

Adaptive Contrast for Image Regression in Computer-Aided Disease Assessment

Weiham Dai¹, Xiaomeng Li², *Member, IEEE*, Wan Hang Keith Chiu,
Michael D. Kuo, and Kwang-Ting Cheng³, *Fellow, IEEE*

Abstract—Image regression tasks for medical applications, such as bone mineral density (BMD) estimation and left-ventricular ejection fraction (LVEF) prediction, play an important role in computer-aided disease assessment. Most deep regression methods train the neural network with a single regression loss function like MSE or L1 loss. In this paper, we propose the first contrastive learning framework for deep image regression, namely AdaCon, which consists of a feature learning branch via a novel adaptive-margin contrastive loss and a regression prediction branch. Our method incorporates label distance relationships as part of the learned feature representations, which allows for better performance in downstream regression tasks. Moreover, it can be used as a plug-and-play module to improve performance of existing regression methods. We demonstrate the effectiveness of AdaCon on two medical image regression tasks, *i.e.*, bone mineral density estimation from X-ray images and left-ventricular ejection fraction prediction from echocardiogram videos. AdaCon leads to relative improvements of 3.3% and 5.9% in MAE over state-of-the-art BMD estimation and LVEF prediction methods, respectively.

Index Terms—Bone mineral density estimation, contrastive learning, ejection fraction prediction, image regression.

I. INTRODUCTION

MEDICAL image classification and segmentation with convolutional neural networks (CNNs) have seen huge

adoption in various applications for computer-aided diagnosis, such as disease grading [1], [2], organ segmentation [3], [4], and tumor segmentation [5]–[8]. Comparatively, medical image regression tasks based on CNNs have received less attention. Unlike classification and segmentation problems, regression tasks generate continuous value predictions that are important indicators for various disease assessments in clinical practice. Two examples of this are bone mineral density (BMD) and left-ventricular ejection fraction (LVEF). BMD is used to identify osteoporosis, determine the risk of fractures, and measure patient response to osteoporosis treatment [9]. Typical BMD values for the Asian population range between 0.5–0.9, with a low reading indicating loss of bone mass and higher risks of fractures. LVEF is a percentage measurement of how much blood the left ventricle pumps out with each contraction, serving as an essential indicator to diagnose and track heart failure. Normal LVEF ranges from 55% to 70%, with LVEF higher than 75% or lower than 40% indicating a potential for heart failure or cardiomyopathy [10]. Early detection and accurate assessment of BMD and LVEF are important for early intervention in clinical practice.

Early works for image regression relied on hand-crafted filters as feature inputs in a multivariate regression model. For example, Pulkkinen *et al.* [11] estimate BMD by applying a gradient filter to patches of femur-neck X-ray images and calculating summary statistics as feature inputs. Similar techniques have also been applied for computer vision tasks using natural images, such as facial age estimation [12], [13] and single-image depth estimation [14]. Recent progress in image regression has increasingly featured deep learning techniques [15]–[17]. For instance, Luo *et al.* [15] used a VGG backbone [18] on cardiac magnetic resonance volumes to directly regress ventricular volumes. Ouyang *et al.* [16] employed the R2+1D ResNet network [19] to directly regress LVEF values from echocardiogram videos. Among these works, mean squared error (MSE) is used as the loss function to measure prediction error and supervise training. Alternative loss functions such as L1 loss [20], Huber Loss [21], [22], and Tukey loss [23] have also been proposed for image regression as ways to reduce the effect of sample outliers.

Although different regression loss functions have been explored, model performance from training with a singular regression loss function can still be limited. Recently, contrastive learning has shown great promise for feature representation learning in a variety of computer vision tasks, such as unsupervised pretraining [24]–[27], image classification [28],

Manuscript received November 23, 2021; accepted December 15, 2021. Date of publication December 23, 2021; date of current version May 2, 2022. This work was supported in part by the Research Grant from Shenzhen Municipal Central Government Guides Local Science and Technology Development Special Funded Projects under Grant 2021Szzup139 and in part by the Research Grant from the Hong Kong Research Grants Council (RGC) through the National Nature Science Foundation of China/Research Grants Council Joint Research Scheme under Grant N_HKUST627/20. (Corresponding author: Xiaomeng Li.)

