Artificial Intelligence for Science, Industry, and Society

AISIS - 2019

Variational Autoencoders for Medical Imaging

Juan J. Cerrolaza, PhD

accenture technology

artificial intelligence IBERIA

BioMedIA Imperial College London

Imperial College London

KING'S College LONDON

Children's National

Sheikh Zayed Institute for Pediatric Surgical Innovation

Part of the Children's National Health System

accenture

20th – 25th October 2019

Mexico City - UNAM

Variational Autoencoders - Background

Neural Network Perspective

Dimensionality Reduction Method

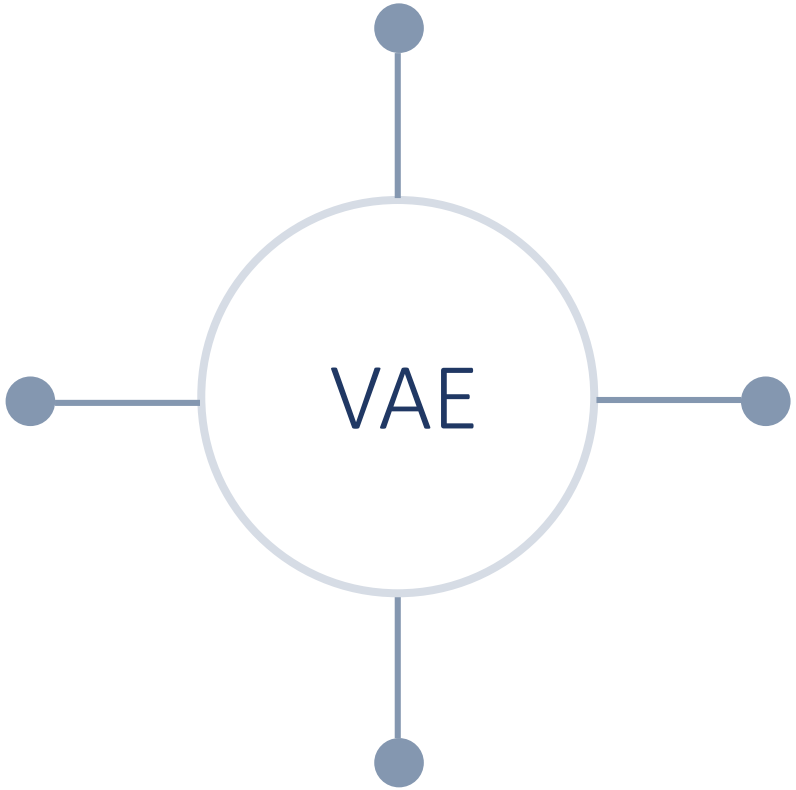

Generative Model

Probability Model Perspective

Variational Autoencoders - Background

Dimensionality Reduction Method

Generative Model

Variational Autoencoders - Background

Neural Network Perspective

Probability Model Perspective

Variational Autoencoders - Background

Neural Network Perspective

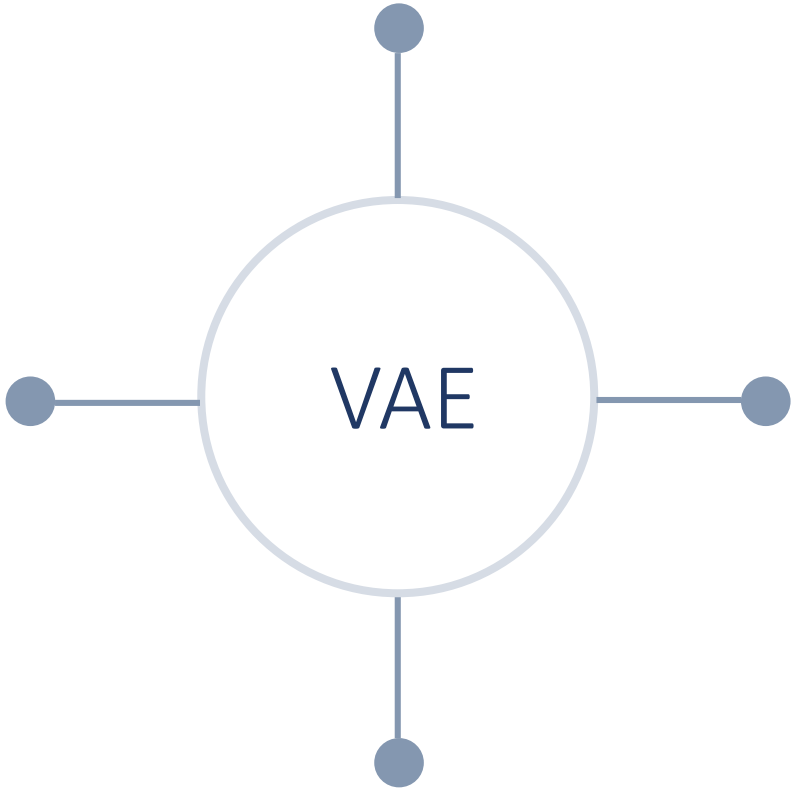

Dimensionality Reduction Method

Generative Model

Probability Model Perspective

Variational Autoencoders - Background

Autoencoder Architecture

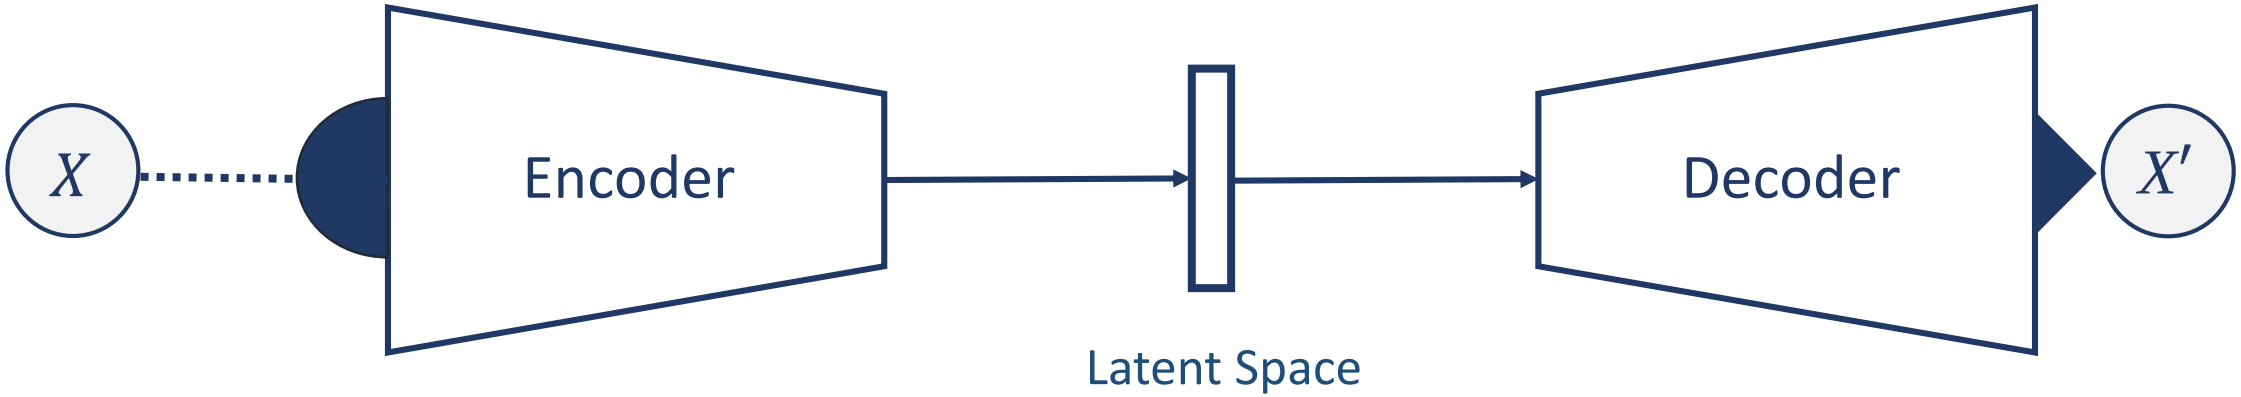

Loss Function	$\|X - X'\|$

Variational Autoencoders - Background

Variational Autoencoder – Autonecoder with a “twist”

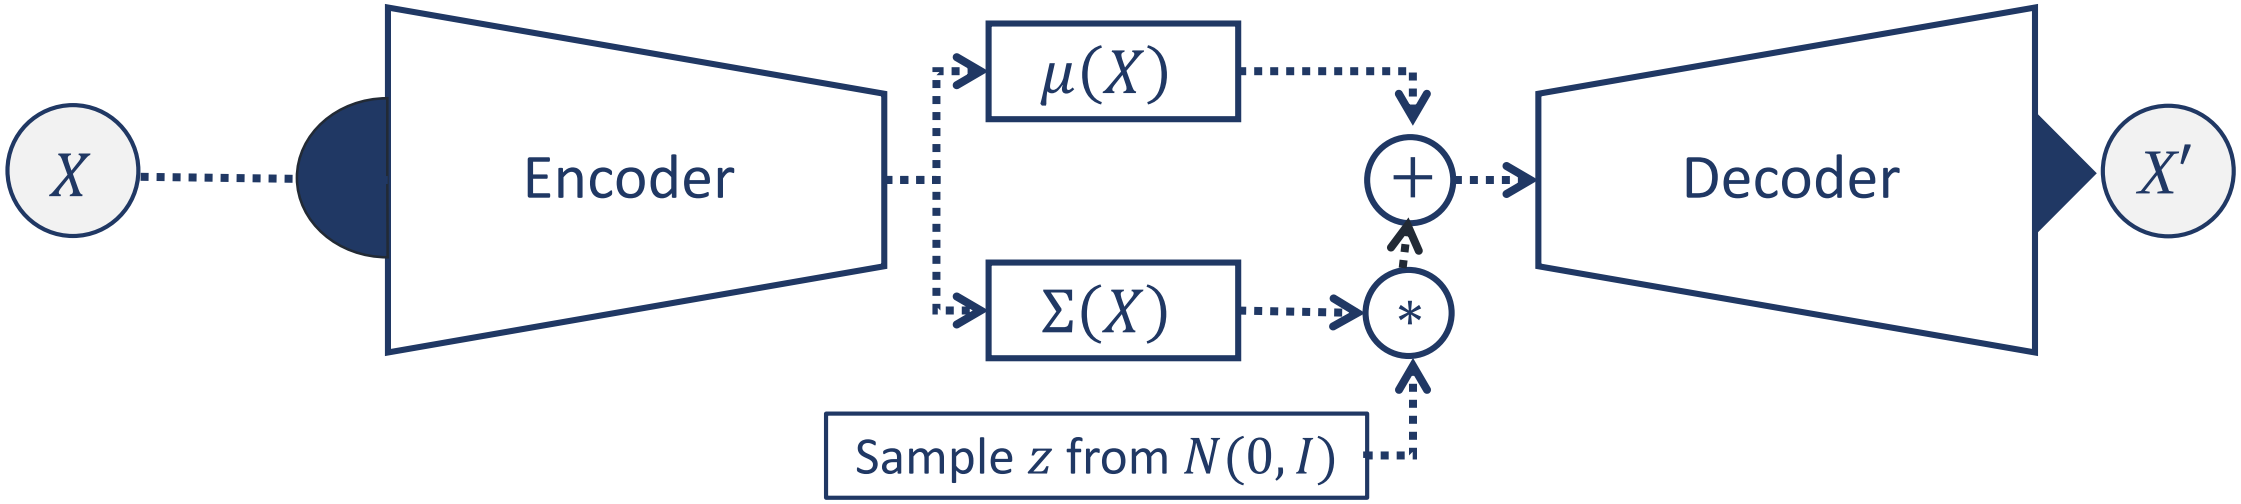

Loss Function $\|X - X'\| + \lambda \cdot KL\left[N\left(\mu(X), \Sigma(X)\right) \| N(0, I)\right]$

Variational Autoencoders - Background

Generative Model

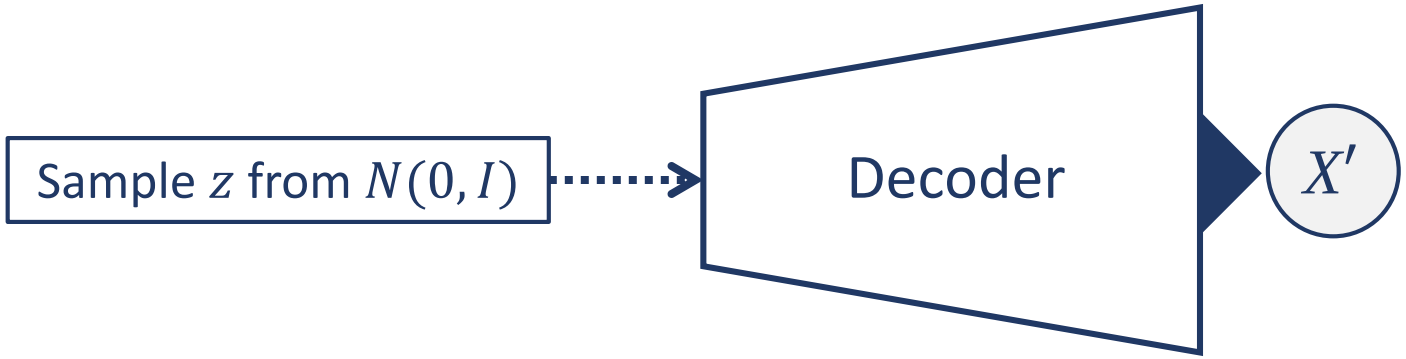

Loss Function $\|X - X'\| + \lambda \cdot KL\left[N\left(\mu(X), \Sigma(X)\right) \| N(0, I)\right]$

