

Anomaly Detection in Echocardiograms with Dynamic Variational Trajectory Models

Alain Ryser¹, Laura Manduchi¹, Fabian Laumer¹, Holger Michel², Sven Wellmann², and Julia E. Vogt¹

¹Department of Computer Science, ETH Zurich

²Department of Neonatology, University Children’s Hospital Regensburg (KUNO), University of Regensburg, Germany

Abstract

We propose a novel anomaly detection method for echocardiogram videos. The introduced method takes advantage of the periodic nature of the heart cycle to learn three variants of a *variational latent trajectory* model (TVAE). While the first two variants (TVAE-C and TVA-E-R) model strict periodic movements of the heart, the third (TVAE-S) is more general and allows shifts in the spatial representation throughout the video. All models are trained on the healthy samples of a novel in-house dataset of infant echocardiogram videos consisting of multiple chamber views to learn a normative prior of the healthy population. During inference, maximum a posteriori (MAP) based anomaly detection is performed to detect out-of-distribution samples in our dataset. The proposed method reliably identifies severe congenital heart defects, such as Ebstein’s Anomaly or Shone-complex. Moreover, it achieves superior performance over MAP-based anomaly detection with standard variational autoencoders when detecting pulmonary hypertension and right ventricular dilation. Finally, we demonstrate that the proposed method enables interpretable explanations of its output through heatmaps highlighting the regions corresponding to anomalous heart structures.

1. Introduction

Congenital heart defects (CHDs) account for about 28% of all congenital defects worldwide (Van Der Linde et al., 2011). CHDs manifest in several different heart diseases with various degrees of frequency and severity and are usually diagnosed primarily with echocardiography. Echocardiography is one of the most common non-invasive screening tools due to the rapid data acquisition, low cost, portability, and measurement without ionizing radiation. Early screening of heart defects in newborns is crucial to ensure the long-term health of the patient (Buskens et al., 1996; Singh and McGeoch, 2016; Van Velzen et al., 2016). However, due to the subtlety of various heart defects and the inherently noisy nature of echocardiogram video (echo) data, a thorough examination of the heart and the diagnosis of CHD remains challenging and time-consuming, raising the need for an automated approach. Still, collecting real-world datasets from large populations to apply state-of-the-art supervised deep learning methods is often infeasible. The reason is that many CHDs like Ebstein’s Anomaly, Shone-complex, or complete atrioventricular septal defect (cAVSD) rarely occur, making the dataset extremely imbalanced. On the other hand, we have access

to an abundance of echos from healthy infant hearts generated during standard screening procedures, often performed on infants shortly after birth. In this work, we leverage the healthy population and propose a novel anomaly detection method to identify a variety of CHDs. The proposed approach learns a structured normative prior of healthy newborn echos using a periodic variational latent trajectory model. At test time, the method can detect out-of-distribution samples corresponding to CHDs. The advantage of this approach is that the model is trained purely on healthy samples, eliminating the need to collect large amounts of often rarely occurring CHDs.

In anomaly detection, we assume that all data is drawn from a space \mathcal{X} with some probability density p_H . We define anomalies as samples drawn from low probability regions of \mathcal{X} under p_H . More formally, the space of anomalies $\mathcal{A} \subset \mathcal{X}$ under density p_H and anomaly threshold $\tau \geq 0$ is defined by

$$\mathcal{A} = \{x \in \mathcal{X}; p_H(x) \leq \tau\}$$

Note that τ is a task-specific measure, as the definition of anomaly can vary drastically over different problem settings. Consequently, most anomaly detection algorithms assign anomaly scores rather than discriminating between normal and anomalous samples.

In this work, we focus on reconstruction-based approaches, which encompass some of the most widespread methods for anomaly detection (Chalapathy and Chawla, 2019; Ruff et al., 2021; Pang et al., 2021). This family of methods aims to learn generative models that can reconstruct normal samples well but decrease in performance for anomalous inputs. A given measure $\alpha_f(x)$ that quantifies the reconstruction quality achieved by model f when given sample x can then be interpreted as the anomaly score of x . The models are commonly trained on healthy samples, and during inference, an anomalous sample x_a is assumed to get projected into the learned normal latent space. This leads to high reconstruction errors, resulting in high anomaly scores $\alpha_f(x_a)$. More recently, Chen et al. (2020) proposed a variation of the reconstruction-based approach that allows us to incorporate prior knowledge on anomalies during inference by detecting anomalies using a *maximum a posteriori* (MAP) based approach. However, this approach requires an estimate of the log-likelihood, which restricts model choice to generative models such as *variational autoencoders* (VAE Kingma and Welling (2013)).

Although various generative architectures have been proposed in the literature, little effort has been directed toward echocardiogram videos. One exception is the work of Laumer et al. (2020), where the authors introduced a model that specifically targets the periodicity of heartbeats for ejection fraction prediction and arrhythmia classification. However, the model enforces somewhat restrictive assumptions on the heart dynamics and is purely deterministic. In contrast, we propose a variational latent trajectory model that overcomes the simplistic assumptions of previous approaches and learns a distribution over dynamic trajectories, enabling the detection of different types of CHDs in echocardiograms using the MAP approach. Furthermore, the proposed algorithm produces incredibly accurate video reconstructions from surprisingly few latent factors.

To allow the community to reuse our methods in future work, we provide code for the TVAE and pretrained models used in our experiments on Github¹.

1. <https://github.com/alain-ryser/tvae>

Generalizable Insights about Machine Learning in the Context of Healthcare

This work introduces a novel generative model designed explicitly for echocardiogram videos. By leveraging key assumptions on this data modality, our model allows a flexible encoding of videos in a small number of latent dimensions from which accurate reconstructions can be retrieved. We demonstrate how this method outperforms previous approaches on an anomaly detection task using a new in-house echo video dataset containing samples of newborns and infants with various forms of congenital heart defects. Our model learns an accurate normative prior on healthy echo data and then performs maximum a posteriori-based anomaly detection to detect CHDs. Additionally, we demonstrate how our model produces interpretable outputs by showcasing decision heatmaps that highlight regions that drive anomaly scores. To summarize, the contributions of this paper are the following:

1. We propose a novel variational latent trajectory model (TVAE) for reconstruction-based anomaly detection on echocardiogram videos.
2. We perform extensive evaluation of the proposed method on the challenging task of CHD detection in a real-world dataset.
3. We complement our predictions with decision heatmaps, highlighting the echocardiogram regions corresponding to anomalous heart structures.

2. Related Work

The rapid data acquisition, the high observer variation in their interpretation, and the non-invasive technology have made echocardiography a suitable data modality for many machine learning algorithms. In recent years, a variety of algorithms for *Segmentation* (Dong et al., 2016; Moradi et al., 2019; Leclerc et al., 2019), *View Classification* (Gao et al., 2017; Vaseli et al., 2019) or *Disease Prediction* (Madani et al., 2018; Kwon et al., 2019) have been proposed. However, their performance often relies on the assumption that a large *labeled* dataset can be collected. This assumption does not hold for rare diseases, where the amount of collected data is often too scarce to train a supervised algorithm. Hence, reconstruction-based anomaly detection algorithms could be used in such a setting, but their application to echocardiography is, to the best of our knowledge, left unexplored.

Previous work on reconstruction-based anomaly detection are often based on generative models, such as *autoencoders* (AE) (Chen et al., 2017; Principi et al., 2017; Chen and Konukoglu, 2018; Pawlowski et al., 2018) or *variational autoencoders* (VAE Kingma and Welling (2013)) (An and Cho, 2015; Park et al., 2018; Xu et al., 2018; Cerri et al., 2019; You et al., 2019). Their application to the medical domain is mostly limited to disease detection in MRI, (Chen and Konukoglu, 2018; Baur et al., 2018, 2020; Chen et al., 2020; Baur et al., 2021; Pinaya et al., 2021) where anomalies are often easily detectable as they are clearly defined by regions of tissue that contain lesions. On the other hand, pathologies of CHDs in echos are largely heterogeneous and can usually not be described by unique structural differences from healthy echos. Identifying them is often challenging, as they can be caused by small perturbations of ventricles (ventricular dilation) or subtle malfunctions like pressure differences between chambers in certain phases of the cardiac cycle (pulmonary

hypertension). Detecting certain CHDs thus requires the inclusion of temporal structures in addition to the spatial information leveraged in MRI anomaly detection.

Different extensions to AE/VAE have been proposed to perform reconstruction-based anomaly detection methods on video data (Xu et al., 2015; Hasan et al., 2016; Yan et al., 2018). However, these methods are often mainly designed for abnormal event detection, where anomalies can arise and disappear throughout the video. On the other hand, we are interested in whether a given video represents a healthy or anomalous heart. Another method for video anomaly detection is *future frame prediction* (Liu et al., 2018). This approach trains models to predict a video frame from one or more previous ones. During inference, it is then assumed that such a model achieves better performance on normal than on anomalous frames. Recently, Yu et al. (2020) proposed a method that combines reconstruction and future frame prediction-based approaches in one framework. Though achieving good performance on videos with varying scenes, future frame prediction does not seem suitable for echos as just returning any input frame will always lead to good prediction scores due to the periodic nature of the cardiac cycle. An entirely different approach to anomaly detection is given by *One-Class Classification* (Moya and Hush, 1996). In contrast to the previous approaches, the latter relies on discriminating anomalies from normal samples instead of assigning an anomaly score. This is usually achieved by learning a high-dimensional manifold that encloses most or all normal data. The surface of this manifold then serves as a decision boundary that discriminates anomalies from normal samples. One of the more prominent methods of that family is the so-called *Support Vector Data Description* (SVDD) (Tax and Duin, 2004) model. The SVDD learns parameters of a hypersphere that encloses the training data. Similar to SVMs, it provides a way to introduce some slack into the estimation process, allowing certain normal samples to lie outside the decision boundary. A similar approach is given by the *One-Class SVMs* (OC-SVM) (Schölkopf et al., 2001), where anomalies are discriminated from normal samples by learning a hyperplane instead of a hypersphere. Like with SVMs, the expressivity of SVDD and OC-SVM can be drastically improved by introducing kernelized versions (Ratsch et al., 2002; Ghasemi et al., 2012; Dufrenois, 2014; Gautam et al., 2019). More recently, deep neural networks have been proposed to perform anomaly detection based on similar principles (Sabokrou et al., 2018; Ruff et al., 2018, 2020; Ghafoori and Leckie, 2020). While conceptually interesting, One-Class Classification methods often require large amounts of data to work accurately, making them unsuitable in many clinical applications.

