

Echo-SyncNet: Self-Supervised Cardiac View Synchronization in Echocardiography

Fatemeh Taheri Dezaki¹, Christina Luong, Tom Ginsberg, Robert Rohling², *Fellow, IEEE*,
Ken Gin³, Purang Abolmaesumi⁴, and Teresa Tsang⁵

Abstract—In echocardiography (echo), an electrocardiogram (ECG) is conventionally used to temporally align different cardiac views for assessing critical measurements. However, in emergencies or point-of-care situations, acquiring an ECG is often not an option, hence motivating the need for alternative temporal synchronization methods. Here, we propose Echo-SyncNet, a self-supervised learning framework to synchronize various cross-sectional 2D echo series without any human supervision or external inputs. The proposed framework takes advantage of two types of supervisory signals derived from the input data: spatiotemporal patterns found between the frames of a single cine (intra-view self-supervision) and interdependencies between multiple cines (inter-view self-supervision). The combined supervisory signals are used to learn a feature-rich and low dimensional embedding space where multiple echo cines can be temporally synchronized. Two intra-view self-supervisions are used, the first is based on the information encoded by the temporal ordering of

a cine (temporal intra-view) and the second on the spatial similarities between nearby frames (spatial intra-view). The inter-view self-supervision is used to promote the learning of similar embeddings for frames captured from the same cardiac phase in different echo views. We evaluate the framework with multiple experiments: 1) Using data from 998 patients, Echo-SyncNet shows promising results for synchronizing Apical 2 chamber and Apical 4 chamber cardiac views, which are acquired spatially perpendicular to each other; 2) Using data from 3070 patients, our experiments reveal that the learned representations of Echo-SyncNet outperform a supervised deep learning method that is optimized for automatic detection of fine-grained cardiac cycle phase; 3) We go one step further and show the usefulness of the learned representations in a one-shot learning scenario of cardiac key-frame detection. Without any fine-tuning, key frames in 1188 validation patient studies are identified by synchronizing them with only one labeled reference cine. We do not make any prior assumption about what specific cardiac views are used for training, and hence we show that Echo-SyncNet can accurately generalize to views not present in its training set. Project repository: github.com/fatemehtd/Echo-SyncNet.

Index Terms—Echocardiography, fine-grained phase classification, self-supervised learning, synchronization.

I. INTRODUCTION

ECHOCARDIOGRAPHY (echo) plays an important role in cardiac imaging and provides a non-invasive, low-cost, and widely available diagnostic tool for the comprehensive evaluation of cardiac structure and function. A cine is an ultrasound video that captures a 2D cross-section of the heart from a specific orientation. Several common cine views, characterized by the plane in which they are captured, are frequently taken by cardiac specialists and used for qualitative and/or quantitative analysis of cardiac function.

To evaluate important cardiac measurements such as left ventricular ejection fraction (EF), detection of wall motion abnormality (WMA) or severity of aortic stenosis (AS) clinicians often use several cines from various views. When analysing multiple views, temporal synchronization or simply cine synchronization, becomes critical for making sure the cardiac phase is the same at every instant in all views. Cine synchronization is a challenging task considering the large quality variations in typical echo and the lack of obvious structural correspondences between views. Normally, synchronization of multiple cardiac views is based on concurrently

Manuscript received January 20, 2021; revised April 2, 2021; accepted April 6, 2021. Date of publication April 9, 2021; date of current version July 30, 2021. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) and in part by the Canadian Institutes of Health Research (CIHR). (Fatemeh Taheri Dezaki and Christina Luong contributed equally to this work as first authors.) (Purang Abolmaesumi and Teresa Tsang contributed equally to this work as senior authors.) (Corresponding author: Fatemeh Taheri Dezaki.)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by The University of British Columbia Clinical Research Ethics Board, approval number H19-00132.

Fatemeh Taheri Dezaki and Purang Abolmaesumi are with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: fatemeht@ece.ubc.ca; purang@ece.ubc.ca).

Christina Luong and Ken Gin are with the Vancouver General Hospital Echocardiography Laboratory, Division of Cardiology, Department of Medicine, The University of British Columbia, Vancouver, BC V5Z 1M9, Canada (e-mail: christina.luong@ubc.ca; kenneth.gin@vch.ca).

Tom Ginsberg is with the Department of Engineering Physics, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: tom.ginsberg@alumni.ubc.ca).

Robert Rohling is with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada, and also with the Department of Mechanical Engineering, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: rohling@ece.ubc.ca).

Teresa Tsang is with the Vancouver General Hospital and University of British Columbia Echocardiography Laboratories, Division of Cardiology, Department of Medicine, The University of British Columbia, Vancouver, BC V5Z 1M9, Canada (e-mail: t.tsang@ubc.ca).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TMI.2021.3071951>, provided by the authors.

Digital Object Identifier 10.1109/TMI.2021.3071951

recorded electrocardiogram (ECG) signals. ECGs can be easily synchronized allowing clinicians to view multiple cines play together with near-perfect temporal correspondance [1]. Cine synchronization allows for better visualization of the cardiac structure, more accurate measurements and hence, a simpler and more reliable cardiac assessment.

In machine learning literature, several research groups have taken advantage of synchronized echo views in developing their models [2]–[4]. In echo based EF prediction, Behnami *et al.* [2] showed that a prediction model performs better when two synchronized views, apical 4 chamber (AP4) and apical 2 chamber (AP2), were fed to the model.

In [3], Li *et al.* applied their proposed segmentation method on synchronized sequences from three views, Apical 3 chamber (AP3), AP4, and AP2. To produce synchronized echo cines, it is required to have either expert annotations of the cardiac phase or the presence of an ECG. Unfortunately, expert annotations or ECG-based cine synchronization is often not feasible in emergencies or point-of-care [5]. The use of ECG sensors can be also prohibited when there is sterilization barriers in clinics, such as those we are facing with the ongoing COVID-19 pandemic.

With the advent of deep learning in computer vision, the medical imaging community has rapidly adopted and developed algorithms to assist physicians in clinical routine. Nevertheless, the overall performance of fully-supervised learning highly depends on the size of the annotated dataset, resulting in sub-optimal performance for smaller datasets. Generating expert annotations of patient data at scale is non-trivial, expensive, and time-consuming. Consequently, unsupervised and self-supervised representation learning is gaining attention to overcome the problem of learning under a low data regimen and to reduce the required cost for expert annotation. In self-supervised learning, supervision comes from the data itself. The only and most important step is to set the learning objectives properly.

Numerous attempts have been made on self-supervised representation learning with non-medical images and videos. Self-supervised approaches can be generally categorized into encoder and encoder-decoder architectures. In the first category, a convolutional neural network encoder is trained based on developing pretext tasks that learn meaningful representations [6]–[8]. In the encoder-decoder architectures [9]–[11], a model is trained based on reconstruction tasks. Self-supervised learning has been applied to many computer vision problems in image and video representation. To stay aligned with the work in this paper, we only review prior methods related to video learning. In [9], [10], tracking moving objects in videos is used as a tool for learning visual representations. Several other self-supervised tasks rely on learning representations based on the chronological order of frames [6]–[8]. Moreover, video colorization by copying colors from a reference frame to target frames is proposed [11] as a means of self-supervised video representation learning.

Most of the prominent self-supervised methods have been derived in the context of natural images and videos, without considering the unique properties that medical imaging has

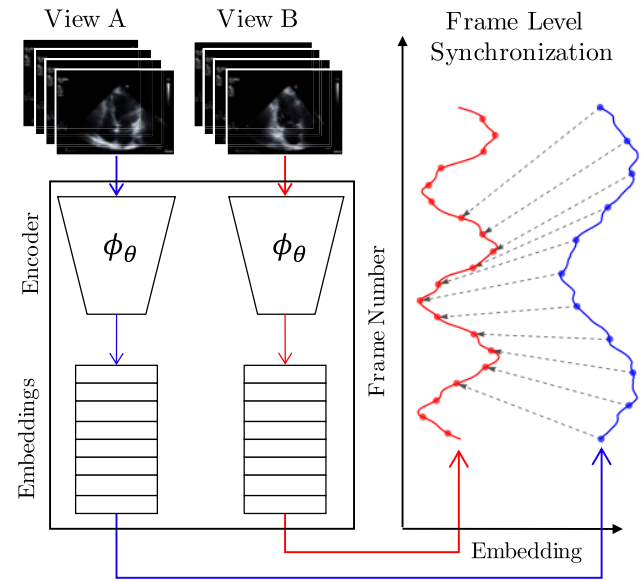


Fig. 1. Overview of the proposed Echo-SyncNet framework. Two cardiac echo views are mapped to embeddings via a shared encoder ϕ_θ . The embeddings are designed to enable a robust temporal synchronization.

to offer. In echo, with the lack of ECG or other external annotations related to synchronization, we propose a self-supervised framework, Echo-SyncNet (Figure 1), to temporally synchronize cardiac echo cines acquired from different views. We hypothesize that consistent temporal and structural information across cardiac views can be leveraged to establish a robust learning objective. Echo-SyncNet exploits two types of supervisory signals derived from the input data: spatiotemporal patterns found between the frames of a single cine (intra-view self-supervision), and interdependencies between multiple cines (inter-view self-supervision). The combined supervisory signals are used to learn a feature-rich and low-dimensional embedding space, where multiple echo cines can be temporally synchronized. Using data from 998 patients, our proposed framework shows promising results for synchronizing AP2 and AP4 cardiac views, which are acquired spatially perpendicular to each other. Furthermore, using data from 3070 patients, we demonstrate that the learned representations of our proposed framework outperform a supervised deep learning method that is optimized for automatic detection of fine-grained cardiac phase-detection. We do not make any prior assumption about specific cardiac views used for training, and hence we show that Echo-SyncNet can accurately generalize to views not present in its training set. Moreover, in an experiment in a one-shot learning scenario using only one cine with labeled key frames, we show the usefulness of the learned embeddings in identifying key frames in unseen validation cines, significantly reducing the burden of manual annotation.