Variational Autoencoders - Background

Neural Network Perspective

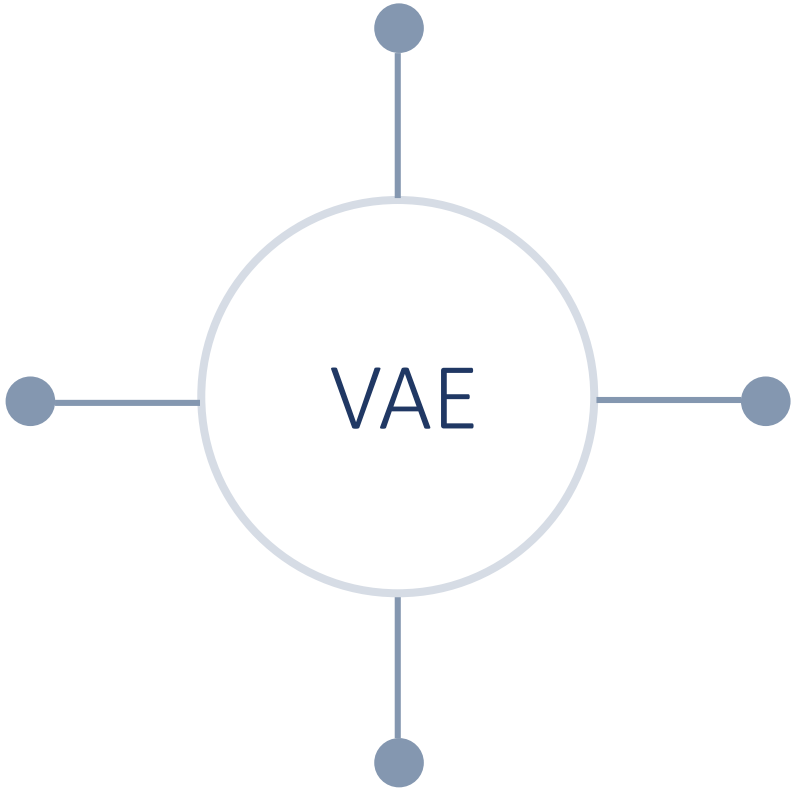

Dimensionality Reduction Method

Generative Model

Probability Model Perspective

Variational Autoencoders - Background

X: training dataset

z: latent variable (w. known prob. distribution, $z \sim N(0,I)$)

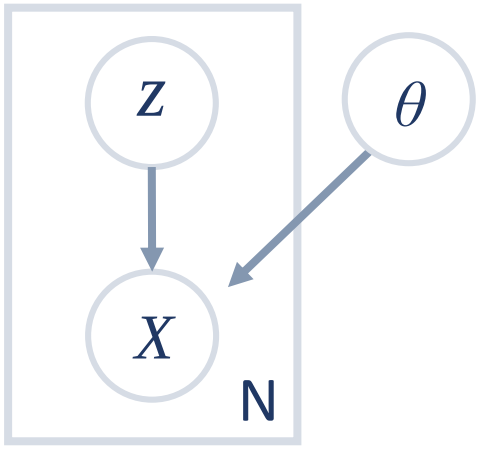

$$P(X) = \int P(X|z;\theta)P(z)dz \qquad (1)$$

Two problems:

1. How to define the latent variables (what information from X they represent)?
2. How to solve *(1)*? → Sampling over every possible z´s is not practical!

Variational Autoencoders - Background

1. How to define the latent variables (what information from X they represent)?

- Assuming a powerful (high-capacity) mapping function θ, z can be drawn from a simple distribution ($N(0,I)$) → $P(z) = N(z|0, I)$

- $f(z; \theta)$ is a deep multi-layer neural network which will map z to the corresponding X.

Variational Autoencoders - Background

2. How solve $P(X)$

- $P(X) = E_{z \sim P(z)} P(X|z)$ Difficult in practice to infer $P(X|z)$ without sampling a large number of z values.

To learn a new function Q, that takes X and gives a distribution over z that are likely to produce X.

- $E_{z \sim Q(z)} P(X|z)$ More practical

$$E_{z \sim Q(z)}[\log P(X|z)] \overset{\text{Bayes' rule}}{=} E_{z \sim Q(z)}[\log P(z|X) + \log P(X) - \log P(z)]$$

2. How solve $P(X)$

$$E_{z\sim Q(z)}[\log P(X|z)] = E_{z\sim Q(z)}[\log P(z|X) + \log P(X) - \log P(z)]$$

Rearranging and subtracting $E_{z\sim Q(z)}[\log Q(z|X)]$ on both sides…

$$\log P(X) - E_{z\sim Q(z)}[\log Q(z|X) - \log P(z|X)] =$$

$$E_{z\sim Q(z)}[\log P(X|z)] - E_{z\sim Q(z)}[\log Q(z|X) - \log P(z)]$$

Definition of KL divergence $E_{z\sim Q}[\log a - \log b] = KL[a||b]$

$$\log P(X) - KL[Q(z|X)\|P(z|X)] = E_{z\sim Q(z)}[\log P(X|z)] - KL[Q(z|X)\|P(z)]$$

Variational Autoencoders - Background

2. How solve $P(X)$

$$\log P(X) - KL[Q(z|X)\|P(z|X)] = E_{z\sim Q(z)}[\log P(X|z)] - KL[Q(z|X)\|P(z)]$$

★ $\log P(X)$

Lower bound (KL > 0): $KL[Q(z|X)\|P(z|X)]$

Reconstruction error (Decoder): $E_{z\sim Q(z)}[\log P(X|z)]$

$KL[Q(z|X)\|P(z)]$

$Q(z|X)$ Gaussian $\mathcal{N}(\mu, \Sigma)$

KL-div. of two Gaussians → closed form

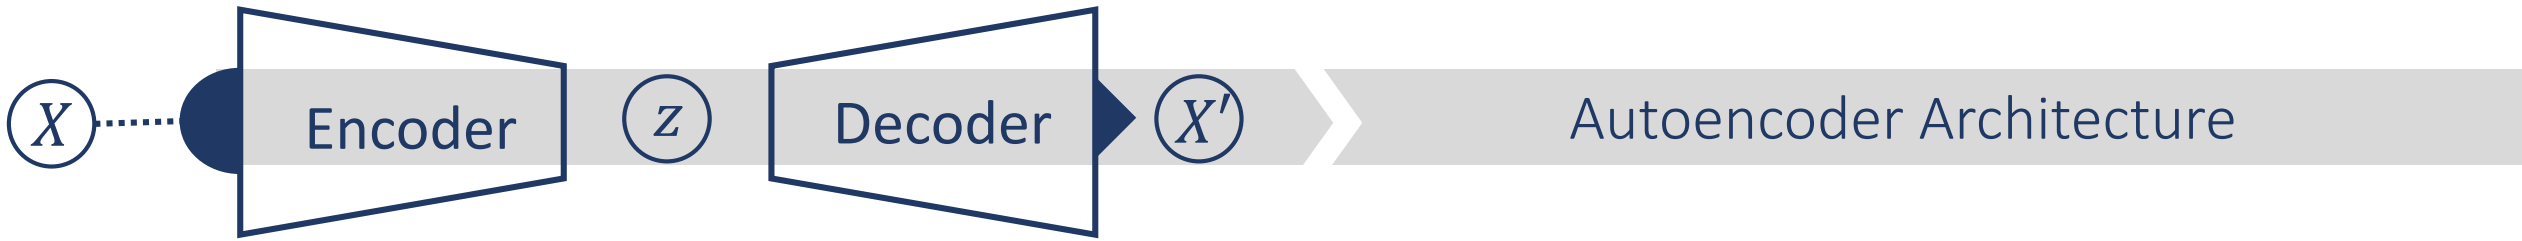

Variational Autoencoders - Background

Reparameterization trick

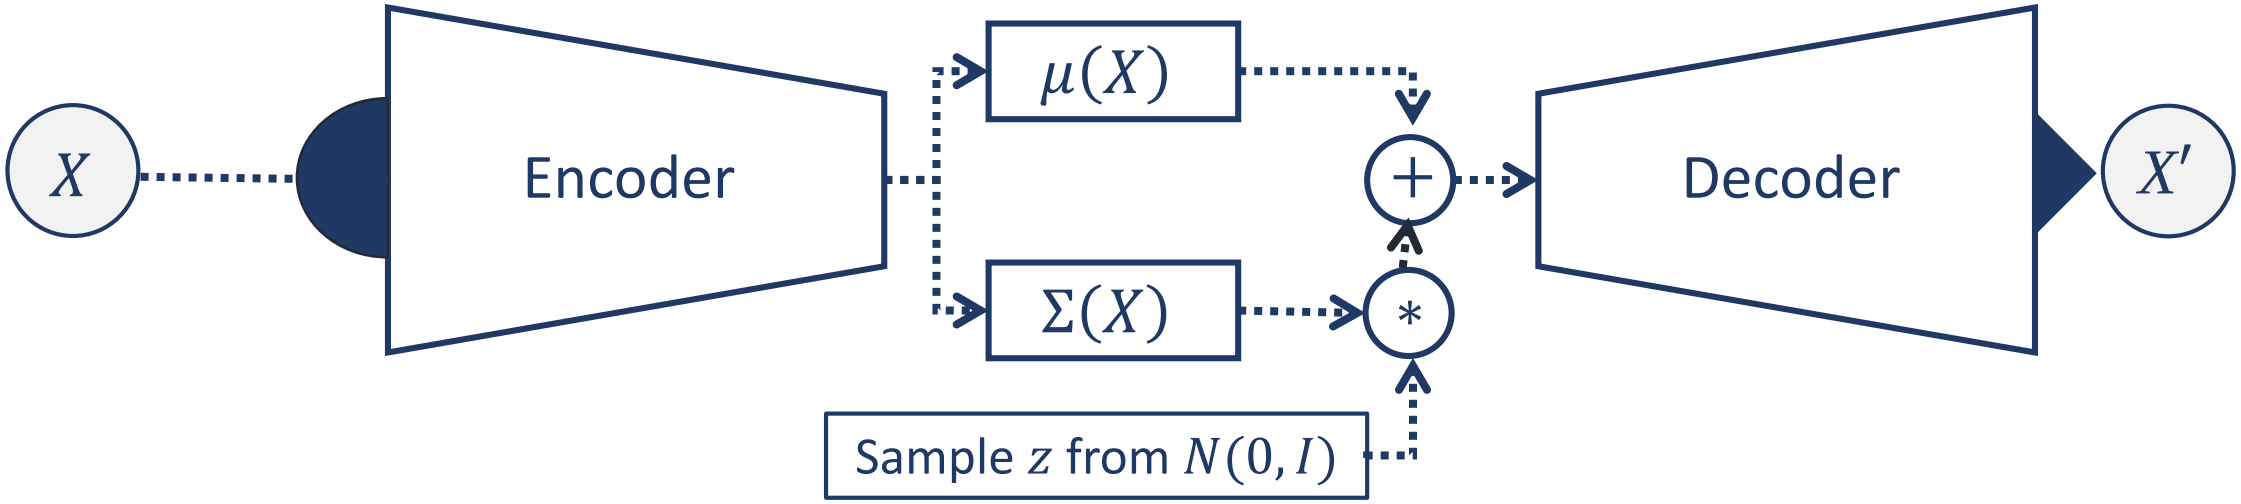

Loss Function $\|X - X'\| + \lambda \cdot KL\left[N\left(\mu(X), \Sigma(X)\right) \| N(0, I)\right]$

Variational Autoencoders - Background

Generative Model

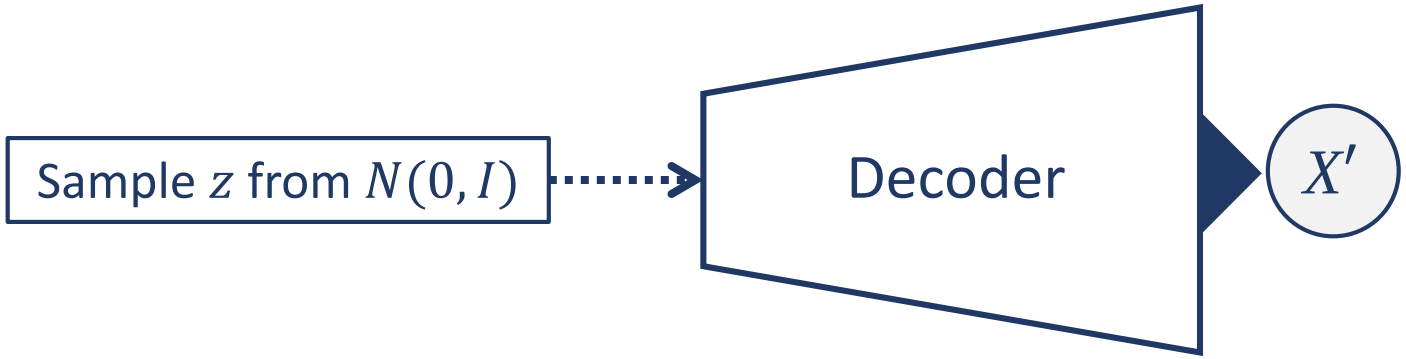

Loss Function $\|X - X'\| + \lambda \cdot KL\left[N\left(\mu(X), \Sigma(X)\right) \| N(0, I)\right]$

Conditional Variational Autoencoders - Background

$$P(X) \rightarrow P(X|Y)$$

$$\log P(X|Y) - KL[Q(z|X,Y)\|P(z|X,Y)] =$$

Lower bound (KL > 0)

$Q(z|X,Y)$ Gaussian $\mathcal{N}(\mu, \Sigma)$

KL-div. of two Gaussians → closed form

$$= E_{z\sim Q(z)}[\log P(X|Y,z)] - KL[Q(z|X,Y)\|P(z|Y)]$$

Reconstruction error (Decoder)