Weiham 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).

Xiaomeng 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, Shenzhen 518057, China (e-mail: eexmli@ust.hk).

Wan Hang Keith Chiu and Michael D. Kuo are with the Department of Diagnostic Radiology, LKS Faculty of Medicine, The University of Hong Kong, Hong Kong, SAR, China (e-mail: kwhchiu@hku.hk; mikedkuo@gmail.com).

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

Digital Object Identifier 10.1109/TMI.2021.3137854

and video detection [29]. For example, Chen *et al.* [24] showed that their unsupervised SimCLR framework achieves comparable performance with state-of-the-art performance on ImageNet classification. Khosla *et al.* [28] further demonstrated that contrastive learning used in a supervised setting could surpass state-of-the-art ImageNet results. Although contrastive learning has demonstrated excellent performance for various tasks, adapting it for image regression problems still remains under-explored.

In this paper, our goal is to explore the feasibility of contrastive learning for image regression tasks. One naïve solution is to treat each continuous value as an individual class and use the supervised contrastive loss function to learn feature representations [28]. However, this method ignores the underlying distance relationships between labels, leading to inferior results; see results of SupCon in Tables II and IX. To this end, we present a simple and effective framework, namely AdaCon, that introduces a novel distance-adaptive contrastive loss function to model distance relationships in the representation space for image regression tasks. AdaCon conducts contrastive learning and regression prediction in a multi-task framework and learns features that are more consistent with the regression task, thus improving the overall prediction performance. We demonstrate our methodology on two medical image regression tasks, *i.e.*, BMD estimation from X-rays images and LVEF prediction from echocardiogram videos.

The key contributions of this study are:

- We present a simple yet effective framework (AdaCon) for image regression tasks, consisting of a contrastive feature learning branch for representation learning and a regression estimation branch for regression. Our method can serve as a plug-and-play component in existing methods to improve image regression performance.
- AdaCon introduces a novel adaptive-margin loss function to preserve distance and order relationships among continuous labels in the representation space. We show that the enhanced feature representation can improve image regression results.
- AdaCon achieves state-of-the-art performance on two medical image regression tasks: BMD estimation from single X-ray images and LVEF prediction from echocardiogram videos.¹ Ablation studies demonstrate AdaCon's superior performance over other methods. Results are also validated using three additional datasets.

II. RELATED WORK

In this section, we review contrastive learning for representation learning and some of its applications. We further discuss related works for BMD estimation and LVEF prediction.

A. Contrastive Learning

1) *Unsupervised Representation Learning*: In the early stages of unsupervised representation learning, researchers relied on dimension-reduction techniques such as PCA [30], and clustering-based methods on image patches [31]–[33] to

extract feature representations. With the popularity of deep learning techniques, architectures like deep auto-encoders, which introduce dimension reducing bottlenecks and are trained by minimizing reconstruction error [34]–[36], have become increasingly used. Studies have also explored pretraining networks on pretext tasks such as image transformation identification, showing that it can be an effective way to learn feature representations for images [37]–[39]. The pretrained networks can then be applied to downstream tasks by training a classifier on top of the feature extractor.

In recent years, however, contrastive learning methods have dominated this field. Chen *et al.* [24] proposed the SimCLR framework, in which a network is pretrained to identify positive sample pairs, defined as samples augmented from the same source image. They demonstrated that when fine-tuning the pretrained network to perform downstream image classification tasks, they can achieve state-of-the-art performance, sometimes even exceeding results from supervised learning. Similarly, MoCo identifies augmented images from the same source by maximizing feature similarity between positive pairs using a momentum encoder, thus reducing the reliance on large batch sizes [40]. Subsequent variations adopt improvements from SimCLR and vision transformers into Moco_v2 [27] and Moco_v3 [41]. Additional studies have also looked at improving performance through stronger augmentations [42], adapting underlying concepts to different architectures [43], or using contrastive learning for unsupervised tasks like clustering [25].

