Cyclical Self-Supervision for Semi-Supervised Ejection Fraction Prediction from Echocardiogram Videos

Weihang Dai, Xiaomeng Li, Member, IEEE, Xinpeng Ding, Kwang-Ting Cheng, Fellow, IEEE

Abstract-Left-ventricular ejection fraction (LVEF) is an important indicator of heart failure. Existing methods for LVEF estimation from video require large amounts of annotated data to achieve high performance, e.g. using 10,030 labeled echocardiogram videos to achieve mean absolute error (MAE) of 4.10. Labeling these videos is time-consuming however and limits potential downstream applications to other heart diseases. This paper presents the first semi-supervised approach for LVEF prediction. Unlike general video prediction tasks, LVEF prediction is specifically related to changes in the left ventricle (LV) in echocardiogram videos. By incorporating knowledge learned from predicting LV segmentations into LVEF regression, we can provide additional context to the model for better predictions. To this end, we propose a novel Cyclical Self-Supervision (CSS) method for learning video-based LV segmentation, which is motivated by the observation that the heartbeat is a cyclical process with temporal repetition. Prediction masks from our segmentation model can then be used as additional input for LVEF regression to provide spatial context for the LV region. We also introduce teacher-student distillation to distill the information from LV segmentation masks into an end-to-end LVEF regression model that only requires video inputs. Results show our method outperforms alternative semi-supervised methods and can achieve MAE of 4.17, which is competitive with stateof-the-art supervised performance, using half the number of labels. Validation on an external dataset also shows improved generalization ability from using our method.

Index Terms—Ejection Fraction, Echocardiogram Video, Semisupervised Learning, Video Regression

I. INTRODUCTION

Left ventricular ejection fraction (LVEF) is one of the most commonly reported measures of left ventricular (LV) systolic

Manuscript received XX XX, 2022.

W. Dai is with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail:wdai03@gmail.com).

X. Li is with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China, and also with The Hong Kong University of Science and Technology Shenzhen Research Institute (e-mails: eexmli@ust.hk) (Corresponding author: Xiaomeng Li.)

X. Ding is with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail:xdingaf@connect.ust.hk).

K.T. Cheng is with the Department of Computer Science and Engineering and the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail: timcheng@ust.hk).

Copyright (c) 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, including reprinting/republishing this material for advertising or promotional purposes, collecting new collected works for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.



1



(b) Problem setting for semi-supervised LVEF prediction

Fig. 1: (a) The labeling process for left-ventricular ejection fraction (LVEF) involves first identifying representative enddiastole (ED) and end-systole (ES) frames from the echocardiogram video sequence. The left ventricle (LV) is then segmented from these two frames to predict end-diastole volume (EDV) and end-systole volume (ESV) using method of disks [1]. LVEF is calculated as their percentage difference. (b) For semi-supervised LVEF prediction, we aim to use labeled sequences together with completely unlabeled video sequences for model training to reduce annotation requirements.

function and is used to diagnose cardiac disease [2–4]. It is calculated as the percentage volume change between the enddiastole (ED) and end-systole (ES) phases of the heartbeat, which corresponds to maximum and minimum volumes of the LV respectively. Values outside the normal range of 40% to 70% indicate potential heart failure or cardiomyopathy [2]. Manual measurement of LVEF involves identifying representative ED and ES frames from the echocardiogram sequence, tracing the LV chamber to estimate volume, and calculating the percentage difference [1, 5]; see Fig. 1a. This process is timeconsuming and suffers from large inter-observer variation, thus



Fig. 2: We approach semi-supervised LVEF prediction in two steps. (a) Step 1: We perform semi-supervised LV segmentation to obtain LV mask predictions for multi-input LVEF regression. We sample a batch of ED and ES frames (red), which have dimensions $3 \times 112 \times 112$, to calculate segmentation loss (\mathcal{L}_{seg}) for supervised segmentation. We *separately* sample an input clip of 40 frames (blue), which covers a minimum of two cardiac cycles, for enforcing cyclical consistency through minimizing Cyclical Self-Supervision loss (\mathcal{L}_{css}). This is based on the observation that the heartbeat is a cyclically repeating process, and features should be cyclically consistent. The two tasks share the same model encoder and are jointly trained. (b) Step 2: We train a multi-input model f_m (red) using mean squared error (MSE) loss (\mathcal{L}_m) on video concatenated with segmentation predictions, which has 4 input channels. The spatial context from LV mask predictions is then distilled into an end-to-end model for 3 channel video inputs, f_e (blue), by using f_m as a teacher model in teacher-student distillation. Distillation loss (\mathcal{L}_{dst}) is calculated based on MSE of labeled (\mathcal{L}_{lb}) and unlabeled samples (\mathcal{L}_{ulb}) with unlabeled loss weighting of w_{ulb} .

motivating the development of automated solutions [6, 7].

One approach for LVEF prediction is to estimate enddiastole volume (EDV) and end-systole volume (ESV) separately based on predicted LV segmentations [8–10]. These methods require accurate identification of ED and ES frames however since errors from frame selection and volume estimation can propagate to the final LVEF estimate [10, 11]. Current state-of-the-art methods achieve decent results by directly regressing LVEF from video inputs, an end-to-end approach. However, video understanding requires large amounts of labeled data. For example, Ouyang et al. [6] and Dai et al. [12] develop LVEF regression models using 10,030 annotated echocardiogram videos, which are costly to label. Moreover, the requirement for extensive annotations makes it challenging to generalize the methods for diagnosing related diseases, such as elevated B-type natriuretic peptide (BNP) levels or LV hypertrophy [4, 13]. To this end, we propose a semi-supervised method to predict LVEF from echocardiogram videos using labeled and unlabeled data. Fig. 1b shows our problem setting.

To our best knowledge, this is the first paper exploring semisupervised LVEF prediction from echocardiogram videos. We observe that, unlike natural video understanding, LVEF is directly related to changes in LV volume. We can therefore provide additional context for LVEF video regression by concatenating LV segmentation masks with video frames, as this helps highlight the relevant region. Based on this insight, we approach semi-supervised LVEF prediction in two steps: (1) we train a semi-supervised video segmentation model to generate LV segmentation predictions for individual frames; (2) we then distill the information provided by segmentation masks into an end-to-end LVEF regression model that directly performs inference using raw video inputs only.

For semi-supervised video-based LV segmentation, we introduce a novel cyclical self-supervision (CSS) method that enforces feature similarity based on the cyclical nature of the heartbeat. This is based on the observation that the heart is visually similar during the same phases of a cycle, such as at the start of contraction, and therefore should have similar features. Experiments show that segmentation masks generated using CSS improve LVEF video regression when concatenated with video as additional input and outperform alternative segmentation methods (see Table V). For model distillation, we propose teacher-student distillation to distill information captured by segmentation predictions into an endto-end model. Our method uses pseudo-labels on unlabeled videos for distillation, which also allows additional data to be used in training and further improves predictions (see Table IV). An illustration of our workflow is shown in Fig 2.

Our method outperforms existing semi-supervised approaches for video learning, achieving 4.90 mean absolute error (MAE) with one-eighth of available labels (see Table I). We also achieve 4.17 MAE using half the available labels, which is competitive with 4.10 MAE achieved using the supervised approach by Ouyang *et al.* [6] on a fully labeled dataset. External validation further shows that our semi-supervised method has better generalization ability and outperforms state-of-the-art fully supervised models on data from a different hospital (see Table IX). Our key contributions are:

- This paper is the first work to propose an annotationefficient approach to LVEF prediction from echocardiogram videos using labeled and unlabeled video sequences.
- · Given the cyclical nature of echocardiograms, we intro-

duce a novel cyclical self-supervision (CSS) method for semi-supervised video segmentation that enforces feature similarity based on temporal cyclicality.

- Because LVEF is determined by changes in the volume of the LV, we introduce a novel teacher-student distillation method to distill spatial context from LV segmentation masks into an end-to-end LVEF video regression model.
- Our method outperforms existing semi-supervised video learning techniques, achieves results competitive with fully supervised methodologies, and demonstrates good generalization through external validation ¹.

II. RELATED WORK

In this section, we review related works on automated LVEF prediction and semi-supervised learning for images and videos.

A. Automated LVEF Prediction

Segmentation-based approach. LVEF can be predicted from ED and ES frames by estimating EDV and ESV separately from segmentation masks and calculating the percentage difference [8–10]. Volumes can be directly calculated for 3D segmentation masks and approximated from 2D masks by using method of disks [1]. Leclerc *et al.* [14] estimate LVEF by first performing LV segmentation on pre-selected ED and ES frames with a standard U-Net architecture. Liu *et al.* [15] use pyramid attention modules to learn different feature resolutions for segmentation before calculating LVEF.

Segmentation approaches for video inputs are more challenging because ED and ES frames must be identified from the input sequence. What's more, LV segmentation labels are typically available only for the ED and ES frame because these are the only labels collected during manual measurement of LVEF [6, 14] (see Fig. 1a). Segmentation errors on non-ED and non-ES frames can produce incorrect volume approximations, and these can propagate to the final LVEF value, which is taken as the percentage difference between maximum and minimum volume predictions [10, 16, 11, 17]. Jafari et al. [8] use a Bayesian U-Net for frame-wise segmentation and use top and bottom percentile estimates to obtain EDV and ESV. Painchaud et al. [16] apply temporal smoothing to volume predictions to reduce outliers. Other works train separate models to first identify ED and ES frames from video before performing segmentation [18, 9].

Regression-based approach. State-of-the-art LVEF prediction directly regresses LVEF from video inputs using spatial-temporal models, which avoids the need to estimate EDV and ESV separately. Ouyang *et al.* [6] use the R2+1D ResNet [19] to perform regression, achieving an MAE of 4.10. Dai *et al.* [12] further improve predictions by proposing AdaCon, a novel contrastive learning framework for regression problems. These end-to-end regression methods produce more reliable estimates by avoiding error propagation from separate frame identification and volume approximation stages.

Video regression requires large amounts of data for training however. Ouyang and Dai use 10,030 echocardiogram video sequences in their works for example [6, 12], which requires extensive labeling effort. What's more, current regression approaches do not make use of the LV segmentation labels that are collected as part of the labeling process. Segmentation labels still provide valuable information because LVEF is directly related to changes in the LV. The spatial context from segmentation masks contains location as well as volume information of the LV chamber [1]. Works by Ouyang *et al.* [13] also support the fact that spatial-temporal networks attend to the LV region when learning LVEF regression. Models trained to perform LV segmentation have specialized knowledge of the LV region, and *this can be used to provide additional guidance for LVEF video regression* if used effectively.