3. Methods

In this work, we propose a probabilistic latent trajectory model to perform reconstruction-based anomaly detection on echocardiogram videos. To that end, we take inspiration from latent trajectory models (Louis et al., 2019; Laumer et al., 2020) and introduce a variational autoencoder that learns a structured normative distribution of the heart’s shape and dynamic. In particular, the model encodes the echos into stochastic trajectories in the latent space of a VAE, enabling us to accurately generate high-quality reconstructions while maintaining a low dimensional latent bottleneck. The learned approximate distribution of healthy hearts allows us to detect anomalies post-hoc using a maximum a posteriori (MAP)

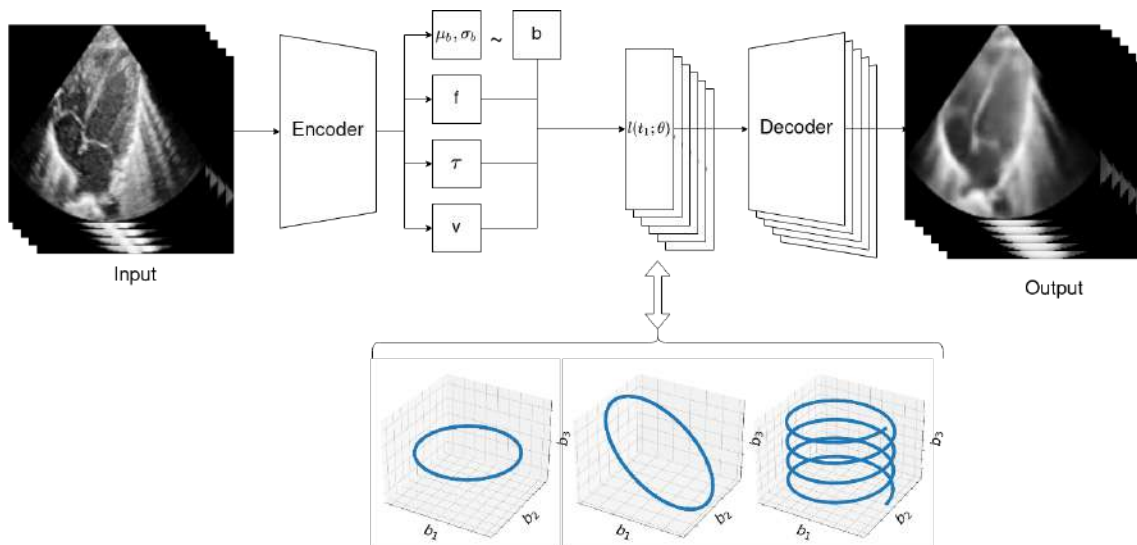


Figure 1: Overview of the model architecture with $\ell_{circular}$ (left), ℓ_{rot} (middle) and ℓ_{spiral} (right).

approach (Chen et al., 2020). High-quality normative reconstructions and informative latent representations are essential for correctly detecting out-of-distribution echos.

3.1. Latent Trajectory Model

The latent trajectory model (Laumer et al., 2020) is an autoencoder that is designed to learn latent representations from periodic sequences of the heart, i.e. echos in this case. The main idea is to capture the *periodic nature* of the observed data by learning an encoder ϕ that maps an echo $X := (\mathbf{x}^{(j)}, t^{(j)})_{j=1}^T$ with frames $\mathbf{x}^{(j)} \in \mathbb{R}^{w \times h}$ at time points $t^{(j)}$ to a prototypical function $\ell_{circular}(t; \phi(X))$ whose parameters contain information about the heart’s shape and dynamic. The decoder ψ reconstructs the original video frame by frame from the latent embedding $\ell_{circular}$ via

$$\tilde{\mathbf{x}}^{(j)} = \psi(\ell_{circular}(t^{(j)}; \phi(X)))$$

Here, $\ell_{circular}$ corresponds to the following cyclic trajectory:

$$\ell_{circular}(t; f, \omega, \mathbf{b}) = \begin{pmatrix} \cos(2\pi ft - \omega) + b_1 \\ \sin(2\pi ft - \omega) + b_2 \\ b_3 \\ \vdots \\ b_d \end{pmatrix},$$

where the frequency parameter, $f > 0$, corresponds to the number of cycles per time unit, and the offset parameter $\omega \in [0, 2\pi]$ allows the sequence to start at an arbitrary point within the (cardiac) cycle. The parameter $\mathbf{b} \in \mathbb{R}^d$ characterizes the *spatial information* of

the signal. See Figure 1 a) for an illustration of $\ell_{circular}$. This model thus describes a simple tool to learn the disentanglement of temporal components (f, ω) from a common spatial representation (\mathbf{b}) for a given echo. On the other hand, the assumptions made may be too simplistic to result in good reconstructions. We will address this issue in the following sections.

3.2. Dynamic Trajectories

The above formulation, $\ell_{circular}$, allows modeling time-related information only through the first two latent dimensions, thereby limiting the amount of time-dependent information that can be encoded in the latent space. The reduced flexibility results in insufficient reconstruction quality, impairing the reconstruction-based anomaly detection performance. To circumvent this problem, we distribute time-dependent components over each dimension of the latent space while retaining the periodicity. We thus define the rotated trajectory function ℓ_{rot} as

$$\ell_{rot}(t; f, \omega, \mathbf{b}) = \begin{pmatrix} \cos(2\pi ft - \omega) - \sin(2\pi ft - \omega) + b^{(1)} \\ \cos(2\pi ft - \omega) + \sin(2\pi ft - \omega) + b^{(2)} \\ \vdots \\ \cos(2\pi ft - \omega) + \sin(2\pi ft - \omega) + b^{(d)} \end{pmatrix}.$$

See Figure 1 b) for an illustration of ℓ_{rot} .

Furthermore, in real-world applications, it is often the case that doctors may either change the settings of the echocardiogram machine during screening or zoom in or out to get better views of specific cardiac structures. Additionally, some patients might slightly move while scans are performed, which leads to a heart displacement with respect to the transducer position throughout an echo recording. This is particularly prominent in our in-house dataset, which consists of echocardiograms of newborn children. Such echocardiograms are not necessarily well represented with a simple periodic trajectory, as over multiple cycles, the spatial structure of a sample shifts and looks different than in the beginning, even though temporal information like the frequency or phase shift is preserved. Thus, the current trajectory model fails in such scenarios, which can manifest in two ways: either the model gets stuck in a local optimum with high reconstruction error, or the model tries to reconstruct the video from one long cycle, hence not leveraging the heart cycle periodicity. Hence, to account for movements of the recording device, we extend ℓ_{rot} with a velocity parameter $v \in \mathbb{R}$ that allows the model to learn gradual shifts of the latent trajectory over time, resulting in a trajectory that is no longer circular but a spiral embedded in high dimensional space. More formally, we define the spiral trajectory function as

$$\ell_{spiral}(t; f, \omega, v, \mathbf{b})_i = \ell_{rot}(t; f, \omega, \mathbf{b})_i + tv$$

See Figure 1 c) for an illustration of the spiral model.

3.2.1. VARIATIONAL FORMULATION

Previous work often applied VAEs to anomaly detection, as its generative nature enables more sophisticated variants of reconstruction-based anomaly detection (Baur et al., 2018;

Chen et al., 2020; Xu et al., 2018). However, the current latent trajectory model is purely deterministic. Thus, we introduce the variational latent trajectory model and perform a post-hoc MAP estimation to detect anomalies similar to Chen et al. (2020).

We modify the encoder $\phi(X; \theta)$ such that it outputs trajectory parameters $v, f, \omega \in \mathbb{R}$ and $\boldsymbol{\mu}_b, \boldsymbol{\sigma}_b \in \mathbb{R}^d$. The model is then extended with a stochastic layer by defining $\mathbf{b} \sim q_\theta(\mathbf{b}|X) := \mathcal{N}(\boldsymbol{\mu}_b, \text{diag}(\boldsymbol{\sigma}_b))$. While we aim to learn a distribution over heart shapes, we would also like to accurately identify the frequency f , phase shift ω , and spatial shift v given an echo video X , instead of sampling them from a latent distribution. We thus leave those parameters deterministic. Next, we define an isotropic Gaussian prior $p(\mathbf{b}) := \mathcal{N}(0, \mathbb{I})$ on \mathbf{b} and assume that $x^{(i)} \sim p_\eta(X|\mathbf{b}, f, \omega, v) := \mathcal{N}(\psi(\ell_{\text{spiral}}(t^{(i)}; f, \omega, v, \mathbf{b}); \eta), \sigma \mathbb{I})$, where ψ is our decoder with weights η and σ is some fixed constant. Given these assumptions, we can derive the following *evidence lower bound* (ELBO):

$$ELBO(X) := E_{q_\theta(\mathbf{b}|X)}[\log(p_\eta(X|\mathbf{b}, \phi_f(X), \phi_\omega(X), \phi_v(X)))] - KL[q_\theta(\mathbf{b}|X)||p(\mathbf{b})]$$

Here, $\phi_f(X)$, $\phi_\omega(X)$ and $\phi_v(X)$ are the trajectory parameter outputs of the encoder ϕ for f, ω and v , respectively. Note that VAEs on ℓ_{circular} and ℓ_{rot} are defined in a similar fashion. A derivation of this ELBO can be found in Appendix A.

