



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

CoReEcho: Continuous Representation Learning for 2D+time Echocardiography Analysis

Fadillah Adamsyah Maani^{✉1}, Numan Saeed¹, Aleksandr Matsun¹, and Mohammad Yaqub¹

Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE
{firstname.lastname}@mbzuai.ac.ae

Abstract. Deep learning (DL) models have been advancing automatic medical image analysis on various modalities, including echocardiography, by offering a comprehensive end-to-end training pipeline. This approach enables DL models to regress ejection fraction (EF) directly from 2D+time echocardiograms, resulting in superior performance. However, the end-to-end training pipeline makes the learned representations less explainable. The representations may also fail to capture the continuous relation among echocardiogram clips, indicating the existence of spurious correlations, which can negatively affect the generalization. To mitigate this issue, we propose **CoReEcho**, a novel training framework emphasizing continuous representations tailored for direct EF regression. Our extensive experiments demonstrate that **CoReEcho**: 1) outperforms the current state-of-the-art (SOTA) on the largest echocardiography dataset (EchoNet-Dynamic) with MAE of 3.90 & R2 of 82.44, and 2) provides robust and generalizable features that transfer more effectively in related downstream tasks. The code is publicly available at <https://github.com/BioMedIA-MBZUAI/CoReEcho>.

Keywords: Echocardiography · Ejection Fraction · Direct Diagnosis · Continuous Representation · Model Transferability.

1 Introduction

Heart failure, a leading cause of mortality in clinical practice, affects 1-2% of individuals over 40 and 10% of those aged 60-70, representing the fastest-growing cardiac disease [2]. Echocardiography is the preferred option for efficiently and accurately diagnosing heart conditions. From an echocardiogram video, clinicians usually assess the heart's condition by measuring biomarkers such as the ejection fraction (EF) [23,11,19], which represents the percentage of blood that can be pumped out of the left ventricle (LV). Echocardiography also allows clinicians to analyze LV characteristics to detect early myocardial infarction (MI) or heart attack [10,4]. However, the natural characteristics of echocardiograms (e.g., speckle, operator dependence, and signal attenuation) make diagnosis challenging. Thus, recent initiatives exist to automate echocardiogram analysis through

the creation of properly labeled datasets, such as [11,19,4]. These publicly accessible datasets have been instrumental in advancing 2D+time echocardiography analysis, ultimately leading to a more accurate and faster diagnosis.

Early studies on EF estimation relied on the LV segmentation [11,14,12]. Ouyang et al. [19] proposed a two-stage approach in which they first segmented the LV frame-by-frame to find heartbeat cycle(s) and then used a deep learning (DL) model to regress the EF. The main limitation of these methods is the temporal inconsistency in the segmentation. These methods use 2D frames to segment the LV independently because most echocardiograms are sparsely labeled. To address this, several studies have proposed temporally-aware approaches [28,15] to enable video-based segmentation models for echocardiography.

Despite the promising performance of segmentation-based EF estimation, it fails on several challenging occasions where the model struggles to segment the LV properly, leading to an inability to detect any heartbeat cycle [26]. This challenge has led researchers to develop methods for measuring EF without relying on LV segmentation [20,26,17]. For instance, EchoCoTr [17] utilized a hybrid CNN-transformer video network [13] to estimate the EF from a sampled echocardiogram clip. This approach achieved the best EF estimation in the largest echocardiography dataset [19], demonstrating that the direct estimation of EF from raw video clips is more robust than the segmentation-based method.

Similar to EF estimation, there are two primary methods for classifying MI: segmentation-based and direct classification. Degerli et al. [4] segmented LV walls to extract 2D+time features, followed by SVM to classify MI. Nguyen et al. [18] utilized two segmentation models to extract the features for feature ensembling followed by logistic regression, improving the MI classification. However, the improvement is inconsistent, as they perform worse when using three feature extractors. Conversely, Saeed et al. [21] proposed an end-to-end method and achieved a competitive performance [18] without complex feature engineering.

