ACQuA: Anomaly Classification with Quasi-Attractors

William Rudman, Jack Merullo, Laura Mercurio, Carsten Eickhoff

BrownBAI

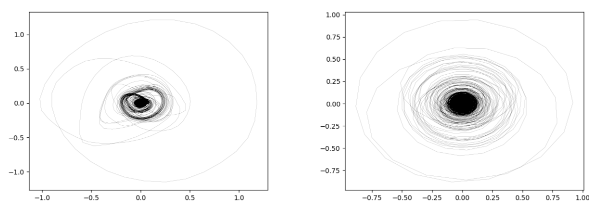

Figure 1: Sample normal and abnormal sample attractor images taken from an aortic valve point.

In recent years, deep learning has redefined algorithms for detecting cardiac abnormalities. However, many state of the art algorithms still rely on calculating handcrafted features from a given heart signal that are then fed into a shallow 1D convolutional neural network or transformer architecture. We propose ACQuA (Anomaly Classification with Quasi Attractors), a task agnostic algorithm that can be used in a wide variety of cardiac settings, from classifying cardiac arrhythmias from ECG signals to detecting the presence of heart murmurs from PCG signals.

Using theorems from dynamical analysis and topological data analysis (TDA), we create informative attractor images that 1) are human distinguishable and 2) can be used to train small, off the shelf deep neural networks for anomaly classification. We can approximate the attractor of a dynamical system (i.e. the set of points for which the system converges to over time) by using time-delay embeddings that project a time series onto an M dimensional manifold[1]. Quasi-attractors have been previously used to perform signal classification in the field of topological data analysis, however, TDA-based approaches for extracting informative features of the quasi-attractor are extremely computationally expensive. We overcome this constraint by performing principal component analysis to compress our M dimensional attractors into a 2D images. These images can then be fed into a convolutional neural network and trained to differentiate cardiac anomalies with an image classification objective. Using an off-the-shelf ResNet-18 architecture, we are able to automatically extract the most salient features from the quasi-attractor for signal classification, obviating the need to perform expensive post-processing. We receive an official challenge score of 1764 without any pretraining or hyperparameter tuning.

[1] According to *Taken's Theorem*, if the embedding dimension is sufficiently large, the time-delay embedding will be diffeomorphic to the true attractor of the dynamical system. Studies typically make the naïve choice of setting $M = 2$. In contrast, we set $M = 36$ when creating attractors from PCG signals.