3.2.2. ANOMALY DETECTION

The variational formulation of the latent trajectory model allows us to perform anomaly detection by *Maximum a Posteriori* (MAP) inference as proposed in Chen et al. (2020). They suggest to model anomalies as an additive perturbation of a healthy sample. Following their reasoning we define a healthy sample $X_H := (\mathbf{x}_H^{(j)}, t^{(j)})_{j=1}^T \sim \mathcal{H}$, (anomalous) sample $Y := (\mathbf{y}^{(j)}, t^{(j)})_{j=1}^T \sim \mathcal{D}$, and anomaly perturbation $A := (\mathbf{a}^{(j)}, t^{(j)})_{j=1}^T$, where \mathcal{H} is the healthy data distribution and \mathcal{D} the overall data distribution, and assume that

$$\mathbf{y}^{(j)} = \mathbf{x}_H^{(j)} + \mathbf{a}^{(j)}$$

In the case of CHD, A could, e.g., remove walls between heart chambers or produce holes in the myocardium for certain frames. The anomaly score α can then be defined as $\alpha(Y) := \frac{1}{T} \sum_{j=1}^T \|\mathbf{a}^{(j)}\|_2^2$. When training a VAE on healthy samples only, i.e. $\mathbf{a}^{(j)} = 0$ for all $j \in \{1, \dots, T\}$, the variational latent trajectory model learns to approximate $P(X_H)$ by maximizing $ELBO(X_H)$. The usual MAP estimation maximizes the posterior distribution of X_H given Y . By Bayes' theorem

$$P(X_H|Y) \propto P(Y|X_H)P(X_H),$$

the concavity of the logarithm, as well as the fact that $\log(P(X_H)) \geq ELBO(X_H)$ it is then possible to estimate X_H by

$$\tilde{X}_H = \arg \max_{X_H} (\log(P(Y|X_H)) + ELBO(X_H))$$

To compute the anomaly score we compute $\tilde{\mathbf{a}}^{(j)} = \mathbf{y}^{(j)} - \tilde{\mathbf{x}}_H^{(j)}$ and arrive at $\alpha(Y) := \frac{1}{T} \sum_{t=1}^T \|\tilde{\mathbf{a}}^{(t)}\|_2^2$. Similar to Chen et al. (2020), we choose $\log P(Y|X) = \|(\mathbf{x}^{(j)} - \mathbf{y}^{(j)})_{j=1}^T\|_{TV}$,

where $\|\cdot\|_{TV}$ denotes the *Total Variation Norm* in ℓ_1 , as this incorporates the assumption that anomalies should consist of contiguous regions rather than single pixel perturbations. Note that since we have a temporal model, we can incorporate temporal gradients into the TV norm, i.e.

$$\|X\|_{TV} := \sum_{i=1}^w \sum_{j=1}^h \sum_{k=1}^T \|\nabla \mathbf{x}_{ij}^{(k)}\|_1$$

In our experiments, we approximate gradients by

$$\nabla \mathbf{x}_{ij}^{(k)} \approx \begin{pmatrix} \mathbf{x}_{(i+1)j}^{(k)} - \mathbf{x}_{(i-1)j}^{(k)} \\ \mathbf{x}_{i(j+1)}^{(k)} - \mathbf{x}_{i(j-1)}^{(k)} \\ \mathbf{x}_{ij}^{(k+1)} - \mathbf{x}_{ij}^{(k-1)} \end{pmatrix}$$

4. Cohort

Table 1: Cohort Statistics

Feature	Statistic
No. of Patients	192
No. of Patients with no CHD	69
No. of Patients with SSD	5
No. of Patients with PH	73
No. of Patients with RVDIL	73
Age (Days) (Mean± SD)	34 ± 48
111 Time until birth (Days) (Mean± SD)	232 ± 46
Weight (Gramms) (Mean± SD)	2774 ± 1227
Manufacturer (Ultrasound Machine / Transducer)	GE Logic S8 / S4-10 at 6 MHz
Original Video Size (pixel×pixels)	1440 × 866
Video length (frames) (Mean± SD)	122 ± 7
Video FPS	25 fps

The dataset for this study consists of echos of 192 newborns and infants up to one year of age collected between 2019 and 2020 at a single center by a single pediatric cardiologist. All examinations were performed with the GE Logic S8 ultrasound machine and contain 2D video sequences of at least 2 standard echo views, i.e., apical 4-chamber view (4CV) and parasternal long-axis view (PLAX). Of the 192 patients, 123 suffer from, potentially multiple, CHDs, and 69 are healthy. See Table 1 for more details.

In order to evaluate anomaly detection performance, a pediatric cardiologist labeled the dataset into three categories. These include *Pulmonary Hypertension* (PH), *Right Ventricular Dilation* (RVDil) and *Severe Structural Defects* (SSD). While PH and RVDil are well-defined pathologies, SSD was defined as a category of multiple rare but severe CHDs, including Ebsteins anomaly, anomalous left coronary artery origin from pulmonary artery (ALCAPA), atrio-ventricular discordance, and ventricular-artery concordance (AVD-VAC),

Shone-complex, total anomalous pulmonary venous drainage (TAPVD), tetralogy of Fallot (ToF) and complete atrioventricular septal defect (cAVSD). We illustrate examples for healthy, SSD, PH, and RVDil echos of both 4CV and PLAX views in Appendix B.

All collected echocardiograms were preprocessed by resizing them to 128×128 pixels. Additionally, histogram equalization was performed to increase the contrast of the frames, and pixel values were normalized to the range $[0, 1]$. Consequently, models in the experiments in the following section are trained and evaluated on the preprocessed videos.

5. Experiments

In addition to the *variational latent trajectory* (TVAE) model with the *circular* (TVAE-C), *rotated* (TVAE-R) and *spiral* (TVAE-S) trajectories described in Section 3.1, as a baseline, we train a standard variational autoencoder (Kingma and Welling, 2013) model on the individual video frames of the dataset.

We run experiments for each of the three CHD categories described in Section 4 by training the models exclusively on samples that do not exhibit these pathologies. Each experiment is evaluated on 10 separate data splits, leaving out 30 healthy patients for evaluation of PH and RVDil and 7 for SSD, respectively. Additionally, every experiment is performed on both the *apical four chamber* (4CV) and *parasternal long axis* (PLAX) views.


5.1. Implementation Details

We assume that any anomaly of the heart should always be visible for a certain period over the heart cycle. It thus suffices to have a model that reconstructs only a fixed number of video frames, as long as at least one heart cycle is present in the video. The collected videos are recorded with 24 frames per second (FPS), and we assume that a heart beats at least 30 times a minute. Therefore, we decided to subsample the video frequency to 12 FPS and reconstruct videos with a fixed length of 25 frames, which is enough to capture at least one cycle in every video.


Reconstructing a fixed number of frames enables us to implement efficient architectures to aggregate echo frames and predict the trajectory parameters. More specifically, we implement the encoder by concatenating all input frames of the video, hence treating them like different channels of an image, and passing them to a *residual* (He et al. (2016)) encoder backbone. Each frame $(x^{(i)}, t^{(i)})$ is then individually decoded by passing $\ell_{circular}(t^{(i)})$, $\ell_{rot}(t^{(i)})$ or $\ell_{spiral}(t^{(i)})$ to a *deconvolution* (Zeiler et al. (2010)) based decoder. To train the VAE, we used identical encoder and decoder architectures, only changing the first layer to take a single grayscale channel instead of 25 frames and adapting latent fully connected layers to match dimensions. For more detailed schematics of the architecture and an overview of the chosen hyperparameters like latent dimension, batch size, or learning rate, we refer to Appendix C.

We pretrained all models on the EchoDynamic dataset to speed up training convergence (Ouyang et al., 2020). As in most clinical applications, the scarcity of the data makes optimized models prone to overfitting. To prevent this, we apply data augmentation during training by transforming samples with random affine transformations, brightness adjustments, gamma corrections, blurring, and adding Salt and Pepper noise before performing the forward pass.

5.2. Reconstruction



(a) Healthy reconstructions



(b) SSD reconstructions

Figure 2: Examples of healthy (a) and SSD (b) samples (first and third rows) and their reconstructions (second and fourth rows) using the TVAE-S model. We sample 3 frames for each echo from the 25 frame long sequences.

Reconstruction quality is directly related to reconstruction-based anomaly detection performance, as we rely on the *manifold* and *prototype* assumptions formalized in [Ruff et al. \(2021\)](#). The manifold assumption is often used in many machine learning-based applications and states that \mathcal{X} , the space of healthy echos, can be generated from some latent space \mathcal{Z} by a decoding function ψ and that it is possible to learn a function ϕ that encodes \mathcal{X} into \mathcal{Z} . The *better* a learned function $f(x) := \psi(\phi(x))$ reconstructs x on a test set, the better

Table 2: Apical 4-chamber view reconstruction performance on test data of the proposed approaches (TVAE-C, TVAER and TVAES) compared with the baseline (VAE). Means and standard deviations are computed across 10 data splits.

		VAE	TVAE-C	TVAE-R	TVAE-S
SSD	MSE	0.013 \pm 0.0	0.014 \pm 0.0	0.014 \pm 0.0	0.013 \pm 0.0
	PSNR	19.008 \pm 0.11	18.574 \pm 0.15	18.58 \pm 0.13	18.774 \pm 0.09
	SSIM	0.545 \pm 0.01	0.544 \pm 0.01	0.545 \pm 0.01	0.552 \pm 0.01
RVDil	MSE	0.012 \pm 0.0	0.014 \pm 0.0	0.013 \pm 0.0	0.013 \pm 0.0
	PSNR	19.146 \pm 0.05	18.7 \pm 0.07	18.82 \pm 0.08	18.803 \pm 0.04
	SSIM	0.555 \pm 0.0	0.549 \pm 0.0	0.554 \pm 0.0	0.555 \pm 0.0
PH	MSE	0.012 \pm 0.0	0.014 \pm 0.0	0.014 \pm 0.0	0.014 \pm 0.0
	PSNR	19.084 \pm 0.07	18.66 \pm 0.07	18.723 \pm 0.08	18.727 \pm 0.07
	SSIM	0.552 \pm 0.0	0.55 \pm 0.0	0.551 \pm 0.0	0.553 \pm 0.0

we meet the manifold assumption. The prototype assumption, on the other hand, assumes that there is some set of prototypes that characterizes the healthy distribution well. In our case, the prototypes would be echos corresponding to healthy hearts, i.e., a subset of \mathcal{X} . Under the prototype assumption, our model f must be able to assign a given sample to one of the learned prototypes, i.e., project anomalies to the closest healthy echo.