Our proposed method is one of the first to incorporate knowledge from segmentation models into LVEF video regression. By distilling spatial context from LV segmentation predictions into an end-to-end video regression model, we can achieve results competitive with state-of-the-art performance with significantly fewer labeling requirements.

B. Semi-Supervised Learning

Segmentation. Semi-supervised image segmentation methods combine augmentations with consistency regularization or pseudo-labels to learn from unlabeled data [20-28]. The meanteacher approach [20] uses a teacher model to generate pseudolabels for the student model. CCT [23] applies augmentations to intermediate features and enforces consistency on predicted outputs. CPS [25] uses predictions from co-trained segmentation models for cross-supervision. Other approaches using contrastive learning [29, 30] and class activation maps [31] have also been explored. Video segmentation methods makes additional use of temporal information either through optical flow for pseudo-label generation [32] and multi-task learning [33, 34], or through multi-frame attention [35-37]. These methods achieve strong results on benchmark tasks but are typically designed for natural videos. They are not able to utilize the unique characteristics of echocardiogram sequences.

Semi-supervised methods have also been applied to LV segmentation by making use of unlabeled non-ED and non-ES frames in echocardiogram sequences. Works by Zhang et al. [38] and Pedrosa et al. [39] use different label propagation techniques to generate pseudo-labels for LV segmentation. Others propose jointly learning motion flow and segmentation from video and enforcing consistency between the two tasks [40, 41, 10]. These methods share similarities with natural video segmentation and do not make use of the special characteristics of echocardiogram sequences. Moreover, they require the ED and ES frames in a sequence to be labeled in order to apply motion propagation and *cannot use completely* unlabeled sequences. Our proposed CCS framework uses the cyclical nature of the heart as a prior constraint and can be applied to labeled and unlabeled sequences for improved segmentation and downstream LVEF regression performance (see Table V).

Video Understanding. Semi-supervised learning has also been applied to tasks such as video classification, action recognition, and temporal action proposal. State-of-the-art methods

¹Code is available at https://github.com/xmed-lab/CSS-SemiVideo.

share similarities with those for images and use consistency regularization and pseudo-labels on unlabeled data. Xiong et al. [42] generate pseudo-labels by ensembling multi-view predictions. Jing et al. [43] train additional image classifiers to enforce consistency on unlabeled videos for classification. Xu et al. [44] use co-trained models to generate pseudo-labels for each other to perform action recognition. These methods target classification tasks and cannot be applied directly to LVEF regression without modification however.

Semi-supervised video regression tasks receive less attention compared to classification tasks and have fewer established methods. Ding et al. [45] propose feature perturbation strategies for semi-supervised temporal action localization. Ji et al. [46] use mean-teacher learning by generating pseudolabels on unlabeled inputs with a teacher model. These generalized techniques achieve decent performance on natural video but are not well-tailored to echocardiograms. We show in Table I that our novel approach outperforms existing stateof-the-art methods by effectively making use of the unique properties of the problem domain.

III. METHODOLOGY

Our proposed methodology consists of two steps: (1) we use a novel cyclical self-supervision (CSS) method to perform semi-supervised LV segmentation using labeled and unlabeled echocardiogram sequences (2) we distill the context provided by LV segmentation predictions into an end-to-end video regression model with teacher-student distillation. Fig. 2 illustrates the overall flow. We expand on the two steps in the subsections below.

We introduce notations and denote $\mathcal{D} := \{(X_i, Y_i)\}_{i=1}^N$ as the labeled dataset with N samples, where $X_i :=$ $\{x_i^1, x_i^2, ..., x_i^{T_i}\}$ is an echocardiogram video sequence and x_i^t is the frame at time-step t. $Y_i := \{EF_i, LV_i^{ED}, LV_i^{ES}\}$ represents the LVEF label and LV segmentation masks for ED and ES frames, respectively. We denote the unlabeled dataset as $\mathcal{D}' \coloneqq \{X'_{i'}\}_{i'=1}^{N'}$ which has no corresponding labels.

The segmentation model consists of a feature encoder and decoder, which we denote $f_{seg} = f_{dec} \circ f_{enc}$. The encoder generates feature embeddings $z_i^t = f_{enc}(x_i^t) \in \mathbb{R}^d$ where d is the feature dimension. The decoder f_{dec} generates LV segmentation predictions based on embeddings, $\hat{LV}_i^t = f_{dec}(z_i^t)$.

Two deep regression models are used for LVEF prediction. We denote f_m as the multi-input LVEF regression model that accepts inputs of two types: raw echocardiogram videos concatenated with the LV segmentation mask prediction. We denote f_e as the end-to-end model which only accepts raw echocardiogram video inputs.

A. Cyclical Self-Supervision (CSS) for Video-based Semi-Supervised LV Segmentation

1) Motivation: The heart pumps blood in a cyclical process and is visually similar during matching phases of the cardiac cycle, such as at the start or end of contraction. This is evident from observing echocardiogram videos as they typically feature repeating sequences of cardiac motion. Based on this observation, we aim to train an encoder on labeled

and unlabeled echocardiogram sequences by learning feature embeddings that are consistent with the cyclical nature of the heartbeat. Enforcing cyclicality allows us to utilize additional unlabeled data for training, which also helps to reduce overfitting to labeled ED and ES frames.

Identifying and enforcing cyclicality within echocardiograms is challenging however because there are typically no labels associated with cycle duration. The number of frames per cycle varies depending on the patient's heart rate and the video frame rate, and different sequences have different numbers of cycles. To formulate an unsupervised methodology, we observe that humans can recognize cyclical repetition without cycle labels by performing the following actions: 1) We first scan ahead of a reference point and find a time-point where the contents are similar to our reference phase; 2) We then skip a few frames ahead of the reference and matching timepoints to see if the contents continue to be aligned. If this is true for all sampled references then the sequence is cyclically repeating. We can also express this process rigorously through the following relation: if a phase at time-point p^* of a cyclical sequence is the same as time-point q^* , then the phase at timepoint $q^* + c$ is the same as time-point $p^* + c$, where c is some temporal offset (see Fig. 3a for illustration). We design CSS as a differentiable loss function to enforce this relationship on features such that they exhibit temporal cyclicality.

Labels for p^* and q^* are unavailable for echocardiogram sequences, which means we cannot rely on supervised methods. Instead, we can use feature similarity to calculate the probability a phase matches with our reference for different time-points. This mimics the process of visually searching for similar content to find a matching phase (see Fig. 3b). We can then use the probabilities to calculate an expected value for the phase at $q^* + c$ to match with $p^* + c$, which mimics the process of checking for continued alignment (see Fig. 3c). By enforcing a close match between the two, we ensure the features learned by our model encoder are cyclically consistent. To ensure features are also relevant to the LV segmentation task, we train CSS jointly with supervised segmentation using a shared encoder. We design our CSS loss function based on these considerations.

2) Loss Formulation: Given some input sequence X_i that covers at least two cardiac cycles, we first divide it into template region P, search region Q, and a buffer region Bused to accommodate frame selection. We analytically define a phase, *i.e.* a given stage in the cardiac process, by E_i^t , an array of s feature embeddings starting at time-point t:

$$E_i^t = [z_i^t, z_i^{t+1}, ..., z_i^{t+s-1}].$$
(1)

Our design for E_i^t captures both spatial and temporal characteristics, which reflects how a person visually identifies video phases based on sequences of key features. Given some template phase $E_i^{p^*}$ where p^* is randomly selected from P, we can enforce cyclical consistency by finding a matching phase in Q and ensuring that the phases continue to match after an equal temporal offset. In other words, we find $q^* \in Q$ and ensure that $E_i^{q^*+c}$ matches with $E_i^{p^*+c}$ for offset c. First, we find the probability α_i^q that the phase starting at

time-point $q \in Q$ is the same phase as our template $E_i^{p^*}$:



(c) Offset matching in the template region

Fig. 3: (a) We tend to perform the following actions when we visually check for cyclicality: 1) We search ahead of our reference point for a matching time-point where the contents are highly similar 2) We skip ahead a number of frames and check that the contents remain aligned. Expressed rigorously: if the phase at time-point p^* of a cyclical process is the same as time-point q^* , then the phase at time-point q^*+c is the same as time-point $p^* + c$. We enforce this relationship on feature embeddings to reflect cyclicality of cardiac motion. (b) Using template phase $E_i^{p^*}$ selected from P, we calculate matching probability α_i^q for all time-points $q \in Q$. This is similar to visually searching for a matching phase. (c) We calculate a soft match $\tilde{E}_i^{q^*+c}$ by weighting embeddings by α_i^q and find the closest matching phase in P, which should be at $p^* + c$. This is similar to visually checking for continued alignment after a temporal offset of c frames.

$$\alpha_i^q = \frac{\exp(\gamma(E_i^{p^*}, E_i^q) \cdot \tau)}{\sum_{j \in Q} \exp(\gamma(E_i^{p^*}, E_i^j) \cdot \tau)},$$
(2)

where τ is a temperature parameter and γ is a similarity function based on inverse squared distance of feature embeddings:

$$\gamma(E_i^{p^*}, E_i^j) = -\frac{1}{d \times s} \|E_i^{p^*} - E_i^j\|^2$$

= $-\frac{1}{d \times s} \sum_{k=0}^{s-1} \|z_i^{p^*+k} - z_i^{j+k}\|^2.$ (3)

As defined earlier, d is the feature dimension and s is the number of embeddings. Fig. 3b illustrates this step, which is similar to visually finding a matching phase.

