CISTIB Center for Computational Imaging & Simulation Technologies in Biomedicine

UNIVERSITY OF LEEDS

MANCHESTER 1824 The University of Manchester

The Alan Turing Institute

Haoran Dou

Demo

A Conditional Flow Variational Autoencoder for Controllable Synthesis of Virtual Populations of Anatomy

Haoran Dou[1], Nishant Ravikumar[1], Alejandro F. Frangi[1,2,3,4]

[1]Center for Computational Imaging & Simulation Technologies in Biomedicine (CISTIB), University of Leeds, UK
[2]Division of Informatics, Imaging and Data Science, Schools of Computer Science and Health Sciences, University of Manchester, UK
[3]Medical Imaging Research Center (MIRC), Electrical Engineering and Cardiovascular Sciences Departments, KU Leuven, Belgium
[4]Alan Turing Institute, UK

Introduction

In-silico trials (ISTs) use computational modelling and simulation techniques with virtual twin or patient models of anatomy and physiology to evaluate the safety and efficacy of medical devices virtually. Virtual patient populations (VPs), distinct from virtual twin populations, comprise plausible instances of anatomy and physiology that do not represent any specific real patient's data (as in the case of the latter, viz. virtual twins). In other words, VPs comprise synthetic data that help expand/enrich the diversity of anatomical and physiological characteristics that can be investigated within an IST for a given medical device. A key aspect of patient recruitment in real clinical trials used to assess device performance and generate regulatory evidence for device approval is the clear definition of inclusion and exclusion criteria for the trial. These criteria define the target patient population considered appropriate/safe to assess the performance of the device of interest. Consequently, it is desirable to enable the controlled synthesis of VPs that may be used for device ISTs, in a manner that emulates the imposition of trial inclusion and exclusion criteria.

Highlights

- We introduce normalising flows (**NF**) to help the generative model capture greater anatomical variability from the observed real population, leading to the synthesis of more diverse VPs.
- We condition the flow-based VAE on patient demographic data and clinical measurements to enable conditional synthesis of plausible VPs.