In summary, our contributions are:

- a self-supervised framework for fine-grained temporal alignment of cardiac echo cines captured from different echo views;

- evaluation of the proposed method on a large echocardiography data set, and demonstration of its effectiveness to synchronize various cardiac views;
- a thorough analysis of the use of our framework for the detection of fine-grained cardiac phases, specifically end-diastolic (ED) and end-systolic (ES) cardiac events, without the need for significant labeled data; and,
- showing the extension of our framework to the synchronization of multiple types of cardiac views—even views not present in the training set—with empirical evidence.

Additionally our source code is made available¹ included with an example usage of Echo-SyncNet on the Stanford EchoNet-Dynamic [12] public dataset and the supplementary materials.

II. METHODS

In this section, we describe the proposed Echo-SyncNet approach and its components.

A. Echo-SyncNet

In practice, 2D echo cines are obtained from different imaging windows, capturing various cross-sections of the heart. For example, AP4 and AP2 are two perpendicular cardiac views that are most commonly used to study the left ventricle and quantify its systolic function [1].

We intend to learn an embedding space where multiple echo cines from different views can be synchronized temporally. Each cine by itself or jointly with other cines can contribute to the learning of a view point-invariant embedding space. By learning from multi-view cines, the generated representations effectively disentangle functional attributes such as contraction or expansion of the left ventricle, and opening or closing of the mitral valve.

For a given cine, every frame f_i with height, width, and channels H , W , and C , respectively, is encoded as $p_i = \phi_\theta(f_i)$, where $\phi_\theta : \mathbb{R}^{H \times W \times C} \mapsto \mathbb{R}^d$ is the encoder network parameterized by θ which maps video frames to d -dimensional feature vectors. Echo-SyncNet learns embeddings by a weighted sum of three losses as in (1): temporal and spatial intra-view and inter-view losses, where \mathcal{L}_i represents a loss function and λ_i is an associated weight:

$$\mathcal{L}_{\text{Echo-SyncNet}} = \lambda_1 \mathcal{L}_{\text{intra-spatial}} + \lambda_2 \mathcal{L}_{\text{intra-temporal}} + \lambda_3 \mathcal{L}_{\text{inter-view}}. \quad (1)$$

The loss functions that constitute the total loss leverage a shared encoder network architecture and are combined in a single multitask training framework outlined in subsection III-B. Therefore, the combined losses backpropagate through the same network parameters. The weights for the relative contribution of each loss were discovered through a hyperparameter search.

A visual description of the entire process of Echo-SyncNet is shown in Figure 2 and a detailed explanation of each component is provided in the following subsections.

B. Temporal and Spatial Intra-View Self-Supervision

The goal of the Echo-SyncNet encoder, ϕ_θ , is to produce a latent representation of a cine that is fine-tuned for robust temporal alignment. To ensure that ϕ_θ possess an adequate understanding of individual cines, we formulate two self-supervised learning objectives to take advantage of hidden spatial and temporal features.

1) Temporal Intra-View Self-Supervision: In sequential data types, such as echo cines, the ordering of elements encodes an easily understandable notion of the temporal structure. To exploit this property, it is natural to consider the self-supervised task of verifying if cines are arranged in correct chronological order. To frame this problem as a direct supervision over the quality of the embeddings, we implement a binary classifier $\mathcal{C} : \mathbb{R}^{3,d} \mapsto \mathbb{R}$ that attempts to classify sequences of three unique d -dimensional feature vectors (three-tuples) output by ϕ_θ as either appearing in sorted or unsorted order in the original cine. Inspired by [6] to make sure that the classification of the tuples is not ambiguous, we only sample tuples from temporal windows with high motion. As shown in Figure 3, a coarse optical-flow is used as a measure of motion between frames, and the training data for \mathcal{C} is chosen by considering the motion as a sampling weight—frames with higher motion have a higher probability to be selected. Within a selected temporal window, all possible sorted three-tuples are selected as positive training examples and unsorted three-tuples as negatives. A sorted tuple may follow increasing or decreasing temporal order while an unsorted tuple follows neither. This simple sequential verification task encourages the embeddings to encode information about the ordering of events in a cardiac cycle.

To function as a direct supervision over the embeddings, the joint network $(\phi_\theta \circ \mathcal{C})$ is trained end-to-end. The parameters of $\phi_\theta \circ \mathcal{C}$ are learned by minimizing the binary cross-entropy loss of the predictions on each tuple. For a training sample n this loss is defined as:

$$\mathcal{L}_{\text{temporal-intra}} = (1 - o_n) \log(1 - \hat{o}_n) + o_n \log(\hat{o}_n), \quad (2)$$

where $o_n \in \{0, 1\}$ is a binary label indicating if sample n is sorted and $\hat{o}_n \in [0, 1]$ is the predicted probability that sample n is sorted. The average value of this loss is computed over a given batch of data during training.

2) Spatial Intra-View Self-Supervision: We now formulate a self-supervised learning objective to help Echo-SyncNet understand the spatial similarities between nearby frames. We adopt a method originally proposed in [13] in the context of distance metric learning, which is the process of learning embeddings that preserve the distance between similar data points.

In our method, we partition the embeddings of an n frame cine $\mathcal{Q} = \{q_1, q_2, \dots, q_{n-1}, q_n\}$ into positive and negative inner product pairs. For a given embedding q_i , there is a single positive pair $q_i^\top q_{i+1}$. The negatives are the inner products of all embeddings pairs separated by a minimum window radius α —the window radius is a hyperparameter of our experiments further described in subsection III-B. The objective is then formulated as minimizing the mean