The aforementioned DL-based direct diagnosis methods [17,21] offer straightforward solutions and superior results on large-scale data yet have reduced explainability due to their end-to-end training. The resulting representations often fail to capture the underlying relations among samples [29]. In other words, the model fails to learn that hearts with similar EF values have similar features. Instead, it may learn clinically irrelevant relations, as indicated in Figure 1, that can negatively affect performance and transferability. This highlights a critical need for methodologies that can effectively leverage the intrinsic relations among data samples. Incorporating techniques such as contrastive learning to learn from the relationships among samples can improve model performance, generalizability, and explainability [9,22,29]. For example, [29] demonstrated that modeling continuous representations boosts regression performance and transferability, an essential factor in achieving enhanced performance on small datasets [21,5,27].

Contributions. We propose **CoReEcho**, a novel framework that enables a video analysis model to infer continuous representations when predicting a clinical outcome, achieving SOTA performance on EF regression in the largest echocardiog-




Fig. 1. UMAP of learned feature embeddings on the EchoNet [19] test set. Unlike EchoCoTr, CoReEcho assembles continuous representations in the embedding space with respect to true EF, thus mitigating clinically irrelevant relations.

Algorithm 1 CoReEcho: First Stage

Require: $\mathcal{F}_E, \mathcal{F}_R, \mathcal{A}, \mathcal{C}, \tau$, batch size N .
1: **for** sampled minibatch $\{V^n, y_r^n\}_{n=1}^N$ **do**
2: **for** $n \in \{1, \dots, N\}$ **do**
3: Assign labels $y^{2n-1}, y^{2n} = y_r^n$
4: Sample clips c^{2n-1}, c^{2n} from $\mathcal{A}(\mathcal{C}(V^n))$
5: Extract embeddings E^{2n-1}, E^{2n} using $\mathcal{F}_E(\cdot)$
6: Predict $\hat{y}^{2n-1}, \hat{y}^{2n}$ using $\mathcal{F}_R(\cdot)$
7: **end for**
8: Compute \mathcal{L}_{RnC} and \mathcal{L}_{L1} as in Eq. (1)
9: Calculate $\mathcal{L} = \mathcal{L}_{RnC} + \mathcal{L}_{L1}$
10: update \mathcal{F}_E and \mathcal{F}_R to minimize \mathcal{L}
11: **end for**
12: **Return** $\mathcal{F}_E, \mathcal{F}_R$

raphy dataset [19]. We further demonstrate how the CoReEcho-pretrained model exhibits higher transferability than the current SOTA echocardiography model [17] when probing and finetuning on other small datasets. We demonstrate how our work serves as the SOTA in MI classification for a single model benchmark. Our qualitative analysis of three public datasets shows that CoReEcho produces better explainable representations and visualizations than the baselines.

2 Methodology

2.1 Problem Definition

Given an echocardiography dataset $D := \{V^n, y_r^n\}_{n=1}^{N_V}$ where $V^n \in \mathbb{R}^{F^n \times H \times W \times 3}$ is an echocardiogram video containing $F^n \in \mathbb{Z}^+$ frames with a frame size of $H \times W$, a corresponding EF value $y_r^n \in (0, 100)$, and N_V is the number of samples. We aim to train a DL-based model to predict EF value \hat{y}^n from a given V^n . Our model is composed of two parts: a *feature extractor*, \mathcal{F}_E , extracting a feature embedding E^n from an input echocardiogram and a *shallow MLP*, \mathcal{F}_R , predicting EF (\hat{y}^n) from the feature embedding E^n .

Feature extractor. Our feature extractor \mathcal{F}_E is based on UniFormer-S [13], driven by the fact that it achieves SOTA performance in echocardiography [17]. The CNNs in the first two blocks of UniFormer provide an inductive bias that helps the model learn effectively from a small dataset, followed by two global Multi-Head Relation Aggregator blocks that effectively aggregate spatiotemporal information. Thus, the embedding produced by \mathcal{F}_E can be written as $E^n \in \mathbb{R}^{C_E}$.