Although we do not know the exact value of $E_i^{q^*+c}$ without knowing the true value of q^* , we can calculate a soft expectation for $E_i^{q^*+c}$ by using the matching probabilities α_i^q . By finding the probability-weighted average of feature embeddings across time-steps, we can get the expected value of embeddings for $E_i^{q^*+c}$. We define a probability-weighted feature vector \tilde{z}_i^{θ} as:

$$\tilde{z}_i^{\theta} = \sum_{j \in Q} \alpha_i^j z_i^{j+\theta} \,. \tag{4}$$

where θ denotes a temporal offset. The expected value for $E_i^{q^*+c}$, which we denote $\tilde{E}_i^{q^*+c}$, can then be formulated as:

$$\tilde{E}_{i}^{q^{*}+c} = \left[\tilde{z}_{i}^{c}, \tilde{z}_{i}^{c+1}, ..., \tilde{z}_{i}^{c+s-1}\right],$$
(5)

where the c + s - 1 frames that lie outside of region Q are taken from buffer region B. This soft expectation serves as a differentiable approximation of $E_i^{q^*+c}$ and can be used to perform matching in the template region. We find the probability that the phase at $p \in P$ matches with $\tilde{E}_i^{q^*+c}$ by calculating:

$$\beta_i^p = \frac{\exp(\gamma(\tilde{E}_i^{q^*+c}, E_i^p) \cdot \tau)}{\sum_{k \in P} \exp(\gamma(\tilde{E}_i^{q^*+c}, E_i^k) \cdot \tau)} \,. \tag{6}$$

Fig. 3c illustrates this step, which is similar to visually checking if the contents remain aligned after a temporal offset. Given the cyclical properties of echocardiograms, β_i^p should be highest for $p^* + c$. We use cross-entropy loss with ground truth $p^* + c$ to obtain the final CSS loss value \mathcal{L}_{css} . Taken across labeled and unlabeled sequences, we have:

$$\mathcal{L}_{css} = -\frac{1}{N+N'} \sum_{i \in N, N'} \sum_{k \in P} \mathbb{1}_{k=p*+c} \log(\beta_i^k), \quad (7)$$

where N and N' refer to the number of labeled and unlabeled samples respectively. We note that CSS can be performed even with multiple matching phases present in search region Q, which we illustrate in Appendix A. This is vital because it is not possible to guarantee the number of cycles present in the input clip due to variations in cycle rate. CSS does not require any labels regarding cycle length or duration and is the first to perform this task unsupervised.

Our objective for enforcing cyclical feature consistency within sequences is to perform semi-supervised segmentation using unlabeled frames. Thus, CSS is trained jointly with supervised LV segmentation through a shared encoder to ensure learned features are relevant to the segmentation task. We obtain predictions on labeled ED and ES frames by applying the segmentation decoder to their feature embeddings:

$$\hat{LV}_{i}^{ED} = f_{dec}(z_{i}^{ED}); \ \hat{LV}_{i}^{ES} = f_{dec}(z_{i}^{ES})$$
 (8)

and use the corresponding ground-truth labels LV_i^{ED} and LV_i^{ES} for supervision. We apply standard pixel-wise crossentropy loss to obtain the segmentation loss \mathcal{L}_{seg} . The total

5

loss, \mathcal{L}_{vol} , for our semi-supervised LV segmentation method with CSS weight of w_{css} is then:

$$\mathcal{L}_{vol} = \mathcal{L}_{seg} + w_{css} \mathcal{L}_{css} . \tag{9}$$

Because CSS is an unsupervised constraint and does not require labels, *sampling for the two tasks is done independently*.

We also note that CSS is applied on intermediate features instead of segmentation predictions directly to avoid over constraining the model. This is because motion noise can cause slight changes in the true LV position, and enforcing cyclically equivalent segmentations will force the model to learn wrong predictions. By enforcing cyclical consistency only on feature embeddings, which is an abstract representation that captures key characteristics [47, 48], the model is encouraged to learn similar key features for matching phases. Slight variations in LV appearance can then be accurately reflected in predictions by the model decoder.

B. Teacher-Student Distillation for LVEF Regression

LV segmentation masks provide additional context that can supplement LVEF video regression. This is because LVEF is specifically related to changes in the LV and overlaying segmentation masks gives the regression model additional spatial information about the relevant region. We first generate LV prediction masks for every frame using our trained segmentation model and concatenate them with raw video as input for multi-input regression. We use a spatial-temporal model for f_m and train it using available LVEF labels from the labeled dataset with mean squared error (MSE) loss:

$$\mathcal{L}_m = \frac{1}{N} \sum_{i \in N} (\hat{EF}_{m,i} - EF_i)^2 , \qquad (10)$$

where $\hat{EF}_{m,i}$ is the prediction from model f_m for sample *i*. Results in Table IV show that LVEF prediction with additional segmentation masks consistently improves performance.

There are still limitations to the multi-input model however, as inference requires pre-computing segmentation masks for every frame which is impractical. Also, f_m is trained using only labeled data and ignores the unlabeled dataset, which can also provide useful information. We propose teacher-student distillation, which simultaneously performs semi-supervised regression using unlabeled videos, whilst also distilling the multi-input model f_m into an end-to-end model f_e . Our approach is novel in that *it is the first to distill spatial knowledge* from a LV segmentation model into a spatial-temporal video regression model for LVEF prediction.

To achieve this, we train f_e on the labeled dataset using LVEF labels and on the unlabeled dataset using pseudo-labels from f_m . We use the same architecture as f_m but with a reduced input channel. Using f_m as a teacher model for pseudo-labels serves two purposes. Firstly, f_e learns to predict like the teacher model by using pseudo-labels for guidance, thereby distilling information from segmentation inputs into its own network. Secondly, training with pseudo-labels on unlabeled data is a form of semi-supervised regression, which leads to further improvements over f_m (see ablation results in

Table IV). The overall result is a more efficient and accurate distilled model, f_e . We calculate the loss for labeled (\mathcal{L}_{lb}) and unlabeled (\mathcal{L}_{ulb}) samples using MSE:

$$\mathcal{L}_{lb} = \frac{1}{N} \sum_{i \in N} (\hat{EF}_{e,i} - EF_i)^2 , \qquad (11)$$

$$\mathcal{L}_{ulb} = \frac{1}{N'} \sum_{i' \in N'} (\hat{EF}_{e,i'} - \hat{EF}_{m,i'})^2 \,. \tag{12}$$

where $\hat{EF}_{e,i}$ and $\hat{EF}_{e,i'}$ are the predictions from f_e for sample *i* and *i'* respectively. Training is done by minimizing the loss function \mathcal{L}_{dst} :

$$\mathcal{L}_{dst} = \mathcal{L}_{lb} + w_{ulb} \mathcal{L}_{ulb} , \qquad (13)$$

where w_{ulb} is the weight value for unlabeled samples. Step 2 of Fig. 2 illustrates this step. The final trained model f_e can then be used directly on raw video input for inference without pre-computing segmentation masks.

IV. EXPERIMENTS

A. Implementation Details

EchoNet-Dynamic Dataset: We conduct experiments using EchoNet-Dynamic [5], which consists of 10,030 apical-fourchamber echocardiogram videos labeled with LVEF values and LV segmentation tracings for reference ED and ES frames (Fig. 1a shows an example). The echocardiogram videos are collected from Stanford University Hospital using five different ultrasound machines with an average length of 175 frames, covering multiple cardiac cycles. The videos are saved using three color channels and have been rescaled to 112×112 pixels with auxiliary text removed. Summary statistics are shown in Table A2. We also show frame similarity plots in Fig. A2 of the Appendix to highlight cyclical repetition. The dataset has been pre-divided into 7,465 videos for training, 1,288 videos for validation, and 1,277 videos for testing. We maintain the original data splits but treat subsets of the training data as labeled data and the remainder as unlabeled data.

CAMUS dataset: We perform external validation using the four-chamber echocardiogram sequences in CAMUS [14], which are taken from 500 patients at the University Hospital of St. Etienne in France with GE Vivid E95 ultrasound scanners. LVEF labels are available for 450 patients. Sequences begin with the ED frame and terminate with the ES frame or vice-versa. The sequences have an average length of 20 frames and capture only part of the cardiac cycle. They are saved in varying sizes and have been further classified as good, medium, or poor quality as determined by the labeling cardiologist [14]. Summary statistics are shown in Table A3.

Training CSS: We use the DeepLabV3 [49] architecture for our segmentation model, which consists of a ResNet-50 [50] encoder for f_{enc} and an Atrous Spatial Pyramid Pooling (ASPP) segmentation decoder for f_{dec} . The model is trained using SGD with a learning rate of 10^{-5} and momentum of 0.9. We jointly train the model with CSS and segmentation loss and use mini-batch random sampling for optimization.

7

 TABLE I

 Comparison with state-of-the-art supervised and semi-supervised methodologies using 1/8, 1/4, 1/2, and all training labels on EchoNet-Dynamic.

 Note that "Supervised" methods use only labeled data while "Semi-supervised" methods use both labeled and unlabeled data. We report mean results ± standard deviation of five separate runs.

			MAE Values ↓	·		
Туре	Method	Backbone	1/8 labels	1/4 labels	1/2 labels	All labels
Supervised	Ouyang et al. [6]	R2+1D	5.64 ± 0.08	4.80 ± 0.05	4.35 ± 0.04	4.10 ± 0.04
Supervised	Dai <i>et al</i> . [12]	R2+1D	5.47 ± 0.07	4.65 ± 0.06	4.28 ± 0.04	$\textbf{3.86} \pm \textbf{0.03}$
Sami	Ji et al. [46]	R2+1D	5.61 ± 0.07	4.76 ± 0.06	4.31 ± 0.04	-
Supervised	Xu et al. [44]	R2+1D	5.08 ± 0.04	4.49 ± 0.04	4.28 ± 0.04	-
Supervised	Ours	R2+1D	$\textbf{4.90} \pm \textbf{0.04}$	$\textbf{4.34} \pm \textbf{0.03}$	$\textbf{4.17} \pm \textbf{0.03}$	-
			\mathbf{R}^2 Values \uparrow			
Туре	Method	Backbone	1/8 labels	1/4 labels	1/2 labels	All labels
Supervised	Ouyang et al. [6]	R2+1D	$60.6\% \pm 0.8$	$71.2\% \pm 0.6$	$76.7\% \pm 0.4$	$80.5\% \pm 0.2$
Supervised	Dai et al. [12]	R2+1D	$62.8\% \pm 0.6$	$73.4\% \pm 0.5$	$78.1\% \pm 0.3$	$82.8\%~\pm~0.2$
Sami	Ji et al. [46]	R2+1D	$61.6\% \pm 0.6$	$71.4\% \pm 0.6$	$77.2\% \pm 0.4$	-
Supervised	Xu et al. [44]	R2+1D	$66.8\% \pm 0.4$	$75.8\% \pm 0.3$	$77.6\% \pm 0.3$	-
Supervised	Ours	R2+1D	$71.1\%~\pm~0.4$	77.0% \pm 0.4	$80.1\%~\pm~0.3$	-

Every iteration, we sample 20 echocardiogram sequences from the labeled dataset and use only their labeled ED and ES frames to minimize \mathcal{L}_{seq} . We separately sample an input clip of 40 frames from either the labeled or unlabeled dataset at a rate of 1 in every 3 frames to minimize \mathcal{L}_{css} . This clip length sufficiently covers two cardiac cycles based on the range of human pulse rates and the frame rate used for EchoNet-Dynamic (see Appendix B for calculations). We set frames 1-15 as template region P, 16-36 as search region Q, and 37-40 as buffer region B. We choose s = 3 and c = 2empirically based on the validation set and set $w_{css} = 0.01$ and $\tau = 10$ after performing a coarse parameter search (see Section IV-F). p^* is randomly sampled every iteration from the range 1-13. The frames are passed through the encoder f_{enc} to calculate \mathcal{L}_{css} , which is jointly optimized with \mathcal{L}_{seq} . We train for 25 epochs, where an epoch is when all labeled data in \mathcal{D} has been used once for supervised segmentation.