¹Code availability: github.com/fatemehtd/Echo-SyncNet

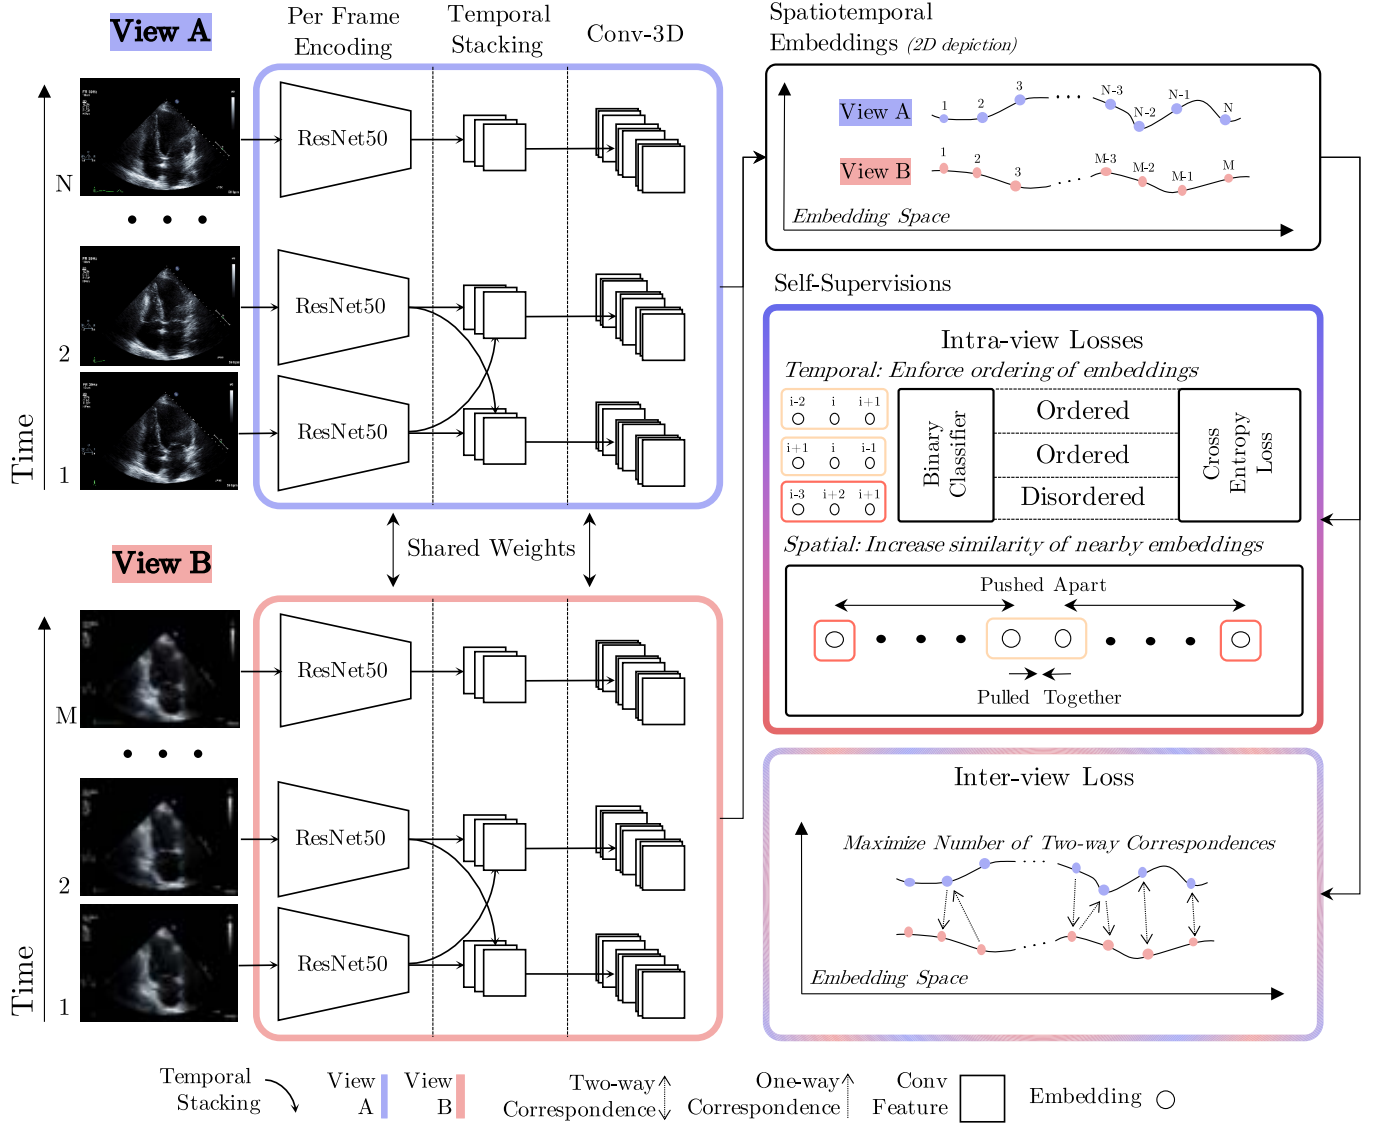


Fig. 2. The full training framework for Echo-SyncNet. The base 2D network gives an embedding for each frame after which k context frames are temporally stacked and passed through the 3D convolutional network. Losses from three self-supervisions on the encodings are linearly combined and used to train the encoder network to produce embeddings optimized for temporal synchronization. The two intra-view supervisions are used to capture temporal and spatial features found within individual cardiac echos. The inter-view supervision (synchronization loss) is used to maximize the similarity of embedded features across different echo views. It should be noted that this framework may be trained with any two echo views; here, we show the AP2 and AP4 views as an example.

cross-entropy loss on all sequences of inner products of positive and negative pairs with the positive for each embedding as the target. This loss function, further illustrated in Figure 4, is adapted from the n -tuple loss originally formulated in [13] and defined for this work in Equation (3). The intuitive objective of this loss is to maximize the inner product between adjacent elements of \mathcal{Q} while minimizing it for distant elements:

$$\begin{aligned} \mathcal{L}_{\text{spatial}}^{\text{inter}} &= \frac{1}{n-1} \sum_{i=1}^{n-1} -\log \left(\frac{e^{v_i}}{e^{v_i} + \sum_k e^{\mu_i^k}} \right) \\ &= \frac{1}{n-1} \sum_{i=1}^{n-1} -v_i + \log \left(e^{v_i} + \sum_k e^{\mu_i^k} \right), \end{aligned} \quad (3)$$

where v_i and μ_i^k are the positive and negative inner product pairs shown in Equation (4) and n is the number of frames in the cine for which this loss is computed on:

$$v_i = q_i^\top q_{i+1} \quad \mu_i^k = q_i^\top q_{i+k} \quad \text{s.t. } \text{abs}(k) \geq \alpha \text{ and } 1 \leq i+k \leq n. \quad (4)$$

This learning objective enforces that nearby and hence spatially similar frames map to similar embeddings and frames from different points of the cardiac cycle map to distant embeddings. For this loss to function properly, it is required that the training data consists of no more than one cardiac cycle. The reasoning for this restriction is to avoid pushing apart embeddings that correspond to the same cardiac event but are captured from different cycles.

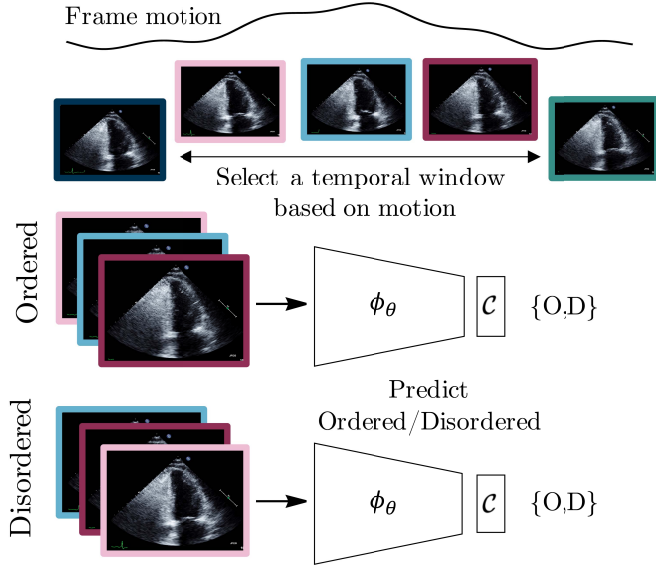


Fig. 3. Temporal Intra-view Self-supervision: three-Tuples of frames are sampled from high motion windows in a cardiac cycle. A small binary classifier \mathcal{C} is trained end-to-end with the encoder ϕ_θ to identify sequences as being ordered or disordered. It should be noted that this task is applied to all cine views separately; here, we show the AP2 view as an example.

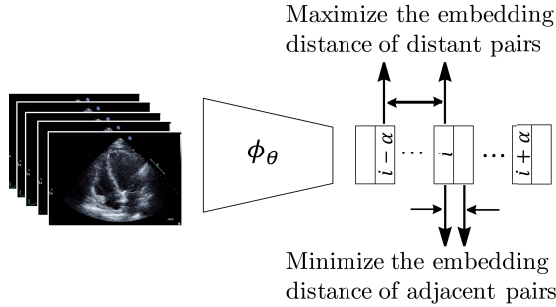


Fig. 4. Spatial Intra-view Self-supervision: Adjacent frames are selected for their embedding distance to be minimized, while distant frames (separated by at least α) are selected for their embedding distance to be maximized. It should be noted that this loss function is applied to both cine views separately; here, we show the AP4 view as an example.

C. Inter-View Self-Supervision

In this section, we describe the main framework of Echo-SyncNet for synchronizing multiple cardiac view cines. Given two echo cines from AP4 and AP2 views with N and M frames, respectively, their embeddings are noted as $\mathcal{P} = \{p_1, p_2, \dots, p_N\}$ and $\mathcal{Q} = \{q_1, q_2, \dots, q_M\}$. To formulate a self-supervised learning objective over a pair of embeddings $(\mathcal{P}, \mathcal{Q})$, we exploit their inherent consistency. Namely, the consistency property is that for each element $p_i \in \mathcal{P}$, its most similar element $q_j \in \mathcal{Q}$ should itself be most similar to p_i out of all possible choices of \mathcal{P} . More plainly put, each distinct point in a cardiac cycle from one cine should correspond directly to the same point of the cardiac cycle in another cine. If the embedding sequences \mathcal{P} and \mathcal{Q} both come from the same patient and are each sampled over a single cardiac cycle, then upholding the consistency property is a

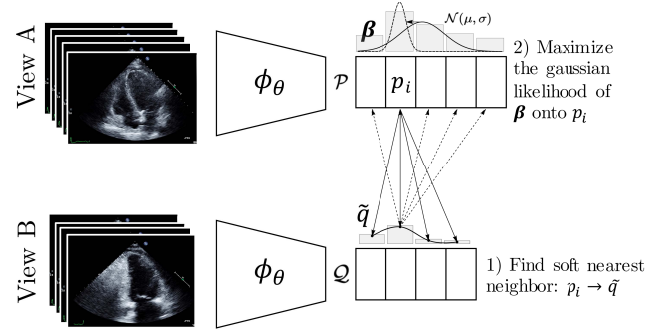


Fig. 5. Inter-view Self-supervision: Two echo cines taken in different views are encoded as two embedding sequences \mathcal{P} and \mathcal{Q} . For each embedding in $p_i \in \mathcal{P}$, the soft nearest neighbor \tilde{q}_i from \mathcal{Q} is computed and compared with each element of \mathcal{P} to make the discrete distribution β_i . The encoder ϕ_θ is trained to minimize the empirical negative log likelihood of the Gaussian prior $\mathcal{N}(\mu, \sigma)$ on β_i to the target position i .

valid objective for promoting the sharing of common features between embeddings.

