Conditional generative modeling of cardiac anatomy with flow-based architectures

MASTER THESIS PROJECT

TU Delft

Supervisors:

Konstantinos Kevopoulos
Dr. Alexander Heinlein
Dr. ir. Mathias Peirlinck

October 6, 2025

Delft University of Technology

1 Problem description

There is considerable natural variation in the shape of the human heart among healthy individuals. This variation is influenced by demographic and anthropometric factors such as age, sex, body mass index (BMI), and body surface area (BSA). Understanding these shape differences is important since cardiac geometry directly affects heart function, and deviations from expected shape patterns may indicate early stages of disease. To study the shape variations of large cohorts in a data-driven manner, data-driven generative models should be developed, that are able to learn the underlying anthropometric and demographic-conditioned shape distributions.

Several different approaches have been proposed to develop generative models of cardiac anatomy, including architectures based on variational autoencoders (VAEs) [2, 3, 4, 8, 9], or generative adversarial networks (GANs) [1, 5, 6, 7]. However, both modeling methods have intrinsic limitations that complicate the training and generative process. VAEs usually constrain the autoencoder latent space to a standard normal distribution, limiting the expressivity of the architectures, while the training process in GANs can be highly unstable, due to the balance to be maintained between the generator and discriminator.

To circumvent the aforementioned issues, more recently, we have proposed a generative model based on normalizing flows (paper in preparation). The approximation results of this alternative formulation for generation have been promising, highlighting an interesting research direction towards this kind of methods for the generation of cardiac anatomies.

2 Research objective

In the developed normalizing flow model, we aim to learn the shape distribution, according to patient-specific characteristics such as age, sex, BMI. We call these characteristics the metadata. After training the framework, the goal is to generate synthetic cardiac anatomies by only inputting metadata information to the model.

This project aims to extend the developed normalizing flow framework, to make it potentially simpler, more efficient, or more explainable. Depending on the experience and motivation of the student, several research directions can be explored, such as:

- Leverage the capability of normalizing flows to compute the exact likelihood of data, to derive a data-driven biomarker for early identification of cardiac disease, and anatomical abnormalities. Investigate the efficiency of the biomarker in identifying diseased patients.
- Make the generative model more parameter-efficient, or further enhance its ability to approximate "reference" distributions obtained from data.
- Examples of extending the current architecture would be to introduce attention-based mechanisms in the normalizing flow model, or substitute the normalizing flow with flow matching. Other extensions could also be explored by the student.
- Analyse the new framework in terms of efficiency and approximation abilities. Investigate what the features are that make the model better. Investigate if there are any additional benefits in the explainability of the new architecture.

3 Candidate requirements

- Hands-on experience with building, training, and testing deep learning models.
- Experience in performing statistical analyses, and understanding basic statistics
- Experience with developing and training (generative) deep learning models (VAEs, GANs, denoising diffusion, normalizing flows, flow matching), or motivation to get hands-on experience with such models.

4 Envisioned tasks

- Literature review regarding conditional generative models in cardiac anatomy, and frameworks that could extend the current model.
- Familiarise with the developed normalizing flow model. Formulate some concrete ideas to improve the method.

- Implement your ideas. Evaluate the new model (or the derived biomarker) in terms of efficiency and explainability
- Perform more experiments to illustrate the framework performance, and write the dissertation.

References

[1] Amirrajab, S., Abbasi-Sureshjani, S., Al Khalil, Y., Lorenz, C., Weese, J., Pluim, J., and Breeuwer, M. (2020). Xcat-gan for synthesizing 3d consistent labeled cardiac mr images on anatomically variable xcat phantoms. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part IV 23*, pages 128–137. Springer.

[2] Beetz, M., Banerjee, A., and Grau, V. (2022a). Multi-domain variational autoencoders for combined modeling of mri-based biventricular anatomy and ecg-based cardiac electrophysiology. *Frontiers in physiology*, 13:886723.

[3] Beetz, M., Corral Acero, J., Banerjee, A., Eitel, I., Zacur, E., Lange, T., Stiermaier, T., Evertz, R., Backhaus, S. J., Thiele, H., et al. (2022b). Interpretable cardiac anatomy modeling using variational mesh autoencoders. *Frontiers in Cardiovascular Medicine*, 9:983868.

[4] Dou, H., Ravikumar, N., and Frangi, A. F. (2023). A conditional flow variational autoencoder for controllable synthesis of virtual populations of anatomy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 143–152. Springer.

[5] Liao, H., Tang, Y., Funka-Lea, G., Luo, J., and Zhou, S. K. (2018). More knowledge is better: Cross-modality volume completion and 3d+ 2d segmentation for intracardiac echocardiography contouring. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11*, pages 535–543. Springer.

[6] Liu, Y., Dwivedi, G., Boussaid, F., Sanfilippo, F., Yamada, M., and Bennamoun, M. (2023). Inflating 2d convolution weights for efficient generation of 3d medical images. *Computer Methods and Programs in Biomedicine*, 240:107685.

[7] Qiao, M., Basaran, B. D., Qiu, H., Wang, S., Guo, Y., Wang, Y., Matthews, P. M., Rueckert, D., and Bai, W. (2022). Generative modelling of the ageing heart with cross-sectional imaging and clinical data. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 3–12. Springer.

[8] Qiao, M., McGurk, K. A., Wang, S., Matthews, P. M., O'Regan, D. P., and Bai, W. (2025). A personalized time-resolved 3d mesh generative model for unveiling normal heart dynamics. *Nature Machine Intelligence*, pages 1–12.

[9] Qiao, M., Wang, S., Qiu, H., De Marvao, A., O'Regan, D. P., Rueckert, D., and Bai, W. (2023). Cheart: A conditional spatio-temporal generative model for cardiac anatomy. *IEEE transactions on medical imaging*, 43(3):1259–1269.