Table 2 contains the scores of the VAE, TVAEC, TVAER, and TVAES with respect to the *Mean Squared Error* (MSE), *Peak Signal to Noise Ratio* (PSNR) and *Structural Similarity Index Measure* (SSIM). We observe how TVAEC has consistently higher MSE and SSIM errors and lower PSNR than both TVAER and TVAES. Upon inspection of the reconstructed test videos, we notice that, for most seeds, TVAEC converges to a local optimum where the model learns mean representations of the input videos, thus ignoring the latent dimensions containing temporal information, as described in Section 3. On the other hand, we did not observe this behavior in TVAER and TVAES, suggesting that these models indeed capture dynamic properties of echos through the learned latent representations. Additionally, TVAES achieves good echo reconstructions even for samples with transducer position displacement, improving upon TVAER and achieving similar performance as VAE despite having a smaller information bottleneck. The proposed approaches, TVAEC, TVAER, and TVAES, encode videos into $d + 2$ or $d + 3$ trajectory parameters respectively, while the VAE encodes each frame in \mathbb{R}^d , resulting in a total of $25 \times d$ latent parameters. In conclusion, TVAES and the standard VAE fulfill the manifold assumption. Figure 2 presents reconstructed healthy, and SSD samples for the 4CV and PLAX echo views.

In Figure 3, we qualitatively demonstrate that TVAES satisfies the prototype assumption. We observe how the perturbed septum and enlarged/shrunken heart chambers of SSD anomalies are projected to healthy echo reconstructions.

We provide more reconstructions and comprehensive performance comparison of the deterministic and variational models for the 4CV and PLAX echo views in Appendix D.




Figure 3: Projection of 4CV view anomalous echo (top) to healthy prototype (bottom). Projections of right (R) and left (L) ventricle (V) and atrium (A) are highlighted in color. The reconstruction of SSD samples approximates a healthy version of the input, e.g., by normalizing the scale of the right and left ventricles (left), adding the ventricular septum (middle), or fixing the location of the valves (right).

Table 3: The area under the curve and average precision of the proposed approaches (TVAE-C, TVAE-R, and TVAE-S) compared with the baseline (VAE) on the four-chamber view and long-axis view for the three different CHD labels. Means and standard deviations are computed across 10 data splits on the test sets. We defined positive labels to correspond to anomalous echos to compute the scores. AP scores of a random classifier are 0.58 (SSD), 0.72 (RVDil), and 0.72 (PH).

		SSD		RVDil		PH	
		AUROC	AP	AUROC	AP	AUROC	AP
4CV	VAE	0.645±0.08	0.667±0.08	0.477±0.05	0.715±0.05	0.498±0.05	0.722±0.03
	TVAE-C	0.913±0.09	0.916±0.11	0.6 ±0.05	0.762±0.04	0.612±0.05	0.786±0.04
	TVAE-R	0.917 ±0.05	0.928 ±0.05	0.594±0.07	0.771±0.04	0.629±0.08	0.797 ±0.06
	TVAE-S	0.868±0.05	0.892±0.05	0.595±0.03	0.774 ±0.02	0.649 ±0.06	0.794±0.05
PLAX	VAE	0.628±0.14	0.457±0.07	0.455±0.05	0.702±0.03	0.432±0.04	0.695±0.03
	TVAE-C	0.87±0.1	0.811±0.15	0.599±0.07	0.794 ±0.04	0.631±0.05	0.818±0.03
	TVAE-R	0.877±0.08	0.826±0.1	0.61 ±0.04	0.794 ±0.02	0.629±0.06	0.817±0.03
	TVAE-S	0.914 ±0.09	0.876 ±0.14	0.592±0.05	0.791±0.03	0.636 ±0.05	0.821 ±0.02

5.3. Anomaly Detection

As described in Section 3.2.2, we detect anomalies by MAP estimation:

$$\tilde{X}_H = \arg \max_{X_H} (\log(P(Y|X_H)) + ELBO(X_H))$$

Due to the reconstruction loss in the ELBO, this optimization problem requires us to back-propagate through the whole model in every step. As a result, inference with the standard MAP formulation is inefficient and proved to be infeasible for our experiments. To circumvent this problem, we assumed the reconstruction part of the ELBO to be constant and solely balanced the posterior with the KL-Divergence of the encoded \mathbf{b} , i.e., how well X_H is mapped to a standard Gaussian, thus computing

$$\tilde{X}_H = \arg \max_{X_H} (P(Y|X_H) - KL[q(\mathbf{b}|X_H)||p(\mathbf{b})])$$

Solving this optimization procedure results in only backpropagating through the encoder instead of the whole model, which leads to a significant speedup while performance was not affected.

To optimize this objective we initialize \tilde{X}_H with the reconstructions computed by the respective model, i.e. $\tilde{X}_H^{(0)} = f(Y)$ for model f and input Y . We then solve the inference problem with the Adam optimizer, incorporating a learning rate of 0.01 and taking 100 optimizer steps per sample. Additionally, we weight the TV norm with a factor of 0.001. For each sample Y , we define the anomaly score $\alpha(Y) := \frac{1}{T} \sum_t \|\tilde{\mathbf{a}}^{(j)}\|_2^2$ as described in Section 3.2.2. Anomaly detection performance is then evaluated in terms of the *Area Under the Receiver Operator Curve* (AUROC) and *Average Precision* (AP) when considering the anomalies as the positive class. In Table 3, we provide a complete overview of the results of the anomaly detection experiments over both views.

We observe that the proposed approaches outperform the VAE in all experiments. This holds especially true when detecting SSD, where our models, TVAE-C, TVAE-R, and TVAE-S, have significantly better performance and can reliably detect such anomalies. Despite outperforming TVAE-C and TVAE-R in terms of reconstruction quality, we also note that TVAE-S does not always perform better in the anomaly detection task. We explain the score discrepancies between SSD and RVDil/PH because SSD deviates considerably from the healthy distribution. RVDil and PH, on the other hand, are more subtle and require expert knowledge and several echocardiogram views to be detected in practice.

Additionally, we argue that we achieve superior performance over VAE since TVAE-R and TVAE-S have considerably smaller latent spaces ($d + 2/d + 3$) than the VAE ($25 \cdot d$), but similar performance regarding reconstruction quality as demonstrated in Section 5.2. This gives the optimizer more flexibility when solving the MAP problem since the frames of \tilde{X}_H can be updated independently to encode them on Gaussian parameters close to $\mathcal{N}(0, \mathbb{I})$, which may result in overfitting during MAP estimation.

Another reconstruction-based inference method approach where we simply define $\alpha_f(X)$ over the MSE, i.e. $\alpha_f(X) = \frac{1}{T} \sum_{j=1}^T \|(x^{(j)} - (f(X))^{(j)})\|_2^2$, is presented in Appendix E.

5.4. Decision Heatmaps

This experiment presents how the estimated anomaly perturbation \tilde{A} can be applied to highlight anomalous regions. Intuitively, anomalous regions in input echos Y differ more substantially from its healthy projection X_H than healthy regions. Consequently, this leads to higher magnitude values in the corresponding locations in the frames of \tilde{A} . In turn, we are able to compute an anomaly heatmap by temporally averaging the estimated

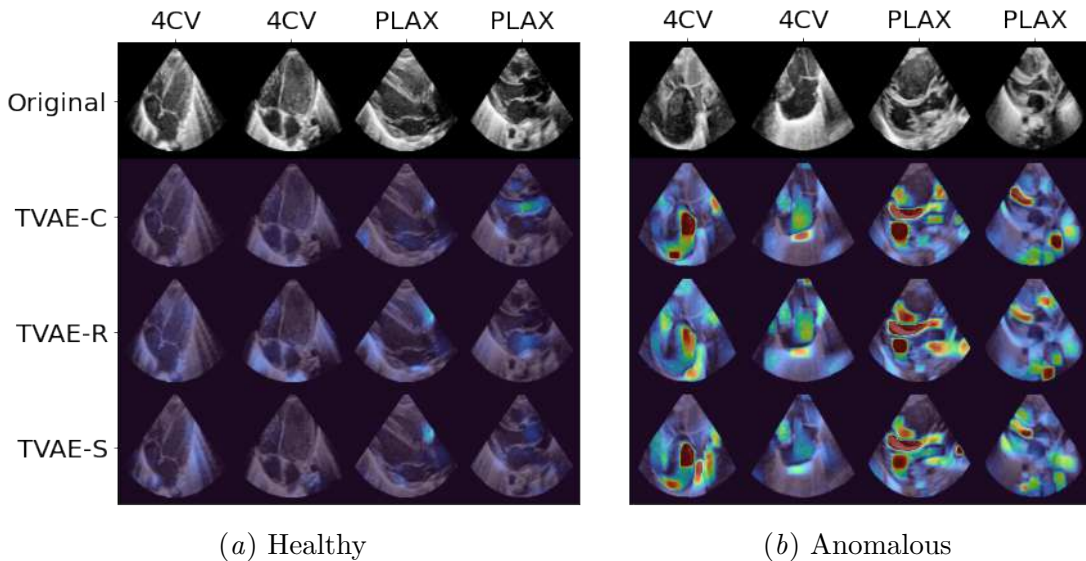


Figure 4: Anomaly response maps of TVAE-R and TVAE-S for healthy samples (a) and echos with CHDs (b). Note how healthy heatmaps are mostly constant, while anomalous maps contain regions with high responses in anomalous regions, corresponding to enlarged ventricles (first/second) or perturbed septums (third/fourth).

anomaly perturbation with $\frac{1}{T} \sum_{j=1}^T \tilde{\mathbf{a}}^{(j)}$. Figure 4 presents examples of such maps for each TVAE variation. There we can verify that not only do we have consistently low magnitude responses for healthy echos, but regions corresponding to, e.g., enlarged chambers, are well highlighted for echos with CHDs. These heatmaps provide TVAE with an additional layer of interpretability and could make our method feasible in a clinical setting, as the reason for decisions made by our method can easily be followed by clinicians. This helps practitioners build trust in model decisions and provides a more intuitive explanation of the outputs of our method. More examples of decision heatmaps are provided in Appendix F.

