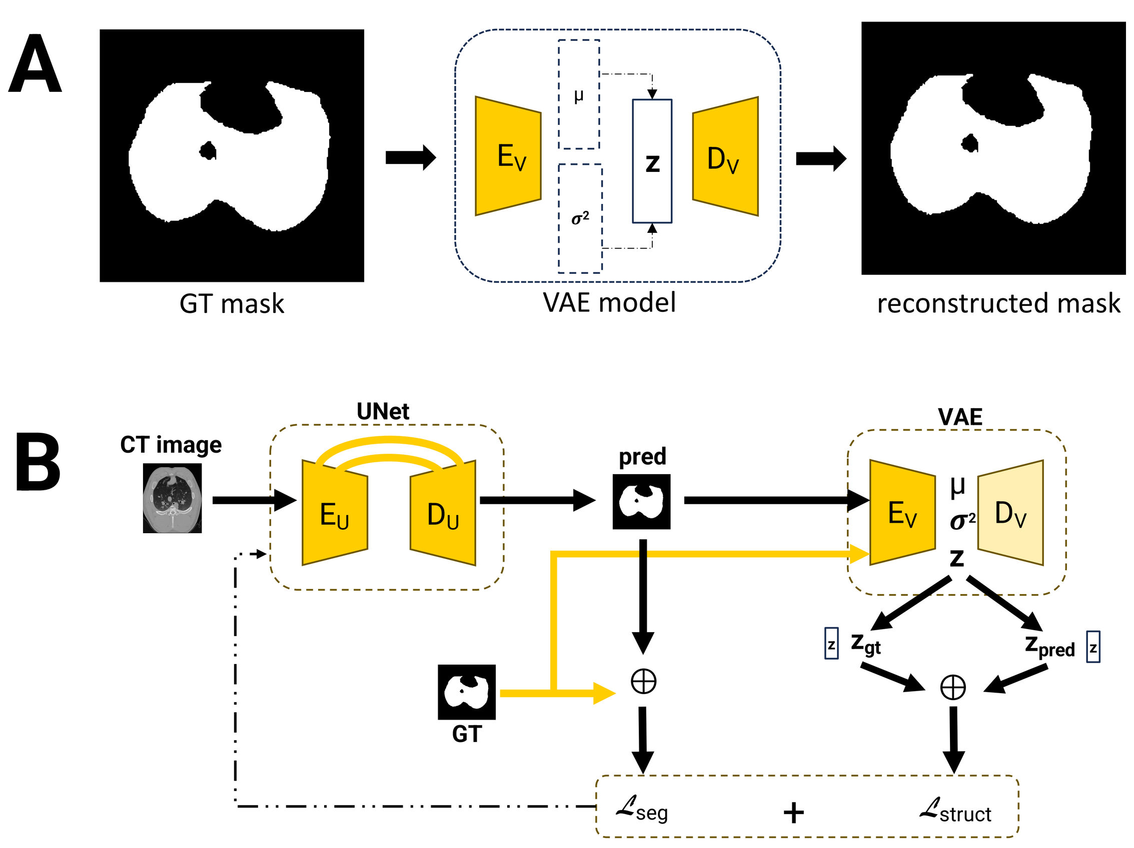

Schematic diagram of the proposed method

Dataset: For model training and testing, we used 508 and 117 CT images, respectively. Half the test dataset were considered severely injured cases defined by percent of non-aerated lung volume.

Pre-processing: CT images were spatially resampled to 1 mm isotropic voxels for high resolution (HR) models and 3 mm for low resolution (LR) models. Intensities were clipped in the range -1024 to 200 Hounsfield units (HU). During training, images were randomly cropped to 96 x 96 x 96 for LR and 96 x 192 x 192 for HR.