Conditional Variational Autoencoders - Background

Reparameterization trick

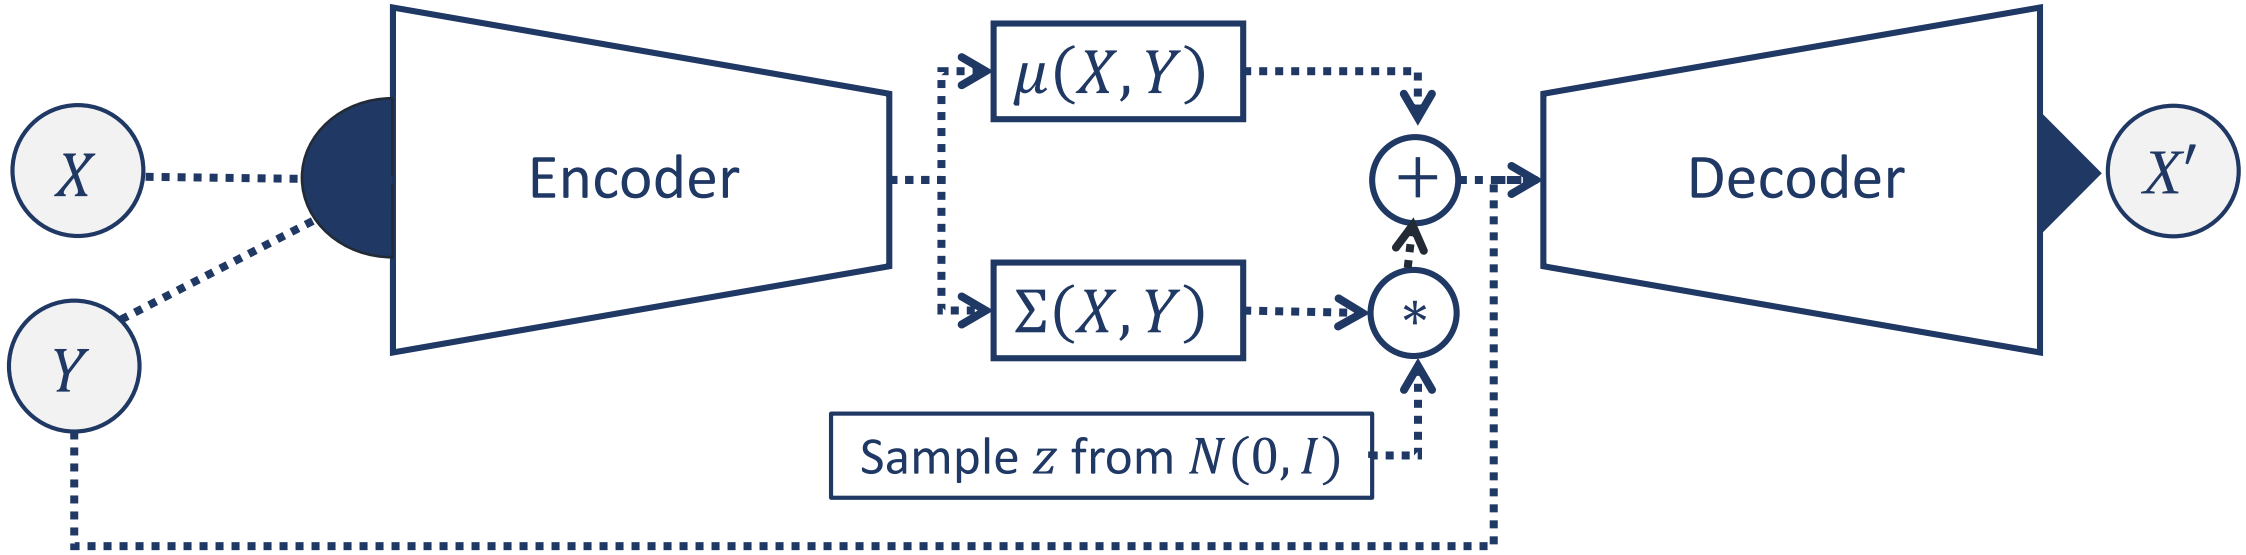

Loss Function $\|X - X'\| + \lambda \cdot KL\left[N\left(\mu(X,Y), \Sigma(X,Y)\right) \| N(0,I)\right]$

Conditional Variational Autoencoders - Background

Reparameterization trick

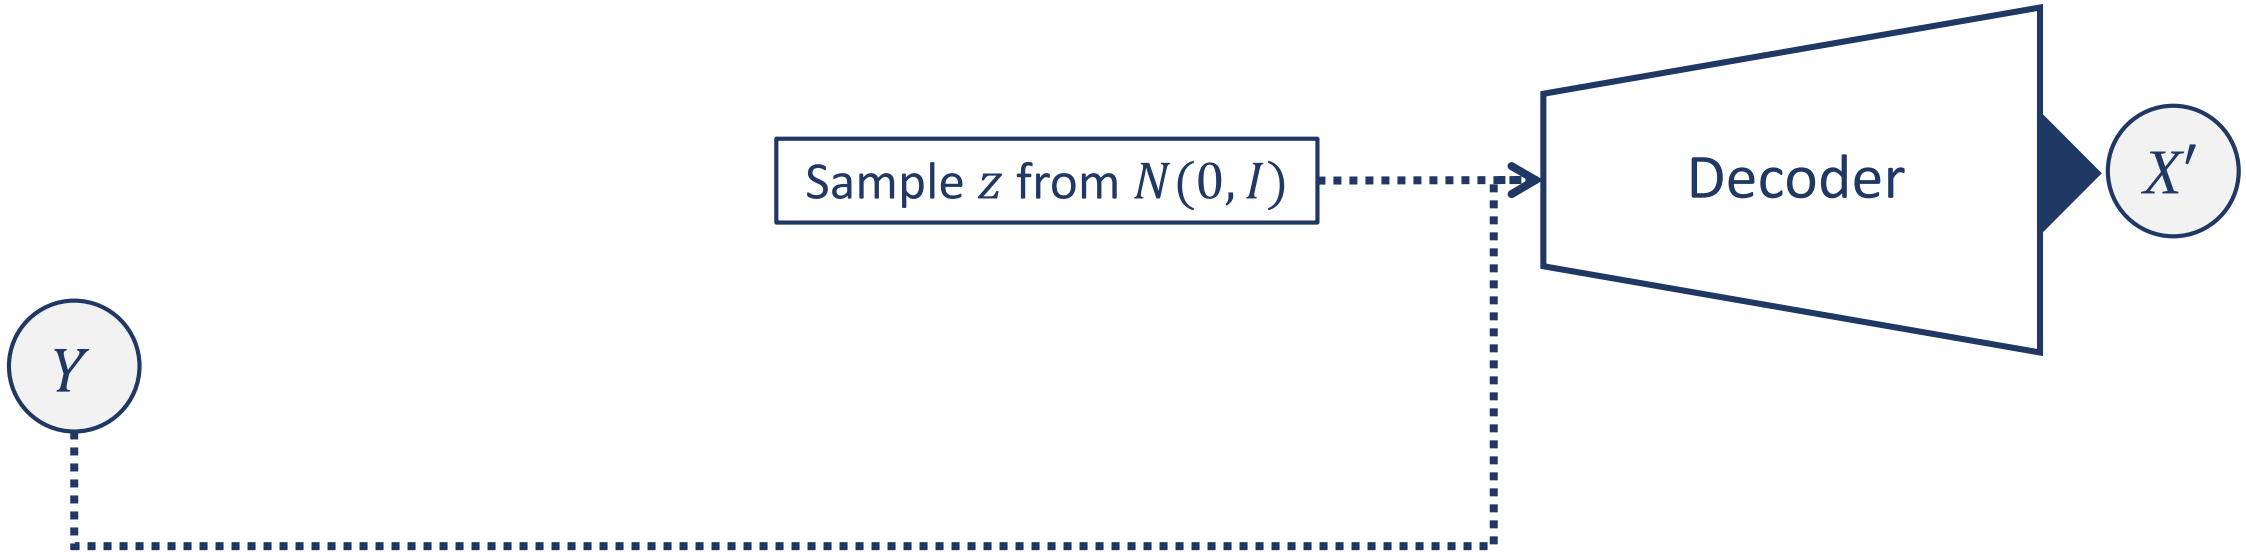

Loss Function $\|X - X'\| + \lambda \cdot KL\left[N\left(\mu(X, Y), \Sigma(X, Y)\right) \| N(0, I)\right]$

Variational Autoencoders - Background

Neural Network Perspective

Dimensionality Reduction Method

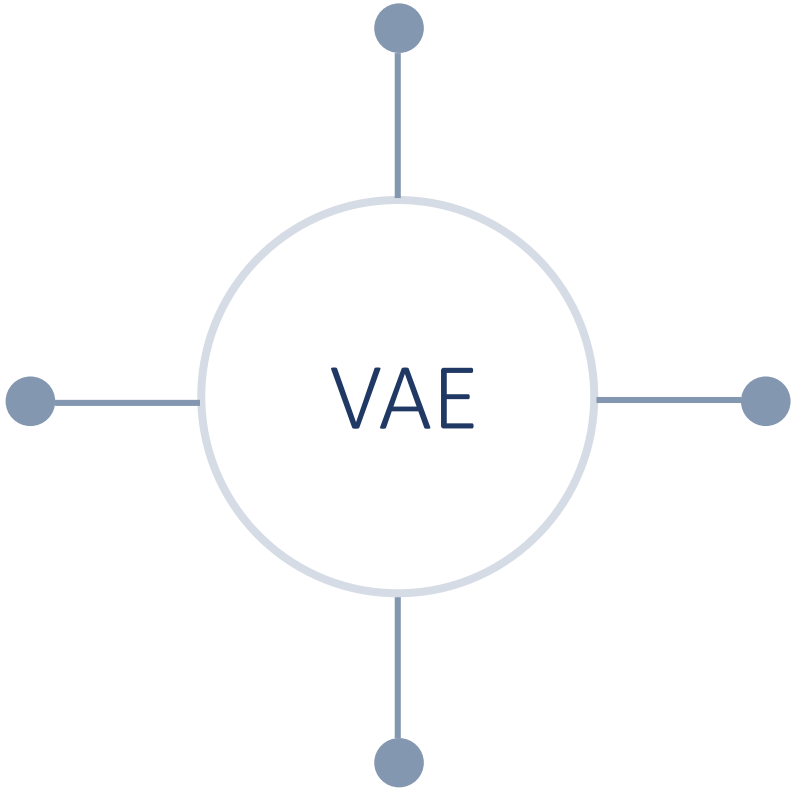

Generative Model

Probability Model Perspective

3D Anatomy Reconstruction from 2DUS

sagittal

coronal

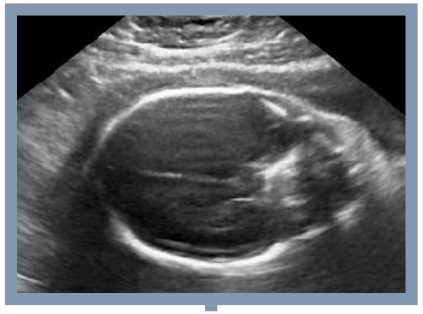

axial (transvent.)

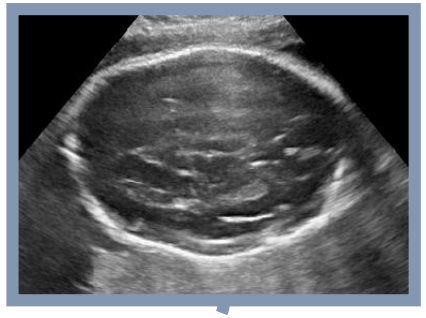

$P(X|Y_{1,2,3})$ where X: 3D skull and $Y_{1,2,3} = \{Y_1, Y_2, Y_3\}$

Y_1

Y_2

Y_3

Sagittal

Coronal

Transcerebellar

Transventricular

Transthalamic

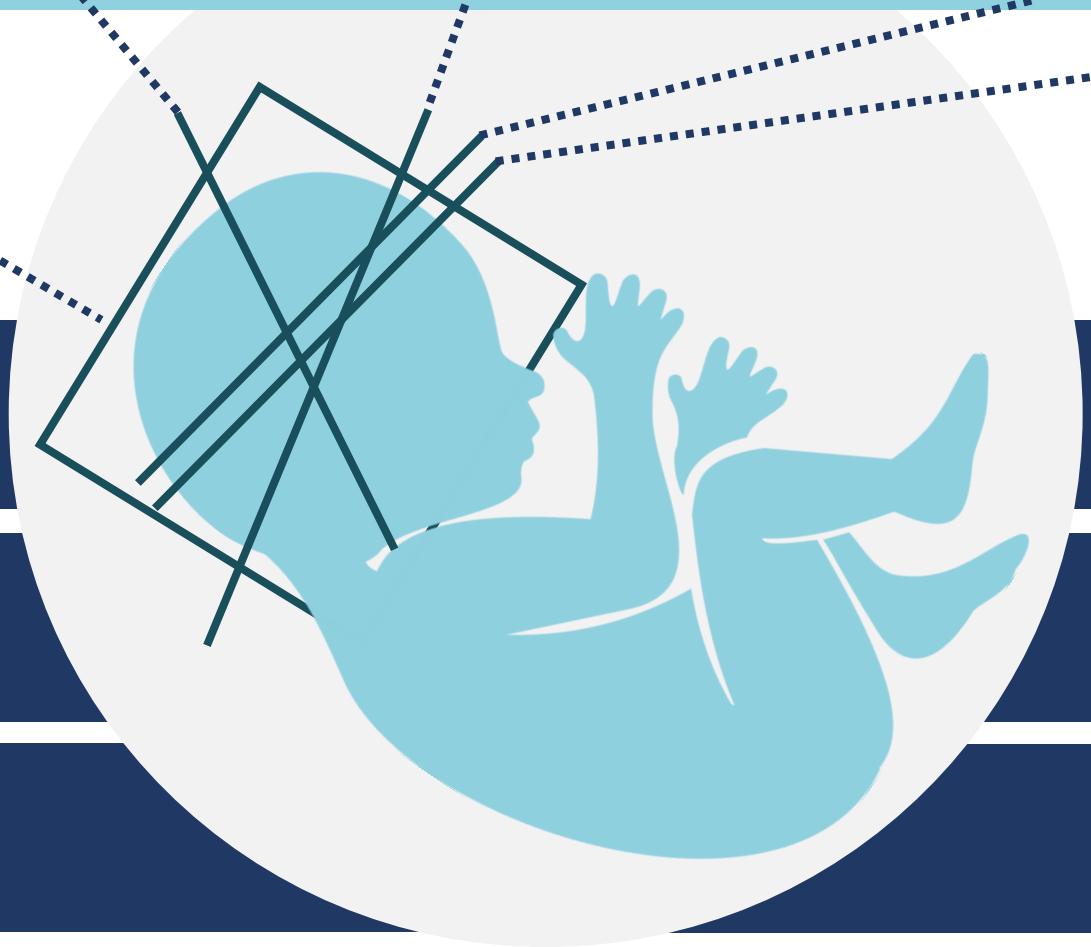

Standard planes manually acquired.

