DIFF·3: A latent diffusion model for the generation of synthetic 3D echocardiographic images and corresponding labels

Edward Ferdian[1,2*], Debbie Zhao[1*], Gonzalo Maso Talou[1], Gina Quill[1], Malcom Legget[3], Robert Doughty[3,4], Martyn Nash[1,5**] and Alistair Young[6**]

[1]Auckland Bioengineering Institute, University of Auckland, New Zealand, [2]Faculty of Informatics, Telkom University, Indonesia, [3]Department of Medicine, University of Auckland, New Zealand, [4]Green Lane Cardiovascular Service, Auckland City Hospital, New Zealand, [5]Department of Engineering Science and Biomedical Engineering, University of Auckland, New Zealand, [6]School of Biomedical Engineering & Imaging Sciences, King's College London, United Kingdom.

*Corresponding authors: e.ferdian@auckland.ac.nz, debbie.zhao@auckland.ac.nz

BACKGROUND

Echocardiography is the most widely used modality for cardiac assessment, however, issues related to the high costs of acquisition, time-consuming analysis, and patient privacy, have limited the number of publicly available datasets. Deep learning models for automated analysis are being actively developed where large amounts of labelled data are needed.

Here, we present **DIFF·3**, a latent diffusion model (LDM) for the synthesis of realistic 3D echocardiograms paired with left-ventricular segmentation labels informed by cardiac magnetic resonance (CMR) imaging to provide a more efficient means of generating realistic and high-quality datasets.

METHODS

Network design

DIFF·3 is a Latent Diffusion Model, which consists of two separate networks:

1. Variational Auto Encoder

2D LPIPS used on the grid images to minimize **perceptual loss**

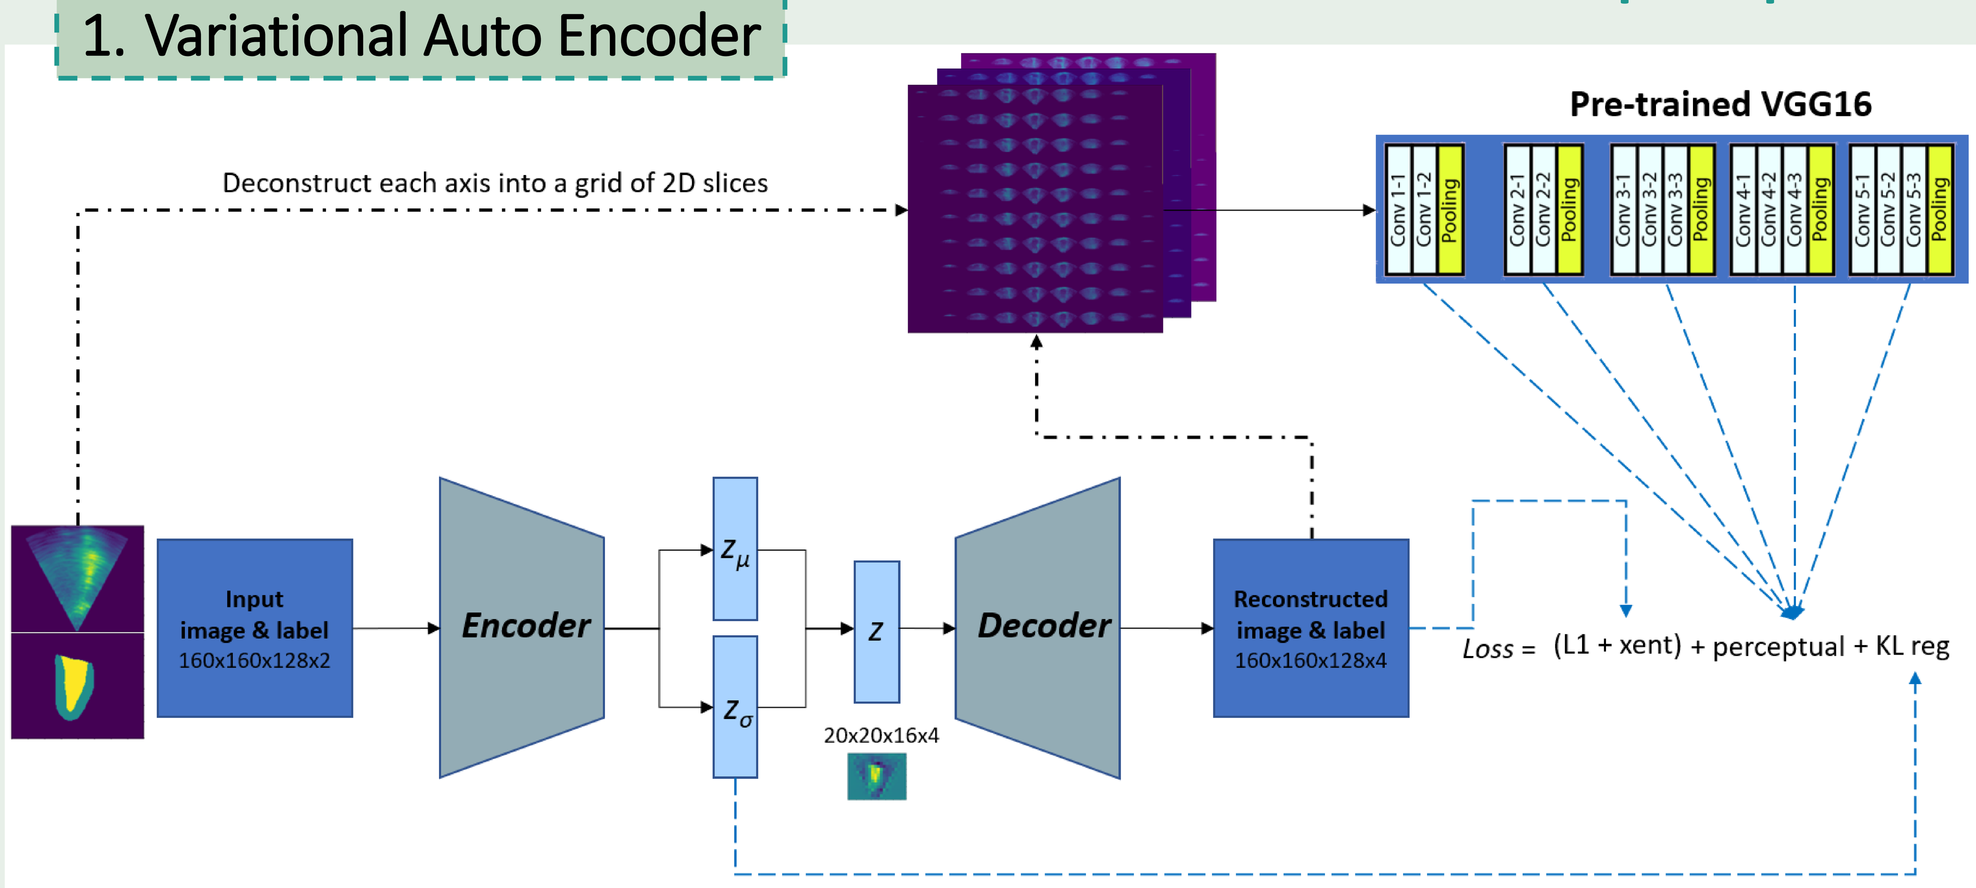

*VAE accepts **image + label** stacked in channel dimension*

2. Diffusion Probabilistic Model (DDPM)

- DDPM trained on the latent samples of the reference dataset
- DDPM was adapted for 3D input data
- Small latent input data allow high number of batches (b=64)
- Trained over 50000 iterations with a pixel-loss (L1) objective function

Reference dataset

We leverage **MITEA**, a publicly available 3D echocardiography (3DE) dataset with subject-specific labels (of the left ventricular myocardium and cavity) derived from higher-resolution CMR (train=428, test=108 cases).