Shallow MLP. We employ a shallow MLP consisting of two layers to transform the extracted features into an EF value. More specifically, \hat{y}^n is given by $\hat{y}^n = W_2[g(\text{BN}_2(W_1[\text{BN}_1(E^n); 1])); 1]$, where BN denotes a batch norm layer [7] and $g(\cdot)$ is GELU activation [6]. $W_1 \in \mathbb{R}^{C_E \times (C_E+1)}$, and $W_2 \in \mathbb{R}^{1 \times (C_E+1)}$ are linear layers. Additionally, we apply dropout [25] before every linear layer.

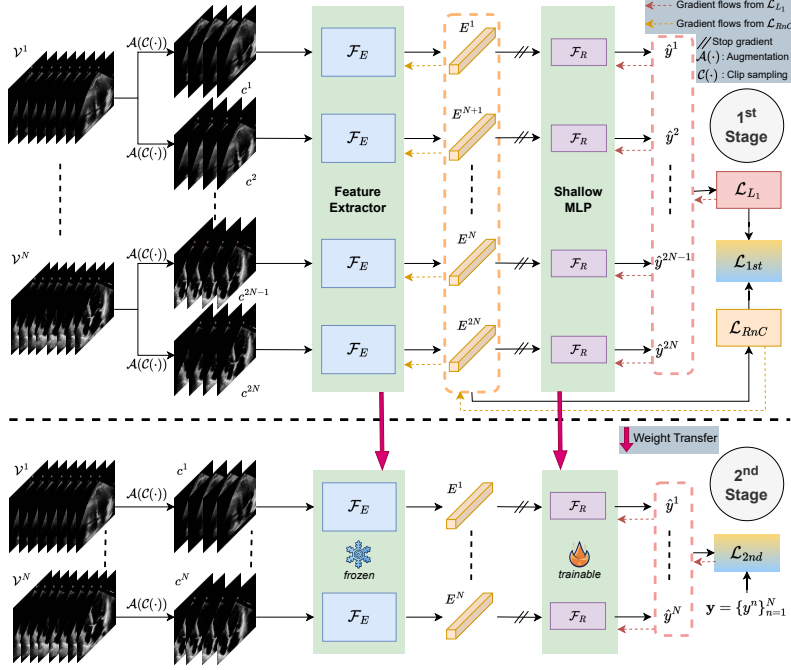


Fig. 2. CoReEcho consists of two training stages: 1) Optimize \mathcal{F}_E to learn continuous relation between echocardiogram clips while simultaneously teaching \mathcal{F}_R to estimate EF. 2) Perform MLP probing to further optimize \mathcal{F}_R while \mathcal{F}_E is frozen.

Clip sampling. DL models typically require uniform-shaped inputs for efficient processing, whereas echocardiogram videos vary in frame count. We apply clip sampling $\mathcal{C}(\cdot)$ to unify the input shapes and remove redundancy between adjacent frames before passing to \mathcal{F}_E . We randomly sample a clip $c^n \in \mathbb{R}^{F_C \times H \times W \times 3}$ from V^n with a fixed number of frames, F_C , and a sampling stride of T .

2.2 CoReEcho Training Pipeline

As illustrated in Figure 2, the training pipeline consists of two stages. In the *first stage*, we focus on learning the continuous relationship between echocardiogram clips while teaching the model to estimate EF. We then freeze the encoder in the *second stage* to refine the regression head further.

First stage. This stage is summarized in Alg. 1. We randomly select N samples from every batch ($\{V^n, y_r^n\}_{n=1}^N$). Let $\mathcal{C}(\cdot)$ and c denote the clip sampling and an echocardiogram clip, respectively. Then, we sample two clips for every echocardiogram video sample on the batch, i.e. $c^{2n-1} \triangleq \mathcal{C}(V^n)$, $c^{2n} \triangleq \mathcal{C}(V^n)$, and $y^{2n-1} \equiv y^{2n} \triangleq y_r^n$. A set of clip augmentation $\mathcal{A}(\cdot)$ is then applied to each c^n . Thus, the training data in each batch can be written as $\{\mathcal{A}(c^n), y^n\}_{n=1}^{2N}$. Through these processes, we expose the model to diverse augmented clips, which