Training teacher-student distillation: We use the R2+1D ResNet architecture [19], pretrained on Kinetics 400 [51], for f_m and f_e . This is the same baseline architecture used by previous works [5, 6, 12], although we also include backbone comparisons in Appendix C. The first layer of f_m is modified to accept an additional input channel for multi-input regression. The models are trained with SGD using learning rate 10^{-4} and momentum of 0.9. For minimizing \mathcal{L}_m , we sample a batch of 20 video clips taken from labeled echocardiogram sequences for every iteration. The clips are 32 frames in length and sampled at a rate of 1 in every 2 frames. For minimizing \mathcal{L}_{dst} , we sample 20 clips from labeled sequences and 10 clips from unlabeled sequences every iteration. We set $w_{ulb} = 5$ after performing a coarse search to determine stable weights (see Section IV-F). Both models are trained for 25 epochs.

We use PyTorch and run experiments on four Titan-RTX GPUs. We perform evaluation using MAE and R^2 . Each experiment is run five times using the same data splits with different random seed initializations. Mean results with standard deviation are reported. P-Values where reported are calculated based on the Student's t-test to indicate level of statistical significance.

B. Comparison with State-of-the-art Methodologies

We demonstrate that our proposed method outperforms state-of-the-art supervised approaches by Ouyang *et al.* [6] and Dai *et al.* [12], as well as state-of-the-art semi-supervised video learning techniques by Ji *et al.* [46] and Xu *et al.* [44] for LVEF prediction. We apply different semi-supervised settings where 1/8, 1/4, and 1/2 of available labels are used for training and the rest is treated as unlabeled data. Supervised methods use both labeled and unlabeled data. We also include results using a fully labeled dataset for supervised methods as reference.

We directly follow the implementation in [6, 12] for supervised methods. Because there are no existing semi-supervised video regression approaches used for LVEF prediction, we adapt the methods by Ji *et al.* [46] and Xu *et al.* [44] such that they can be applied to our task, which we briefly detail below. The same backbone architecture, batch size, clip length, and frame sampling rates are used by all methods for fairness.

Ji *et al.* [46]: The authors use a mean-teacher approach together with temporal and frame-wise augmentations for semi-supervised temporal localization. We use the same mean-teacher method but introduce a warm-up period of 10 epochs before using pseudo-labels from the teacher model, which we find leads to better results. We only use random frame jittering and pixel noise as our augmentation operations as we find excessive augmentations handicaps performance.

Xu *et al.* [44]:. The authors propose a cross-model pseudolabeling scheme, where predictions for weakly augmented inputs are used as pseudo-labels for strongly augmented inputs. We change the classification output to a single regression prediction for our task. We use frame jittering and pixel noise as our strong augmentations and use original frames for weak operations. We use the same R2+1D architectures for the primary and auxiliary backbones to ensure fairness.

We show results of our comparison in Table I. Our method gives the best overall performance across all levels of label availability, achieving MAE of 4.90, 4.34, and 4.17 using one-eighth, one-quarter, and one-half of available labels respectively. \mathbf{R}^2 is also significantly higher than alternative approaches for all semi-supervised settings (p-Value < 0.05). Furthermore, we achieve MAE of 4.17 and \mathbf{R}^2 of 80.1% using

This article has been accepted for publication in IEEE Transactions on Medical Imaging. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TMI.2022.3229136

IEEE Transactions on Medical Imaging, VOL. XX, NO. XX, XXXX 2022



0X678B0D95D2EF1F37, ED Frame

Fig. 4: Qualitative samples of attention heatmaps on EchoNet-Dynamic using different methods. Models are trained using one-eighth of labeled data, with remaining samples treated as unlabeled data. Models attend better to the LV region when trained using our method than alternative approaches.

TABLE II

Dice score between LV segmentation labels and pixels with the top 5% highest absolute gradient values for different methods. Results are shown for the EchoNet-Dynamic dataset. Pixel gradients are calculated using SmoothGrad [52]. We use models trained with only one-eighth of available labels.

Туре	Method	Dice↑ ED	Dice↑ ES	Dice↑ Overall
Supervised	Ouyang et al. [6]	19.2%	19.5%	19.4%
Supervised	Dai <i>et al</i> . [12]	25.7%	39.1%	34.0%
Sami	Ji et al. [46]	29.0%	30.0%	29.6%
Supervised	Xu et al. [44]	31.7%	47.2%	40.9%
	Ours	41.5%	56.1%	50.7%

only half the available labels, which is competitive with MAE of 4.10 and \mathbf{R}^2 of 80.5% achieved by Ouyang *et al.* [6] when trained on a fully labeled dataset.

To interpret model performance, we visualize the attention heatmaps generated by the different methods. We use Smooth-Grad [52] to calculate the absolute gradient of each input pixel relative to the prediction, which indicates its relative importance to the model, and overlay the normalized values over the frame to generate the heatmap. This is the same approach used by Ouyang *et al.* [6]. We use models trained with one-eighth of available labels for comparison. Because our method provides additional spatial context through mask predictions to guide video regression, we expect to see higher absolute gradients for pixels in the LV region compared to alternative approaches. We confirm that this is observed for most samples in the test set and show examples in Fig. 4.

To quantify model attention to the LV region, we calculate a Dice score between segmentation labels and the most important pixels, which we determine based on the top 5% of pixels with the largest absolute gradients. We show results on labeled ED and ES frames in Table II. We can see that the overall Dice value of our method (50.7%) is significantly higher than alternative methods, indicating greater attention by the model

TABLE III

LVEF prediction results using TCCL compared with CSS (Ours) on EchoNet-Dynamic. One-eighth of labels are used as the labeled dataset, and remaining samples are used as unlabeled data. Results are shown for multi-input model f_m and end-to-end model f_e . We report mean results \pm standard deviation of five separate runs.



Fig. 5: t-SNE plot of features learned using different methods on EchoNet-Dynamic. Each point represents a single frame of a sequence. Different colors represent different video sequences. (a) Features learned using TCCL do not exhibit any cyclicality. Features from different video sequences are also not well separated. (b) Features learned using CSS are clustered in circular formations, demonstrating cyclical similarity. Features from different sequences are well separated.

to the LV. Overall, our method is better designed for processing echocardiogram videos and successfully reduces the number of labels required for training.

C. Comparing CSS with Unsupervised Temporal Alignment Methods

CSS makes use of cyclical repetition within echocardiograms as a prior for semi-supervised segmentation. We compare with unsupervised feature alignment methods that *do not* enforce cyclical consistency to highlight its effectiveness for LVEF prediction. One such approach proposed by Dwibedi *et al.* [53] is TCCL, which performs unsupervised temporal alignment by enforcing feature consistency between the same actions within different videos. Dezaki *et al.* [54] also use a similar method to synchronize echocardiograms of different views. Unlike CSS, these methods enforce feature consistency *between different sequences* for temporal alignment instead of *within the same sequence* for ensuring temporal cyclicality.

We replace CSS with TCCL within our semi-supervised framework and compare results for LVEF prediction. Instead of sampling template region P and search region Q from the same sequence, we sample Q from a separate sequence to perform phase matching and consistency enforcement. This allows the model to learn representations for key phases *across sequences*. For fairness, we use the same hyper-parameters as our method except for temporal offset c, which we do not need for TCCL. We use one-eighth of available labels and treat

TABLE IV

Ablation study using 3/32, 4/32, 6/32, 8/32, 16/32, and all training labels on EchoNet-Dynamic. "Seg." refers to the use of segmentation prediction masks as additional input. "CSS" refers to the use of CSS loss for semi-supervised segmentation. "Dist." refers to the use of teacher-student distillation. Methods that do not use "Dist." report results from f_m . We report mean results \pm standard deviation of five separate runs.

MAE Values ↓									
Method	Seg.	CSS	Dist.	3/32 labels	4/32 labels	6/32 labels	8/32 labels	16/32 labels	All labels
Video Only				5.82 ± 0.09	5.64 ± 0.08	5.22 ± 0.07	4.80 ± 0.05	4.35 ± 0.05	4.10 ± 0.04
Video+Seg.	\checkmark			5.61 ± 0.06	5.42 ± 0.07	5.03 ± 0.05	4.70 ± 0.05	4.29 ± 0.04	4.13 ± 0.04
Video+Seg.+CSS	\checkmark	\checkmark		5.47 ± 0.05	5.13 ± 0.05	4.87 ± 0.03	4.67 ± 0.04	4.29 ± 0.03	4.09 ± 0.03
Video+Seg.+Dist.	\checkmark		\checkmark	5.41 ± 0.07	5.20 ± 0.06	4.72 ± 0.05	4.41 ± 0.04	4.17 ± 0.03	-
Video+Seg.+CSS+Dist. (Ours)	\checkmark	\checkmark	\checkmark	$\textbf{5.18} \pm \textbf{0.06}$	$\textbf{4.90} \pm \textbf{0.04}$	$\textbf{4.55} \pm \textbf{0.04}$	$\textbf{4.34} \pm \textbf{0.03}$	$\textbf{4.14} \pm \textbf{0.03}$	-
					\mathbf{R}^2 Values \uparrow				
Method	Seg.	CSS	Dist.	3/32 labels	4/32 labels	6/32 labels	8/32 labels	16/32 labels	All labels
Video Only				$57.2\% \pm 0.9$	$60.6\% \pm 0.8$	$66.9\% \pm 0.6$	$71.2\% \pm 0.6$	$76.7\% \pm 0.4$	$80.5\% \pm 0.2$
Video+Seg.	\checkmark			$60.4\% \pm 0.7$	$63.0\% \pm 0.7$	$68.7\% \pm 0.6$	$73.5\% \pm 0.3$	$78.3\% \pm 0.4$	$80.3\% \pm 0.3$
Video+Seg.+CSS	\checkmark	\checkmark		$62.6\% \pm 0.6$	$67.6\% \pm 0.5$	$71.8\% \pm 0.3$	$74.0\% \pm 0.3$	$78.2\% \pm 0.3$	$80.4\% \pm 0.3$
Video+Seg.+Dist.	\checkmark		\checkmark	$62.8\% \pm 0.8$	$65.2\% \pm 0.5$	$72.3\% \pm 0.6$	$76.4\% \pm 0.4$	$79.8\% \pm 0.4$	-
Video+Seg.+CSS+Dist. (Ours)	\checkmark	\checkmark	\checkmark	$66.3\% \pm 0.6$	$71.1\% \pm 0.4$	$74.6\%~\pm~0.4$	$77.0\% \pm 0.4$	$80.1\%~\pm~0.3$	-



