Abstract

Diffusion models have recently gained immense interest for their generative capabilities, specifically the high quality and diversity of the produced data. However, examples of their applications in 3D medical imaging are still scarce, especially in the field of cardiology. The ability to generate diverse realistic cardiac anatomies is crucial for applications such as in silico trials - electrophysiological computer simulations or data augmentation for machine learning solutions (e.g. cardiac disease classifiers). In this work, we investigate the application of Latent Diffusion Models (LDMs) for generating 3D meshes of human cardiac anatomies. To this end, we propose a novel Latent Diffusion Model architecture - Mesh LDM. We evaluate the proposed model on a dataset of 3D meshes of left ventricle cardiac anatomies from patients with acute myocardial infarction, using clinical metrics and visuals. The proposed model successfully captures the characteristics of the cardiac shapes at the end-systolic (contraction) and end-diastolic (relaxation) cardiac phases, generating meshes with a 2% difference in population mean compared to the gold standard. We also analyse the limitations of our model and the utilised variational mesh autoencoder. Finally, we propose solutions to resolve them and discuss potential future extensions of our work.