2) *Contrastive Learning for Downstream Applications*: Given the stunning improvements seen in unsupervised representation learning, researchers started to adapt contrastive methods to various downstream tasks such as image classification [28], [44], segmentation [45], [46], and video detection [29]. Khosla *et al.* [28] showed that by using label information to identify additional positive-pair samples through supervised contrastive learning, classification accuracy on the benchmark ImageNet classification can be improved by 0.8%. Wang *et al.* [44] also demonstrated that contrastive learning can be used to improve classification performance for highly imbalanced datasets. Wang and Zhao *et al.* [45], [46] demonstrated that contrastive learning can be used to improve performance on image segmentation tasks by performing contrastive learning on pixel-level features instead of global features. By considering temporal information, Dave *et al.* [29] also showed that contrastive learning can also be used efficiently to learn feature representations for videos.

Although contrastive learning has been applied to various tasks, adapting it to regression remains an open and non-trivial problem. This study proposes a simple adaptive contrastive learning framework, AdaCon, that allows the same principles to be applied to continuous value prediction. We demonstrate that AdaCon can be used as a general technique to improve performance for different image regression tasks.

3) *Comparison With Deep Metric Learning*: Contrastive learning is closely related to traditional distance metric learning, which uses margins to impose separation between different classes in the feature space [47]–[50]. The main difference between the two is that contrastive learning generalizes to an

¹Code is available at <https://github.com/XMed-Lab/AdaCon>

arbitrary number of positive and negative sample pairs. In contrast, metric learning loss functions, such as triplet loss [48] and N-pair loss [50], are restricted in the number of pairs they use. Triplet loss, in particular, is known to be unstable and highly dependent on hard-sample mining [49]–[51]. The number of possible triplet combinations also scales at a cubic rate relative to sample size, making it difficult to sample well for large batch sizes. Khosla *et al.* [28] show analytically that the SupCon loss function is a more general form of metric learning loss that reduces to triplet loss and N-pair loss under certain conditions and is a superior alternative. They also show empirically that SupCon outperforms these losses due to the inherent hard-mining properties of its loss function.

Distance metric learning has been successfully used in regression problems by making the separation margins adaptive. The adaptive triplet loss, for example, imposes larger margins on samples pairs with larger differences in values [17], [52]. However, these functions still rely on triplet sampling and only applies constraints within each sampled triplet group. Unlike this method, the optimization of our method is performed over an entire batch with constraints applied simultaneously across all samples, contributing to a more stable training process and better feature representations. We also empirically show in Tables II and IX that AdaCon outperforms adaptive triplet loss.

B. Bone Mineral Density Estimation From Plain X-Ray Films

Medical imaging studies on bone X-rays have traditionally been focused on fracture detection [53]–[55], disease diagnosis [56], [57], or segmentation [57], [58]. Little attention has been given to bone mineral density (BMD) estimation, which is important for diagnosing osteoporosis, a condition characterized by decreased levels of bone density. BMD is calculated as bone-mass/coverage-area (g/cm^2) over an region of interest, typically the femur-neck or spine, and the gold standard is measurement by dual-energy X-ray absorptiometry (DXA). However, due to the limited availability of DXA devices, especially in developing countries, osteoporosis is often under-diagnosed and under-treated. Compared to DXA, plain X-ray films are easier and cheaper to obtain. Therefore, alternative lower-cost BMD evaluation methods using more accessible mediums like X-ray films are highly beneficial.

Early works explored using handcrafted features from X-rays to correlate BMD. Pulkkinen and Chappard [11], [59] for example both calculate texture-based summary statistics on X-ray patches to regress on BMD using multi-variate regression. Deep learning methodologies have also been used recently following their popularity in computer vision problems. Chu *et al.* [60] use a siamese network on patches from dental radiographs to perform osteoporosis diagnosis. Wang *et al.* [61] use a GCN on chest X-rays to first perform ROI detection before regressing image crops on BMD with a VGG-16 network and MSE loss. Hsieh *et al.* [20] use a VGG-11 network and L1 loss to regress BMD from femur-neck and spine vertebrae crops after performing ROI detection. Zheng *et al.* [17] use a VGG architecture for BMD

regression with additional adaptive triplet loss and semi-supervised learning.