Methods

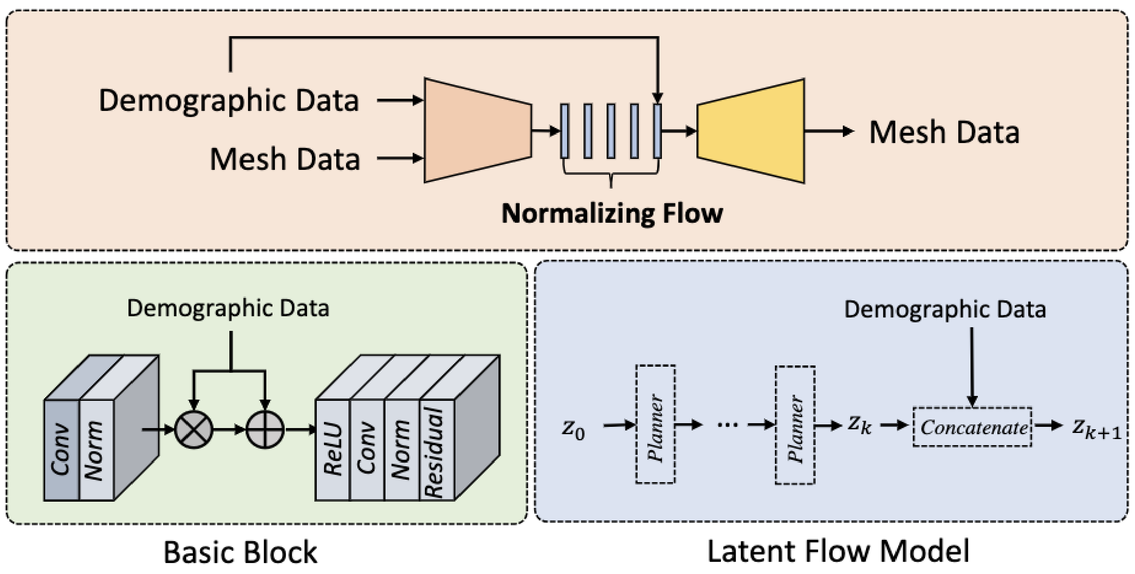

Fig. 1. Schematic illustration of our proposed conditional flow VAE

As shown in Fig.1, our conditional flow VAE is a graph-convolutional network that takes as input a triangular surface mesh representation of an anatomical structure of interest and its associated covariates/conditioning variables, i.e., the patient demographic data and clinical measurements and outputs the reconstructed surface mesh. The network is built based on CoMA equipped with the NFs in the bottleneck of the autoencoder. The NFs construct a flexible multi-modal latent posterior distribution by applying a series of differentiable, invertible/diffeomorphic transformations iteratively to the initial simple unimodal latent distribution (see Fig.2).

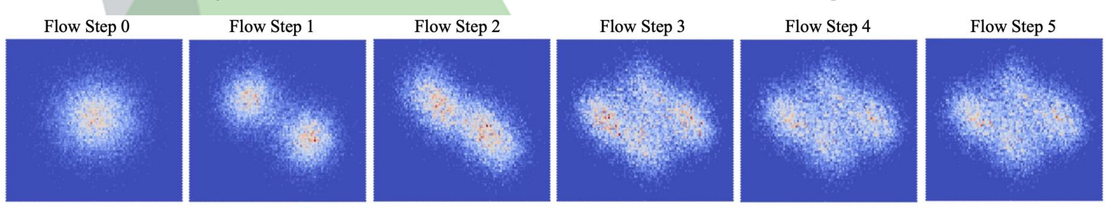

Fig. 2. Effect of normalising flow on Gaussian distribution. Step 0 is the initial two-dimensional Gaussian distribution, and step 1-5 represents the distribution of latent variables transformed by the normalising flow layers (i.e., planar flow).

Results

Data: We created a cohort of 2360 triangular meshes of the left ventricle (LV) based on a subset of the UK Biobank. We randomly split the data set into 422/59/1879 for training, validation, and testing, respectively.

Quantitative Analysis:

The following table illustrates the quantitative results of the investigated methods in a hold-out test dataset in terms of reconstruction error, specificity error, and volume variability.

Methods	Reconstruction Error↓	Specificity Error↓	Volume Variability↑
cVAE	1.43±0.26	**1.32±0.21**	28.39
Ours	**1.23±0.23**	1.38±0.20	**29.91**

Exploring Scenario:

Two representative examples of the reconstructed shapes and their variations through manipulation over two demographic attributes, i.e., weight and age. MyoVol and BPVol are shown in the bottom right corner.