increases the model’s robustness. Next, we extract the feature embeddings of every augmented clip, i.e., $E^n \triangleq \mathcal{F}_E(\mathcal{A}(c^n))$. Inspired by a recent work [29], we develop a representation-aware regression loss to enable \mathcal{F}_E to model the underlying continuous relation among clips in the embedding space and \mathcal{F}_R to regress \hat{y}^n from E^n . Our loss is composed of two losses, i.e., RnC loss [29] to optimize \mathcal{F}_E and L1 loss to optimize \mathcal{F}_R . We apply a stop-gradient operator $\mathcal{SG}(\cdot)$ to prevent the gradients from L1 loss to \mathcal{F}_E so that the regression loss does not disturb the continuous representation learning, i.e., $\hat{y}^n \triangleq \mathcal{F}_R(\mathcal{SG}(E^n))$. Let $s(A, B) \triangleq \exp(-L_2(A, B)/\tau)$, $\tau \in \mathbb{R}$, and $E^l \in \mathcal{S}^{n,m}$, our loss is given by

$$\mathcal{L}_{1st} = \underbrace{\frac{1}{2N(2N-1)} \sum_{n=1}^{2N} \sum_{m=1, m \neq n}^{2N} -\log \frac{s(E^n, E^m)}{\sum_{E^l} s(E^n, E^l)}}_{\mathcal{L}_{RnC}} + \underbrace{\frac{1}{2N} \sum_{n=1}^{2N} |y^n - \hat{y}^n|}_{\mathcal{L}_{L1}} \quad (1)$$

\mathcal{L}_{RnC} allows \mathcal{F}_E to learn continuous representations by iteratively contrasting an anchor n with a positive pair m and the corresponding set of negative pairs $\mathcal{S}^{n,m}$. The negative pairs $\mathcal{S}^{n,m}$ are samples within a batch where the distance from the anchor n to the negative pair is greater than the distance from the anchor n to the positive pair m . Thus, for an (n, m) pair, the corresponding negative pairs are $\mathcal{S}^{n,m} \triangleq \{E^l | l \in \{1, 2, \dots, 2N\} \setminus \{n\}, L_1(y^n, y^l) \geq L_1(y^n, y^m)\}$. Decreasing \mathcal{L}_{RnC} implies pulling a positive pair to an anchor while pushing negative pairs away. This optimization process is analogous to a ranking mechanism, fostering continuous representations that reflect the hierarchical structure of samples in a batch with respect to the true labels. Although we only force continuity at the batch level, Markov’s inequality guarantees continuity across all partitions if the \mathcal{L}_{RnC} is sufficiently low and batches are sampled randomly. **Second stage.** We continue training for several epochs further to refine \mathcal{F}_R while *freezing* \mathcal{F}_E , considering \mathcal{F}_E is now capable of extracting valuable features that resemble continuous relations among samples concerning EF. Here, we only sample a clip from every video and apply the L1 loss to update \mathcal{F}_R :

$$\mathcal{L}_{2nd} = \frac{1}{N} \sum_{n=1}^N |y^n - \mathcal{F}_R(\mathcal{SG}(\mathcal{F}_E(\mathcal{A}(v^n))))| \quad (2)$$

Compared to [29], CoReEcho is more efficient since \mathcal{F}_R is also optimized with additional cheap forward and backward passes in the 1st stage, requiring only a few epochs of refinement in the 2nd stage.

2.3 Model Transferability

We experiment with two prevalent transfer learning strategies, i.e., *MLP probing* (MLP-P) and *fine-tuning* (FT), to evaluate the transferability of a CoReEcho-pre-trained model and to show how continuous latent representations can help in generalization. The probing involves attaching randomly initialized *MLP* layers, \mathcal{F}_R , on top of a frozen pretrained CoReEcho, \mathcal{F}_E , to evaluate the learned

latent features on a target task. In contrast, FT allows \mathcal{F}_E to be updated during training for the target task, enabling \mathcal{F}_E to refine its knowledge based on the nuances of the new task.