Fig. 6: Plot of Table IV showing LVEF prediction results using different proposed components. Training with CSS loss consistently leads to better results compared to without CSS loss (solid lines *vs.* dashed lines). Training with teacher-student distillation also leads to consistent improvements across different semi-supervised settings (blue lines *vs.* red lines).

remaining samples as unlabeled data. Results for multi-input model f_m and end-to-end model f_e are shown in Table III.

We can see that the MAE of our method is 0.19 lower than TCCL (4.90 vs. 5.09, p-Value < 0.05), demonstrating the effectiveness of cyclical consistency. We compare t-SNE plots of segmentation features to further interpret differences between the two methods, which we show in Fig. 5. Each point of the scatter plot represents the feature embedding of a single frame within a sequence. Points with the same colors belong to the same sequence. When using CSS, features from the same video sequence are clustered in circular formations, which is consistent with cyclical repetition. Features from different videos are clearly separated, which also helps with LV segmentation since different sequences have largely different visual appearances. In contrast, features from different sequences are close together when using TCCL, since key phases are encouraged to have similar features for all sequences. There are also few signs of cyclicality, which means cyclicality of cardiac motion is not adequately reflected on a feature level. Features learned using CSS are therefore more consistent with our task of semi-supervised segmentation based on cyclical consistency.

D. Ablation Study on Proposed Components

We analyze the performance contribution from the different proposed components of our method, namely the use of LV segmentations (Seg.), CSS (CSS), and teacher-student distillation (Dist.), by separately combining these modules and comparing their results. This is done using 3/32, 4/32, 6/32, 8/32, and 16/32 of available labels to comprehensively show results under different semi-supervised settings. We also conduct experiments with a fully labeled dataset for methods that do not require unlabeled data for distillation. We briefly detail the combination of modules used in each of the experiments below. We show results in Table IV and visually plot them in Fig. 6.

Video Only: We perform naive regression on raw video input by directly regressing the available LVEF labels. This is the same methodology used by Ouyang *et al.* [6].

Video+Seg.: We use segmentation masks as additional input for multi-input regression but do not use the proposed CSS and teacher-student distillation methods. A supervised segmentation model is trained using only labeled ED and ES frames and predictions are concatenated with video for input. Only the multi-input model f_m is used.

Video+Seg.+Dist .: We perform the same procedure as the

"Video + Seg." methodology but use additional teacher-student distillation by introducing f_e and unlabeled sequences.

Video+Seg.+CSS.: We perform the same procedure as the "Video+Seg." methodology but use CSS loss when training the segmentation model. Teacher-student distillation is not used.

We can see from Fig. 6 that training using only video inputs gives the worst performance (black line). Applying CSS loss boosts performance consistently across different settings (solid lines *vs.* dashed lines) and significantly decreases MAE (p-Value < 0.05 for 4/32 and 6/32 of labels, p-Value < 0.10 for 3/32 of labels). Teacher-student distillation further improves predictions (blue lines *vs.* red lines), especially when used together with CSS (p-Value < 0.05 for MAE reduction). Consistency across different settings provides further evidence of statistical significance. Overall, each proposed component has important contributions to performance.

E. Comparing CSS with Existing Semi-Supervised Segmentation Methods

1) Evaluating performance using LVEF video regression: We can evaluate the quality of segmentation predictions by concatenating them with video for LVEF video regression. Higher quality, temporally consistent segmentations provide better spatial guidance to the regression network, which leads to more accurate predictions. This evaluation method is also consistent with the clinical objective of obtaining accurate LVEF measurements through automated methods. We demonstrate the effectiveness of CSS by comparing with segmentation approaches using CLAS [10], multi-frame attention [55], and CPS [25], which represent state-of-the-art methods for semi-supervised image and video LV segmentation.

The trained models are used to generate LV segmentation mask predictions, which are then concatenated with raw video as input for regression. We use one-eighth of available training labels and treat remaining samples as unlabeled data. A supervised segmentation model trained on labeled data only is also included for comparison. We briefly describe implementation details below.

CLAS [10]: Wei *et al.* simultaneously learn motion flow and LV segmentation from video inputs by enforcing consistency between the two tasks using motion propagation [10]. We follow the CLAS implementation by Chen *et al.* [17] which extends the original implementation to EchoNet-Dynamic.

Multi-frame attention [55]: Ahn *et al.* concatenate feature embeddings from frames surrounding target ED and ES frames to provide additional context for segmentation [55]. We take the three frames before and after the target ED and ES frames as input for the forward and backward directions respectively. We use the ResNet-50 encoder [50] for f_{enc} and ASPP decoder [49] for f_{dec} as our segmentation model and use batch sizes of 20 ED and ES frames for fair comparison. The forward and backward models are trained separately and the average pixel-wise probability is used for prediction.

CPS [25]: Chen *et al.* propose two co-trained segmentation models that generate pseudo-labels for each other [25]. We use the DeepLabV3 architecture [49] and jointly train batches of 20 labeled ED and 20 labeled ES frames together with

TABLE V

LVEF prediction results on EchoNet-Dynamic using different state-ofthe-art methods to generate segmentation masks for input. One-eighth of labels are used as the labeled dataset, and remaining samples are used as unlabeled data. Results are shown for multi-input model f_m and end-to-end model f_e . We report mean results \pm standard deviation of five separate runs.

Method	j	f m	.	f_e
Wiethou	MAE↓	$\mathbf{R}^{2}\uparrow$	MAE↓	$\mathbf{R}^{2}\uparrow$
CLAS [10]	5.72 ± 0.07	$59.3\% \pm 0.9$	5.57 ± 0.07	$61.5\% \pm 0.8$
Supervised	5.42 ± 0.07	$63.0\% \pm 0.7$	5.20 ± 0.06	$65.2\% \pm 0.5$
Multi-frame [55]	5.42 ± 0.06	$63.3\% \pm 0.8$	5.19 ± 0.07	$66.8\% \pm 0.7$
CPS [25]	5.29 ± 0.06	$65.1\% \pm 0.7$	5.01 ± 0.04	$68.9\% \pm 0.5$
Ours	5.13 ± 0.05	$67.6\%~\pm~0.5$	$4.90~\pm~0.04$	$71.1\%~\pm~0.4$

TABLE VI

Dice score for LV segmentation predictions on EchoNet-Dynamic. We use only one-eighth of available labels and treat remaining samples as unlabeled data. Results are shown for labeled ED and ES frames and randomly sampled unlabeled frames. We report mean results \pm standard deviation of five separate runs.

Method	Dice↑				
Wiethou	ED Frame	ES Frame	Unlabeled Frame		
CLAS [10]	$87.2\% \pm 0.3$	$90.1\% \pm 0.3$	$90.2\% \pm 0.3$		
Supervised	$88.8\% \pm 0.3$	$91.4\% \pm 0.2$	$90.8\% \pm 0.2$		
Multi-frame [55]	$88.5\% \pm 0.3$	$91.5\% \pm 0.3$	$91.0\% \pm 0.3$		
CPS [25]	$89.2\% \pm 0.2$	$91.9\% \pm 0.2$	$91.3\% \pm 0.2$		
Ours	$89.0\% \pm 0.2$	$92.2\%~\pm~0.2$	$91.9\%~\pm~0.2$		

40 unlabeled frames for every iteration. This uses the same number of labeled and unlabeled frames per iteration as CSS for fairness.

Supervised: We train a supervised LV segmentation model using available labels for ED and ES frames. We use DeepLabV3 and a batch size of 20 ED and ES frames for fair comparison.

We show in Table V results from multi-input model f_m and distilled model f_e . We can see predictions using CSS (Ours) give the lowest MAE, achieving 5.13 and 4.90 for f_m and f_e respectively. Improved performance on the multi-input model generally leads to more accurate pseudo-labels, which consequently results in a better distilled end-to-end model.

2) Evaluating performance using Dice score: We can also evaluate the performance of CSS by comparing segmentation prediction results directly. This is challenging however since segmentation labels are only available for reference ED and ES frames, whereas semi-supervised approaches also aim to improve predictions for unlabeled non-ED and non-ES frames [10, 41]. To address this, we also assess performance on unlabeled frames by randomly sampling a single frame from each of the 1,277 test sequences and manually segmenting the LV. Our segmentation masks were reviewed by a cardiologist to ensure correctness and used as the label for evaluating different segmentation methods. We calculate Dice over ED, ES, and unlabeled frames and present results in Table VI. We also show qualitative samples in Fig. A3 of the Appendix.

We can see from the results that our method, CSS, achieves the highest Dice score for ES frames (92.2%) and unlabeled frames (91.9%). CSS outperforms alternatives by a larger margin on unlabeled frames because ED and ES frames are trained with labels by all methods, which leads to similar predictions for these phases. Results are more differentiated on

TABLE VII

LVEF prediction results using different values for w_{css} on EchoNet-Dynamic. One-eighth of labels are used as the labeled dataset, and remaining samples are used as unlabeled data. Results are shown for multi-input model f_m and end-to-end model f_e . We report mean results \pm standard deviation of five separate runs.