Low reproducibility

2D Biometrics.

Subjectivity

Manual meassurements.

Prone to errors.

Towards 3D Fetal Analysis

3D US has the potential to mitigate these effects, provinding a more realistic, objective, and reproducible analysis tool.

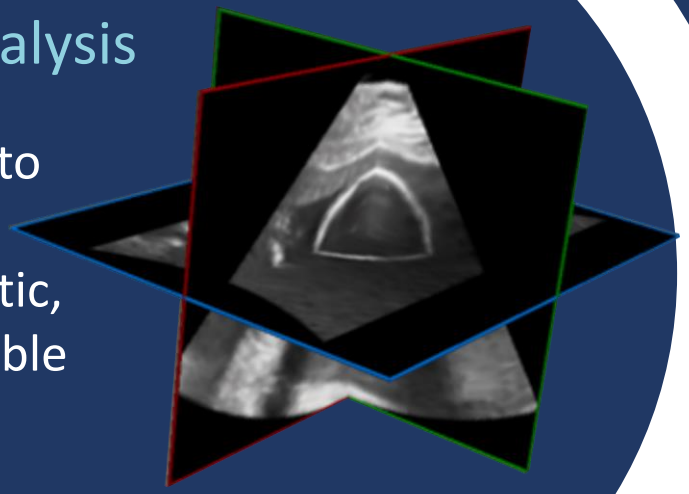

3D vs 2D

Dyson et al. “Three-dimensional ultrasound in the evaluation of fetal anomalies” Ultrasound Obstet. Gynecol. 2000, 16. ☺

Basgul et al. “Evaluation of fetal anomalies by two and three-dimensional ultrasound” Gynaecol Perinatol 2007, 16(2). ☺

Lee et al. “A review of three-dimensional ultrasound applications in fetal growth restriction” Journal of Medical Ultrasound 2012, 20. ☺

Goncalves et al. “Three- and 4-dimensional ultrasound in obstetric practice: Does it help?” J Ultrasound Med. 2005, 24 ☺

Goncalves et al. “What does 2-dimensional imaging add to 3- and 4-dimensional obstetric ultrasonography?” J Ultrasound Med 2006, 25. 😐

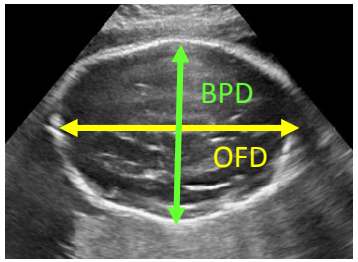

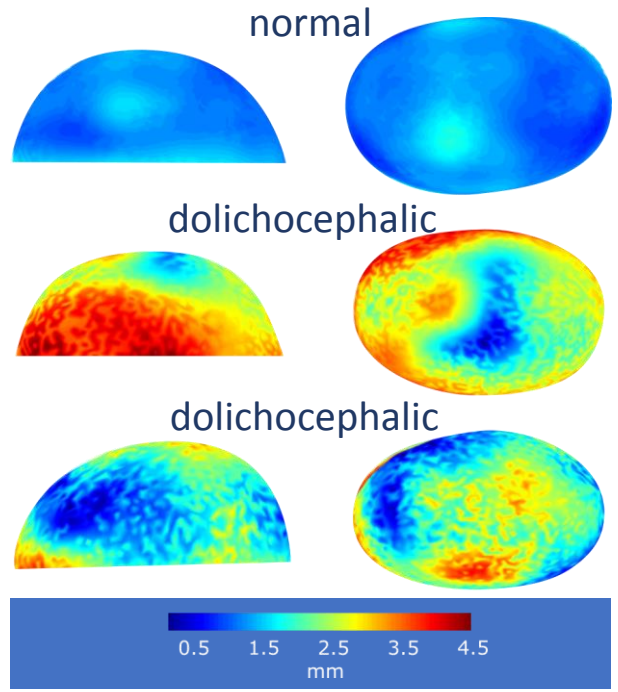

3D Biometry for Fetal Skull Assessment

J. Matthew et al. “Novel 3D-based metric to assess the fetal skull: A Pilot Study. BMUS 2017.

Towards 3D Fetal Analysis

2D → 3D

Analysis and Visualization

Automatic Tools for Segmentation and Analysis

Cerrolaza et al. “Deep learning with ultrasound physics for fetal skull segmentation” ISBI 2018

Namburete et al. “Fully-automated alignment of 3D fetal brain ultrasound to a canonical reference space using multi-task learning” Med. Image Anal. 2018

Limited Experience

Automatic Detection of Anatomical Structures and Standard Planes in 3DUS

Li et al. “Standard plane detection in 3D fetal ultrasound using an iterative transformation network” MICCAI 2018

Huang et al. “VP-Nets: Efficient automatic localization of key brain structures in 3D fetal neurosonography” Med. Image Anal. 2018,47

Lack of Large Dataset

Need of large international/ multi-ethnic studies to create new reference charts for 3D biometry

?

Not all the standard views are always routinely acquired

Hierarchical CVAE

Axial (Transvent.)

Always acquired.
Classic 2D biometrics.

Sagittal

Frequently.
Face / head shape.

Coronal

Sometimes.
Normally acquired as part of a dedicated scan.

3D Anatomy Reconstruction from 2DUS

Conditional Variational Autoencoders for 3D Fetal Anatomy Reconstruction

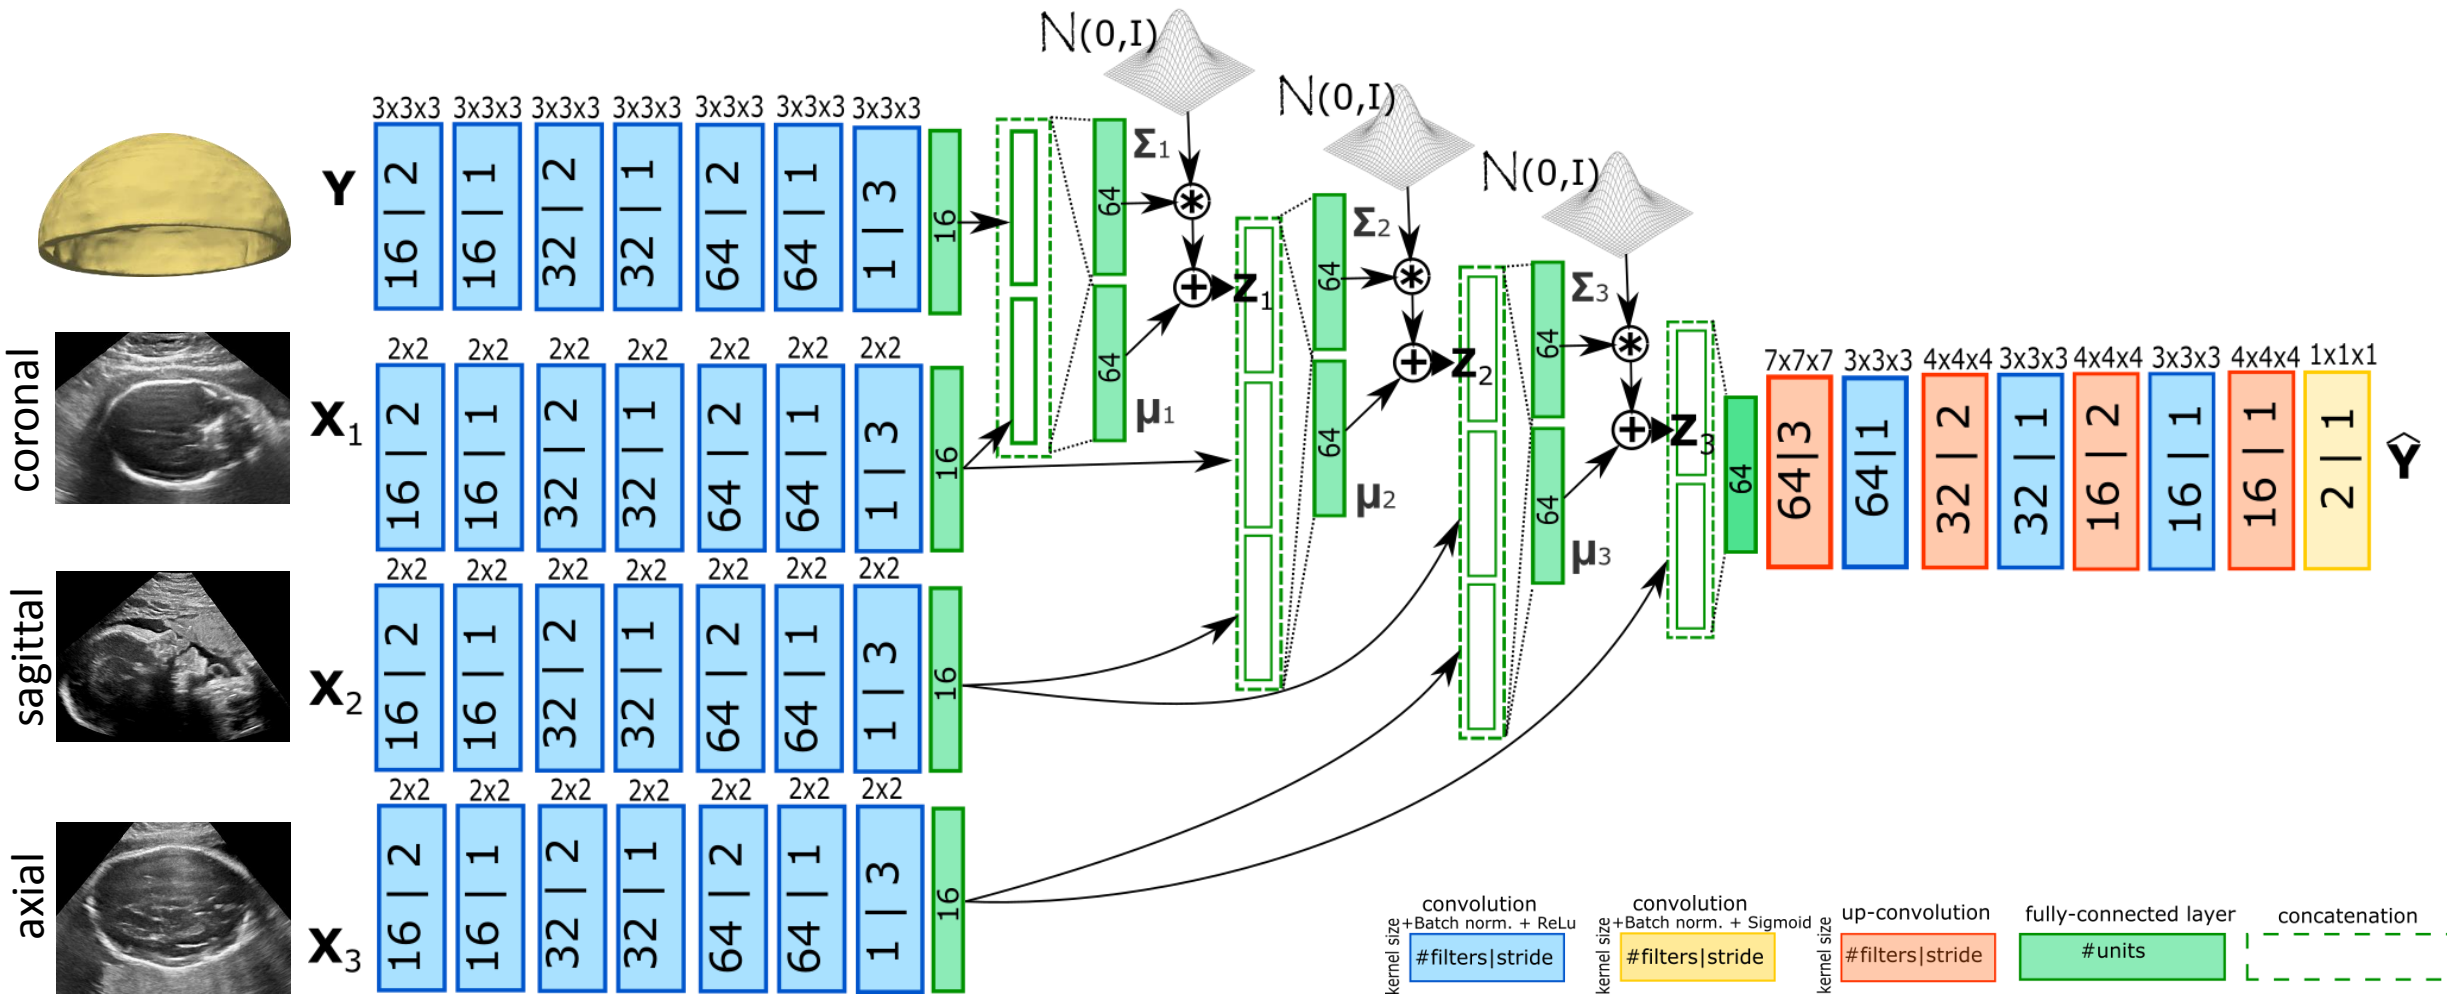

Experiments

- Data: 72 cases (IBD approved); avg. gest. age 24.7 (20 to 36 weeks)
  58 training / 14 testing (3-folds)
- Preprocessing: Resized to 96 x 96 x 96 vox. (96 x 96 pixels); isotropic 0.50 mm
- 3D skull manually segmentation supervised by expert radiologist.
- Random anisotropic scaling, rotations (± 10°) and translations (± 7 pix.).
- Adam (l.r. = 0.001, β_1=0.9, β_2=0.995); 1000 epochs.