6. Discussion

In this work, we introduce the TVAE; a new generative model designed explicitly for echocardiogram data. We propose three variants of the model, TVAE-C, and TVAE-R, which make strong assumptions about the data, and the TVAE-S, which can handle more dynamic inputs. Throughout this work, we compared the proposed approach to the VAE in terms of its reconstruction performance and anomaly detection capabilities in a new in-house echo dataset consisting of two different echo views of healthy patients and patients suffering from various CHD. In exhaustive experiments, we demonstrated how TVAE can achieve reconstruction quality comparable to VAE while having a significantly smaller information bottleneck. Additionally, we verified that the proposed model can project out-of-distribution samples, i.e., patients suffering from CHD, into the subspace of healthy echos

when learning normative priors and concluded that TVAE fulfills crucial assumptions for reconstruction-based anomaly detection. Consequently, we evaluated the CHD detection performance of our model, where we found that it leads to a considerable improvement over frame-wise VAE with MAP-based anomaly detection. Furthermore, we demonstrated how TVAE can separate SSD anomalies almost perfectly from healthy echos. Finally, we present the ability of this model to not only detect but also localize anomalies with heatmaps generated from the MAP output, which could help clinicians with the diagnosis of CHDs.

Limitations and Future Work Even though we observe convincing results for SSD, performance for the detection of RVDil and PH is still insufficient for clinical application. The learned normative prior may not be strong enough for these samples, making it hard to detect them as outliers conclusively. This is expected given that these defects are rather subtle and our in-house dataset is relatively small. It would thus be interesting to apply the proposed approach to different and larger cohorts. In the future, we plan to collect more samples for our in-house dataset. With a more extensive dataset, we look forward to exploring methods that would allow combinations of TVAE with one class classification or future frame prediction methods to achieve more robust anomaly detection in echocardiography-based disease detection.

The spiral trajectory of the TVAE-S model assumes continuous movement over the video and might thus still be limiting in situations where sudden movement occurs. In practice, we did not observe this to be a problem as TVAE-S learned good reconstructions for such samples. Still, investigating accelerating trajectories could be an exciting direction. Further, we want to extend the TVAE to multiple modalities such that it is possible to train a model that learns a coherent latent trajectory of multiple echo views of the same heart. In the future, we are interested in introducing TVAE to modalities in other medical fields by designing trajectory functions that leverage modality-specific characteristics similar to what we did for echos.

References

- Jinwon An and Sungzoon Cho. Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, 2(1):1–18, 2015.
- Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Deep autoencoding models for unsupervised anomaly segmentation in brain mr images. In *International MICCAI brainlesion workshop*, pages 161–169. Springer, 2018.
- Christoph Baur, Robert Graf, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Steganomaly: inhibiting cyclegan steganography for unsupervised anomaly detection in brain mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 718–727. Springer, 2020.
- Christoph Baur, Benedikt Wiestler, Mark Muehlau, Claus Zimmer, Nassir Navab, and Shadi Albarqouni. Modeling healthy anatomy with artificial intelligence for unsupervised anomaly detection in brain mri. *Radiology: Artificial Intelligence*, 3(3):e190169, 2021.

- Erik Buskens, PA Stewart, Jakob Hess, DE Grobbee, and JW Wladimiroff. Efficacy of fetal echocardiography and yield by risk category. *Obstetrics & Gynecology*, 87(3):423–428, 1996.
- Olmo Cerri, Thong Q Nguyen, Maurizio Pierini, Maria Spiropulu, and Jean-Roch Vlimant. Variational autoencoders for new physics mining at the large hadron collider. *Journal of High Energy Physics*, 2019(5):1–29, 2019.
- Raghavendra Chalapathy and Sanjay Chawla. Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407*, 2019.
- Jinghui Chen, Saket Sathe, Charu Aggarwal, and Deepak Turaga. Outlier detection with autoencoder ensembles. In *Proceedings of the 2017 SIAM international conference on data mining*, pages 90–98. SIAM, 2017.
- Xiaoran Chen and Ender Konukoglu. Unsupervised detection of lesions in brain mri using constrained adversarial auto-encoders. *arXiv preprint arXiv:1806.04972*, 2018.
- Xiaoran Chen, Suhang You, Kerem Can Tezcan, and Ender Konukoglu. Unsupervised lesion detection via image restoration with a normative prior. *Medical image analysis*, 64:101713, 2020.
- Suyu Dong, Gongning Luo, Guanxiong Sun, Kuanquan Wang, and Henggui Zhang. A left ventricular segmentation method on 3d echocardiography using deep learning and snake. In *2016 Computing in Cardiology Conference (CinC)*, pages 473–476. IEEE, 2016.
- Franck Dufrenois. A one-class kernel fisher criterion for outlier detection. *IEEE transactions on neural networks and learning systems*, 26(5):982–994, 2014.
- Xiaohong Gao, Wei Li, Martin Loomes, and Lianyi Wang. A fused deep learning architecture for viewpoint classification of echocardiography. *Information Fusion*, 36:103–113, 2017.
- Chandan Gautam, Ramesh Balaji, K Sudharsan, Aruna Tiwari, and Kapil Ahuja. Localized multiple kernel learning for anomaly detection: One-class classification. *Knowledge-Based Systems*, 165:241–252, 2019.
- Zahra Ghafoori and Christopher Leckie. Deep multi-sphere support vector data description. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 109–117. SIAM, 2020.
- Alireza Ghasemi, Hamid R Rabiee, Mohammad Taghi Manzuri, and Mohammad Hossein Rohban. A bayesian approach to the data description problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, pages 907–913, 2012.
- Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis. Learning temporal regularity in video sequences. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 733–742, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.

- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Joon-myung Kwon, Kyung-Hee Kim, Ki-Hyun Jeon, and Jinsik Park. Deep learning for predicting in-hospital mortality among heart disease patients based on echocardiography. *Echocardiography*, 36(2):213–218, 2019.
- Fabian Laumer, Gabriel Fringeli, Alina Dubatovka, Laura Manduchi, and Joachim M Buhmann. Deepheartbeat: Latent trajectory learning of cardiac cycles using cardiac ultrasounds. In *Machine Learning for Health*, pages 194–212. PMLR, 2020.
- Sarah Leclerc, Erik Smistad, Joao Pedrosa, Andreas Østvik, Frederic Cervenansky, Florian Espinosa, Torvald Espeland, Erik Andreas Rye Berg, Pierre-Marc Jodoin, Thomas Grenier, et al. Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE transactions on medical imaging*, 38(9):2198–2210, 2019.
- Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection—a new baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6536–6545, 2018.
- Maxime Louis, Raphaël Couronné, Igor Koval, Benjamin Charlier, and Stanley Durrleman. Riemannian Geometry Learning for Disease Progression Modelling. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11492 LNCS:542–553, 2019. ISSN 16113349. doi: 10.1007/978-3-030-20351-1_42.
- Ali Madani, Jia Rui Ong, Anshul Tibrewal, and Mohammad RK Mofrad. Deep echocardiography: data-efficient supervised and semi-supervised deep learning towards automated diagnosis of cardiac disease. *NPJ digital medicine*, 1(1):1–11, 2018.
- Shakiba Moradi, Mostafa Ghelich Oghli, Azin Alizadehasl, Isaac Shiri, Niki Oveisi, Mehrdad Oveisi, Majid Maleki, and Jan Dhooge. Mfp-unet: A novel deep learning based approach for left ventricle segmentation in echocardiography. *Physica Medica*, 67:58–69, 2019.
- Mary M Moya and Don R Hush. Network constraints and multi-objective optimization for one-class classification. *Neural networks*, 9(3):463–474, 1996.
- David Ouyang, Bryan He, Amirata Ghorbani, Neal Yuan, Joseph Ebinger, Curtis P Langlotz, Paul A Heidenreich, Robert A Harrington, David H Liang, Euan A Ashley, et al. Video-based ai for beat-to-beat assessment of cardiac function. *Nature*, 580(7802):252–256, 2020.
- Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. Deep learning for anomaly detection: A review. *ACM Computing Surveys (CSUR)*, 54(2):1–38, 2021.
- Daehyung Park, Yuuna Hoshi, and Charles C Kemp. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. *IEEE Robotics and Automation Letters*, 3(3):1544–1551, 2018.

- Nick Pawlowski, Matthew CH Lee, Martin Rajchl, Steven McDonagh, Enzo Ferrante, Konstantinos Kamnitsas, Sam Cooke, Susan Stevenson, Aneesh Khetani, Tom Newman, et al. Unsupervised lesion detection in brain ct using bayesian convolutional autoencoders. 2018.
- Walter Hugo Lopez Pinaya, Petru-Daniel Tudosiu, Robert Gray, Geraint Rees, Parashkev Nachev, Sébastien Ourselin, and M Jorge Cardoso. Unsupervised brain anomaly detection and segmentation with transformers. *arXiv preprint arXiv:2102.11650*, 2021.
- Emanuele Principi, Fabio Vesperini, Stefano Squartini, and Francesco Piazza. Acoustic novelty detection with adversarial autoencoders. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 3324–3330. IEEE, 2017.
- Gunnar Ratsch, Sebastian Mika, Bernhard Scholkopf, and K-R Muller. Constructing boosting algorithms from svms: An application to one-class classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(9):1184–1199, 2002.
- Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International conference on machine learning*, pages 4393–4402. PMLR, 2018.
- Lukas Ruff, Robert A Vandermeulen, Billy Joe Franks, Klaus-Robert Müller, and Marius Kloft. Rethinking assumptions in deep anomaly detection. *arXiv preprint arXiv:2006.00339*, 2020.
- Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 2021.
- Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3379–3388, 2018.
- Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- Y Singh and L McGeoch. Fetal anomaly screening for detection of congenital heart defects. *J. Neonatal Biol.*, 5(2):100–115, 2016.
- David MJ Tax and Robert PW Duin. Support vector data description. *Machine learning*, 54(1):45–66, 2004.
- Denise Van Der Linde, Elisabeth EM Konings, Maarten A Slager, Maarten Witsenburg, Willem A Helbing, Johanna JM Takkenberg, and Jolien W Roos-Hesselink. Birth prevalence of congenital heart disease worldwide: a systematic review and meta-analysis. *Journal of the American College of Cardiology*, 58(21):2241–2247, 2011.