Most of the existing techniques for BMD estimation only use a single regression loss to supervise the training process with no further optimization considered on the feature space. In contrast to these methodologies, we present a novel plug-and-play framework, AdaCon, that can be used for different regression tasks. By performing BMD regression and feature representation learning with AdaCon in a multi-task approach, we can outperform alternative methods by ensuring better features are learnt, as seen in Table II. Notably, Zheng *et al.* [17] use the adaptive triplet loss function in addition to regression loss for BMD estimation. We show analytically and empirically that contrastive learning with AdaCon is a better alternative to their methodology.

C. Ejection Fraction Prediction From Echocardiogram Videos

Ejection fraction is the percentage change of a patient's left-ventricular volume between the end-diastole phase (EDV) and the end-systole phase (ESV). It is an indicator for the heart's blood-pumping capability and is commonly used to diagnose heart failure [62]. Measurement is normally done manually by outlining the left ventricle in echocardiograms during different stages of a heart beat and estimating the corresponding volume. This process is laborious and can also lead to significant inter-observer variation [16], [63], therefore driving the need for automated techniques.

Automated LVEF prediction has been studied using both still imaging and video inputs. The most common approach is to first estimate EDV and ESV volumes from still frames and then calculate LVEF based on these values. Luo *et al.* [15] use cardiac magnetic resonance data to perform deep regression with MSE loss for volume estimation after image pre-processing and slice selection. Zhen *et al.* [64] use a regression forest model to jointly regress the left and right ventricle volumes from MRI images with MSE. An alternative technique is to treat volume estimation as a segmentation problem and use the segmented area to estimate volume. Jafari *et al.* [65] use Bayesian estimation for ventricle segmentation from ultra-sound echo cines for volume estimation. Liu *et al.* [66] introduce a pyramid local attention module to capture local feature similarities to improve segmentation accuracy.

Instead of learning from still images, Ouyang *et al.* [16] proposed to predict LVEF values directly from echocardiogram videos using their collected dataset, EchoNet-Dynamic. They regress video inputs directly using a R2+1D ResNet [19] backbone against LVEF values with MSE loss. Reynaud *et al.* [67] also explored LVEF prediction on the EchoNet-Dynamic dataset by using a video transformer network to identify ED and ES frames, which is then used to perform segmentation and volume estimation.

Current deep regression methodologies employ the MSE loss to train the network and do not consider optimizations on the feature space. Unlike these methodologies, we design a new framework, AdaCon, that improves image regression prediction via an improved feature representation learning.

III. METHODOLOGY

We denote $\mathcal{D} := \{(x_i, y_i)\}_{i=1}^N$ as the training dataset consisting of N image-label pairs, where y_i is the label of image x_i . Our goal for contrastive learning is to learn a feature embedding network $f_\theta(\cdot)$ from the training dataset. The network $f_\theta(\cdot)$ maps the input x_i to a L2-normalized d -dimensional embedding, z_i , such that $z_i = f_\theta(x_i) \in \mathbb{S}^{d-1}$ and lies on the unit hypersphere. Let \mathcal{I} denote the sample indices for a randomly sampled batch during training. We denote $\mathcal{P}(i) := \{j \in \mathcal{I} | y_j = y_i, j \neq i\}$ as the set of indices for all positive samples for anchor sample i , i.e., samples with the same label as y_i . The negative pairs in the batch are images with different labels to anchor sample i , the indices for which can be denoted as: $\mathcal{Q}(i) := \{j \in \mathcal{I} | y_j \neq y_i\}$.