Coronal

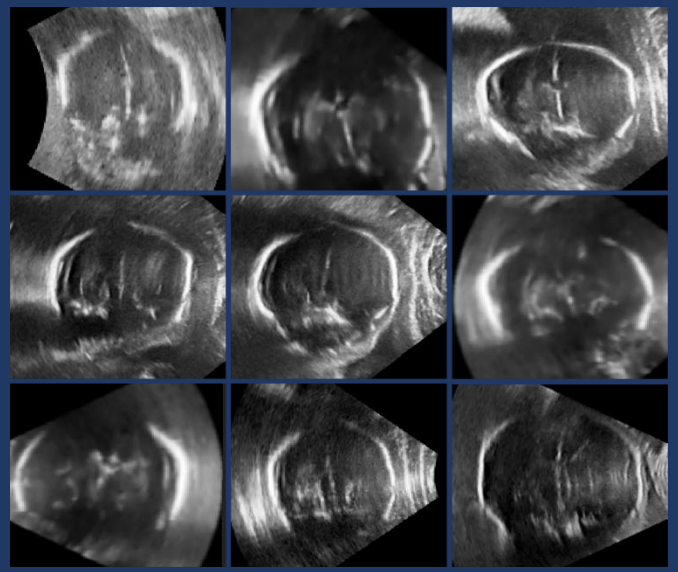

Sagittal

Axial (transvent.)

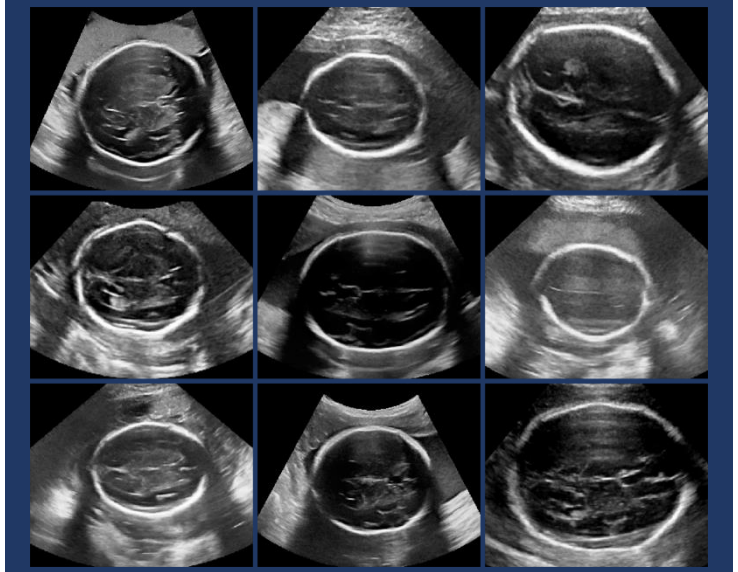

Training progress…

Skull reconstruction

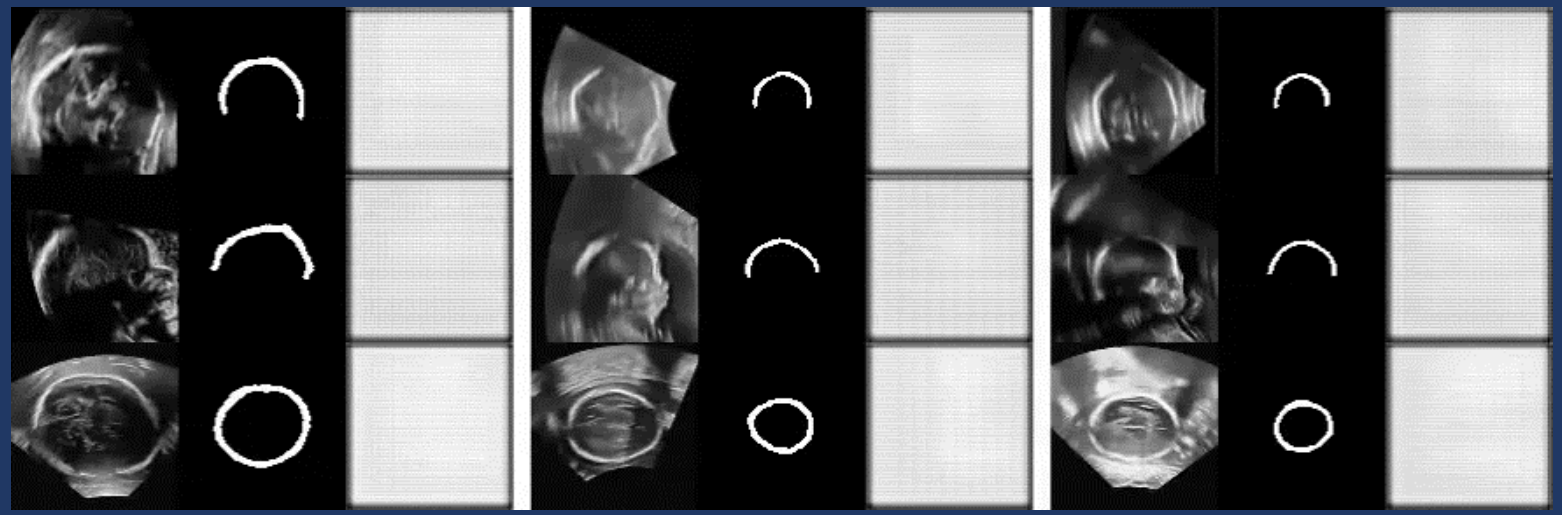

Latent space

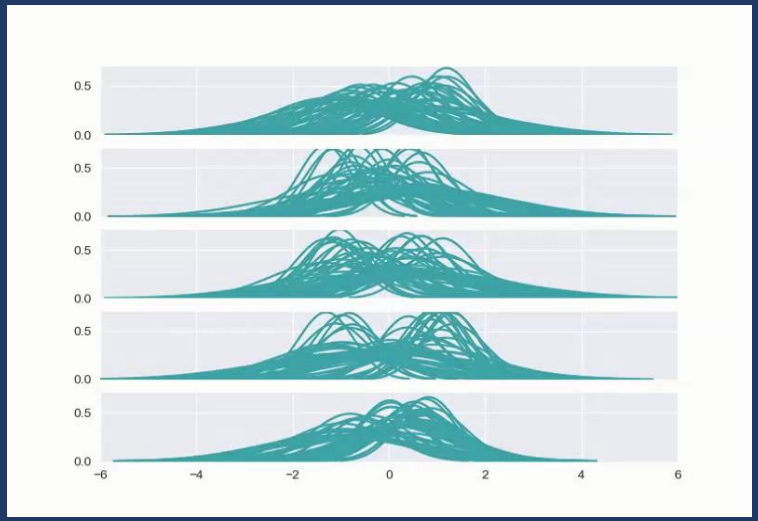

DC: Dice coeff. HD: Hausdorff dist. (mm). RVD: Relative volume diff.

	DC	HD	RVD
CVAE	0.91 ± 0.02	4.33 ± 1.71	0.03 ± 0.12
HCVAE	**0.91 ± 0.04**	**4.12 ± 1.98**	**0.03 ± 0.14**
TL-NET [4]	0.89 ± 0.03	4.79 ± 1.28	0.09 ± 0.17
GAN	0.89 ± 0.04	5.16 ± 1.5	0.03 ± 0.16

axial sagittal coronal

	DC	HD	RVD
CVAE	0.86 ± 0.05	5.43 ± 2.78	0.03 ± 0.30
HCVAE	**0.89 ± 0.05**	**4.81 ± 2.44**	**0.03 ± 0.20**
TL-NET [4]	0.89 ± 0.05	5.32 ± 1.99	0.09 ± 0.19
GAN	0.86 ± 0.07	5.93 ± 2.66	0.05 ± 0.30

axial sagittal

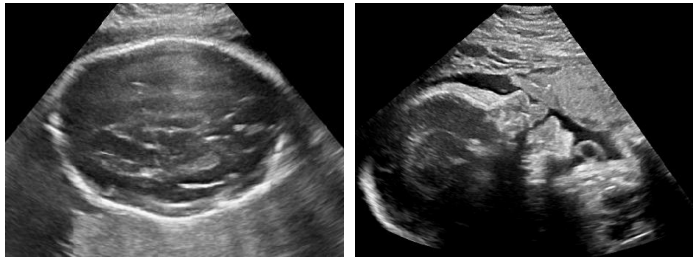

	DC	HD	RVD
CVAE	0.83 ± 0.06	6.23 ± 2.88	0.09 ± 0.33
HCVAE	**0.86 ± 0.05**	**5.04 ± 2.82**	**0.04 ± 0.21**
TL-NET [4]	0.85 ± 0.04	7.04 ± 2.34	0.17 ± 0.20
GAN	0.83 ± 0.08	8.04 ± 3.06	0.17 ± 0.37

axial

Prediction Uncertainty in HCVAE

axial

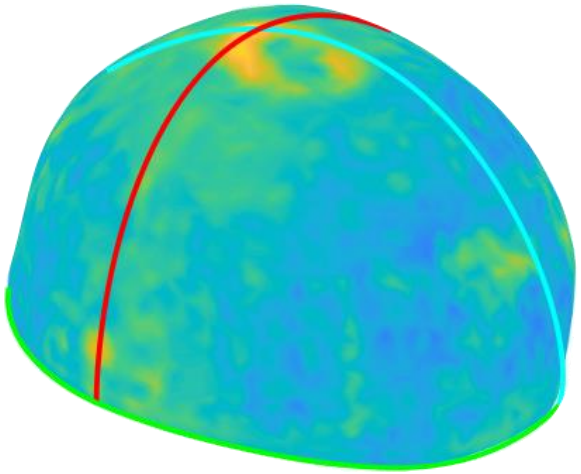

axial + sagittal

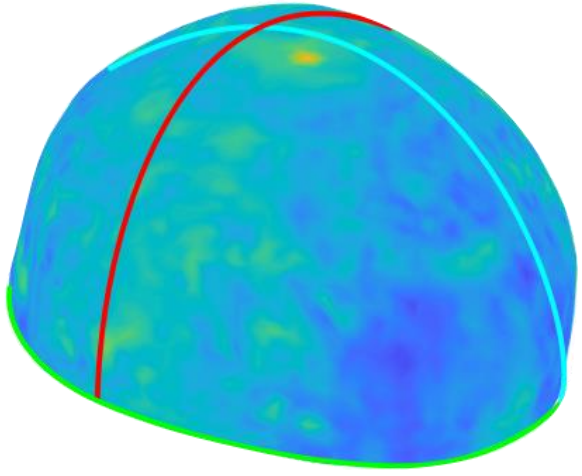

axial + sagittal + coronal

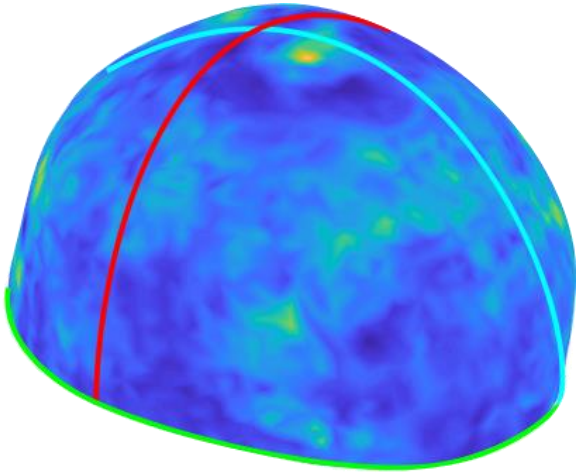

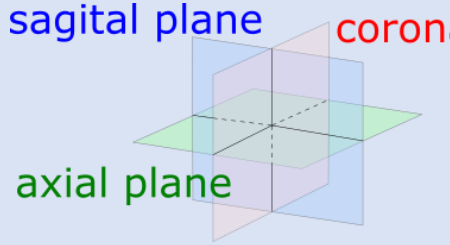

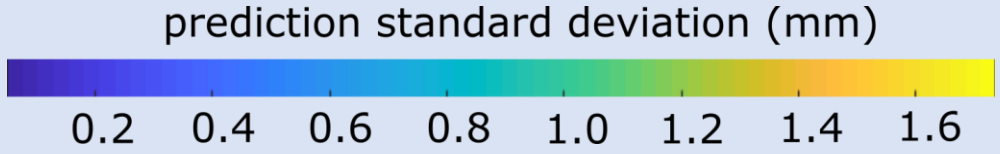

Variational Autoencoders - Background

Neural Network Perspective

Dimensionality Reduction Method

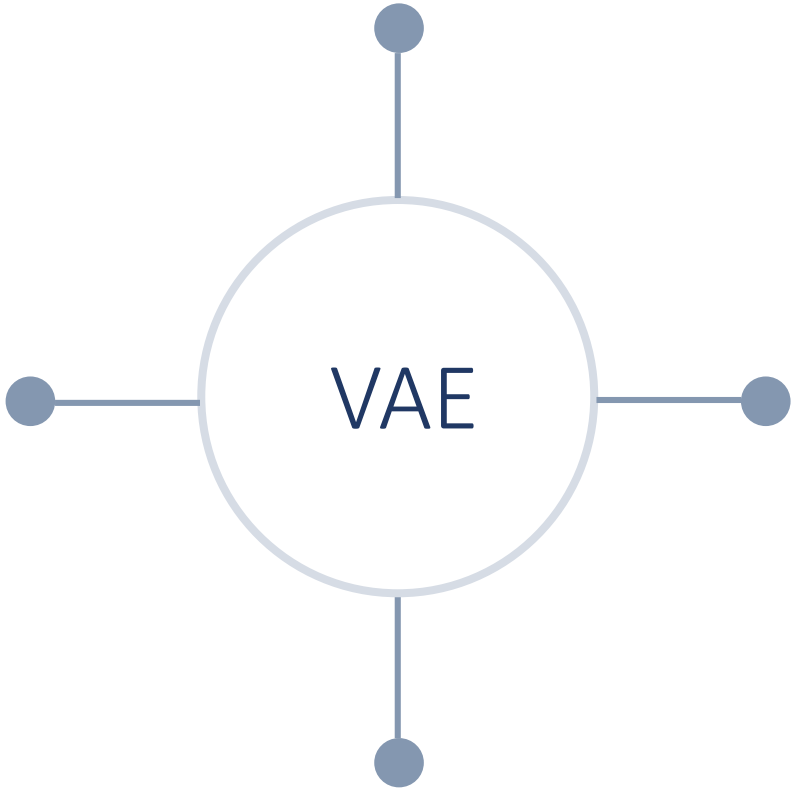

Generative Model

Probability Model Perspective

Explainable Cardiac Remodeling Assessment

Cardiac remodeling

- refers to any change in size and shape of the heart
- strong predictor of survival in cardiac pathologies.