- CL Van Velzen, SA Clur, MEB Rijlaarsdam, E Pajkrt, CJ Bax, J Hruda, CJM de Groot, NA Blom, and MC Haak. Prenatal diagnosis of congenital heart defects: accuracy and discrepancies in a multicenter cohort. *Ultrasound in Obstetrics & Gynecology*, 47(5): 616–622, 2016.
- Hooman Vaseli, Zhibin Liao, Amir H Abdi, Hany Girgis, Delaram Behnami, Christina Luong, Fatemeh Taheri Dezaki, Neeraj Dhungel, Robert Rohling, Ken Gin, et al. Designing lightweight deep learning models for echocardiography view classification. In *Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 10951, pages 93–99. SPIE, 2019.
- Dan Xu, Elisa Ricci, Yan Yan, Jingkuan Song, and Nicu Sebe. Learning deep representations of appearance and motion for anomalous event detection. *arXiv preprint arXiv:1510.01553*, 2015.
- Haowen Xu, Wenxiao Chen, Nengwen Zhao, Zeyan Li, Jiahao Bu, Zhihan Li, Ying Liu, Youjian Zhao, Dan Pei, Yang Feng, et al. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. In *Proceedings of the 2018 world wide web conference*, pages 187–196, 2018.
- Shiyang Yan, Jeremy S Smith, Wenjin Lu, and Bailing Zhang. Abnormal event detection from videos using a two-stream recurrent variational autoencoder. *IEEE Transactions on Cognitive and Developmental Systems*, 12(1):30–42, 2018.
- Suhang You, Kerem C Tezcan, Xiaoran Chen, and Ender Konukoglu. Unsupervised lesion detection via image restoration with a normative prior. In *International Conference on Medical Imaging with Deep Learning*, pages 540–556. PMLR, 2019.
- Guang Yu, Siqi Wang, Zhiping Cai, En Zhu, Chuanfu Xu, Jianping Yin, and Marius Kloft. Cloze test helps: Effective video anomaly detection via learning to complete video events. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 583–591, 2020.
- Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pages 2528–2535. IEEE, 2010.

Appendix A. Variational Trajectory Model ELBO derivation

Recall that we define $\mathbf{b} \sim q_{\theta}(\mathbf{b}|X) := \mathcal{N}(\boldsymbol{\mu}_{\mathbf{b}}, \text{diag}(\boldsymbol{\sigma}_{\mathbf{b}}))$ with prior $p(\mathbf{b}) := \mathcal{N}(0, \mathbb{I})$, while leaving the other trajectory parameters deterministic. Note that this effectively means that we define uniform priors $p(f)$, $p(\omega)$ and $p(v)$ over their support, while having posteriors

$$q_{\theta}(f|X) := \delta_{\phi_f(X)}(f), \quad q_{\theta}(\omega|X) := \delta_{\phi_{\omega}(X)}(\omega), \quad q_{\theta}(v|X) := \delta_{\phi_v(X)}(v)$$

where δ_y is the Dirac Delta spiking at y and $\phi_f(X)$, $\phi_{\omega}(X)$ and $\phi_v(X)$ are the trajectory parameter outputs of the encoder ϕ with weights θ for f , ω and v respectively.

Given input sample x and latent z , recall that VAEs aim to maximize the *Evidence Lower Bound* (ELBO):

$$E_{q_\theta(z|x)}[\log(p_\eta(x|z))] - KL[q_\theta(z|x)||p(z)]$$

Here, x corresponds to the input echocardiogram $X := (\mathbf{x}^{(j)}, t^{(j)})_{j=1}^T$ whereas $z := (\mathbf{b}, f, \omega, v)$. Note that \mathbf{b}, f, ω and v are conditionally independent, i.e.

$$q_\theta(\mathbf{b}, f, \omega, v|X) = q_\theta(\mathbf{b}|X)q_\theta(f|x)q_\theta(\omega|X)q_\theta(v|X)$$

The KL divergence is additive for joint distributions of independent random variables, i.e. for $P = (P_1, P_2)$ and $Q = (Q_1, Q_2)$, where P_1, P_2, Q_1 and Q_2 are independent, it holds that

$$KL(P||Q) = KL(P_1||Q_1) + KL(P_2||Q_2)$$

We can thus rewrite the ELBO as

$$\begin{aligned} & E_{q_\theta(\mathbf{b}, f, \omega, v|X)}[\log(p_\eta(X|\mathbf{b}, f, \omega, v))] \\ & - KL[q_\theta(\mathbf{b}|X)||p(\mathbf{b})] - KL[q_\theta(f|X)||p(f)] \\ & - KL[q_\theta(\omega|X)||p(\omega)] - KL[q_\theta(v|X)||p(v)] \end{aligned}$$

Since we assumed a uniform prior for f, ω and v , their KL-Divergence terms become constant under the Dirac Delta distribution. We can thus ignore the respective terms in the ELBO during optimization as they do not change the result of the *argmax*.

Additionally, since

$$\int \delta_y(x)f(x)dx = f(y)$$

we can rewrite the ELBOs reconstruction term as

$$\begin{aligned} & E_{q_\theta(\mathbf{b}, f, \omega, v|X)}[\log(p_\eta(X|\mathbf{b}, f, \omega, v))] \\ & = \int \delta_{\phi_f(X)}(f)\delta_{\phi_\omega(X)}(\omega)\delta_{\phi_v(X)}(v)q_\theta(\mathbf{b}|X)\log(p_\eta(X|\mathbf{b}, f, \omega, v))d\mathbf{b}dfd\omega dv \\ & = \int q_\theta(\mathbf{b}|X)\log(p_\eta(X|\mathbf{b}, \phi_f(X), \phi_\omega(X), \phi_v(X)))d\mathbf{b} \\ & = E_{q_\theta(\mathbf{b}|X)}[\log(p_\eta(X|\mathbf{b}, \phi_f(X), \phi_\omega(X), \phi_v(X)))] \end{aligned}$$

Finally, this leads to the following reformulation of the ELBO objective:

$$E_{q_\theta(\mathbf{b}|X)}[\log(p_\eta(X|\mathbf{b}, \phi_f(X), \phi_\omega(X), \phi_v(X)))] - KL[q_\theta(\mathbf{b}|X)||p(\mathbf{b})]$$

Appendix B. Cohort Examples

To give some intuition on what CHDs look like in different views, we provide examples in Figure 5 and compare them to healthy samples.

Table 4: Hyperparameters chosen across our experiments.

Hyperparameter	AE/VAE	TAE/TVAE
Latent Dimension	64	66/67 (b:64; f:1; ω :1; v:1)
Batch Size	128	64
Steps	5000	106500
Number of Frames	1	25
Optimizer	Adam	Adam
Learning Rate	10^{-4}	10^{-4}
Reconstruction Loss	MSE	MSE
VAE β	1	1

Appendix C. Architecture

We provide schematics for the building blocks of our architectures in Figure 6 and describe our experiments’ encoder/decoder architecture in Figure 7.

Table 4 contains the hyperparameters used in our experiments. Except for the number of steps, we kept hyperparameters mostly the same for all models. This is because, in contrast to the frame-wise models, TAE and TVAE models required many more steps to converge. We suspect this because the input’s dimensionality is 25 times larger, and the model thus requires more parameter updates to converge to a suitable optimum that results in good reconstructions. The batch size was chosen according to GPU memory capacity.

Appendix D. Further Reconstruction Experiments

In addition to the reconstruction quality experiments provided in Section 5.2, we compared the performance of the variational models to deterministic ones (i.e., standard autoencoder and non-variational trajectory models). As seen in Table 5, the deterministic trajectory models result in similar performance to the variational models and are even slightly better with respect to the structural similarity score. Even though trained on the same architecture and for the same number of steps as the VAE, the autoencoder did not seem to produce very good reconstruction scores in this experiment. We suspect this may be an artifact of overfitting due to the small training set.

We provide more reconstructions of TVAE-S in Figure 8.

Appendix E. Reconstruction error-based anomaly detection and one class classification.

A common alternative to MAP-based anomaly detection is the detection of anomalies based on the model’s reconstruction error. This means, for model f , sample $x \in \mathcal{X}$ and data space \mathcal{X} , we would simply define $\alpha_f(x) = \|x - f(x)\|_2^2$. In order to quantify the performance of the non-variational dynamic trajectory model (TAE) and to have a comparison to a standard autoencoder trained on single frame reconstruction, we performed another ablation on AE, VAE, and the variants of TAE and TVAE. As an additional baseline, we also implemented

the deep one class classification method introduced in [Ruff et al. \(2018\)](#). We present the results of this ablation in [Table 6](#).

Appendix F. More Decision Heatmaps

In addition to the heatmaps presented in [Section 5.4](#), we provide a more extensive collection of TVAE-S decision heatmaps in [Figure 9](#) and [Figure 10](#) and compare them with heatmaps generated by MAP estimation with a standard VAE ([Chen et al., 2020](#)) in [Figure 11](#) and [Figure 12](#).

Appendix G. Generated Videos

The introduced models TVAE-R and TVAE-S are generative models. As such, in addition to producing good reconstructions of existing samples, they allow us to sample from the learned distribution. To qualitatively validate generative performance, we provide random generations of the TVAE-S model in [Figure 13](#) for both 4CV and PLAX views.