One naïve approach to contrastive learning for regression is to apply image augmentations to a batch, treat each continuous label as an individual class, and apply SupCon loss [28] directly. Then, the loss function can be expressed as:

$$\mathcal{L}_{sup} = \sum_{i \in \mathcal{I}} \frac{-1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(s \cos(\theta_{i,p}))}{\sum_{a \in \mathcal{I} \setminus i} \exp(s \cos(\theta_{i,a}))}, \quad (1)$$

where s denotes the temperature scaling factor, usually determined empirically, and $\cos(\theta_{i,j}) = z_i^T z_j$.

Intuitively, for an anchor sample with label 0.5 and negative samples of labels 0.4 and 0.1, the sample with label 0.4 should be closer to the anchor image than the one with label 0.1 in the representation space. However, by using the SupCon loss function directly, *both negative samples are treated the same* during optimization, leading to inferior results. We can see this empirically in [Tables II and IX](#).

To tackle this problem, we must also capture the relationship between the continuous labels in the representation space. For example, given a randomly sampled batch \mathcal{I} with anchor sample i and samples j, k such that $y_i \leq y_j \leq y_k$, the expected relationship between the sample feature projections is:

$$\cos(\theta_{i,j}) \geq \cos(\theta_{i,k}). \quad (2)$$

This is because we expect the sample with label closer to the anchor i to have features that are more similar. Equation (2) should also be true for samples where $y_i \geq y_j \geq y_k$. We can apply adaptive margins at the decision boundaries between different sample pairs to encourage our model to learn this relationship. By incorporating label distance and ranking information into the learned representations, the features learned will be more consistent with the regression task.

A. Adaptive Contrast Method

In contrastive learning for supervised settings, we identify positive sample pairs, defined as pairs with the same labels, amongst a randomly sampled batch. We add an adaptive margin $d_{i,q}$ into our decision boundary to introduce label ordering and distance information. For any positive sample $p \in \mathcal{P}(i)$ and negative sample $q \in \mathcal{Q}(i)$, we would like to achieve the following:

$$\cos(\theta_{i,p}) > \cos(\theta_{i,q}) + d_{i,q}. \quad (3)$$

We want to choose our adaptive margin function, $d_{i,q}$, based on the following characteristics:

- 1) For any $i, j, k \in \mathcal{I}$ where $y_i < y_j < y_k$, or $y_i > y_j > y_k$, we have $d_{i,k} > d_{i,j}$. This means we want a larger margin imposed on the feature similarity of sample pair (i, k) compared to (i, j) because sample k is further away from i in label space.
- 2) Because the cosine function is bounded between $(-1, 1)$, we want $d_{i,q}$ to be dispersed within the range of $(0, 2)$ to encourage feature separation for stronger representation learning.

Based on these characteristics, we define the adaptive margin $d_{i,q}$ as follows:

$$d_{i,q} = 2 \times |\phi(y_i) - \phi(y_q)|. \quad (4)$$

where ϕ is the empirical cumulative distribution function (ECDF) [68], defined by:

$$\phi(y_i) = \frac{1}{N} \sum_{j=1}^N \mathbb{I}_{y_j \leq y_i}. \quad (5)$$

It is easy to see that our adaptive margin function $d_{i,q}$ satisfies both conditions above. Condition one is satisfied because ϕ is a monotonically increasing function and $\phi(y_i) < \phi(y_j) < \phi(y_k)$ if $y_i < y_j < y_k$. Therefore, margin $d_{i,k}$ is greater than $d_{i,j}$. This is also true if $y_i > y_j > y_k$, because we have $\phi(y_i) > \phi(y_j) > \phi(y_k)$. The ECDF also provides a natural mapping of label values onto the $(0,1)$ range by normalizing the label distribution to a uniform distribution [69]. This ensures margin values are robust to different label distributions, even if they are heavily skewed. Condition two is satisfied because the range of $d_{i,q}$ is within $(0,2)$, which encourages feature separation.