Cardiovascular magnetic resonance (CMR) has become the gold-standard for high-resolution imaging of the cardiac structure.

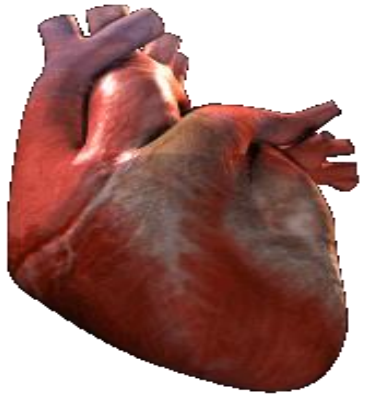

Biological Ground Truth

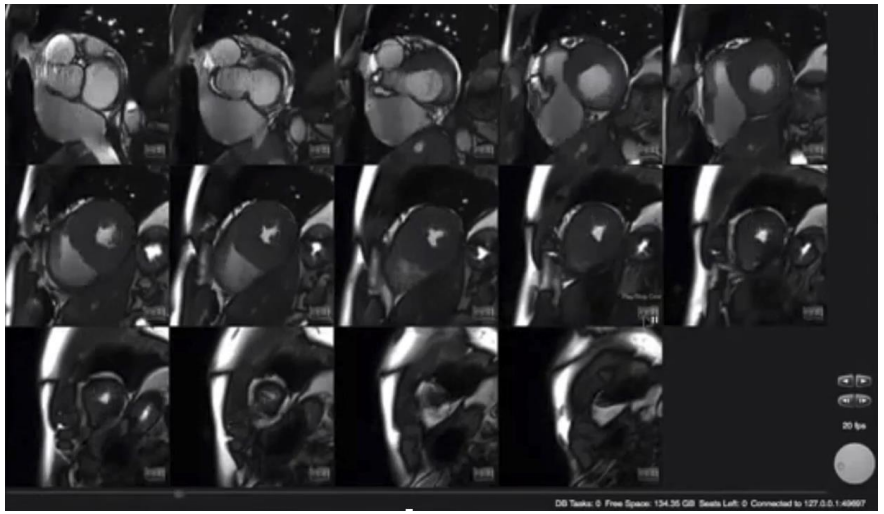

Left Ventricular (LV) short-axis cine MR acquisition

SAX3D Stack LV Function	
EDV:	119 ml
ESV:	42 ml
SV:	78 ml
EF:	65 %
CO:	4805 ml/min
CI:	2 l/min/m²
HR:	62.0/min
Myo Mass (Diast):	96 g
Phase Diastole:	1
Phase Systole:	12
EDV/H:	64 ml/m
EDV/BSA:	59 ml/m²
ESV/H:	22 ml/m
ESV/BSA:	21 ml/m²
SV/H:	42 ml/m
SV/BSA:	38 ml/m²
Myo Mass/H (Diast):	52 g/m
Myo Mass/BSA (Diast):	48 g/m²

Clinical Indexes

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM)

Disease of the heart muscle which manifests clinically with unexplained left ventricular (LV) **hypertrophy** (thickening).

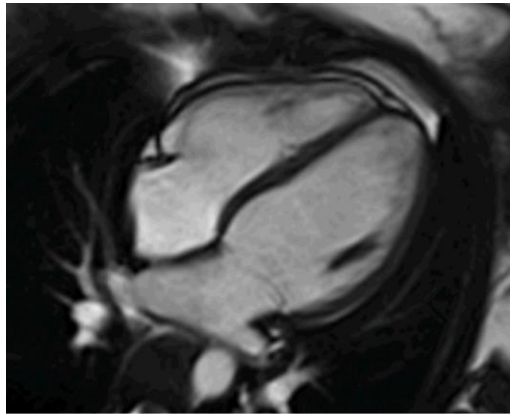

Healthy

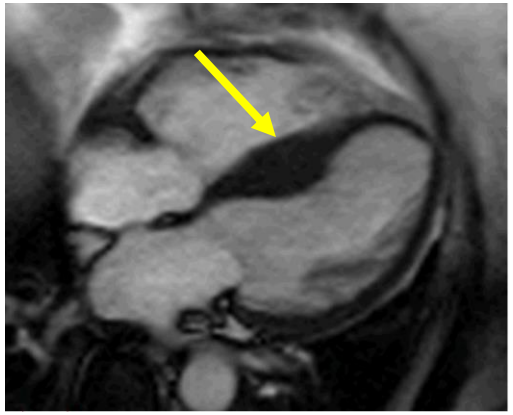

HCM

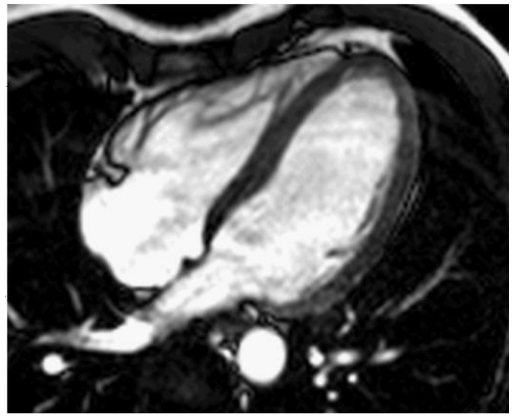

Healthy

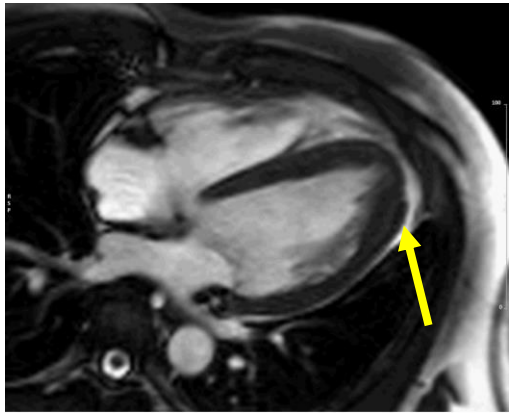

HCM

Conventional clinical indexes are insensitive to **regional** and **asymmetrical** remodeling.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

PCA shape components

☺ Easy to **visualize**.

☹ Just summarizes data variability → they **do not** necessarily encode anatomical features that can **differentiate** between classes.

Machine Learning approaches

☺ Powerful feature extraction.

☹ Lack interpretability

➡ VAE

Learns a set of latent variables *z* that:

1. can **differentiate** clinical conditions *y*.
2. and whose anatomical effect can be **visualized**.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

- End-diastolic (ED) and end-systolic (ES) phases of cardiac cine MR images of healthy (HVols) and HCM subjects were automatically segmented[1].
- 3D segmentations quality was improved by a multi-atlas-aided upsampling scheme and ED and ES frames were registered to a common template.

IMPERIAL COLLEGE DATASET:
TRAINING: 537 (276 HVols, 261 HCMs)
VALIDATION: 150 (75 HVols, 75 HCMs)
TESTING: 200 (200 HVols, 200 HCMs)

ACDC MICCAI 2017 DATASET:
TESTING: 40 (20 HVols, 20 HCMs)

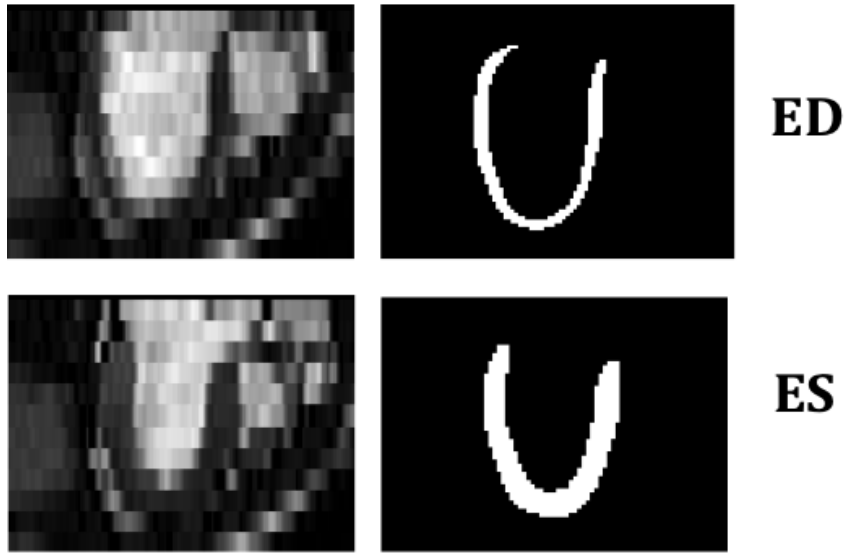

Exemplar 2D LV short-axis views and corresponding LV segmentation at ED and ES.

[1]Bai, W. et al. *A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects.* MedIA 2015 Dec;26(1):133-45.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

VAE Implementation

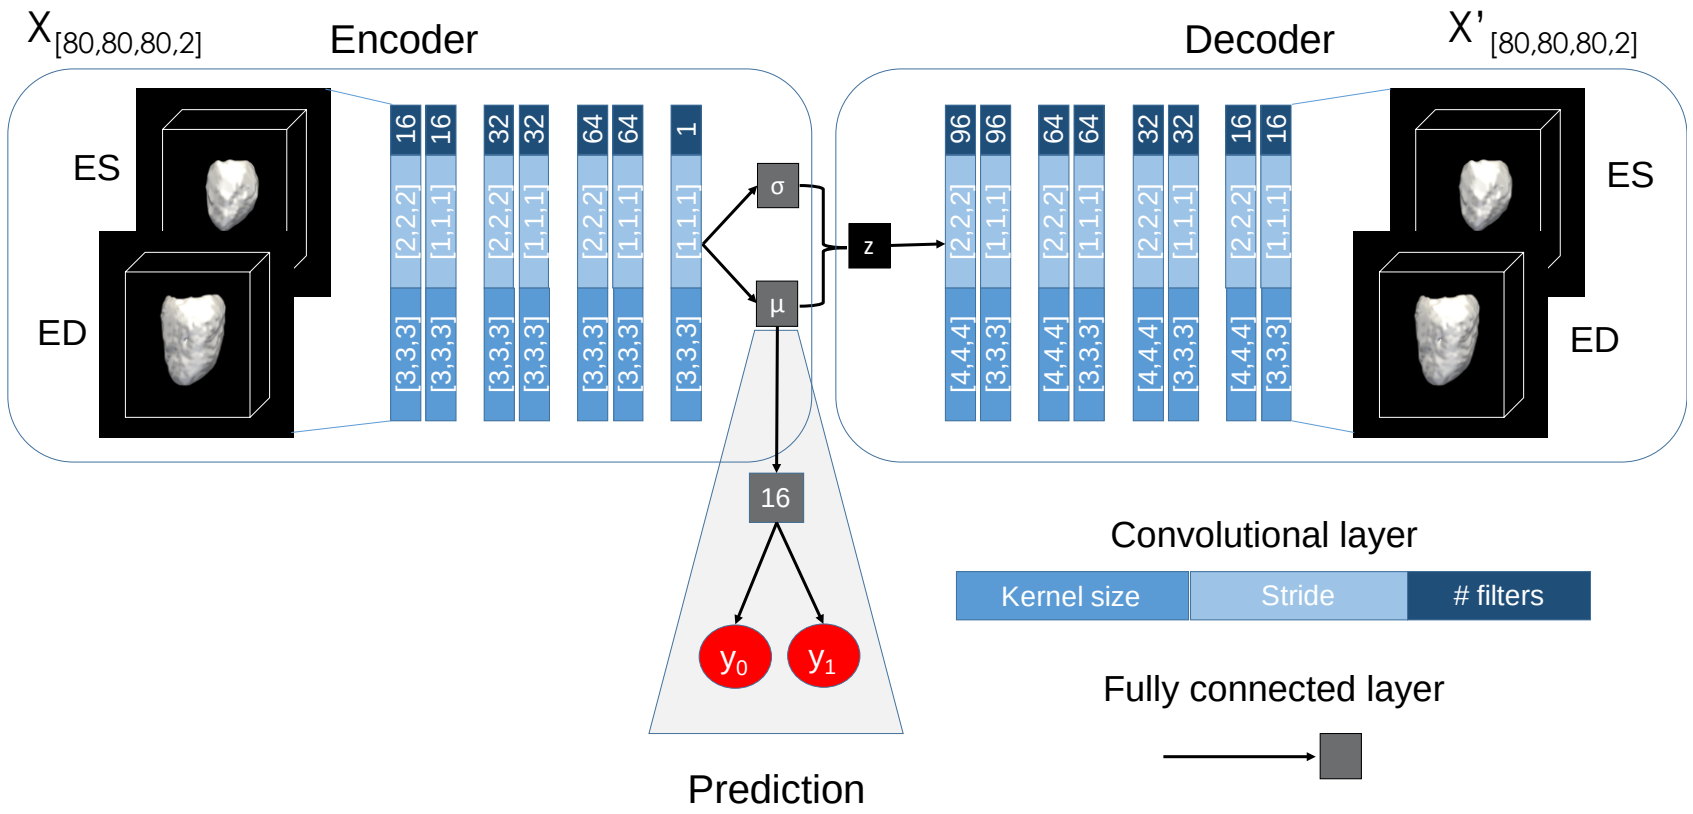

Biffi, C. et al. *Learning Interpretable Anatomical Features Through Deep Generative Models: Application to Cardiac Remodeling*. MICCAI 2018.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