Method	f	m	f_e		
Wiethou	MAE↓	$\mathbf{R}^{2}\uparrow$	MAE↓	$\mathbf{R}^{2}\uparrow$	
$w_{css} = 0.02$	5.24 ± 0.07	$66.4\% \pm 0.6$	4.99 ± 0.06	$70.1\% \pm 0.6$	
$w_{css} = 0.01$	5.13 ± 0.05	$67.6\%~\pm~0.5$	4.90 ± 0.04	$71.1\% \pm 0.4$	
$w_{css} = 0.005$	5.19 ± 0.06	$66.9\% \pm 0.5$	4.96 ± 0.05	$70.3\% \pm 0.5$	

TABLE VIII

LVEF prediction results using different values for w_{ulb} on EchoNet-Dynamic. One-eighth of labels are used as the labeled dataset, and remaining samples are used as unlabeled data. We report mean results \pm standard deviation of five separate runs.

Method	MAE↓	$\mathbf{R}^{2}\uparrow$
$w_{ulb} = 10$	5.13 ± 0.04	$68.2\% \pm 0.4$
$w_{ulb} = 5$	4.90 ± 0.04	$71.1\%~\pm~0.4$
$w_{ulb} = 2$	4.96 ± 0.05	$70.1\% \pm 0.4$

unlabeled frames because different semi-supervision strategies are used. Results on randomly sampled unlabeled frames are likely to be more indicative of performance on full sequences *since most frames are non-ED and non-ES frames*. We also note that the relative performance of different methods is similar to results obtained through LVEF regression. This provides further support that CSS outperforms state-of-the-art alternatives for semi-supervised LV segmentation.

F. Choice of Parameters for CSS

The parameters used for training CSS include the number of feature embeddings in a phase (s), the temporal offset (c), and the softmax temperature (τ). We test the sensitivity of our method to parameter choice by varying parameter values one at a time whilst keeping others constant at our chosen settings s = 3, c = 2, and $\tau = 10$. We perform experiments using one-eighth of available training labels as the labeled dataset and treat the remaining as unlabeled data. Results using values $s \in \{2, 3, 4, 5\}$, $c \in \{1, 2, 3, 4, 5\}$, and $\tau \in \{1, 5, 10, 50, 100\}$ are visually plotted in Fig. 7.

We can see from Fig. 7a that results are generally stable between 4.90 and 4.96 MAE for s >= 3. Results using s =2 are slightly weaker as fewer feature embeddings are used for phase identification, which may lead to incorrect phase matching. Fig. 7b shows that results for different values of care stable between 4.90 and 4.96 MAE. Fig. 7c shows that $\tau = 10$ is the most suitable value for temperature.

We also perform a coarse search for parameters $w_{css} \in \{0.005, 0.01, 0.02\}$ and $w_{ulb} \in \{2, 5, 10\}$ using one-eighth of available training labels to determine stable training weights. For w_{css} , we show results from the multi-input model f_m and the distilled model f_e in Table VII. We can see using $w_{ulb} = 0.01$ gives us the best results, although performance is relatively robust to different values. Results for w_{ulb} are shown in Table VIII, where we can see that using $w_{ulb} = 5$ gives us the best results.



Fig. 7: MAE for LVEF prediction on EchoNet-Dynamic using different values of s, c, and τ for training CSS. (a) s is the number of feature embeddings used to identify a phase for temporal matching. (b) c is the temporal offset used to apply cyclical matching. (c) τ is the softmax temperature parameter controlling the sharpness of the matching probability profile. We use one-eighth of labeled data as the labeled dataset and treat remaining samples as unlabeled data.

G. External Validation of Trained Models

Model generalization is desirable for medical applications as different hospitals may have different equipment and data collection procedures, leading to domain shifts. We perform external validation of the models trained in Section IV-B using the four-chamber sequences from the CAMUS dataset [14] to test for generalized performance. CAMUS sequences have significantly fewer frames on average compared to EchoNet-Dynamic (20 vs. 175) and only capture part of the cardiac cycle. In contrast, sequences for EchoNet-Dynamic cover multiple cycles and to the best of our knowledge is the only publicly available dataset with multi-cycle data. For preprocessing, we double the input sequence by mirroring the data along the temporal dimension at the start and end of the sequence to simulate a full cardiac cycle (see Appendix D for details). Temporal mirroring helps preserve smoothness since the sequence is simply reversed and no discontinuity is introduced. We simulate a full cardiac cycle for testing since this is more consistent with echocardiograms obtained through realistic acquisition procedures [56]. Frames are resized to

12

TABLE IX

External validation of trained models on the CAMUS dataset. MAE is reported based on video quality and overall performance. Models are trained using 1/8, 1/4, 1/2, and all of the available training labels on the EchoNet-Dynamic dataset and tested on CAMUS. Note that "Supervised" methods use only labeled data while "Semi-supervised" methods use both labeled and unlabeled data. We show mean results ± standard deviation of the five training runs.

			MAE Values			
Quality	Туре	Method	1/8 labels	1/4 labels	1/2 labels	All labels
	Sumanniand	Ouyang et al. [6]	7.27 ± 0.19	6.43 ± 0.14	6.37 ± 0.11	6.21 ± 0.08
	Supervised	Dai et al. [12]	7.24 ± 0.18	6.40 ± 0.13	6.31 ± 0.10	5.90 ± 0.07
Good	Sami	Ji et al. [46]	7.87 ± 0.18	6.49 ± 0.13	6.10 ± 0.10	-
	Supervised	Xu et al. [44]	6.63 ± 0.12	6.39 ± 0.13	5.95 ± 0.12	-
	Supervised	Ours	6.51 ± 0.12	$\textbf{6.01} \pm \textbf{0.11}$	$\textbf{5.67} \pm \textbf{0.10}$	-
	Supervised	Ouyang et al. [6]	8.61 ± 0.24	7.96 ± 0.14	8.53 ± 0.15	6.89 ± 0.13
	Supervised	Dai et al. [12]	8.85 ± 0.22	7.49 ± 0.16	7.77 ± 0.15	6.72 ± 0.10
Medium	Sami	Ji et al. [46]	8.54 ± 0.18	7.38 ± 0.17	7.34 ± 0.12	-
	Supervised	Xu et al. [44]	8.41 ± 0.14	7.52 ± 0.14	7.19 ± 0.15	-
		Ours	$\textbf{8.33} \pm \textbf{0.13}$	$\textbf{7.20}\pm\textbf{0.10}$	$\textbf{6.34} \pm \textbf{0.15}$	-
Supervised	Ouyang et al. [6]	11.00 ± 0.26	9.43 ± 0.19	10.64 ± 0.15	10.33 ± 0.15	
	Supervised	Dai et al. [12]	11.31 ± 0.26	9.53 ± 0.20	10.88 ± 0.14	10.82 ± 0.11
Poor	Poor Sami	Ji et al. [46]	9.71 ± 0.25	8.95 ± 0.22	8.51 ± 0.14	-
	Supervised	Xu et al. [44]	10.36 ± 0.18	9.32 ± 0.16	9.51 ± 0.14	-
Supervised	Ours	10.77 ± 0.13	$\textbf{8.78} \pm \textbf{0.15}$	$\textbf{8.00} \pm \textbf{0.16}$	-	
	0 1	Ouyang et al. [6]	8.06 ± 0.22	7.21 ± 0.17	7.48 ± 0.14	6.82 ± 0.12
Supervised	Dai et al. [12]	8.15 ± 0.18	7.05 ± 0.16	7.22 ± 0.13	6.63 ± 0.06	
Overall Semi	Ji et al. [46]	8.26 ± 0.20	7.01 ± 0.14	6.73 ± 0.13	-	
	Supervised	Xu et al. [44]	7.57 ± 0.15	7.04 ± 0.16	6.69 ± 0.11	-
Supervised	Ours	$\textbf{7.51} \pm \textbf{0.14}$	$\textbf{6.66} \pm \textbf{0.09}$	$\textbf{6.11} \pm \textbf{0.12}$	-	



Fig. 8: Qualitative samples of attention heatmaps on CAMUS. Models are trained using EchoNet-Dynamic under different methods and semi-supervised settings. Models trained using our method on half of available labels attend better to the LV region than Dai *et al.* [12] trained on full labels.

 112×112 pixels to fit the trained model. Inference results are shown in Table IX.

We can see our proposed method consistently outperforms alternatives for good and medium-quality sequences and has the best overall performance. Our method trained using half the available labels on EchoNet-Dynamic also outperforms the supervised method by Dai *et al.* [12] trained on full labels, demonstrating superior generalization (MAE 6.11 *vs.* 6.63 with p-Value < 0.05). To understand why this is the case, we compare attention heatmaps to visualize the regions that are important to the models for inference. We use SmoothGrad [52] to calculate pixel gradient contributions, where higher absolute pixel gradients indicate greater sensitivity to changes in that region. For a well generalized model, we expect pixels in the LV to have higher absolute gradients since LVEF is

TABLE X

Dice score between LV segmentation labels and pixels with the top 5% highest absolute gradient values. Results are shown for the CA-MUS dataset with models trained on EchoNet-Dynamic using different methods and different semi-supervised settings. Pixel gradients are calculated using SmoothGrad [52].

Method	Setting	Dice↑ ED	Dice↑ ES	Dice↑ Overall
Dai et al. [12]	All labels	42.9%	47.6%	45.4%
Ours	1/2 labels	46.7%	57.1%	52.3%

primarily determined by changes in this region. We confirm that our method generates heatmaps that are more concentrated in the LV region compared to Dai *et al.* [12] and show qualitative examples in Fig. 8.

To quantify attention to the LV region, we calculate the degree of overlap between segmentation labels and the most important pixels. We calculate the Dice score between LV mask labels and pixels with the top 5% highest absolute gradient values and show results in Table X. We can see our method, even when trained using half the available labels, has a higher Dice score compared to Dai *et al.* [12] trained on a full dataset. This shows that our model can identify important regions more effectively compared to supervised approaches, even when tested on external datasets. Existing works have also shown that training with unlabeled data leads to better model generalization since it reduces over-fitting to the labeled dataset [57, 58]. This further supports the use of our proposed semi-supervised method for LVEF regression.