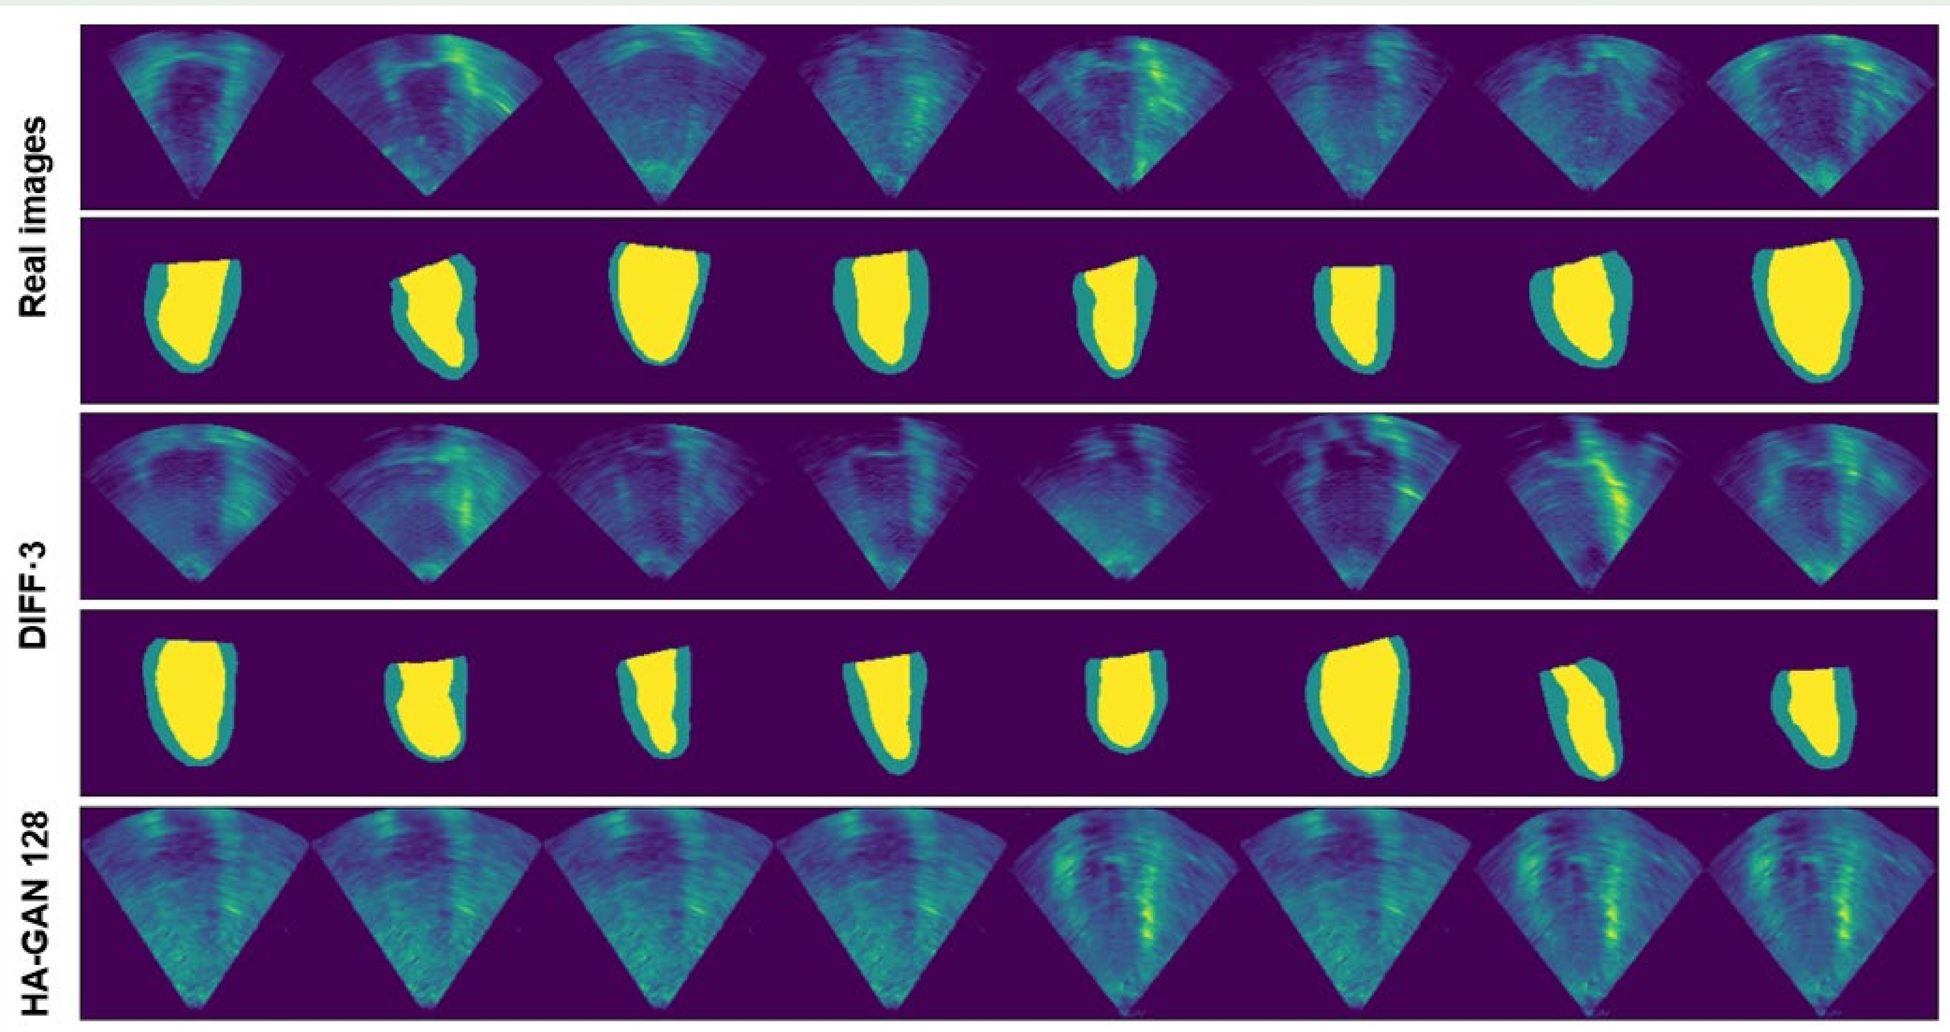

Figure 1. Real and synthetic samples of 3D echocardiograms and their corresponding labels. (Note that HA-GAN does not produce corresponding labels as it was trained unconditionally.)