Training a network to uphold the consistency property has two major pitfalls. First, due to the discrete nature of the definition, a pair of sequences can either satisfy the consistency property or not with no concept of how close a pair is to being consistent. Second, the property demands a bijective mapping between the elements of each sequence hence if the number of elements of the two sequences is not equal, they cannot be consistent. To address both of these problems, we adopt a differential loss function utilising a soft nearest neighbour approach as in [14] that is minimized by making two sequences more consistent with one another.

For each embedding, $p_i \in \mathcal{P}$, its soft nearest neighbour $\tilde{q}_i \in \text{Span}(\mathcal{Q})$ is defined by Equation (5):

$$\tilde{q}_i = \sum_j \frac{\exp(-\|p_i - q_j\|^2)}{\sum_k \exp(-\|p_i - q_k\|^2)} q_j. \quad (5)$$

Using p_i 's soft nearest neighbour \tilde{q}_i , we define β_i^k as the normalized similarity between each $p_k \in \mathcal{P}$ and \tilde{q}_i :

$$\beta_i^k = \frac{\exp(-\|\tilde{q}_i - p_k\|^2)}{\sum_j \exp(-\|\tilde{q}_i - p_j\|^2)}. \quad (6)$$

To satisfy the consistency property, we require that $\beta_i = \{\beta_i^1, \dots, \beta_i^N\}$ is maximized near the i^{th} index. The final expression for the loss is formulated as in [14] by imposing a Gaussian prior $\mathcal{N}(\mu_i, \sigma_i^2)$ defined using the empirical mean and variance of β_i . The objective is then to minimize the negative log likelihood as shown in Equation (7) and Figure 5.

$$\mathcal{L}_{\text{inter-view}} = \frac{|i - \mu_i|^2}{\sigma_i^2} + \frac{\lambda}{2} \log(\sigma_i^2),$$

$$\text{where } \mu_i = \sum_{k=1}^N k \beta_i^k \text{ and } \sigma_i^2 = \sum_{k=1}^N \beta_i^k (k - \mu_i)^2. \quad (7)$$

The regularization λ is an additional hyperparameter to the loss that will be further described in subsection III-B.

III. EXPERIMENTS AND RESULTS

A. Dataset

The echocardiography dataset used in this work was collected from the picture archiving and communication system (PACS) server of the Vancouver General Hospital with ethics approval from the Institutional Medical Research Ethics Board in coordination with the Information Privacy Office. Our first of three sets of experiments are carried out on a dataset consisting of 998 patient studies using a mutually exclusive train and test split of 854 and 144, respectively. As each patient may have multiple echos, this dataset contains a total of 1996 AP4 and AP2 pairs. No explicit temporal correspondences between cines are provided in this dataset.

The second set of experiments focuses on a direct comparison with a supervised learning approach for fine-grained cardiac cycle phase detection in AP4 cines. This set of experiments was conducted using data from 3070 patients, further details of which are explained in subsection III-D under our ablation study.

The third set of experiments focuses on the robustness and generalizability of the learned embeddings. This set of experiments was conducted using PLAX, AP4 and AP2 views taken from 1508, 2355 and 2220 patients respectively. All studies were obtained under the same ethics approval. Further details of which are explained in subsection III-E and subsection III-F. The datasets used in both the second and third sets of experiments contain labeled keyframes for both end-diastole (ED) and end-systole (ES) phases of the cardiac cycle. These labels were assigned by trained sonographers as a part of routine clinical care using both the echo cine and the ECG signal.

For preprocessing, echo cines were cleaned of user interface elements by applying a binary mask around the beam region. The pixels outside of the mask location are set to zero to make sure that the model is not biased to any annotations around the beam. Moreover, echo cines were downsized to a spatial dimension of 224×224 . For the n-tuplet and synchronization loss to function correctly, it is required that each cine in our training set contains approximately one cardiac cycle. To enforce the one cycle constraint, we either used the provided average heart rate and frame rate data to approximate a single period, or in cases where this information was not available, we applied the class agnostic video repetition counting method (RepNet) proposed in [15]. In training studies with more than one cycle, all available cycles are considered as individual training samples.

B. Implementation Details

The encoder network ϕ_θ is composed of two main modules. The first is a base 2D convolutional network that acts on each frame individually. A ResNet-50 [16] backbone truncated after the Conv4c layer is used as the architecture for the base network. Down-sampling in ResNet-50 is done using convolutions with stride 2. The dimension of the last feature map of the base network is $14 \times 14 \times 1024$. For each frame's last feature map, k adjacent maps in a context window are stacked along the temporal dimension and passed to the second part of

the encoder. We empirically chose 5 as the size of the context window. The second part of the encoder is a 3D convolutional network designed to aggregate temporal information [17]. The structure of this network is two 3D convolution layers followed by 3D max-pooling and a fully connected layer to output a 128-dimensional embedding vector. The encoder produces one embedding vector for each frame in the input cine.

For the temporal intra-view self-supervision loss, we sampled eight triplets per echo cine from which six of the triplets were shuffled. Experiments showed that it was important to have a larger percentage of disordered triplets in our training set. The classifier (C), described in subsubsection II-B.1, takes three 128-dimensional vectors as input and consists of a concatenation layer followed by two fully connected layers of output dimension 128 and 64 with ReLU activation and a final linear layer of output dimension two produces logits for binary classification.

For the spatial intra-view self-supervision loss, the window radius α for selecting distant embedding pairs was set to 5. The regularization weight λ in the inter-view loss defined in Equation (7) was set to 0.001.

When evaluating the inter-view loss on a pair of cines we randomly select 20 unique frames for each cine to be used at inputs. When then compute the per frame embeddings of each set of 20 frames using ϕ_θ at which point we can compute the inter-view loss with Equation (7). Using a fixed number of input frames allows for simpler and more efficient forward passes and computations of the loss on large batches of data in parallel. Standardizing the input frame number by randomly selecting frames showed to be superior to resampling video clips using temporal interpolation.

The three self-supervised loss functions were linearly combined to form a single objective. To determine the optimal weights of each loss term, we performed a grid search over the range $[0, 1]$ with grid size of 0.25. The decided weights were 0.25, 0.25, and 0.5 for the temporal intra-view, spatial intra-view and inter-view self-supervisions, respectively. Our models were trained using an Adam optimizer with a batch size of 4. The initial learning rate was set to 10^{-4} with a scheduled decay rate of 0.1 after every 5000 iterations.

C. Qualitative Results

To synchronize two echo cines, we first find their self-supervised representations using the encoder network ϕ_θ . Next, the correspondence of each frame across cines is found using a dynamic time warping method [18], which is an algorithm for time series alignment.

In Figure 6, the synchronization of an AP4 and AP2 echo cine is shown. To examine the details of the synchronization keyframes showing the max contraction of the left ventricle and the opening/closing of the mitral valve are shown. Additionally, we show the dynamic time warping correspondence between the embeddings of the two cines. Further information on the visualization of the cine embeddings in Figure 6 and the supplementary material is explained in Figure 7.