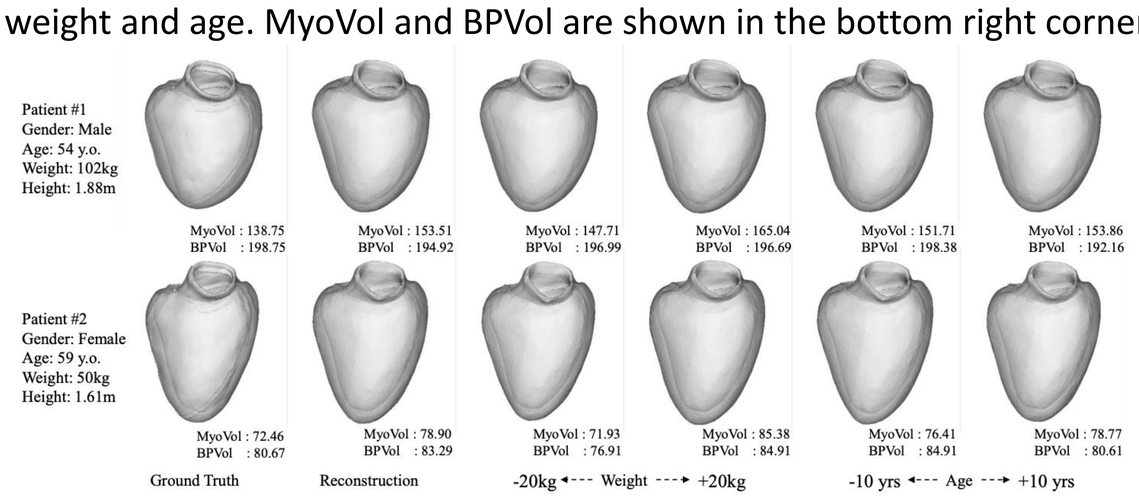

Virtual population Validation

Kernel density plots of 1k patients for BPVol from the VPs generated by our method and the real patient population (UKBB) in terms of gender.

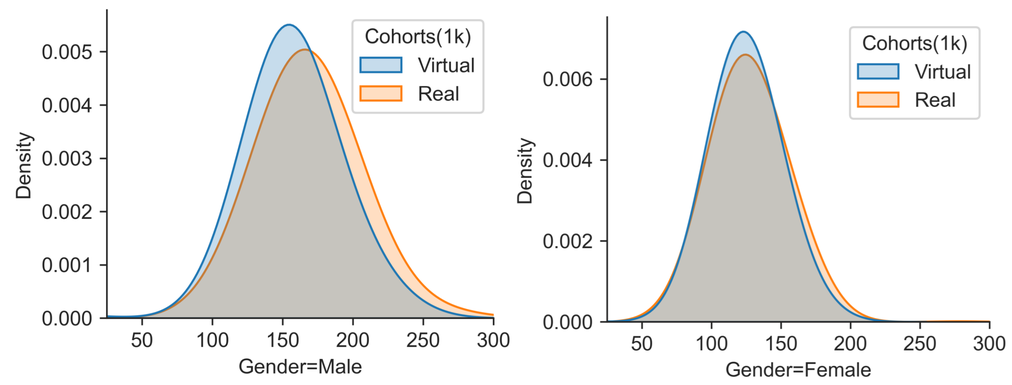

Comparisons of virtual population and real clinical population on correlations between LV/RV Volume and weight.

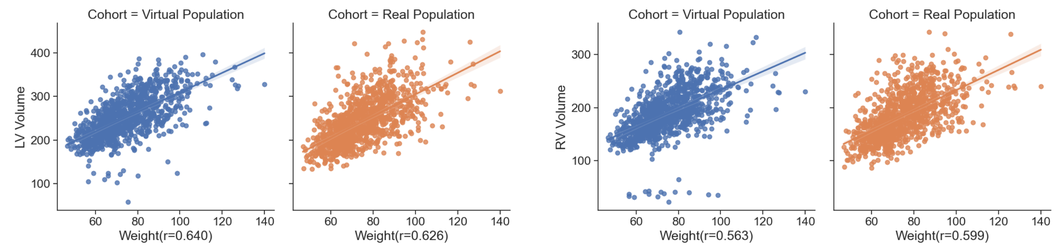

Conclusions

We proposed a conditional flow VAE model for the controllable synthesis of VPs of anatomy. Our approach was demonstrated to increase the flexibility of the learnt latent distribution and reflect the characteristics of the real population accurately. The results suggest that our approach has the potential for the controllable synthesis of diverse, yet plausible, VPs of anatomy. Future work will focus on modelling the whole heart and exploring the impact of individual covariates on VP synthesis in more detail.

ACKNOWLEDGMENT

This research was carried out using data from the UK Biobank (access application 11350). This work was supported by the Royal Academy of Engineering (INSILEX CiET1819/19), Engineering and Physical Sciences Research Council (EPSRC) UKRI Frontier Research Guarantee Programmes (INSILICO, EP/Y030494/1), and the Royal Society Exchange Programme CROSSLINK IES\NSFC\201380.