Results

For **severe** ARDS cases, the UNet-VAE-128 (proposed model) performed significantly better than the high-resolution UNet (UNet-HR), while UNet-LR is better than UNet-HR since the LR images have more global context. For **mild** cases, UNet-HR performs better than both UNet-LR and UNet-VAE-128 models, but the proposed model performs better than UNet-LR.

Conclusions

The results demonstrate VAEs can learn meaningful shape priors. The shape priors can be utilized when training a CNN for segmentation such that the network predicts anatomically correct lung shapes in severe cases of ARDS, where it becomes extremely challenging to discriminate injured lung from surrounding soft tissues.

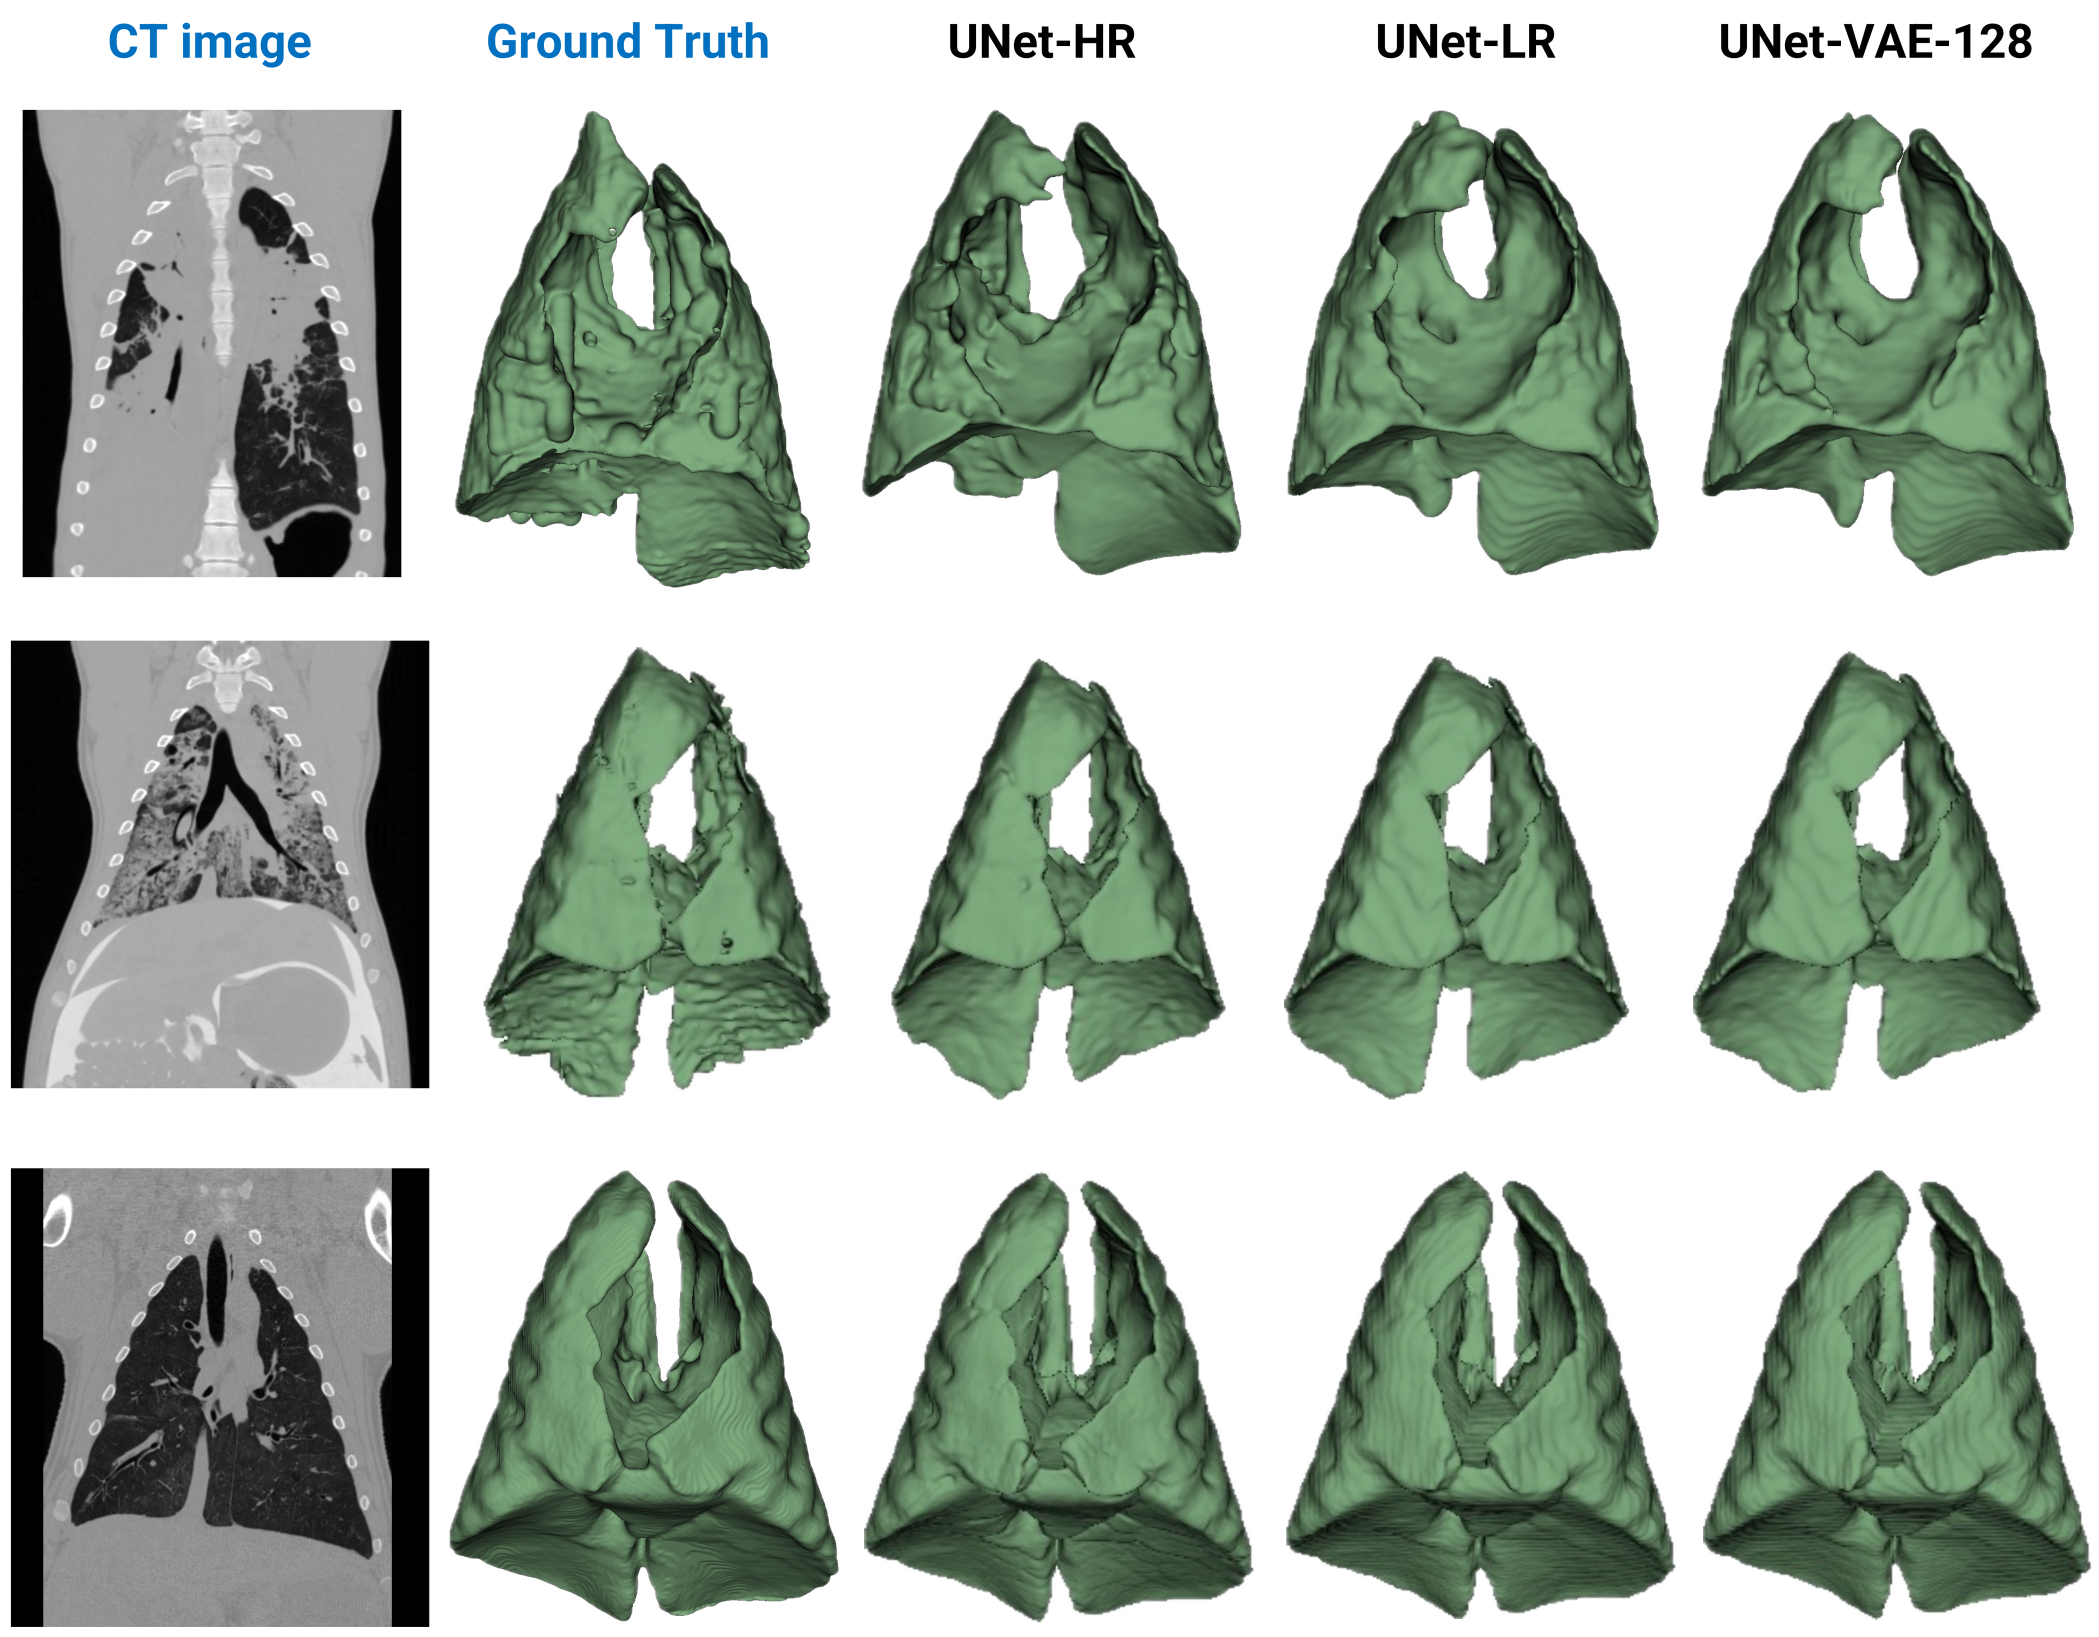

Segmentations for the three models (columns) for three porcine subjects (rows). CT image and ground truth included for reference.