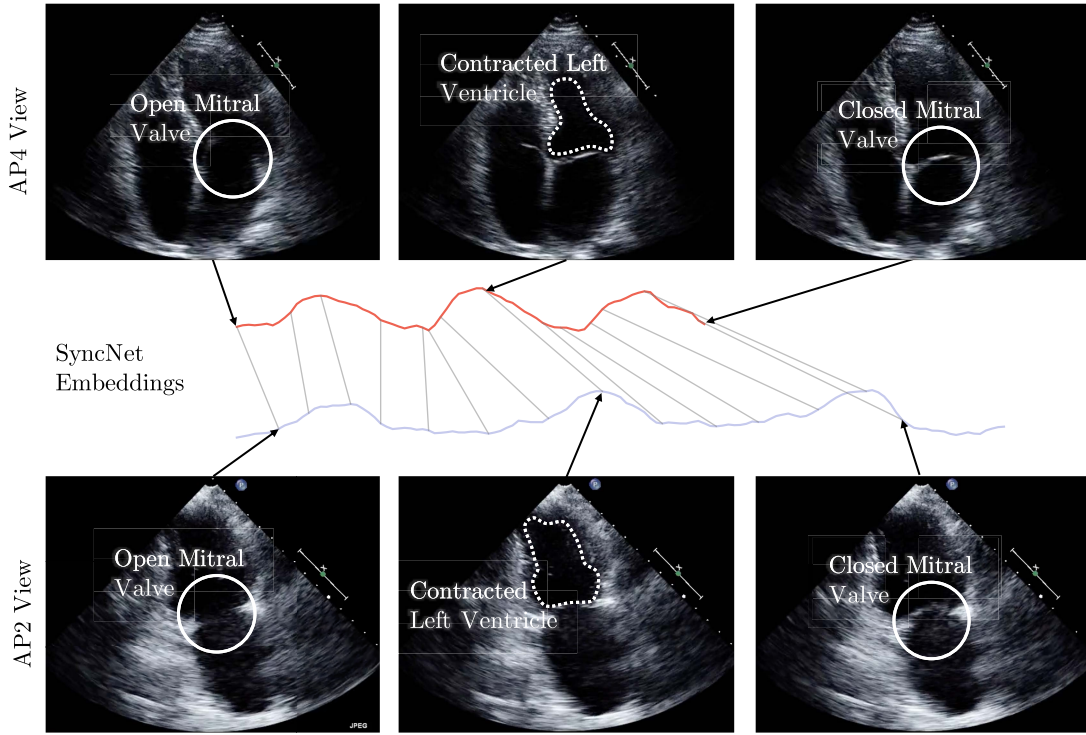


Fig. 6. A sample synchronization produced by Echo-SyncNet. We examine the synchronization of an AP4 and AP2 echo cine by comparing three distinct cardiac events: the earliest opening of the mitral valve, the maximum contraction of the left ventricle and the earliest closing of the mitral valve. We show the correspondence of these three cardiac events on the respective embeddings generated by Echo-SyncNet — AP4 (red), AP2 (blue). Additionally, a full mapping between the AP4 and AP2 embeddings is shown with thin gray lines. Please see the supplementary files for synchronized video results.

D. Empirical Ablation Analysis

We evaluated the learned representations with two frequently used metrics in prior literature to assess temporal synchronization of two videos [5], [14], [19], [20], Kendall’s Tau [21] and the coefficient of determination (R^2 score). To automatically estimate the synchronization quality between cines, we measured Kendall’s Tau for models trained with different combinations of the three self-supervisions. Kendall’s Tau ($K's \tau$) is a statistical metric that measures the ordinal association between two sequences by counting the number of pairwise concordances and discordances between elements. For every pair of frame embeddings in the first cine, (p_i, p_j) where $i < j$, the nearest neighbour correspondence in the second cine (q_α, q_β) is computed. A matching is concordant if and only if $\alpha < \beta$ and is otherwise discordant. For a given pair of cines, Kendall’s Tau is calculated as:

$$K's \tau = \frac{2}{n(n-1)}(\# \text{concordances} - \# \text{discordances}). \quad (8)$$

The highest attainable score is 1 and implies that every matching is concordant, while the lowest score is -1 and is attained only for a completely reversed synchronization. We average Kendall’s Tau over all pairs of cines in the validation set and the results are reported in Table I. Using solely inter-view self-supervision outperforms the joint spatial and temporal intra-view task as it considers two cines together as opposed to just one. However, the best performance is achieved by combining inter and intra-view supervision, which shows the

TABLE I
AP2 AND AP4 SYNCHRONIZATION PERFORMANCE EVALUATION
BASED ON KENDALL’S TAU METRIC

Method	$K's \tau$
Temporal Intra-view (TIV)	0.772
Spatial Intra-view (SIV)	0.694
Inter-view (IV) [14]	0.801
IV and TIV	0.784
IV and SIV	0.713
Echo-SyncNet (Proposed)	0.921

impact of their complementary effects to reduce the model from over-fitting to a particular loss. Given that the inter-view loss is the most significant contributor to the framework, its ability may suffer when the network is overfitting to the spatial or temporal losses, i.e. it makes nearby embeddings very dissimilar or similar. When both intra-view losses are used, the network cannot overfit to one without penalizing the other, which allows the inter-view loss to train effectively with the joint intra-view losses acting partially as regularizers to reduce overfitting. The inter/intra-view combination is our final proposed method and will be used for the remaining experiments.

Embeddings can carry semantic or structural meanings. To test the utility and robustness of our embeddings, we decided to take advantage of them in a practical downstream task, i.e. fine-grained cardiac phase detection. A typical cardiac cycle consists of two major phases, one during which

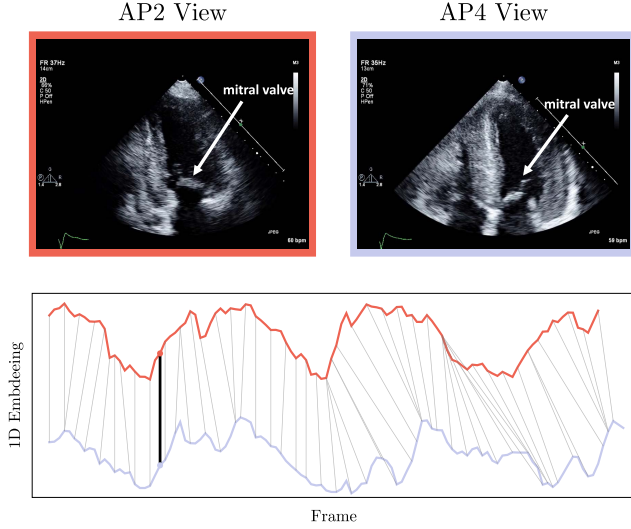


Fig. 7. Visualization of embeddings: For each view, we show trajectories of the cines in the embedding space. Using a principal component analysis reduction approach, we reduce the dimensionality of the embedding from 128 to 1 for visualization. The visualization indicates how the embeddings change as the cardiac cycle progresses and how two cines with a different number of frames can be synced together. The images show a sample correspondence point (shown with black line in the lower sub-figure) for an AP4 and AP2 cardiac echo as well as their respective embeddings. Notice the corresponding position of the mitral valve in both images, and for a more obvious correspondence of cardiac walls and chambers during a cardiac cycle please see the supplementary videos.

all chambers expand and refill with blood, termed diastole, followed by a period of contraction and pumping of blood, termed systole. In [19], authors assigned each frame a fine-grained phase value, which mimics the left ventricle volume using the following equation:

$$y_t = \begin{cases} \left(\frac{|t - T_{ES}^i|}{|T_{ES}^i - T_{ED}^i|} \right)^3, & \text{if } T_{ED}^i < t < T_{ES}^i \\ \left(\frac{|t - T_{ES}^i|}{|T_{ES}^i - T_{ED}^{i+1}|} \right)^{\frac{1}{3}}, & \text{if } T_{ES}^i < t < T_{ED}^{i+1}. \end{cases} \quad (9)$$

where t is the frame index within cardiac cycle i , T_{ED}^i is the index of end-diastolic frame in cardiac cycle i , and T_{ES}^i is the index of end-systolic frame in cardiac cycle i . Using this equation, each frame in a cardiac cycle is assigned to a phase value between 0 and 1. These values show an approximate measure of progress through each phase. To test the representation power of our embeddings, we used a linear regression to try and estimate the phase values directly from the embedding values. We carried out this experiment on a large set of AP4 cines from 3070 individual cases. We partitioned the data into mutually exclusive sets of training (60%), validation (20%) and test (20%) based on unique patients. No patient studies are shared across these sets. All the results shown were computed on unseen test dataset. Ground truth values for T_{ED} and T_{ES} were recorded by sonographers. Using the ground truths, the label for each frame was calculated by (9). The performance of this experiment was measured based on the coefficient of determination (R^2 score) which is

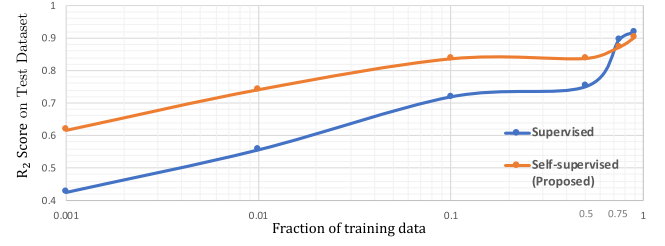


Fig. 8. Comparison of self-supervised learning with supervised learning in fine-grained phase detection with different fractions of training data.

defined by:

$$R^2 = 1 - \frac{\sum (y_t - \hat{y}_t)^2}{\sum (y_t - \bar{y})^2},$$

where \bar{y} is the mean of true labels and \hat{y}_t is the prediction for y_t . The R^2 score of the regression is reported in Figure 8. We also show the advantage of applying self-supervised learning as apposed to supervised training in scenarios with different fractions of labeled training data. Following the experts in this field [14], we train the embedding space either by the proposed self-supervised method or a vanilla supervised method. Then, the aforementioned linear regressor is trained on top of embeddings using different fractions of labeled training data. We observe significant improvement using self-supervised learning in experiments with a small amount of labeled data, from which can be inferred that the embedding space is generalizable for downstream tasks in low-labeled data scenarios. This shows that there is much untapped rich information present in the raw echo cines that can be harvested using self-supervision. A fraction of training data increases, the performance gaps narrow between supervised learning and self-supervised learning, with a slightly better performance in supervised learning as the model is explicitly optimized for the fine-grained phase-detection task. However, self-supervised approach is an effective learning strategy when training target tasks on small annotated datasets. Such strategy has a significant potential for medical imaging analysis, as most of the time, there is a shortage of labeled data. It is not only difficult to get a large dataset labeled because of the manual effort involved, but also the labels may suffer from significant intra/inter-observer variability, specifically amongst less experienced operators [22], [23].