Appendix H. Robustness to Domain Shifts

The proposed approach learns a normative prior on the distribution of the training dataset. Further, it assumes that every outlier of the learned distribution corresponds to an anomaly. Therefore, a trained model is not robust to domain shifts but needs some finetuning on a new dataset to incorporate the new notion of normality into its normative prior. We demonstrate this by combining 100 samples of EchoDynamic, which contains 4CV echocardiograms, with the 4CV views of our in-house dataset for training. We assume that EchoDynamic samples correspond to the healthy class. EchoDynamic exhibits a considerable distribution shift, as these echocardiograms were recorded by a different machine and collected from an adult population. [Table 7](#) contains the results of this experiment. From these results, we can conclude that, for PH and RVDil, learning a new normative prior on this extended dataset does not change the outcome. On the other hand, scores for SSD appear less robust, though they still outperform PH and RVDil in AUROC.

Appendix I. Anomaly oversensitivity

To quantify oversensitivity towards anomalies, we can redefine the task of anomaly detection to healthy instance detection, i.e., assigning the positive label to the healthy instances. Due to its symmetry, the AUROC scores will not change under this new setting; we thus only report AP. [Table 8](#) contains the results of these experiments. In the case of SSD, the AP is stable whereas in the case of PH and RVDil, the APs drop. With now 30 positive versus 73 negative labels, this is expected to happen. Still, results seem to agree with anomaly detection results, as AP scores are still considerably better than random (AP=0.32), suggesting that we are not overly sensitive to anomalies.

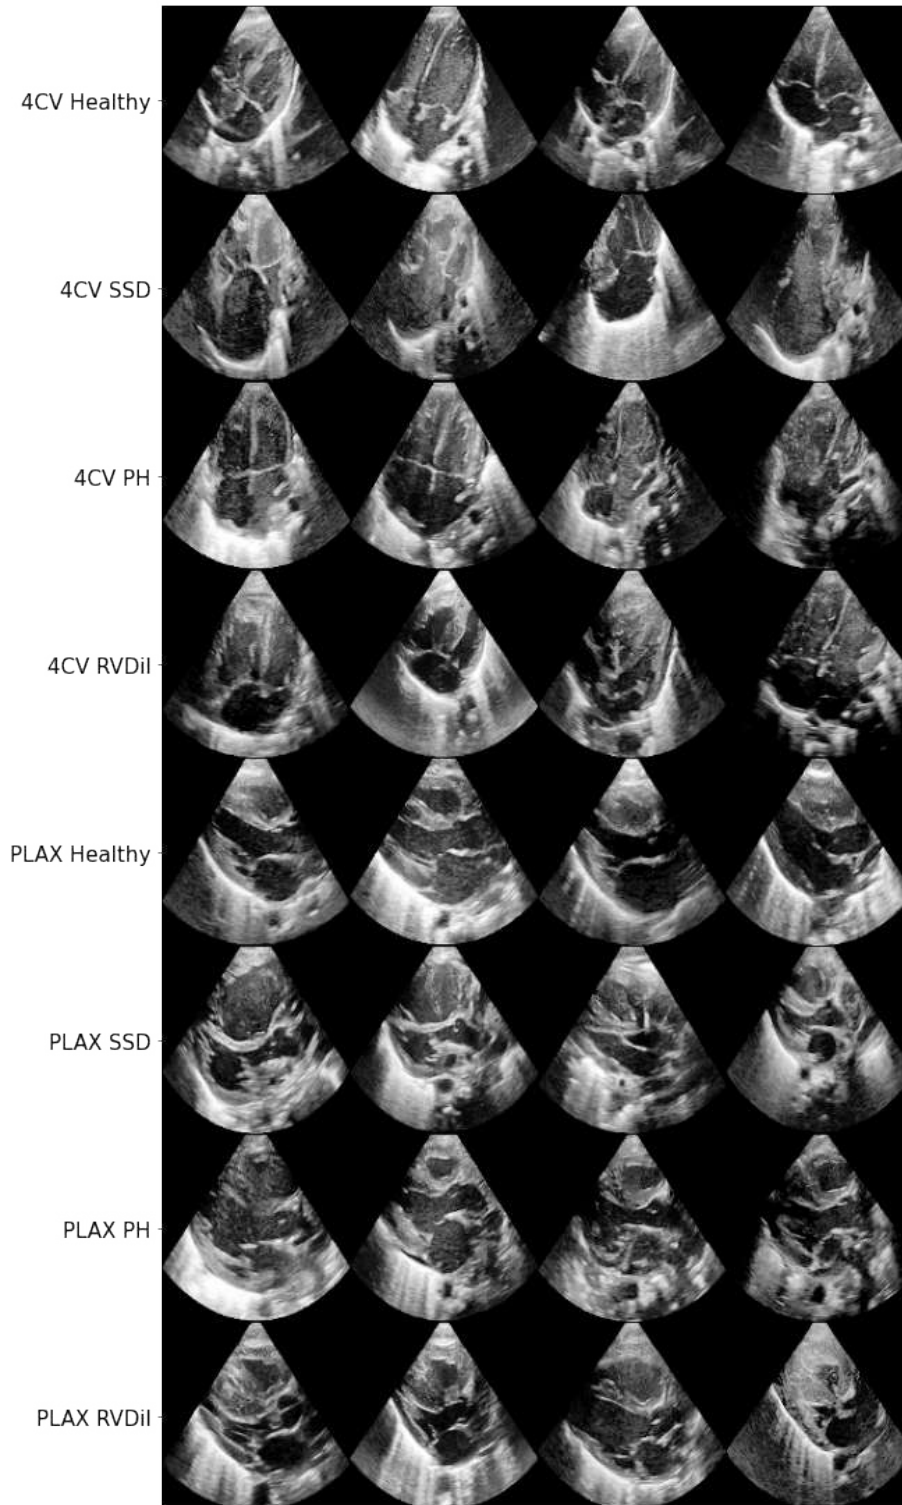


Figure 5: Examples of each label of the cohort in 4CV and PLAX views.

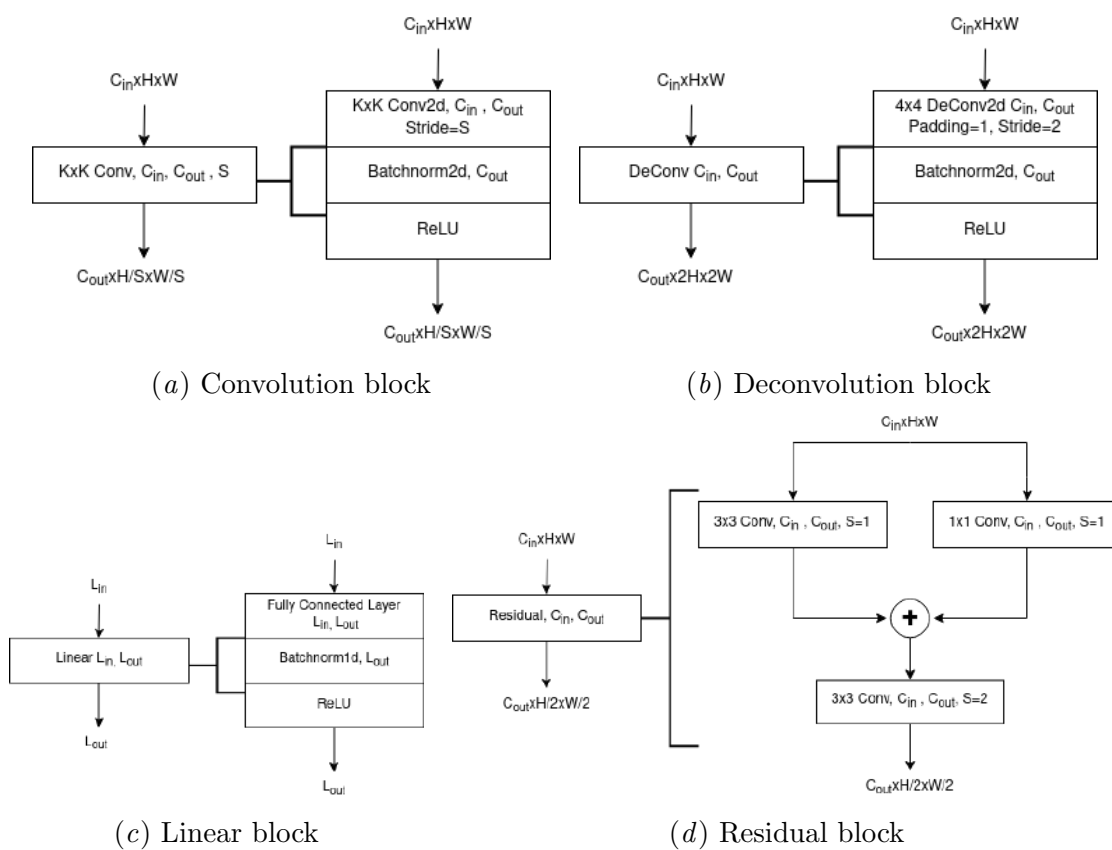


Figure 6: Definitions of the encoder/decoder building blocks.