3 Experiments

We used an RTX 4090 (24 GB) with PyTorch 1.12.1 and CUDA 11.6 throughout our experiments. We conducted experiments using three standard echocardiography benchmarks, namely EchoNet-Dynamic [19], CAMUS [11], and HMC-QU [4]. We tested the CoReEcho training framework on [19] to demonstrate its efficacy and then performed transfer learning experiments on [11] and [4] to investigate the CoReEcho pretrained model transferability. [19] and [11] primarily deal with EF regression, whereas [4] focuses on MI classification. *The preprocessing, augmentation, and hyperparameter details, are provided in the Appendix.*

EchoNet-Dynamic [19] is the largest publicly available echocardiography dataset containing 10,030 echocardiogram A4C videos with 112×112 frame size and a varying number of frames, along with each corresponding EF reference. We follow the train-val-test split in [17] for fair comparisons. We sample 36 frames of each video with $T = 4$ to standardize the clip shape and mitigate redundant information in adjacent frames, following [17]. We train our model in the 1st training stage for 25 epochs ($\tau = 1.0$), followed by the 2nd stage for 5 epochs.

CAMUS [11] is a 2D echocardiogram dataset containing 500 single heartbeat (SH) videos of the A2C and A4C views. It provides the EF value for each video as the main clinical metric. We only consider the A4C view since we focus on a single-view echocardiography analysis. For fair benchmarking, we perform 10-fold cross-validation (CV) following [11] splits. Inspired by [21], we preprocess the dataset to unify the video shape. We utilize MSE loss as the cost function.

HMC-QU [4] includes 109 A4C (SH) videos with apparent myocardial walls (72 MI and 37 normal). We follow the stratified 5-fold split provided by [4,21] and strictly implement the pre-processing by [21]. We also utilize the BCE loss.

4 Results & Discussion

4.1 Ejection Fraction Regression

Comparison with SOTA. In Table 1, we compare CoReEcho with SOTA methods using the largest echocardiography dataset [19]. While assessing both CoReEcho and EchoCoTr [17], we average the EF predictions of three sampled clips for each video to reduce the randomness introduced by the clip sampling process. CoReEcho achieves the lowest MAE of 3.90 and RMSE of 5.13, and the highest R2, indicating superiority over other methods. Additionally, CoReEcho is more efficient than [19,20,26] since it provides a direct EF regression, thus removing the need for the segmentation stage to determine a heartbeat cycle. The training time for CoReEcho is 4 hours, and the inference time is 5.95 ± 0.14 ms per clip, demonstrating our method’s suitability for real-time processing.

Table 1. EchoNet-Dynamic: Comparison with SOTA methods on the test set. **CoReEcho** achieves the best performance. (*): *reproduced performance*. SH: single heartbeat

Method	Clip(s)	MAE	RMSE	R2
[19]	1 x SH	4.22	5.56	79
	Full video	4.05	5.32	81
[20]	1 x SH	5.32	7.23	64
	Full video	5.95	8.38	52
[26]	1 x SH	4.01	5.36	81
	Full video	4.23	5.67	79
[3]	3 x Clips (*)	4.03	5.28	81
[17]	3 x Clips (*)	3.98	5.30	81
CoReEcho	3 x Clips	3.90	5.13	82

Table 2. Ablation study of **CoReEcho** components. **K400**: using pretrained weight on [8]. **LP**: the 2nd stage training.

L1	\mathcal{L}_{1st}	K400	LP	R2
✓				75.11±1.73
✓	✓			77.07±0.38
✓	✓	✓		81.50±0.43
✓	✓	✓	✓	81.59±0.18
EchoCoTr [17]				81.22±0.80

Table 3. Impact of augmentation and clip sampling. $2 \times \mathcal{C}$ means we sample two clips from each video, i.e. $c^{2n-1} \neq c^{2n}$.

\mathcal{L}_{1st}	A	$2 \times \mathcal{C}$	R2
✓	✓	✓	77.07±0.38
✓	✓		76.93±0.44
✓		✓	68.24±1.06

Table 4. CAMUS: EF regression (A4C view) with 10-fold CV. **CoReEcho** provides a stronger pretrained model that can transfer better on CAMUS.