E. Extension of Echo-SyncNet to Other Views

Although we experimented on the AP4 and AP2 views up to this point, the proposed deep learning approach does not make any *a priori* assumption on a specific echo view/views that may be used during training. We conduct a feasibility study with the use of 1508 patient studies on Parasternal Long Axis (PLAX) view, 2220 patient studies on AP2 and 2355 patient studies on AP4 obtained under the same ethics approval. We sample random pairs of cines from these three views and present each pair to the encoder network, ϕ_θ , to learn the inter- and intra-view self-supervisions. The encoded embeddings capture semantic information from this

TABLE II
ONE SHOT LEARNING RESULTS FOR ED/ES DETECTION

View	Validation Size	Units	Mean Absolute Error		Standard Deviation	
			μ_{ED}	μ_{ES}	σ_{ED}	σ_{ES}
AP4	486	frames ms	4.1 138	3.6 121	3.6 120	2.9 98
AP2	408	frames ms	6.8 136	8.0 159	4.2 84	3.5 70
PLAX	294	frames ms	3.1 122	3.9 153	2.4 95	2.6 101

new combination of views and can be used to synchronize them smoothly. Please refer to a sample synched video in the supplementary files to watch its performance. The embeddings can be used to synchronize more than two views, and as a sample, we provide a video in the supplementary files that shows AP2, AP4, AP5 and PLAX cines can be synchronized smoothly. This experiment can have significant promise for echo imaging, as it determines the breadth of network applicability over various captured echo views.

F. One-Shot Learning for Key Frame Detection

We conducted an experiment to test the usefulness of our learned embeddings in the one-shot scenario. Using only one reference cine with multiple cycles and labeled ED/ES keyframes, we test the prediction power of identifying key frames in unseen cines. The reference and a candidate cine are synchronized using Echo-SyncNet. Once the synchronization is found, the ED and ES frames are identified in the candidate cine based on their correspondence to the reference. Using the sonographer labeled ground-truth key frames for the candidate cine, the mean absolute error (μ_{ED} , μ_{ES}) and the standard deviation of the error (σ_{ED} , σ_{ES}) are computed for AP4, AP2 and PLAX views. As cines have been recorded at different frame rates, we report statistics in both the number of frames and milliseconds.

In Table II, the results of key frame detection in three different views can be observed. Please note that for these predictions, only one labeled sample has been used for each view. The errors are similar to reported inter-observer variability for manual detection of ED and ES frames by sonographers (Zolgharni *et al.* [24] demonstrated that the median disagreement among five sonographers for the identification of ED and ES phases is 3 frames).

G. Evaluation of the Proposed Method on an Unseen Public Dataset

We performed an additional validation using a publicly available echo dataset, EchoNet-Dynamic [12]. The dataset includes 10,030 videos of AP4 echocardiography from patients who received imaging as part of routine clinical treatment at Stanford University Hospital between 2016 and 2018. The location of one ED and one ES frame is labelled in each video. We extracted the embeddings of these studies using our trained model on the multiview AP4, AP2 and PLAX

TABLE III
ONE SHOT LEARNING RESULTS FOR ED/ES DETECTION IN PUBLICLY AVAILABLE ECHO DATASET (ECHO NET-DYNAMIC)

View	Validation Size	Units	Mean Absolute Error		Standard Deviation	
			μ_{ED}	μ_{ES}	σ_{ED}	σ_{ES}
AP4	10,030	frames ms	4.0 79.1	3.1 61.1	6.7 128.5	4.6 90.1

dataset. The embeddings are validated in the detection of ED and ES frame location following the same procedure as the one explained in subsection III-F. In Table III, the results of keyframe detection in this public dataset can be found. The errors are similar to the ones we got from our data, which is close to the inter-observer variability of ED and ES frame detection.

H. Limitations

Working with highly diverse datasets, especially in the medical domain, comes with a variety of unintended consequences. Performing a thorough evaluation of methods becomes increasingly costly when external factors such as patient age, ultrasound machine type, image quality, and enumerable health conditions (arrhythmias, valve disease, congenital defects, etc) are taken into account. Both large amounts of data and high-quality annotations of all external factors would be required to identify Echo-SyncNet's major weaknesses before it could be deployed in a real hospital setting. As the focus of this research is to explore the potential of a novel cardiac ultrasound encoding method, and its application to view synchronization and keyframe detection, we must defer a more detailed understanding of generalizability and robustness on unseen data to future work.

IV. CONCLUSION AND FUTURE WORK

In this paper, we proposed a self-supervised representation learning method for multi-view synchronization of echo cines. We showed that such rich yet free semantic representations can be used not only for synchronization, but also for fine-grained cardiac phase detection. The proposed technique does not make any *a priori* assumption on a specific echo view and is hence performance insensitive to various cardiac views.

This work suggests that robust and generalizable feature representations can potentially be improved with self-supervision and opens a new direction towards finding novel self-supervision paradigms for robustness and generalizability rather than developing self-supervision approaches as a mechanism of catching up to fully-supervised performance.

In future, we plan to extend this method and foster such semantic representations from high dimensional cines in analysing the rhythm of the cardiac cycle and detecting abnormalities correlated with cardiac rhythm.

ACKNOWLEDGMENT

The authors would like to thank Dale Hawley for helping with accessing data used in this research.

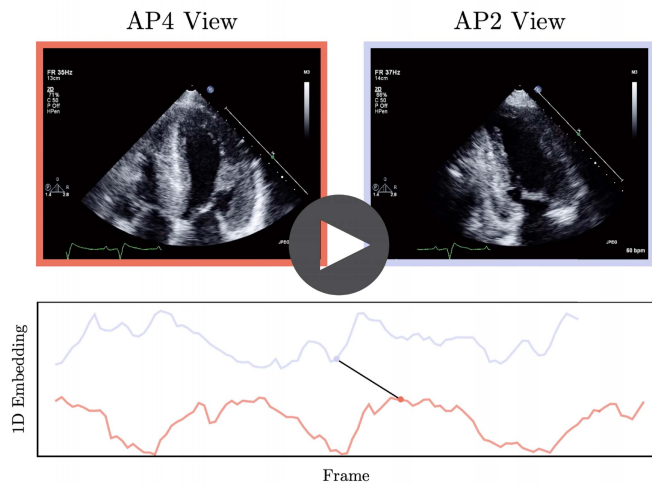


Fig. 9. A screenshot of synched cines captured from AP4 and AP2 views along with their trajectories in the embedding space. It should be noted that we reduced the embeddings' dimensions from 128 to 1 using principal component analysis for better visualization.

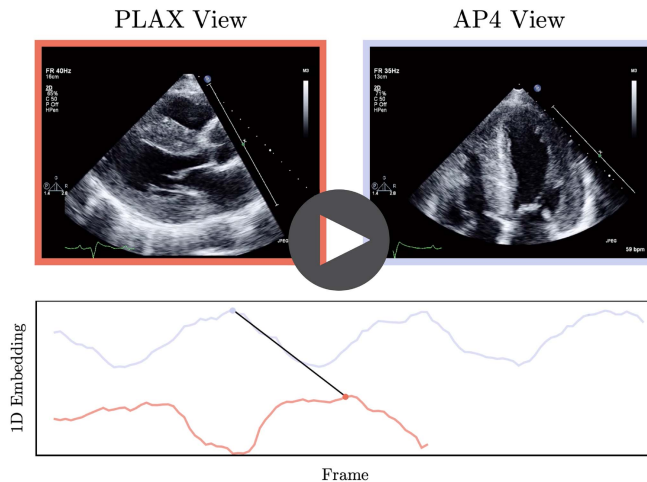


Fig. 10. A screenshot of synched cines captured from AP4 and PLAX views, showing the extensibility and generalizability of the proposed model to other cardiac views. It should be noted that we reduced the embeddings' dimensions from 128 to 1 using principal component analysis for better visualization.