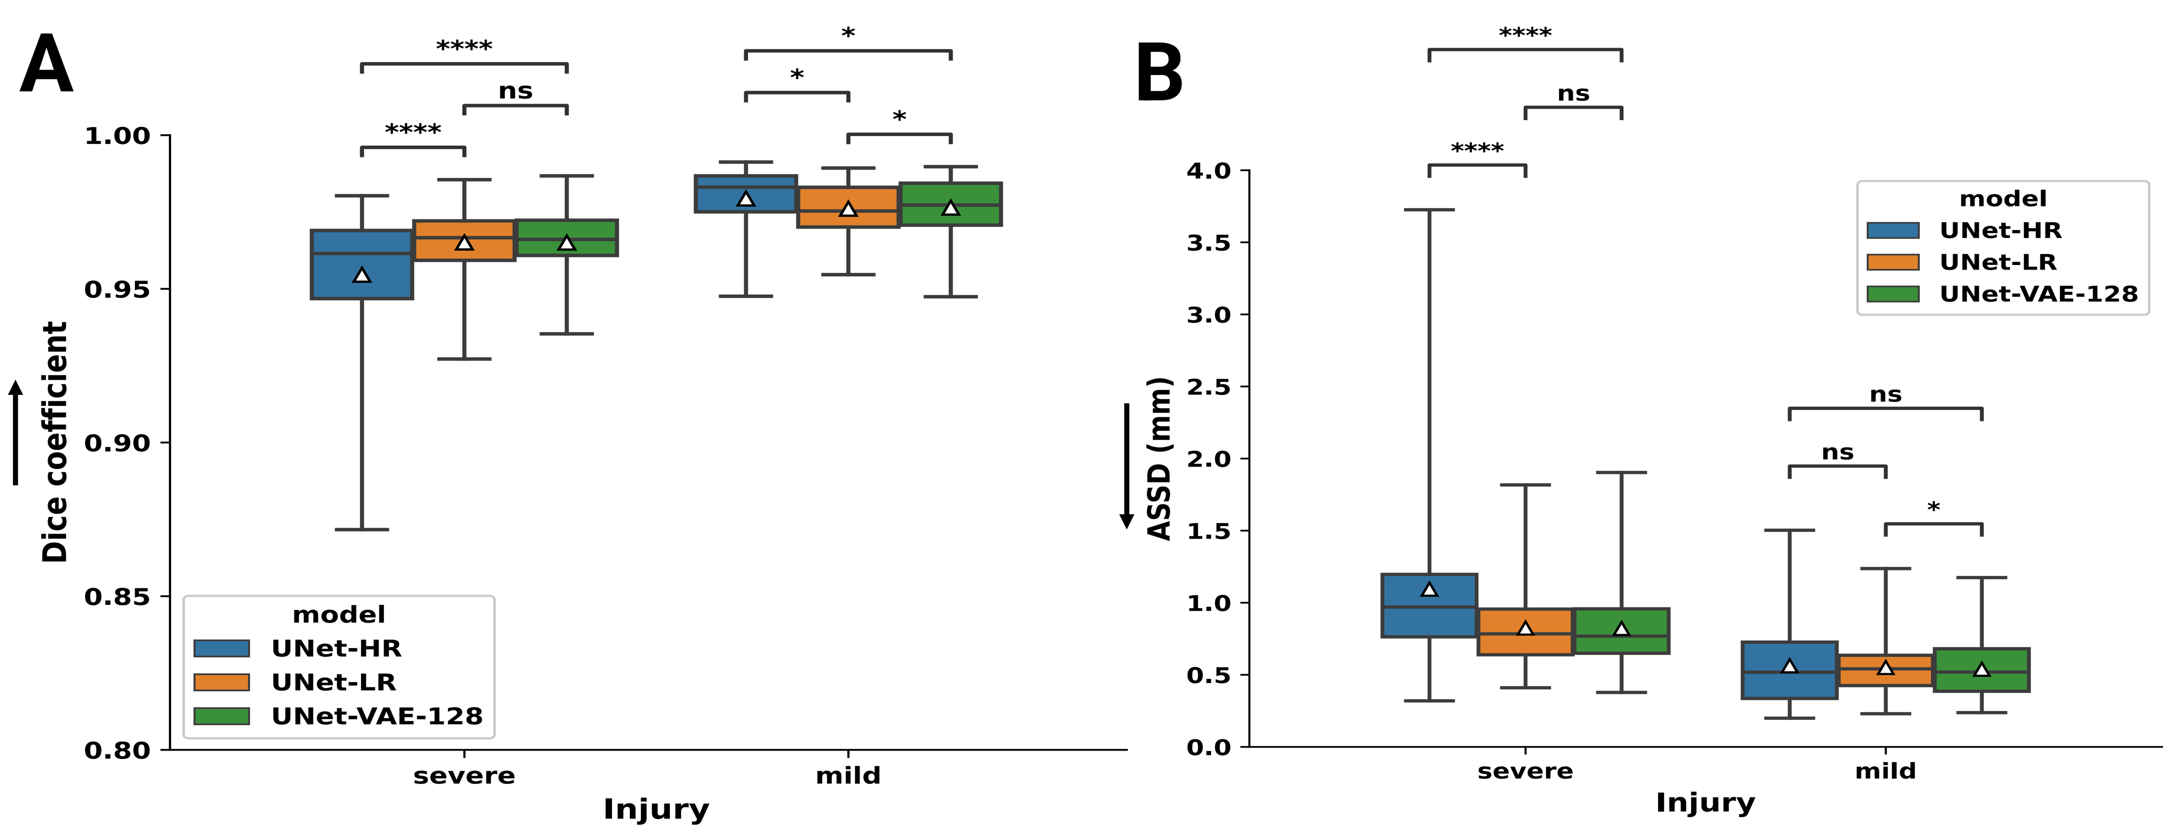

Quantitative results showing distributions of (A) Dice coefficient and (B) average symmetric surface distance (ASSD) on the test dataset, which comprises multiple animal species (porcine, ovine, and canine). Wilcoxon signed rank test is used to compare models. Significant results are annotated with "*" and non-significant results are annotated with "ns".