Work	Method	Pretraining	G&M qual.		P qual.	
			corr	MAE	corr	MAE
Wei et al. [28]	Regression	-	0.725	6.50	<u>0.674</u>	7.53
	Scratch	-	0.714	6.20	0.538	8.28
Our Implementation	MLP	EchoCoTr [17]	0.660	7.00	0.401	9.84
	Probing	CoReEcho	0.730	6.28	0.570	8.42
	Fine-Tuning	EchoCoTr [17]	<u>0.799</u>	<u>5.33</u>	0.599	<u>7.49</u>
		CoReEcho	0.807	5.29	0.693	6.81

Table 5. HMC-QU: MI classification (A4C view) with stratified 5-fold CV. We achieve the highest F1 score and accuracy for the single model benchmark.

Work	Training	Pretraining	Sens.	Spec.	Prec.	F1	Acc.
Degerli et al. [4]	Scratch	-	85.97	70.10	85.52	85.29	80.24
Saeed et al. [21]	Fine-Tuning [19]	1-Heartbeat	86.10	78.21	88.53	87.09	83.44
Nguyen et al. [18]	Scratch	-	83.30	89.30	94.20	88.30	85.10
	Scratch	-	<u>95.81</u>	38.57	75.99	84.36	76.31
	MLP	EchoCoTr [17]	72.19	81.07	88.60	78.90	75.20
Our Implementation	Probing	CoReEcho	93.05	69.64	87.20	89.55	85.18
		K400 [8]	93.14	64.64	85.05	88.25	83.52
	Fine-Tuning	EchoCoTr [17]	98.57	72.50	88.37	92.96	<u>89.77</u>
		CoReEcho	92.95	<u>86.07</u>	<u>93.55</u>	<u>92.91</u>	90.64

Ablation study. Table 2 shows an ablation study in which we use L_1 loss to optimize the entire model as a baseline. We report the mean and standard deviation of the R2 scores from three runs. *Using \mathcal{L}_{1st} to optimize the model encoder and MLP head significantly improves the R2 score by +1.96%*, suggesting that learning continuous representation in the feature embeddings leads to higher performance. Additional linear probing for five epochs (*the 2nd stage*) further improves the performance, demonstrating that **CoReEcho** is more efficient because it only requires a few epochs to refine \mathcal{F}_R compared to [29]. We observe that our augmentation and clip sampling strategy are crucial in learning continuous representations, as indicated in Table 3.




Fig. 3. Grad-CAM [24]. CoReEcho \mathcal{F}_E exhibits reduced focus on backgrounds and more focus on the LV. We provide additional samples in Supplementary Materials.

Qualitative analysis. Figure 1 shows the UMAP [16] of the feature embeddings across all samples in the EchoNet-Dynamic test set. CoReEcho is able to represent the underlying continuous relation between the samples with respect to EF. Furthermore, Figure 3 shows that compared to [17], CoReEcho can exhibit reduced focus on backgrounds and place a greater emphasis on the LV, which is recognized as the most informative area for assessing EF.

4.2 Model Transferability

CAMUS. The results are presented in Table 4. We calculated the performance on both Good & Medium (**G&M**) and Poor (**P**) quality videos following the benchmark [4]. Our experiments demonstrate that in MLP probing, the CoReEcho-pretrained model results in a significantly higher performance than EchoCoTr-pretrained and can even achieve an on-par performance compared to training from scratch, suggesting that *CoReEcho is capable of extracting robust features that can generalize well on out-of-domain (OOD) data*. Our fine-tuning experiment also shows the superiority of a model pretrained with CoReEcho over EchoCoTr. Moreover, we compare our results with a direct regression method from the A4C view proposed by a SOTA method [28] on CAMUS. We outperform the method without requiring segmentation labels for additional supervision.

HMC-QU. We experimented on the A4C view, and the results are presented in Table 5. Similar to MLP probing on CAMUS, the CoReEcho-pretrained model achieves a significantly higher F1 and accuracy than EchoCoTr-pretrained and training from scratch. This indicates that *CoReEcho can extract universal features that are useful for a wide range of downstream echocardiography tasks*. Allowing the encoder to be optimized (fine-tuning) results in better performance. We then compare our result with several SOTA methods on HMC-QU [4,21,18]. Our work outperforms the aforementioned methods for a single model benchmark on F1, accuracy, precision, and sensitivity.