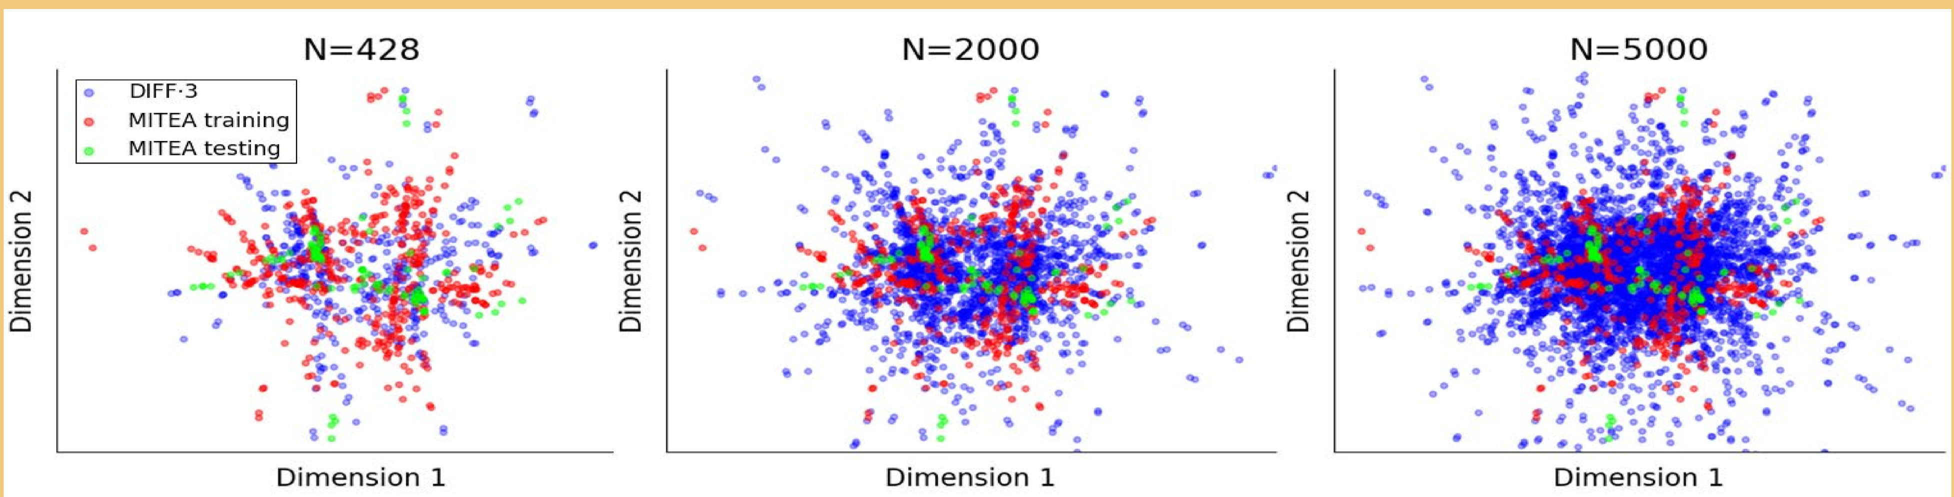

Figure 2. Locally Linear Embedding of the real and generated synthetic images with n=428, 2000, and 5000, respectively.

RESULTS

Table 1. Quantitative evaluation using Fréchet Inception Distance (FID), Improved Precision (IP), Improved Recall (IR), and multi-scale structural similarity (MS-SSIM) for real MITEA data, synthetic data generated using DIFF·3 and HA-GAN against the reference MITEA dataset.

Data	FID↓	IP↑	IR↑	MS-SSIM↓
Real (n=108)	14.0	0.731	0.797	0.781 (n=428)
DIFF·3 (n=428)	13.7	0.888	0.694	0.710
HA-GAN 128 (n=428)	46.7	0.486	0.000	0.999
HA-GAN 256 (n=428)	101.2	0.000	0.000	0.999

Evaluations

- VAE Reconstruction on test data (n=108)
  - SSIM = 0.84 ± 0.03
  - Dice Myocardium = 0.97 ± 0.01
  - Dice cavity = 0.99 ± 0.01
- Generated 5000 synthetic data with paired labels
- Visualized the data distribution using Locally Linear Embedding (LLE)

Table 2. Quantitative evaluation on segmentation task where we retrain nnUnet on real MITEA data + 2000 synthetic data generated by **DIFF·3**

nnUnet	Dice myo↑	Dice cavity↑
MITEA (n=428)	**0.766**	**0.871**
MITEA (n=428) **+ DIFF·3 (n=2000)**	0.762	0.866

CONCLUSIONS

- DIFF·3 provides an effective and efficient means of generating paired 3D echocardiograms and labels to supplement real patient data for deep learning applications.
- Improvement in terms of image quality may improve downstream tasks such as segmentation.

UNIVERSITY OF AUCKLAND Waipapa Taumata Rau NEW ZEALAND

AUCKLAND BIOENGINEERING INSTITUTE

ACKNOWLEDGEMENTS:

This study was funded by the Health Research Council of New Zealand (programme grant 17/608). RD holds the New Zealand Heart Foundation Chair of Heart Health.

Code: https://github.com/EdwardFerdian/diff-3

MITEA: https://www.cardiacatlas.org/mitea/