V. DISCUSSION

A. Impact of Heart Arrhythmia on CSS

Heart arrhythmia refers to when the heart either beats too fast, too slow, or at an irregular pace with altering rates [59]. Our CSS method was designed with heart arrhythmia in

consideration since cases had previously been identified in EchoNet-Dynamic by other studies [6]. CSS does not assume consistent periodic cycles across the entire video sequence and only enforces consistency between adjacent cycles of sampled clips. Because of this, it is relatively robust to irregular heart rates. Although a more detailed study on cardiac arrhythmia cases will be informative, this would require expert labels which are currently unavailable. We leave this to future investigation and note that current experimental results validate the effectiveness of our method even with such cases present in the dataset.

B. Choice of temporal offset

We use a constant temporal offset of c = 2 in our method, which means we apply an offset of two frames from the reference and matching time-points to check for cyclical consistency. Similar results were also achieved using alternative values. Another possible choice is applying an offset value that is relative to cycle length, since different patients have different cycle durations. Doing so will introduce additional complexities however, since labels for cycle lengths are unavailable and can only be inferred indirectly from feature matching probabilities. Furthermore, unlike a constant temporal offset, which can be applied directly to all samples, a relative offset may lead to fractional values and may require temporal interpolation between frames for implementation. Given that our method already outperforms alternative approaches and is the first to make use of unsupervised cyclical consistency, we do not explore relative offsets at this time.

VI. CONCLUSION

In this work, we introduce the first semi-supervised approach to LVEF prediction from echocardiogram videos. We address the problem in two steps by proposing a novel cyclical self-supervision (CSS) method for semi-supervised video LV segmentation, and a teacher-student distillation method for distilling spatial context from segmentation predictions into an end-to-end video regression model. Our method allows echocardiogram videos to be trained using fewer labels, which is important for reducing labeling costs. Experiments show our method outperforms existing semi-supervised alternatives and achieves performance competitive with supervised methods on fully labeled datasets. Furthermore, external validation demonstrates that our method has better generalization when tested on data from a different hospital.

APPENDIX A

CYCLICAL FEATURE MATCHING WITH MULTIPLE CARDIAC CYCLES IN THE SEARCH REGION

One of the challenges in enforcing cyclical feature consistency is that the number of frames per cycle within a video sequence can vary. This means it is impossible to guarantee the number of cycles present in input clip X_i . The illustrations and examples in Fig. 2 and Fig. 3 are based on two cycles being included in the input clip. We show that CSS can be used regardless of the number of cycles present and extend



(a) Cyclical relationship within a sequence for multiple cycles Input Clip X_i



(b) Phase matching in the search region with two cycles present Input Clip X_i



(c) Offset matching in the template region

Fig. A1: (a) We can enforce the same relationship for cyclicality even if three cycles are in the input sequence: if the phase at time-point p^* is the same as time-point q_1^* and q_2^* for a cyclical sequence, then the phase at time-point $q_1^* + c$ and $q_2^* + c$ is the same as time-point $p^* + c$ (b) Using template phase $E_i^{p^*}$, we calculate the matching probability, α_i^q , in the search region for $q \in Q$. For well trained features, the time-points corresponding to true matches will have higher probabilities. (c) $\tilde{E}_i^{q^*+c}$ can still be calculated based on matching probabilities even with multiple cycles. We weigh embeddings by α_i^q and find the closest matching time-point in P, which should be $p^* + c$. The same loss function can still be used.

our examples to three cycles in Fig. A1, showing that the same CSS loss function can still be used.

If there are three cardiac cycles in the input clip and two matching phases in the search region, we can still calculate matching probability at different time-points in Q based on Equation 2. Given that the features are jointly trained with supervised segmentation, the matching probabilities should be higher at the true matching time-points compared with other regions for well trained features. We use the same probability weightings to calculate expected value $\tilde{E}_i^{q^*+c}$ as

per Equation 3. We illustrate this in Fig. A1b.

We can also perform phase matching back in the template region by using Equation 6 regardless of how many cycles are present in the input clip. This is because $\tilde{E}_i^{q^*+c}$ should still match with the phase at p^*+c even if multiple cycles are used to calculate the expectation. Thus, we can use the same CSS loss formulation even for multiple cycles.

One important detail to note is that having multiple phases in the template region P can lead to ambiguous matching for Equation 6. Thus, we use a shorter interval for P compared to Q (15 frames vs. 21) to avoid such cases during implementation. This does not affect other components of our method since we only need P for template sampling and matching, which does not require a full cycle.

TABLE A1

LVEF prediction results on EchoNet-Dynamic using different backbone architectures. We use one-eighth of available labels as the labeled dataset and treat remaining samples as unlabeled data. We report mean results \pm standard deviation of five separate runs.

Туре	Method	Backbone	MAE \downarrow	\mathbf{R}^2 Values \uparrow
Supervised	Ouyang et al. [6]	MC3	6.04 ± 0.06	$58.6\% \pm 0.5$
Supervised	Dai <i>et al</i> . [12]	MC3	5.97 ± 0.05	$60.1\% \pm 0.5$
Sami	Ji et al. [46]	MC3	5.96 ± 0.05	$60.4\% \pm 0.5$
Supervised	Xu et al. [44]	MC3	5.95 ± 0.04	$60.6\% \pm 0.4$
Supervised	Ours	MC3	5.85 ± 0.04	$61.6\%~\pm~0.4$
	Ouyang et al. [6]	R3D	5.90 ± 0.05	$60.0\% \pm 0.5$
Supervised	Dai et al. [12]	R3D	5.84 ± 0.04	$61.1\% \pm 0.5$
Sami	Ji et al. [46]	R3D	5.75 ± 0.05	$61.4\% \pm 0.5$
Supervised	Xu et al. [44]	R3D	5.69 ± 0.04	$63.5\% \pm 0.4$
Supervised	Ours	R3D	5.50 ± 0.04	$66.5\%~\pm~0.4$
Supervised	Ouyang et al. [6]	R2+1D	5.64 ± 0.08	$60.6\% \pm 0.8$
Supervised	Dai et al. [12]	R2+1D	5.47 ± 0.07	$62.8\% \pm 0.6$
Sami	Ji et al. [46]	R2+1D	5.61 ± 0.07	$61.6\% \pm 0.6$
Semi-	Xu et al. [44]	R2+1D	5.08 ± 0.04	$66.8\% \pm 0.4$
Supervised	Ours	R2+1D	$ 4.90 \pm 0.04 $	$71.1\%~\pm~0.4$

APPENDIX B CHOICE OF INPUT LENGTH FOR X_i

For X_i , we use an input clip 40 frames in length, sampled at a rate of 1 in every 3 frames. We show that this length covers at least two cardiac cycles for most sequences in EchoNet-Dynamic.

We note that the human pulse rate is typically between 50 beats-per-minute (bpm) and 90 bpm [60], and that 79% of all video sequences in EchoNet-Dynamic are captured at 50 frames-per-second (fps) [5]. Therefore, the upper bound of frames required for two full cycles is approximately:

$$\frac{60 \text{ seconds}}{50 \text{bpm}} \times 2 \text{ beats} \times 50 \text{ fps} = 120 \text{ frames}.$$
(9)

Sampling 40 frames at a rate of 1 in every 3 frames covers $40 \times 3-2 = 118$ frames from the original sequence, which is most of the two full cardiac cycles. Furthermore, using 40 frames as the input clip length means that the CSS and supervised segmentation task are balanced in the number of frames used. This setting also allows for easier comparison with other semi-supervised methods that use 40 unlabeled frames per iteration.



Fig. A2: Heatmap plot of normalized L2 distance between frames at different time-steps. Diagonal patterns of low L2 distance indicate presence of cyclical repetition since similar frames are repeated at regular intervals.

APPENDIX C CHOICE OF BACKBONE ARCHITECTURE

Ouyang *et al.* compared 3D (R3D), Mixed Convolution (MC3), and R2+1D [19] ResNet architectures for LVEF regression and found that the R2+1D ResNet performed the best [5, 6]. The R2+1D backbone was also used as the baseline architecture in works by Dai *et al.* [12]. We repeat the experiments using one-eighth of available labels in Section IV-B and compare results for different backbone gives the best performance for each method, and therefore use this architecture for experiments. We also see our proposed method (Ours) outperforms alternative state-of-the-art approaches regardless of the backbone used.

TABLE A2 Summary statistics for EchoNet-Dynamic dataset.

Metric	Value (Standard Deviation)
Total number of patients	10,030
Average number of frames	175 (57)
Average frames per second	51.1 (6.2)
Mean LVEF %	55.75 (12.37)

TABLE A3 Summary statistics for CAMUS dataset.

Metric	Value (Standard Deviation)
Total number of patients	450
Average number of frames	20 (4)
Video quality	Good: 260, Medium: 148, Poor: 42
Mean LVEF %	52.03 (12.07)

APPENDIX D Temporal Mirroring for CAMUS sequences

Temporal mirroring is done by reversing the sequences at the starting and ending time-points. For video video sequence V consisting for frames $[v_1, v_2, ..., v_{T-1}, v_T]$, the doubled sequence V' is equal to:

$$V' = [v_{T/2}, v_{T/2-1}, \dots, v_2, v_1, v_2, \dots, v_{T-1}, v_T, v_{T-1}, \dots, v_{T/2+1}, v_{T/2}].$$
(10)

This article has been accepted for publication in IEEE Transactions on Medical Imaging. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TMI.2022.3229136

IEEE Transactions on Medical Imaging, VOL. XX, NO. XX, XXXX 2022



Fig. A3: Qualitative samples of LV mask predictions (red) on unlabeled non-ED and non-ES frames for different methods. Performance is evaluated by comparing with manual LV segmentations (purple). We show ground truth labels on ED (blue) and ES (green) frames for reference. We note that alternative methods tend to have predictions distinctly outside of the true LV region for difficult samples.