5 Conclusion

We propose CoReEcho, a novel training framework for a single-view echocardiography video EF estimation. CoReEcho enables the encoder to represent the

continuity relation across samples in the embedding space and train a shallow MLP head to leverage the extracted features to predict EF. **CoReEcho** outperforms the state-of-the-art methods on the EchoNet-Dynamic test set. Our experiments further show that **CoReEcho** can provide a strong pretrained model which extracts robust and universal features for single-view echocardiography downstream tasks. These observations indicate that allowing a DL-based echocardiography diagnosis model to learn more explainable representation elevates the generalization of in-domain and OOD data. Based on this, we believe **CoReEcho** can serve as an intermediate fine-tuning task for developing an echocardiography foundation model. An intermediate fine-tuning task is critical to tailor foundation model features to be more suitable for related downstream tasks, as evidenced by several works conducted with natural [27] and medical data [1]. In addition, this continuity in representations may be used to detect distribution shifts during deployment, enabling subsequent domain adaptation to maintain performance. Exploring these topics could enhance real-world diagnostic processes, which we leave for future work.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Azizi, S., Culp, L., Freyberg, J., Mustafa, B., Baur, S., Kornblith, S., Chen, T., Tomasev, N., Mitrović, J., Strachan, P., Mahdavi, S.S., Wulczyn, E., Babenko, B., Walker, M., Loh, A., Chen, P.H.C., Liu, Y., Bavishi, P., McKinney, S.M., Winkens, J., Roy, A.G., Beaver, Z., Ryan, F., Krogue, J., Etemadi, M., Telang, U., Liu, Y., Peng, L., Corrado, G.S., Webster, D.R., Fleet, D., Hinton, G., Houlsby, N., Karthikesalingam, A., Norouzi, M., Natarajan, V.: Robust and data-efficient generalization of self-supervised machine learning for diagnostic imaging. *Nature Biomedical Engineering* **7**(6), 756–779 (Jun 2023)
2. Banegas, J.R., Rodríguez-Artalejo, F.: Heart failure and instruments for measuring quality of life. *Revista española de cardiología* **61**(3), 233–235 (2008)
3. Dai, W., Li, X., Chiu, W.H.K., Kuo, M.D., Cheng, K.T.: Adaptive contrast for image regression in computer-aided disease assessment. *IEEE Transactions on Medical Imaging* **41**(5), 1255–1268 (2022). <https://doi.org/10.1109/TMI.2021.3137854>
4. Degerli, A., Zabihi, M., Kiranyaz, S., Hamid, T., Mazhar, R., Hamila, R., Gabbouj, M.: Early detection of myocardial infarction in low-quality echocardiography. *IEEE Access* **9**, 34442–34453 (2021)
5. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *2022 CVPR*. pp. 15979–15988 (2022)
6. Hendrycks, D., Gimpel, K.: Bridging nonlinearities and stochastic regularizers with gaussian error linear units (2016), <http://arxiv.org/abs/1606.08415>
7. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: Bach, F., Blei, D. (eds.) *Proceedings of ICML. Proceedings of Machine Learning Research*, vol. 37, pp. 448–456. PMLR, Lille, France (07–09 Jul 2015)

8. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A.: The kinetics human action video dataset (2017), <http://arxiv.org/abs/1705.06950>
9. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *Advances in NeurIPS*. vol. 33, pp. 18661–18673. Curran Associates, Inc. (2020)
10. Kusunose, K., Haga, A., Abe, T., Sata, M.: Utilization of artificial intelligence in echocardiography. *Circulation Journal* **83**(8), 1623–1629 (2019)
11. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., Lartizien, C., D’hooge, J., Lovstakken, L., Bernard, O.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE Transactions on Medical Imaging* **38**(9), 2198–2210 (2019)
12. Leclerc, S., Smistad, E., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Rye Berg, E.A., Belhamissi, M., Israilov, S., Grenier, T., Lartizien, C., Jodoin, P.M., Lovstakken, L., Bernard, O.: Lu-net: A multistage attention network to improve the robustness of segmentation of left ventricular structures in 2-d echocardiography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **67**(12), 2519–2530 (2020). <https://doi.org/10.1109/TUFFFC.2020.3003403>
13. Li, K., Wang, Y., Gao, P., Song, G., Liu, Y., Li, H., Qiao, Y.: Uniformer: Unified transformer for efficient spatial-temporal representation learning. In: *ICLR (2022)*
14. Ling, H.J., Garcia, D., Bernard, O.: Reaching intra-observer variability in 2-D echocardiographic image segmentation with a simple U-Net architecture. In: *IEEE International Ultrasonics Symposium (IUS)*. Venice, Italy (Oct 2022)
15. Maani, F., Ukaye, A., Saadi, N., Saeed, N., Yaqub, M.: Unilvseg: Unified left ventricular segmentation with sparsely annotated echocardiogram videos through self-supervised temporal masking and weakly supervised training (2023)
16. McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction (2020)
17. Muhtaseb, R., Yaqub, M.: Echocotr: Estimation of the left ventricular ejection fraction from spatiotemporal echocardiography. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) *MICCAI 2022*. pp. 370–379. Springer (2022)
18. Nguyen, T., Nguyen, P., Tran, D., Pham, H., Nguyen, Q., Le, T., Van, H., Do, B., Tran, P., Le, V., et al.: Ensemble learning of myocardial displacements for myocardial infarction detection in echocardiography. *Frontiers in Cardiovascular Medicine* **10** (2023)
19. Ouyang, D., He, B., Ghorbani, A., Yuan, N., Ebinger, J., Langlotz, C.P., Heidenreich, P.A., Harrington, R.A., Liang, D.H., Ashley, E.A., et al.: Video-based ai for beat-to-beat assessment of cardiac function. *Nature* **580**(7802), 252–256 (2020)
20. Reynaud, H., Vlontzos, A., Hou, B., Beqiri, A., Leeson, P., Kainz, B.: Ultrasound video transformers for cardiac ejection fraction estimation. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *MICCAI 2021*. pp. 495–505. Springer (2021)
21. Saeed, M., Yaqub, M.: End-to-end myocardial infarction classification from echocardiographic scans. In: Aylward, S., Noble, J.A., Hu, Y., Lee, S.L., Baum, Z., Min, Z. (eds.) *Simplifying Medical Ultrasound*. pp. 54–63. Springer (2022)
22. Sanjeev, S., Al Khatib, S.K., Shaaban, M.A., Almakky, I., Papineni, V.R., Yaqub, M.: Pecon: Contrastive pretraining to enhance feature alignment between ct

- and ehr data for improved pulmonary embolism diagnosis. In: *Machine Learning in Medical Imaging*. pp. 434–443. Springer (2024)
23. Savarese, G., Stolfo, D., Sinagra, G., Lund, L.H.: Heart failure with mid-range or mildly reduced ejection fraction. *Nature Reviews Cardiology* **19**(2), 100–116 (2022)
 24. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *IEEE ICCV*. pp. 618–626 (2017)
 25. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15**(56), 1929–1958 (2014)
 26. Thomas, S., Gilbert, A., Ben-Yosef, G.: Light-weight spatio-temporal graphs for segmentation and ejection fraction prediction in cardiac ultrasound. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) *MICCAI 2022*. pp. 380–390. Springer (2022)
 27. Wang, L., Huang, B., Zhao, Z., Tong, Z., He, Y., Wang, Y., Wang, Y., Qiao, Y.: Videomae v2: Scaling video masked autoencoders with dual masking. In: *Proceedings of the IEEE/CVF Conference on CVPR*. pp. 14549–14560 (June 2023)
 28. Wei, H., Ma, J., Zhou, Y., Xue, W., Ni, D.: Co-learning of appearance and shape for precise ejection fraction estimation from echocardiographic sequences. *Medical Image Analysis* **84**, 102686 (2023)
 29. Zha, K., Cao, P., Son, J., Yang, Y., Katabi, D.: Rank-n-contrast: Learning continuous representations for regression. In: *Conference on NeurIPS* (2023)