VAE Implementation

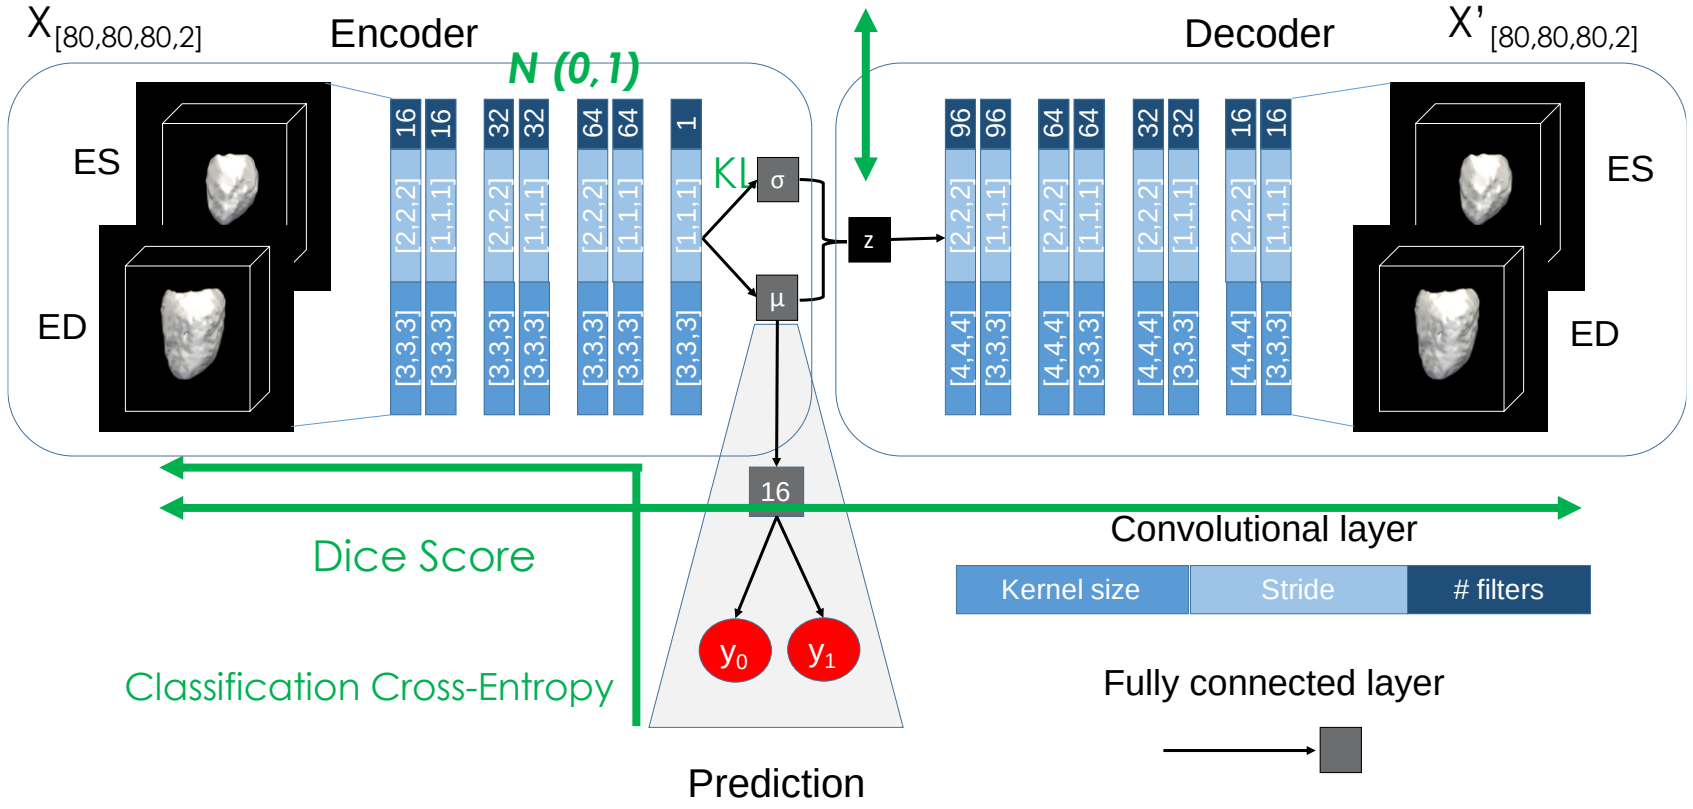

Latent Space Navigation: $\mu_t = \mu_{t-1} + \lambda \frac{\delta y_1}{\delta \mu_{t-1}}$ t = iteration number, λ=0.1

Biffi, C. et al. *Learning Interpretable Anatomical Features Through Deep Generative Models: Application to Cardiac Remodeling*. MICCAI 2018.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

VAE Implementation

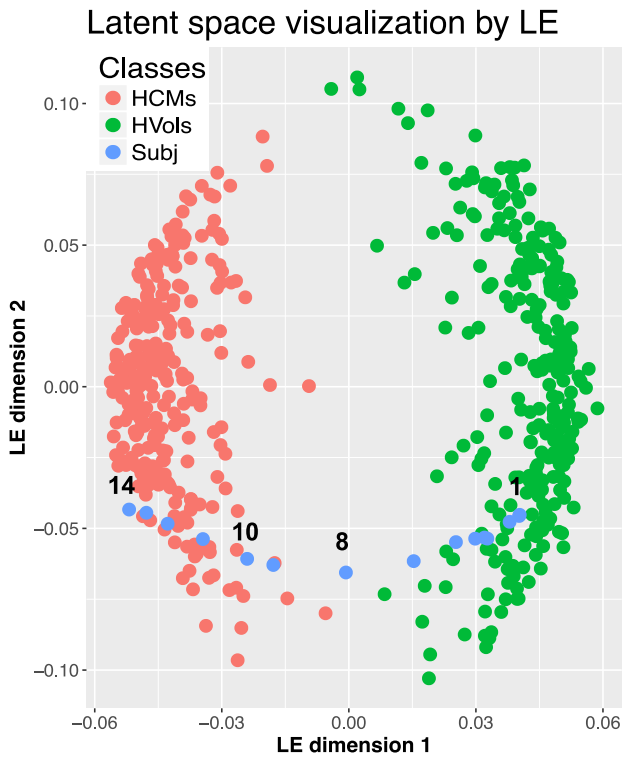

LE = Laplacian Eigenmaps

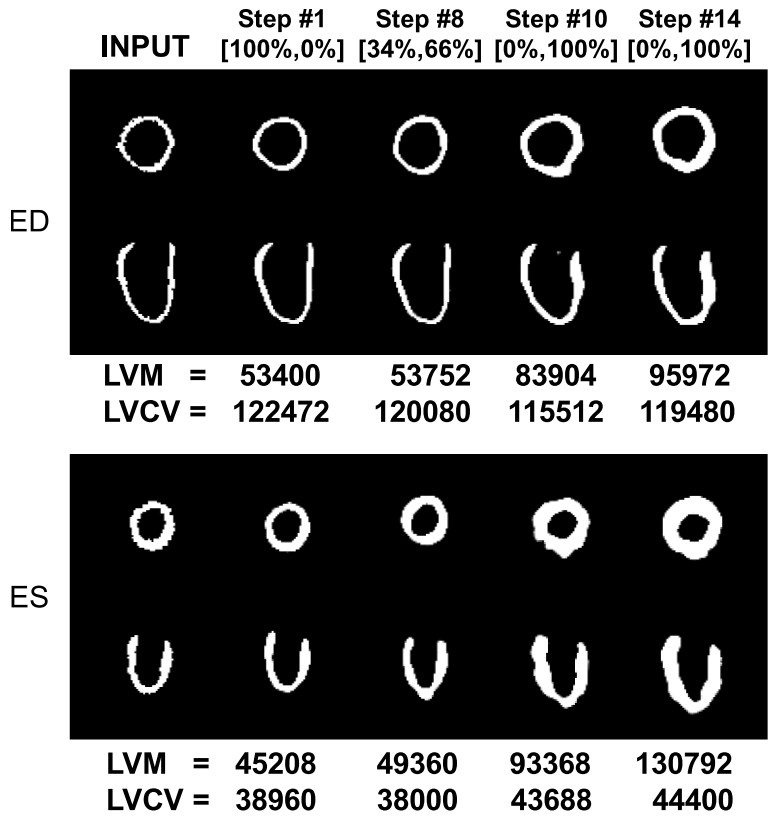

TESTING – IMPERIAL COLLEGE DATASET: 200 (100 HVols, 100 HCMs) – *100% accuracy*
– ACDC MICCAI 2017: 40 (20 HVols, 20 HCMs) – *90% accuracy*

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

Can we do even better?

- Latent space navigation is subject-specific, i.e. no population-based inference.
- Latent space can only visualized with an additional dimensionality reduction technique.

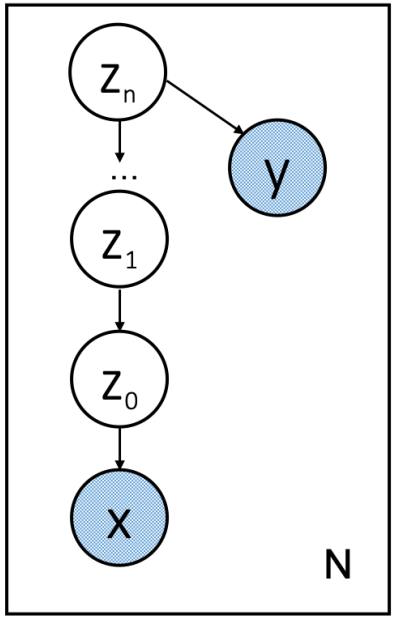

Given a population of N shapes X, we aim at developing a data-driven method that learns a conditional hierarchy of latent variables $\{z_N, \ldots z_1, z_0\}$ where:

1. z_N can **differentiate between** clinical conditions y and is very **low-dimensional**.
2. $\{z_N, \ldots z_1, z_0\}$ anatomical effect can be **visualized**.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

Ladder VAE *Novelty:* Sharing of information between encoder and decoder.

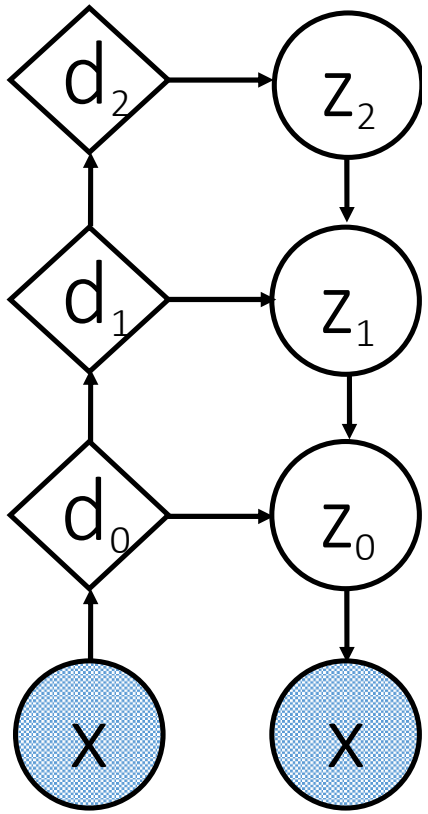

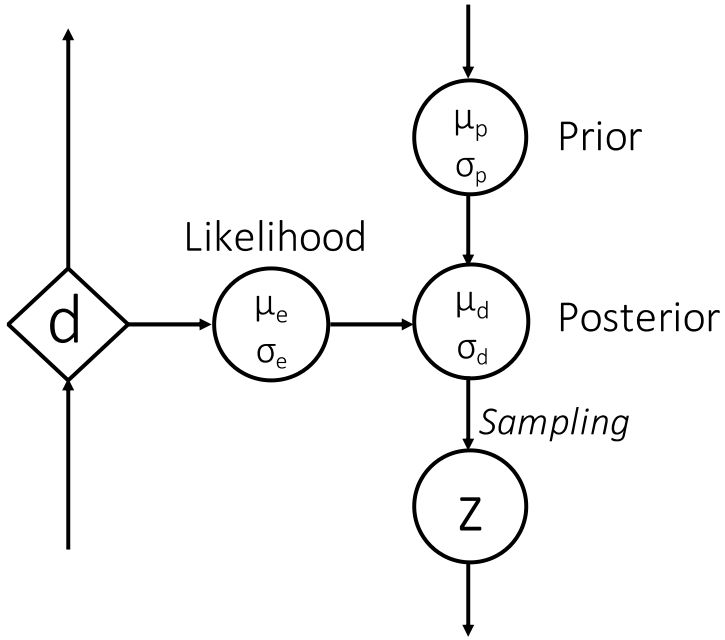

$$\mu_d = \frac{\mu_e \sigma_e^{-2} + \mu \sigma_p^{-2}}{\sigma_e^{-2} + \sigma_p^{-2}} \qquad \sigma_d^2 = \frac{1}{\sigma_e^{-2} + \sigma_p^{-2}}$$

Sønderby, C.K. *et al.,* Ladder variational autoencoders. *Advances in neural information processing systems*. NIPS 2016.