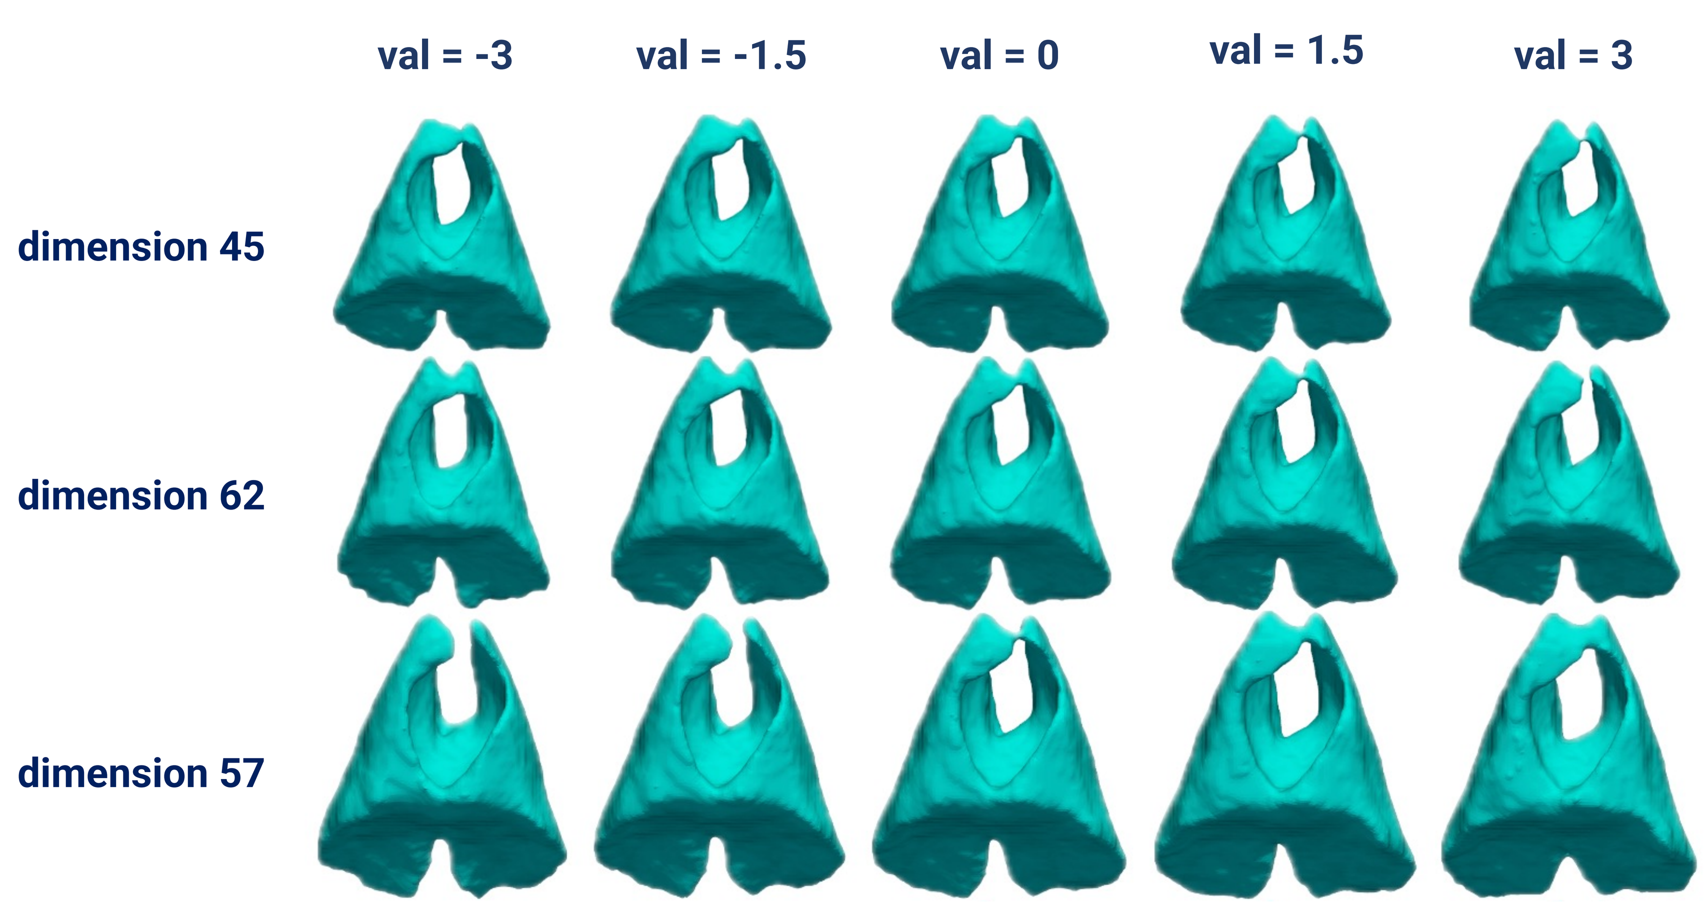

Scan to view latent traversals during VAE model training

Traversing the VAE latent space. Three dimensions out of a 128-dimensional latent space are chosen based on highest to lowest Kullback-Leibler divergence values. The changes observed by traversing different latent dimensions demonstrate that shape and size features as encoded by the VAE.