APPENDIX

DESCRIPTION OF SUPPLEMENTARY MATERIAL

In [Figure 9](#), [Figure 10](#), and [Figure 11](#), we show a screenshot of synched cines captured from different views along with their trajectories in the embedding space. Please refer to the supplementary files to see the videos. In synched videos, note the fine-grained synchronization that in some time points, frames are dropped or repeated to make the movements of valves and walls be aligned between the cines. In [Figure 12](#), a snapshot of a video from four synched cines can be observed. In this example, the generalizability of the method even in the views not found in the training set can be observed.

HYPERPARAMETERS

We tabulate the list of values of the hyperparameters in [Table IV](#).

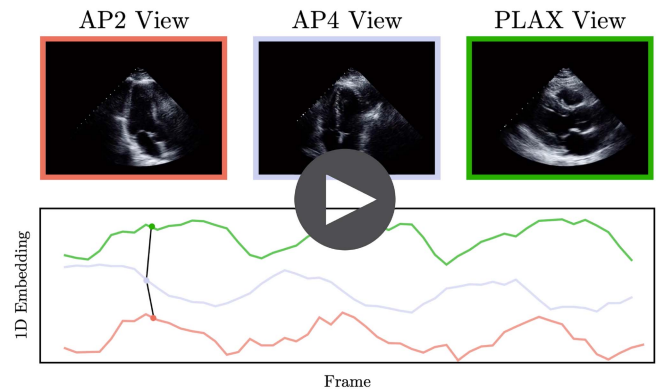


Fig. 11. A screenshot of synched cines captured from AP4, AP2, and PLAX views along with their trajectories in the embedding space. It should be noted that we reduced the embeddings' dimensions from 128 to 1 using principal component analysis for better visualization.

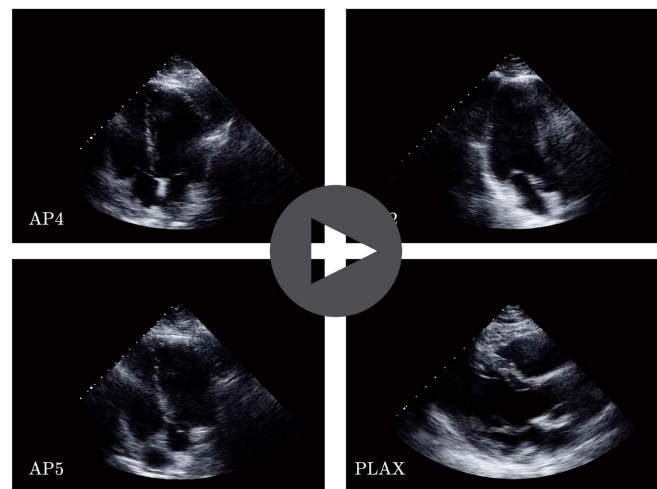


Fig. 12. A screenshot of synched cines captured from AP4, AP2, AP5, and PLAX views.

MORE INFORMATION ON DATASETS

In an effort to reduce selection bias in the data used in this work, we include ultrasound studies from various machine vendors and do not perform filtering of data based on pathology or any other external variables. The unbiased nature of data selection should help the trained models to generalize for translation to real clinical settings. All data were captured as a part of routine clinical care at Vancouver General Hospital. More detailed information of the studies used in each of the experiments can be found in [Table V](#), [Table VI](#) and [Table VII](#).

SENSITIVITY TO FRAME RATE

The dataset used in this study contains a diverse sample of inputs with different frame rates and qualities. We aimed to design and train a network that would be robust to these types of variations. To evaluate this robustness against changes in frame rate, we created embeddings of several studies at original frame rate as well as their temporally down-sampled versions. In [Figure 13](#), a qualitative sample of the sensitivity of embedding quality to frame rate is visualized. In [Figure 14](#),

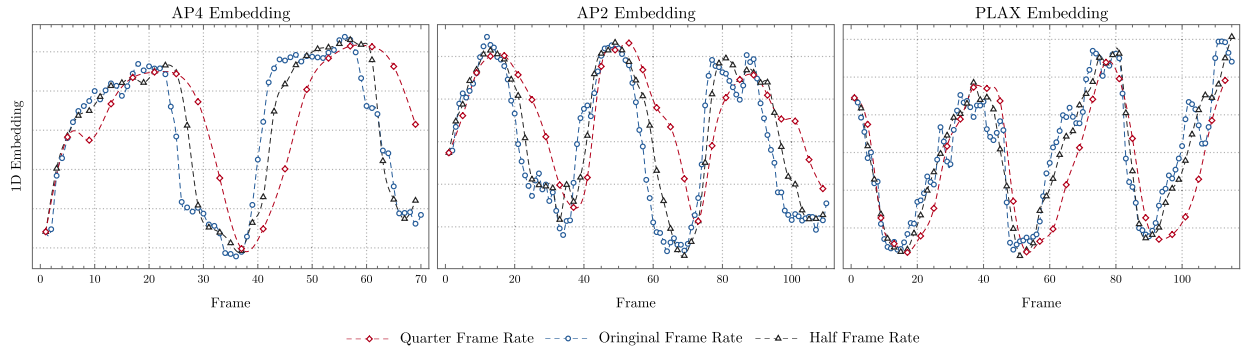


Fig. 13. Qualitative sample of the sensitivity of embedding quality to frame rate. Three sample cines of different cardiac views (AP4, AP2, PLAX) are selected. Each cine is sampled at three different frame rates (original, every second frame, every fourth frame) before being embedded. The 1D PCA reduction of each cine's embeddings is compared for different input sampling methods. The immediate visual conclusion is that even with a significantly lower input frame rate the resulting embeddings have a similar structure and will hence enable robust temporal alignment.

TABLE IV
LIST OF HYPERPARAMETERS USED

Hyperparameter	Value	Hyperparameter	Value
Training Batch Size	4	Temporal Intra-View Loss Weight (λ_1 in Equation (1))	0.25
Learning Rate	10^{-4}	Spatial Intra-View Loss Weight (λ_2 in Equation (1))	0.25
Learning Rate Decay	0.1 / 5000 batches	Inter-View Loss Weight (λ_3 in Equation (1))	0.5
Optimizer	ADAM	Spatial Intra-View Window (α in Equation (4))	5
Number of Input Frames During Training	20	Fraction of Unsorted Samples for Temporal Intra-View Loss	0.75
3D Convolution Context Window Size (k)	5	Inter-View Loss Regularization (λ in Equation (7))	0.001

TABLE V
FIRST EXPERIMENT (AP2-AP4 SYNCHRONIZATION)

	Heart Rate	Frame Time (ms)	Pixel Width (mm)
mean	69.67	25.96	0.385
std	16.49	7.42	0.071
min	28	7.44	0.155
25%	59	20.09	0.332
50%	67	23.29	0.384
75%	78	33.33	0.435
max	164	49.37	0.714

Year	Study Count
2011	62
2012	152
2013	380
2014	404

Machine	AP2-AP4 Study Count
SONOS	9
Vivid i	135
Vivid7	237
iE33	617
Total	998

TABLE VI
SECOND EXPERIMENT (FINE-GRAINED CARDIAC CYCLE PHASE DETECTION IN AP4- CAPTURED WITH iE33)

	Heart Rate	Frame Time (ms)	Pixel Width (mm)
mean	68.40	25.56	0.365
std	12.98	7.21	0.065
min	35	9.06	0.192
25%	60	19.84	0.332
50%	67	22.05	0.358
75%	76	33.33	0.395
max	141	44.60	0.767

Year	Study Count
2010	246
2011	468
2012	304
2013	772
2014	1280

a screenshot of synced cines from AP4 and AP2 views with different frame rates are visualized. Please refer to the supplementary video for a more obvious correspondence of cardiac

walls and chambers during a cardiac cycle with different frame rates. This experiment shows the robustness of the proposed synchronization method to different frame rates.

TABLE VII
THIRD EXPERIMENT (MULTI-VIEW SYNCHRONIZATION)

	Heart Rate	Frame Time (ms)	Pixel Width (mm)
mean	72.02	32.07	0.388
std	16.66	9.73	0.075
min	34	12.14	0.184
25%	61	21.50	0.332
50%	70	40.08	0.395
75%	80	40.69	0.420
max	188	43.64	0.719

Year	Study Count
2011	343
2012	352
2013	880
2014	1304
2015	248
2017	1128
2018	1828

Machine	AP2 Study Count	AP4 Study Count	PLAX Study Count
SONOS	10	9	8
Vivid i	218	235	129
Vivid7	257	394	272
VividE95	0	0	540
iE33	1735	1717	561
Total	2220	2355	1508

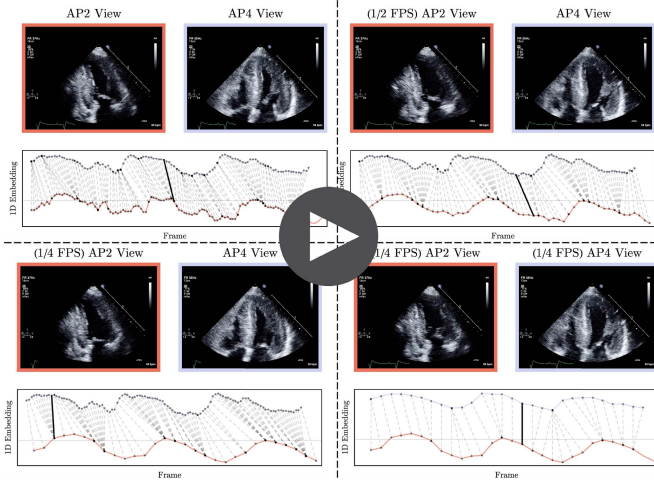


Fig. 14. A screenshot of synched cines captured from AP4 and AP2 views along with their trajectories in the embedding space. This experiment shows the robustness of the proposed synchronization method to different frame rates. We first produce downsampled versions of both cines at one-half and one-quarter of their original frame rate. Next, we create embeddings and perform synchronization using various combinations of original, half sampled and quarter sampled pairs. It should be noted that we reduced the embeddings' dimensions from 128 to 1 using principal component analysis for better visualization.