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

Ladder VAE

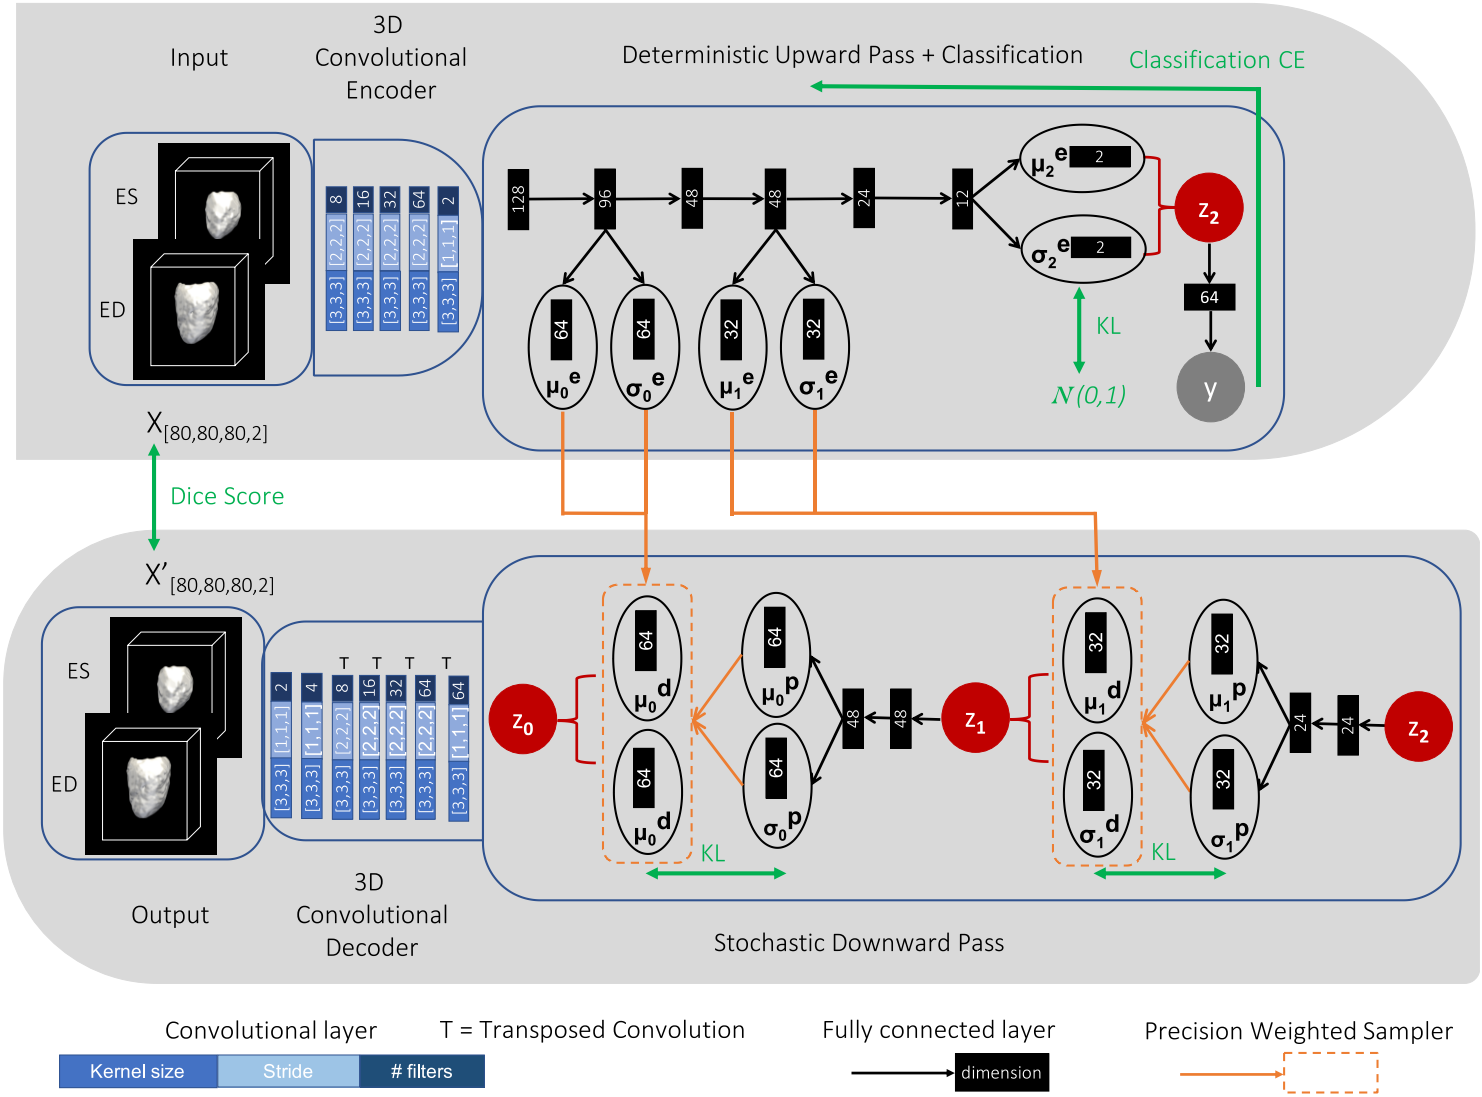

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

Ladder VAE

TESTING – IMPERIAL COLLEGE DATASET: 200 (100 HVols, 100 HCMs) – *100% accuracy*

– ACDC MICCAI 2017: 40 (20 HVols, 20 HCMs) – *90% accuracy*

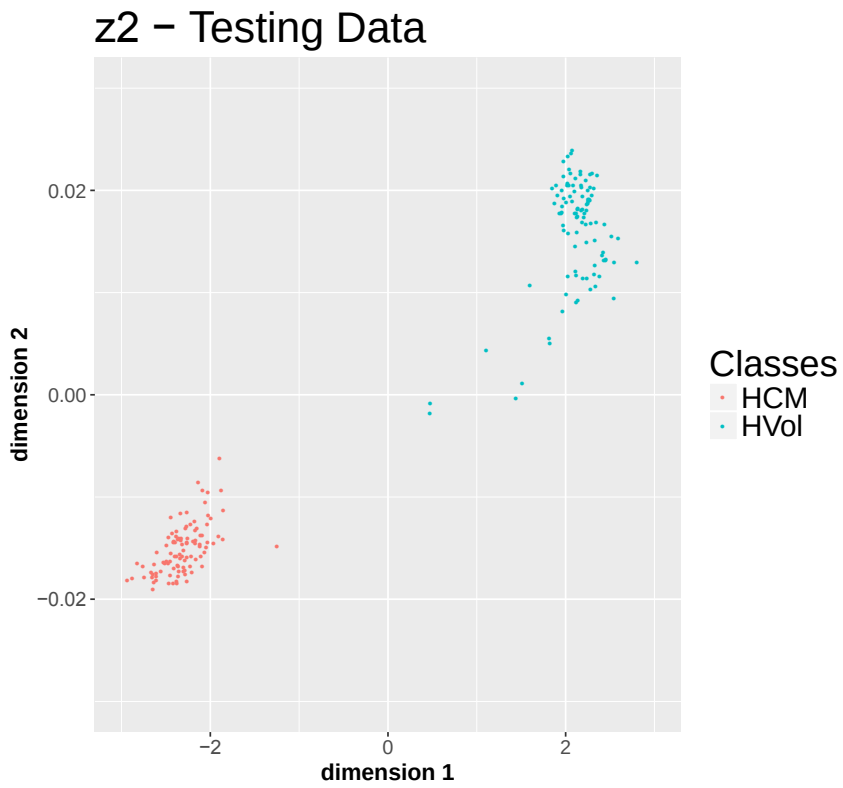

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

Ladder VAE

VAE vs LVAE Reconstruction Accuracy

	DSC_{ED}	DSC_{ES}	$\mathrm{H}_{ED}[mm]$	$\mathrm{H}_{ES}[mm]$
VAE	0.71±0.07	0.78±0.06	4.87±1.23	5.09±1.51
LVAE	0.79±0.06	0.84±0.05	3.92±1.08	4.09±1.36

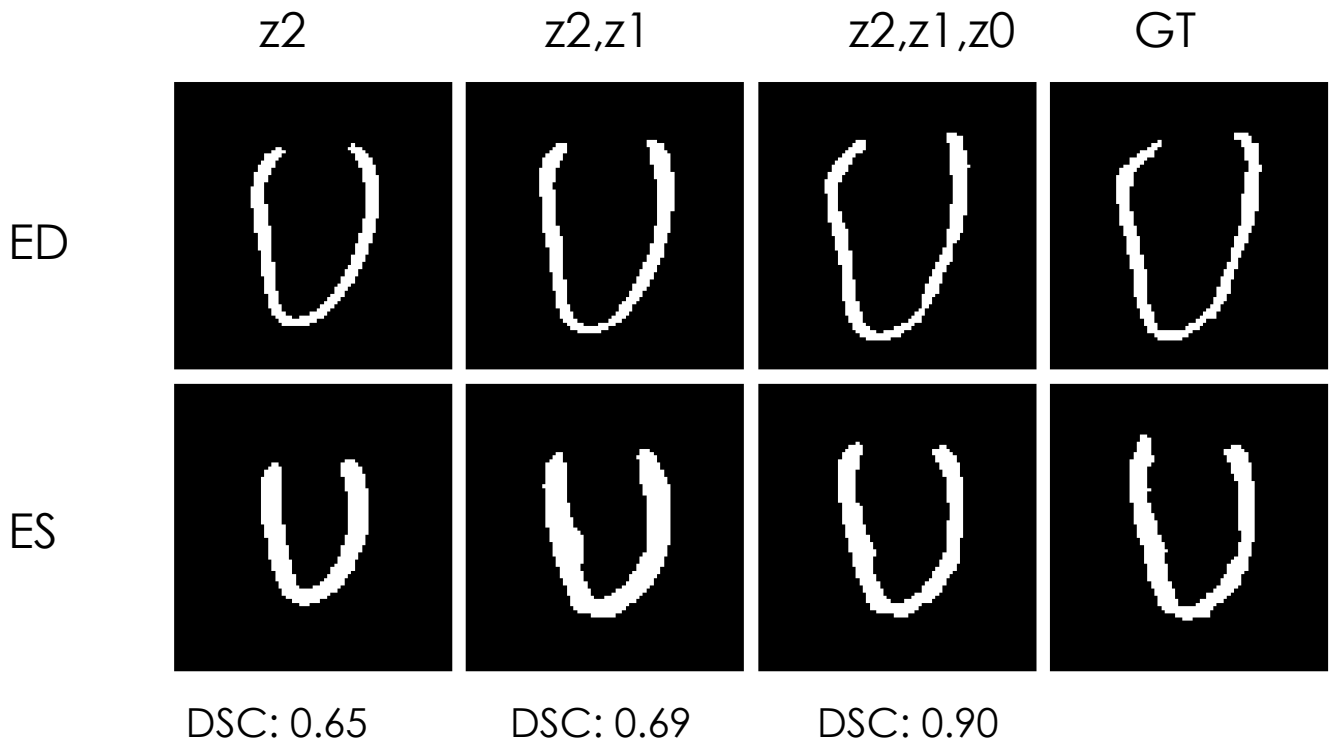

Explainable Cardiac Remodeling Assessment

Hypertrophic Cardiomyopathy (HCM) – Shape Descriptors

Ladder VAE

Mean healthy and HCM shapes generated by sampling from z_2 clusters.
Wall Thickness (WT) values are plotted at each vertex.

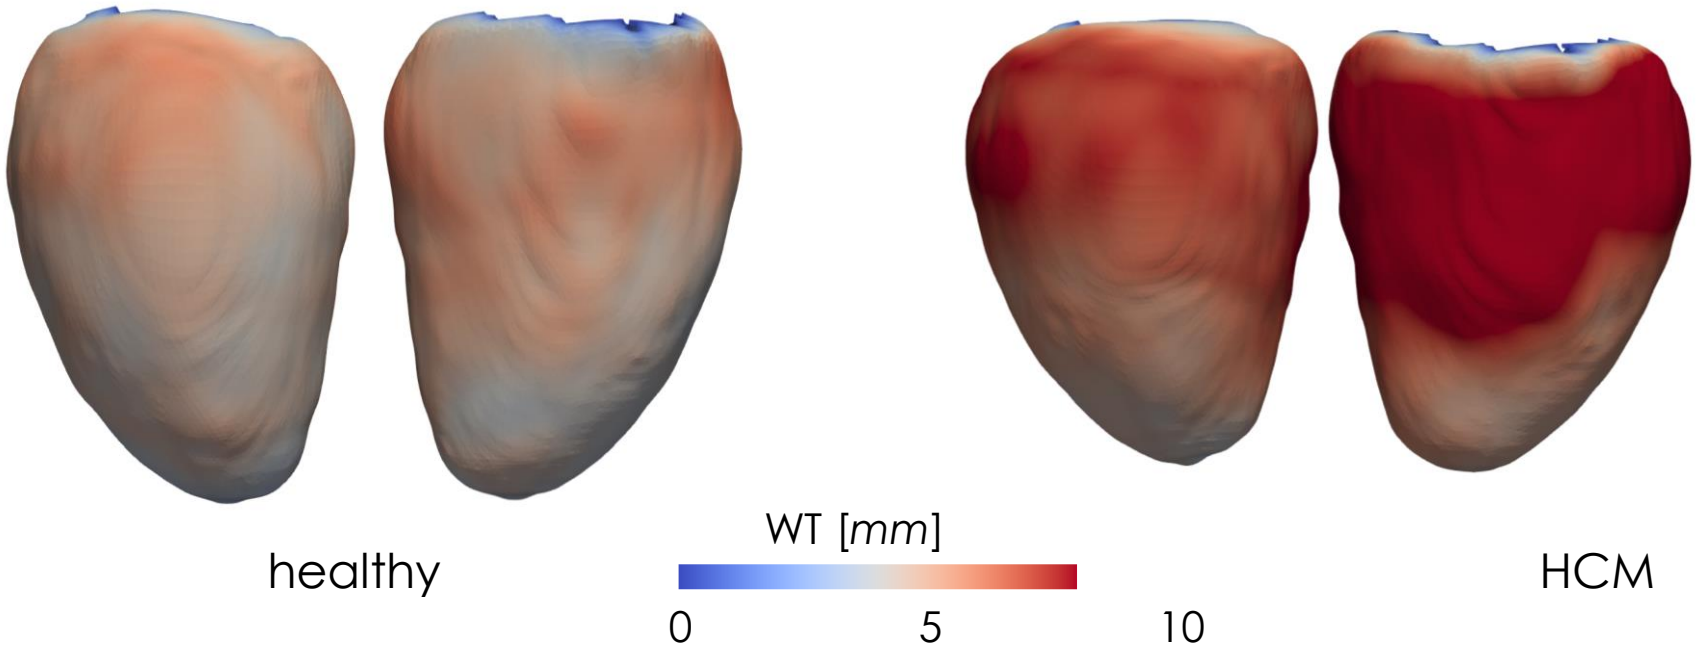

HCM shapes has ~15% greater mass and ~6% bigger cavity volume.

Conclusions

- VAEs are **NOT just a particular subgroup of autoencoders** with a Bayesian “twist”.
- VAEs is a versatile, and flexible family of networks with a solid **foundation in probability theory.**
- Their **intrinsic duality** as **dimensionality reduction** network and **generative models** make them an interesting approach for solving different problems in medical imaging.
- Potential of deep generative networks for the **3D reconstruction of the fetal skull** from non-registered 2DUS standard planes → **New generation of 3D fetal biometry**
- New approach as visualization and classification technique for cardiac pathologies.

Thank You!

Juan J. Cerrolaza, PhD
accenture technology
artificial intelligence IBERIA