REFERENCES

- D. Bamira and M. Picard, "Imaging: Echocardiology—assessment of cardiac structure and function," *Encyclopedia of Cardiovascular Research and Medicine*, pp. 35–54, 2018.
- [2] M. T. Maeder and D. M. Kaye, "Heart failure with normal left ventricular ejection fraction," *Journal of the American College* of Cardiology, vol. 53, no. 11, pp. 905–918, 2009.
- [3] A. H. Abdi et al., "Automatic quality assessment of echocardiograms using convolutional neural networks: feasibility on the apical four-chamber view," *IEEE transactions on medical imaging*, vol. 36, no. 6, pp. 1221–1230, 2017.
- [4] J. W. Hughes *et al.*, "Deep learning evaluation of biomarkers from echocardiogram videos," *EBioMedicine*, vol. 73, p. 103613, 2021.
- [5] D. Ouyang *et al.*, "Echonet-dynamic: a large new cardiac motion video data resource for medical machine learning," in *NeurIPS*, 2019.
- [6] D. Ouyang *et al.*, "Video-based ai for beat-to-beat assessment of cardiac function," *Nature*, vol. 580, no. 7802, pp. 252–256, 2020.
- [7] 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," *The international journal of cardiovascular imaging*, vol. 31, no. 7, pp. 1303–1314, 2015.
- [8] M. H. Jafari, N. Van Woudenberg, C. Luong, P. Abolmaesumi, and T. Tsang, "Deep bayesian image segmentation for a more robust ejection fraction estimation," in *ISBI*. IEEE, 2021, pp. 1264–1268.
- [9] H. Reynaud, A. Vlontzos, B. Hou, A. Beqiri, P. Leeson, and B. Kainz, "Ultrasound video transformers for cardiac ejection fraction estimation," *arXiv preprint arXiv:2107.00977*, 2021.
- [10] H. Wei *et al.*, "Temporal-consistent segmentation of echocardiography with co-learning from appearance and shape," in *MICCAI*. Springer, 2020, pp. 623–632.
- [11] N. Yuan *et al.*, "Systematic quantification of sources of variation in ejection fraction calculation using deep learning," *Cardiovascular Imaging*, vol. 14, no. 11, pp. 2260–2262, 2021.

- [12] W. Dai, X. Li, W. H. K. Chiu, M. D. Kuo, and K.-T. Cheng, "Adaptive contrast for image regression in computer-aided disease assessment," *IEEE Transactions on Medical Imaging*, 2021.
- [13] A. Ghorbani *et al.*, "Deep learning interpretation of echocardiograms," *NPJ digital medicine*, vol. 3, no. 1, pp. 1–10, 2020.
- [14] S. Leclerc *et al.*, "Deep learning for segmentation using an open large-scale dataset in 2d echocardiography," *IEEE transactions* on medical imaging, vol. 38, no. 9, pp. 2198–2210, 2019.
- [15] F. Liu, K. Wang, D. Liu, X. Yang, and J. Tian, "Deep pyramid local attention neural network for cardiac structure segmentation in two-dimensional echocardiography," *Medical Image Analysis*, vol. 67, p. 101873, 2021.
- [16] N. Painchaud, N. Duchateau, O. Bernard, and P.-M. Jodoin, "Echocardiography segmentation with enforced temporal consistency," arXiv preprint arXiv:2112.02102, 2021.
- [17] Y. Chen, X. Zhang, C. M. Haggerty, and J. V. Stough, "Fully automated multi-heartbeat echocardiography video segmentation and motion tracking," in *Medical Imaging 2022: Image Processing*, vol. 12032. SPIE, 2022, pp. 185–192.
- [18] F. T. Dezaki *et al.*, "Cardiac phase detection in echocardiograms with densely gated recurrent neural networks and global extrema loss," *IEEE transactions on medical imaging*, vol. 38, no. 8, pp. 1821–1832, 2018.
- [19] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *CVPR*, 2018, pp. 6450–6459.
- [20] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *arXiv preprint arXiv:1703.01780*, 2017.
- [21] X. Li, L. Yu, H. Chen, C.-W. Fu, and P.-A. Heng, "Semisupervised skin lesion segmentation via transformation consistent self-ensembling model," in *BMVC*, 2018.
- [22] L. Yu, S. Wang, X. Li, C.-W. Fu, and P.-A. Heng, "Uncertaintyaware self-ensembling model for semi-supervised 3d left atrium segmentation," in *MICCAI*. Springer, 2019, pp. 605–613.
- [23] Y. Ouali, C. Hudelot, and M. Tami, "Semi-supervised semantic segmentation with cross-consistency training," in CVPR, 2020,

pp. 12674-12684.

- [24] X. Li, L. Yu, H. Chen, C.-W. Fu, L. Xing, and P.-A. Heng, "Transformation-consistent self-ensembling model for semisupervised medical image segmentation," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 2, pp. 523– 534, 2020.
- [25] X. Chen, Y. Yuan, G. Zeng, and J. Wang, "Semi-supervised semantic segmentation with cross pseudo supervision," in *CVPR*, 2021, pp. 2613–2622.
- [26] H. Yao, X. Hu, and X. Li, "Enhancing pseudo label quality for semi-supervised domain-generalized medical image segmentation," in AAAI, 2022.
- [27] Y. Lin, H. Yao, Z. Li, G. Zheng, and X. Li, "Calibrating label distribution for class-imbalanced barely-supervised knee segmentation," in *MICCAI*, 2022.
- [28] H. Wu, X. Li, and K.-T. Cheng, "Exploring feature representation learning for semi-supervised medical image segmentation," *arXiv preprint arXiv:2111.10989*, 2021.
- [29] I. Alonso, A. Sabater, D. Ferstl, L. Montesano, and A. C. Murillo, "Semi-supervised semantic segmentation with pixellevel contrastive learning from a class-wise memory bank," *arXiv preprint arXiv:2104.13415*, 2021.
- [30] S. Liu, S. Zhi, E. Johns, and A. J. Davison, "Bootstrapping semantic segmentation with regional contrast," *arXiv preprint* arXiv:2104.04465, 2021.
- [31] Y. Zou *et al.*, "Pseudoseg: Designing pseudo labels for semantic segmentation," *arXiv preprint arXiv:2010.09713*, 2020.
- [32] P. Yan et al., "Semi-supervised video salient object detection using pseudo-labels," in CVPR, 2019, pp. 7284–7293.
- [33] S. Varghese *et al.*, "Unsupervised temporal consistency metric for video segmentation in highly-automated driving," in *CVPR*, 2020, pp. 336–337.
- [34] M. Ding *et al.*, "Every frame counts: joint learning of video segmentation and optical flow," in *AAAI*, vol. 34, no. 07, 2020, pp. 10713–10720.
- [35] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung, "Learning video object segmentation from static images," in *CVPR*, 2017, pp. 2663–2672.
- [36] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim, "Video object segmentation using space-time memory networks," in *CVPR*, 2019, pp. 9226–9235.
- [37] H. K. Cheng, Y.-W. Tai, and C.-K. Tang, "Rethinking spacetime networks with improved memory coverage for efficient video object segmentation," *arXiv preprint arXiv:2106.05210*, 2021.
- [38] Y. Zhang *et al.*, "Semi-supervised cardiac image segmentation via label propagation and style transfer," in *STACOM*. Springer, 2020, pp. 219–227.
- [39] J. Pedrosa *et al.*, "Fast and fully automatic left ventricular segmentation and tracking in echocardiography using shapebased b-spline explicit active surfaces," *IEEE transactions on medical imaging*, vol. 36, no. 11, pp. 2287–2296, 2017.
- [40] K. Ta, S. S. Ahn, J. C. Stendahl, A. J. Sinusas, and J. S. Duncan, "A semi-supervised joint network for simultaneous left ventricular motion tracking and segmentation in 4d echocardiography," in *MICCAI*. Springer, 2020, pp. 468–477.
- [41] C. Qin *et al.*, "Joint learning of motion estimation and segmentation for cardiac mr image sequences," in *MICCAI*. Springer, 2018, pp. 472–480.
- [42] B. Xiong, H. Fan, K. Grauman, and C. Feichtenhofer, "Multiview pseudo-labeling for semi-supervised learning from video," *arXiv preprint arXiv:2104.00682*, 2021.
- [43] L. Jing, T. Parag, Z. Wu, Y. Tian, and H. Wang, "Videossl: Semi-supervised learning for video classification," in WACV, 2021, pp. 1110–1119.
- [44] Y. Xu *et al.*, "Cross-model pseudo-labeling for semi-supervised action recognition," *arXiv preprint arXiv:2112.09690*, 2021.
- [45] X. Ding, N. Wang, X. Gao, J. Li, X. Wang, and T. Liu, "Kfc: An efficient framework for semi-supervised temporal action

localization," *IEEE Transactions on Image Processing*, vol. 30, pp. 6869–6878, 2021.

- [46] J. Ji, K. Cao, and J. C. Niebles, "Learning temporal action proposals with fewer labels," in CVPR, 2019, pp. 7073–7082.
- [47] X. Hu, D. Zeng, X. Xu, and Y. Shi, "Semi-supervised contrastive learning for label-efficient medical image segmentation," in *International Conference on Medical Image Computing* and Computer-Assisted Intervention. Springer, 2021, pp. 481– 490.
- [48] W. Ouyang, H. Li, X. Zeng, and X. Wang, "Learning deep representation with large-scale attributes," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1895–1903.
- [49] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," arXiv preprint arXiv:1706.05587, 2017.
- [50] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [51] W. Kay et al., "The kinetics human action video dataset," arXiv preprint arXiv:1705.06950, 2017.
- [52] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, "Smoothgrad: removing noise by adding noise," *arXiv preprint arXiv:1706.03825*, 2017.
- [53] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman, "Temporal cycle-consistency learning," in *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 1801–1810.
- [54] F. T. Dezaki *et al.*, "Echo-syncnet: Self-supervised cardiac view synchronization in echocardiography," *IEEE Transactions on Medical Imaging*, vol. 40, no. 8, pp. 2092–2104, 2021.
- [55] S. S. Ahn, K. Ta, S. Thorn, J. Langdon, A. J. Sinusas, and J. S. Duncan, "Multi-frame attention network for left ventricle segmentation in 3d echocardiography," in *MICCAI*. Springer, 2021, pp. 348–357.
- [56] P. Leeson, D. Augustine, A. R. Mitchell, and H. Becher, *Echocardiography.* Oxford University Press, 2012.
- [57] Y. Zhong, H. Tang, J. Chen, J. Peng, and Y.-X. Wang, "Is self-supervised learning more robust than supervised learning?" arXiv preprint arXiv:2206.05259, 2022.
- [58] D. Hendrycks, M. Mazeika, S. Kadavath, and D. Song, "Using self-supervised learning can improve model robustness and uncertainty," *Advances in neural information processing systems*, vol. 32, 2019.
- [59] G. Tse, "Mechanisms of cardiac arrhythmias," *Journal of ar-rhythmia*, vol. 32, no. 2, pp. 75–81, 2016.
- [60] D. Nanchen, "Resting heart rate: what is normal?" pp. 1048– 1049, 2018.