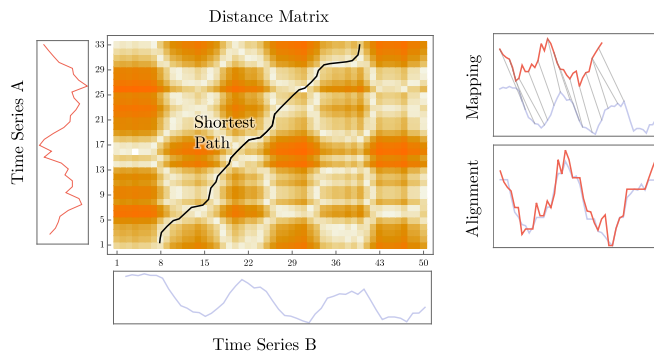


Fig. 15. Visual description of the dynamic time warping alignment algorithm. The pairwise distance matrix is computed for time series A and B. The optimal alignment between the two series is framed as a shortest path problem across the smaller dimension of the distance matrix—this is shown by the black curve. Each point on the black curve represents the mapping between series A and B (top right). The alignment is found by gluing together elements that are mapped to each other (bottom right).

DESCRIPTION OF THE ALIGNMENT ALGORITHM

Echo-SyncNet produces 128-dimensional embeddings vectors for each frame of a given input cine. Two given cines

can be temporally aligned by computing the dynamic time warping (DTW) [18] correspondence between their embedding vectors. DTW is a well known dynamic programming algorithm for time series alignment. The algorithm computes the point-wise correspondence between two temporal sequences by minimizing the sum of the distance between every pair of corresponding points while preserving temporal ordering—a visual description of the algorithm is shown in Figure 15. The complexity of the algorithm is $O(NM)$, where N and M are the lengths of the two sequences, respectively. To align more than two sequences together using DTW, it suffices to choose one as a reference and compute the warping correspondences between it and all targets.

REFERENCES

- [1] R. M. Lang *et al.*, “Recommendations for cardiac chamber quantification by echocardiography in adults: An update from the American society of echocardiography and the European association of cardiovascular imaging,” *J. Amer. Soc. Echocardiogr.*, vol. 28, no. 1, pp. 1–39, Aug. 2016.
- [2] D. Behnami *et al.*, “Dual-view joint estimation of left ventricular ejection fraction with uncertainty modelling in echocardiograms,” in *Medical Image Computing and Computer Assisted Intervention—MICCAI* (Lecture Notes in Computer Science), vol. 11765, D. Shen *et al.*, Eds. Cham, Switzerland: Springer, 2019. [Online]. Available: <https://www.springerprofessional.de/en/dual-view-joint-estimation-of-left-ventricular-ejection-fraction/17255174>, doi: 10.1007/978-3-030-32245-8_77.
- [3] M. Li *et al.*, “Recurrent aggregation learning for multi-view echocardiographic sequences segmentation,” in *Medical Image Computing and Computer Assisted Intervention—MICCAI* (Lecture Notes in Computer Science), vol. 11765D, D. Shen *et al.*, Eds. Cham, Switzerland: Springer, 2019. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-32245-8_75, doi: 10.1007/978-3-030-32245-8_75.
- [4] R. Ge *et al.*, “K-Net: Integrate left ventricle segmentation and direct quantification of paired echo sequence,” *IEEE Trans. Med. Imag.*, vol. 39, no. 5, pp. 1690–1702, May 2020.
- [5] F. T. Dezaki *et al.*, “Cardiac phase detection in echocardiograms with densely gated recurrent neural networks and global extrema loss,” *IEEE Trans. Med. Imag.*, vol. 38, no. 8, pp. 1821–1832, Aug. 2019.
- [6] I. Misra, C. L. Zitnick, and M. Hebert, “Shuffle and learn: Unsupervised learning using temporal order verification,” in *Computer Vision—ECCV* (Lecture Notes in Computer Science), vol. 9905, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham, Switzerland: Springer, 2016. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-46448-0_32, doi: 10.1007/978-3-319-46448-0_32.
- [7] B. Fernando, H. Bilen, E. Gavves, and S. Gould, “Self-supervised video representation learning with odd-one-out networks,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 3636–3645.
- [8] D. Wei, J. Lim, A. Zisserman, and W. T. Freeman, “Learning and using the arrow of time,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8052–8060.
- [9] X. Wang and A. Gupta, “Unsupervised learning of visual representations using videos,” in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2015, pp. 2794–2802.

- [10] X. Wang, A. Jabri, and A. A. Efros, "Learning correspondence from the cycle-consistency of time," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 2566–2576.
- [11] C. Vondrick, A. Shrivastava, A. Fathi, S. Guadarrama, and K. Murphy, "Tracking emerges by colorizing videos," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 391–408.
- [12] D. Ouyang *et al.*, "Video-based AI for beat-to-beat assessment of cardiac function," *Nature*, vol. 580, no. 7802, pp. 252–256, Apr. 2020.
- [13] K. Sohn, "Improved deep metric learning with multi-class n-pair loss objective," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 1857–1865.
- [14] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman, "Temporal cycle-consistency learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 1801–1810.
- [15] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman, "Counting out time: Class agnostic video repetition counting in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10387–10396.
- [16] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [17] F. T. Dezaqi, H. Girgis, R. Rohling, K. Gin, P. Abolmaesumi, and T. Tsang, "Frame rate up-conversion in echocardiography using a conditioned variational autoencoder and generative adversarial model," in *Medical Image Computing and Computer Assisted Intervention—MICCAI* (Lecture Notes in Computer Science), vol. 11765, D. Shen *et al.*, Eds. Springer, Cham, 2019. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-32245-8_78, doi: [10.1007/978-3-030-32245-8_78](https://doi.org/10.1007/978-3-030-32245-8_78).
- [18] Wikipedia Contributors. (2020). *Dynamic Time Warping—Wikipedia, the Free Encyclopedia*. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Dynamic_time_warping&oldid=959272386
- [19] F. T. Dezaqi *et al.*, "Deep residual recurrent neural networks for characterisation of cardiac cycle phase from echocardiograms," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support* (Lecture Notes in Computer Science), vol. 10553, M. Cardoso *et al.*, Eds. Cham, Switzerland: Springer, 2017. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-67558-9_12, doi: [10.1007/978-3-319-67558-9_12](https://doi.org/10.1007/978-3-319-67558-9_12).
- [20] B. Kong, Y. Zhan, M. Shin, T. Denny, and S. Zhang, "Recognizing end-diastole and end-systole frames via deep temporal regression network," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI* (Lecture Notes in Computer Science), vol. 9902, S. Ourselin, L. Joskowicz, M. Sabuncu, G. Unal, and W. Wells, Eds. Cham, Switzerland: Springer, 2016. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-46726-9_31, doi: [10.1007/978-3-319-46726-9_31](https://doi.org/10.1007/978-3-319-46726-9_31).
- [21] Wikipedia Contributors. (2020). *Kendall Rank Correlation Coefficient—Wikipedia, the Free Encyclopedia*. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Kendall_rank_correlation_coefficient&oldid=954618069
- [22] P. Bresser, J. de Beer, and Y. de Wet, "A study investigating variability of left ventricular ejection fraction using manual and automatic processing modes in a single setting," *Radiography*, vol. 21, no. 1, pp. e41–e44, Feb. 2015.
- [23] G. D. Cole *et al.*, "Defining the real-world reproducibility of visual grading of left ventricular function and visual estimation of left ventricular ejection fraction: Impact of image quality, experience and accreditation," *Int. J. Cardiovascular Imag.*, vol. 31, no. 7, pp. 1303–1314, Oct. 2015.
- [24] M. Zolgharni *et al.*, "Automatic detection of end-diastolic and end-systolic frames in 2D echocardiography," *Echocardiography*, vol. 34, no. 7, pp. 956–967, Jul. 2017.