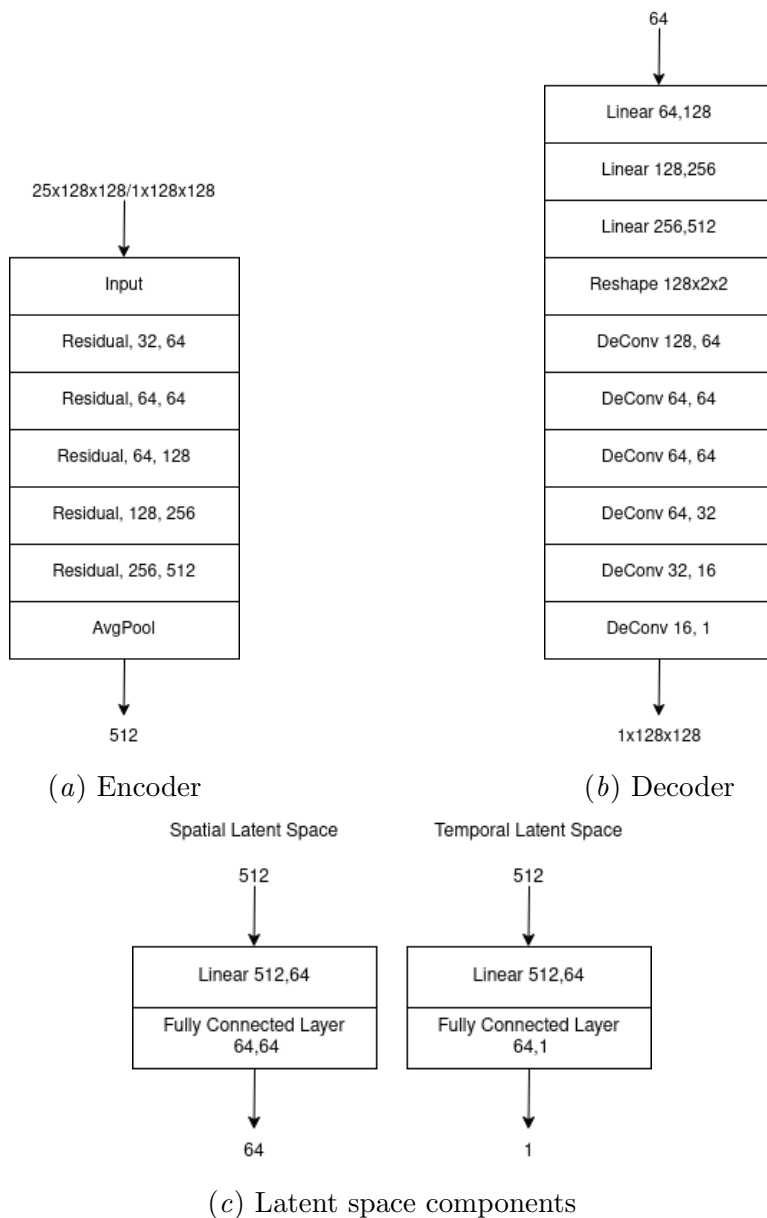


Figure 7: Architectures of encoder (a), decoder (b) and latent space components (c). Spatial latent space components are used to learn z , μ and σ for AE/VAE or b , μ_b and σ_b for TAE/TVAE. Temporal latent space components learn f, ω or v for TAE/TVAE.




Figure 8: More TVAE-S reconstructions.

Table 5: Reconstruction scores across all introduced models and compared to AE/VAE.

	SSD			RVDii			PH			
	MSE	PSNR	SSIM	MSE	PSNR	SSIM	MSE	PSNR	SSIM	
4CV	AE	0.427±0.07	4.692±0.65	0.246±0.01	0.526±0.1	3.722±0.78	0.246±0.0	0.57±0.08	3.248±0.54	0.244±0.0
	VAE	0.013 ±0.0	19.02 ±0.11	0.546±0.01	0.012 ±0.0	19.146 ±0.05	0.555±0.0	0.012 ±0.0	19.093 ±0.07	0.553±0.0
	TAE-R	0.013 ±0.0	18.964±0.15	0.555±0.01	0.013±0.0	19.065±0.05	0.563 ±0.0	0.013±0.0	19.017±0.07	0.562 ±0.0
	TAE-S	0.013 ±0.0	18.949±0.1	0.557 ±0.01	0.013±0.0	18.926±0.06	0.559±0.0	0.013±0.0	18.897±0.07	0.558±0.0
	TVAE-C	0.014±0.0	18.562±0.15	0.544±0.01	0.014±0.0	18.707±0.07	0.549±0.0	0.014±0.0	18.661±0.07	0.55±0.0
	TVAE-R	0.014±0.0	18.568±0.15	0.545±0.01	0.013±0.0	18.813±0.09	0.554±0.0	0.014±0.0	18.73±0.08	0.552±0.0
	TVAE-S	0.013 ±0.0	18.784±0.09	0.551±0.01	0.013±0.0	18.797±0.04	0.554±0.0	0.013±0.0	18.747±0.08	0.554±0.0
PLAX	AE	0.518±0.14	3.724±1.01	0.226±0.0	0.649±0.07	2.545±0.5	0.23±0.0	0.748±0.11	1.935±0.59	0.228±0.0
	VAE	0.017 ±0.0	17.874 ±0.15	0.524±0.0	0.015 ±0.0	18.394 ±0.05	0.541 ±0.0	0.015 ±0.0	18.442 ±0.03	0.542 ±0.0
	TAE-C	0.018±0.0	17.576±0.12	0.518±0.0	0.016±0.0	17.986±0.1	0.532±0.0	0.016±0.0	18.117±0.07	0.536±0.0
	TAE-R	0.017 ±0.0	17.836±0.15	0.528 ±0.01	0.015 ±0.0	18.196±0.13	0.54±0.0	0.015 ±0.0	18.192±0.12	0.54±0.0
	TAE-S	0.017 ±0.0	17.785±0.12	0.525±0.0	0.015 ±0.0	18.184±0.04	0.539±0.0	0.015 ±0.0	18.26±0.04	0.542 ±0.0
	TVAE-C	0.019±0.0	17.264±0.15	0.509±0.01	0.017±0.0	17.795±0.05	0.526±0.0	0.017±0.0	17.778±0.07	0.526±0.0
	TVAE-R	0.018±0.0	17.555±0.18	0.519±0.01	0.016±0.0	17.958±0.08	0.533±0.0	0.016±0.0	17.971±0.12	0.533±0.0
TVAE-S	0.018±0.0	17.633±0.09	0.517±0.0	0.016±0.0	18.115±0.1	0.536±0.0	0.015 ±0.0	18.152±0.08	0.537±0.0	

Table 6: Area under the curve and average precision for experiments performed with one class classification (OCC) or reconstruction-based with anomaly score $\alpha_f(x) = \frac{1}{T} \sum_{t=1}^T \|x^{(t)} - f^{(t)}(x)\|_2^2$.

		SSD		RVDil		PH	
		AUROC	AP	AUROC	AP	AUROC	AP
4CV	OCC	0.51±0.05	0.509±0.04	0.498±0.01	0.719±0.01	0.505±0.01	0.72±0.01
	AE	0.566±0.1	0.602±0.09	0.634±0.05	0.816±0.03	0.612±0.03	0.803±0.02
	VAE	0.699±0.09	0.732±0.08	0.619±0.05	0.803±0.03	0.635±0.04	0.808±0.04
	TAE-C	0.572±0.09	0.651±0.05	0.581±0.04	0.775±0.03	0.609±0.02	0.795±0.01
	TAE-R	0.612±0.06	0.695±0.06	0.594±0.04	0.781±0.03	0.617±0.02	0.8±0.02
	TAE-S	0.558±0.1	0.646±0.08	0.612±0.04	0.794±0.03	0.614±0.03	0.802±0.02
	TVAE-C	0.672±0.06	0.736±0.05	0.6±0.04	0.779±0.03	0.622±0.03	0.803±0.01
	TVAE-R	0.673±0.07	0.745±0.07	0.611±0.03	0.787±0.02	0.621±0.04	0.803±0.03
	TVAE-S	0.616±0.1	0.679±0.05	0.61±0.05	0.786±0.04	0.631±0.03	0.8±0.03
PLAX	OCC	0.529±0.05	0.394±0.07	0.503±0.01	0.727±0.01	0.506±0.01	0.732±0.01
	AE	0.917±0.08	0.889±0.1	0.681±0.03	0.848±0.02	0.637±0.05	0.827±0.03
	VAE	0.918±0.07	0.915±0.05	0.683±0.03	0.851±0.02	0.631±0.05	0.819±0.03
	TAE-C	0.917±0.09	0.885±0.13	0.65±0.03	0.825±0.02	0.646±0.06	0.827±0.03
	TAE-R	0.913±0.08	0.878±0.1	0.68±0.05	0.84±0.03	0.639±0.05	0.823±0.02
	TAE-S	0.927±0.07	0.905±0.09	0.666±0.03	0.836±0.01	0.635±0.04	0.827±0.02
	TVAE-C	0.927±0.07	0.907±0.08	0.65±0.04	0.828±0.02	0.617±0.05	0.814±0.03
	TVAE-R	0.917±0.07	0.883±0.11	0.664±0.05	0.831±0.03	0.622±0.05	0.818±0.02
	TVAE-S	0.935±0.08	0.916±0.11	0.658±0.05	0.828±0.02	0.625±0.04	0.823±0.02

Table 7: Results of the domain shift experiments. We retrained TVAE-S on a modified dataset that combined 100 EchoDynamic samples with our dataset. We aggregated scores across the test sets of 10 different data folds. To compute AUROC and AP, we assigned the positive label to the anomalous samples. AP of a random classifier is 0.085 (SSD), 0.374 (RVDil), and 0.374 (PH) respectively.

SSD		RVDil		PH	
AUROC	AP	AUROC	AP	AUROC	AP
0.688±0.09	0.271±0.06	0.594±0.04	0.638±0.03	0.61±0.02	0.664±0.02



Figure 9: More TVAE-S decision heatmaps for healthy echos.




Figure 10: More TVAE-S decision heatmaps for anomalous echos.




Figure 11: VAE decision heatmaps for healthy echos.



Figure 12: VAE decision heatmaps for anomalous echos.



Figure 13: Random TVAE-S generations of samples in 4CV and PLAX views.

Table 8: Average precision of the proposed approaches (TVAE-C, TVAER, and TVAES) compared to the baseline (VAE) on the four-chamber view and long-axis view for the three different CHD labels. Means and standard deviations are computed across 10 data splits on the test sets. We defined positive labels to correspond to healthy instances to compute the scores. AP scores of a random classifier are 0.535 (SSD), 0.317 (RVDil), and 0.317 (PH).

		SSD	RVDil	PH
4CV	VAE	0.66±0.1	0.273±0.04	0.315±0.06
	TVAE-C	0.91±0.1	0.419±0.05	0.392±0.03
	TVAE-R	0.919 ±0.05	0.423 ±0.05	0.421±0.05
	TVAE-S	0.863±0.07	0.39±0.05	0.458 ±0.07
PLAX	VAE	0.819±0.09	0.269±0.03	0.275±0.03
	TVAE-C	0.921±0.07	0.388±0.07	0.381±0.05
	TVAE-R	0.927±0.05	0.393±0.04	0.387±0.07
	TVAE-S	0.937 ±0.08	0.413 ±0.08	0.392 ±0.06