We can incorporate our adaptive margins into the supervised contrastive loss function in (1) by adding it after the cosine similarity term, giving us:

$$\mathcal{L}_{con} = \sum_{i \in \mathcal{I}} \frac{-1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(s \cos(\theta_{i,p}))}{\sum_{a \in \mathcal{I} \setminus i} \exp(s(\cos(\theta_{i,a}) + d_{i,a}))}, \quad (6)$$

where $d_{i,a}$ follows definition (4). The loss function is minimized when all positive pairs are correctly identified within sample batch \mathcal{I} , i.e. when condition (3) is satisfied, thus enforcing our adaptive decision boundary constraints. A more detailed derivation of (6) is provided in Appendix A. We use this adaptive-margin loss function to supervise contrastive feature learning in our AdaCon framework. [Fig. 1](#) illustrates AdaCon's loss function as compared to SupCon using a simplified example on the 2D plane.

1) Margin Analysis: The ECDF of some value y , $\phi(y)$, is the sample estimate of $Pr(Y < y)$. Using this as part of our adaptive margin function gives us a number of desirable properties.

The first property is that the ECDF reduces the effect of sample outliers on margin values because of distribution normalization [69]. For example, for some outlying value y_j , its relative order ranking will be preserved by the adaptive

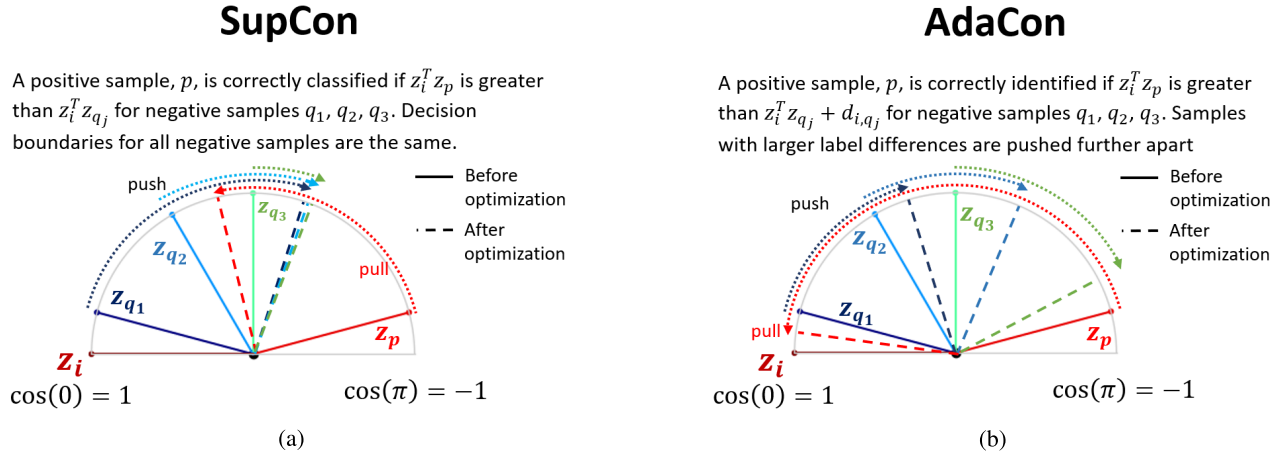


Fig. 1. (a) For classification based supervised contrastive loss, SupCon [28], samples that form positive pairs are pulled closer together in the angular space, whilst negative pairs are pushed further apart. The boundary condition for correct positive-pair identification is $z_i^T z_p > z_i^T z_{q_j}$, i.e. $\cos(\theta_{i,p}) > \cos(\theta_{i,q_j})$, where $p \in \mathcal{P}$ and $q_j \in \mathcal{Q}$; negative pairs are pushed apart indiscriminately. (b) By adding adaptive margins to the boundary conditions for classification, we enforce the model such that the feature projections z_{q_2} is pushed further apart on the angular space compared to z_{q_1} if $y_p < y_{q_1} < y_{q_2}$ or $y_{q_2} > y_{q_1} > y_p$.

margin $d_{i,j}$ after transformation, but the margin will not be excessively large because of large differences in label value. This allows for greater feature separation within the $(-1, 1)$ range for remaining samples.

The second property is that the ECDF transformation gives an intuitive mapping of the adaptive margin function based on sample probability. For example, for samples i, j where $y_i > y_j$, we can substitute $d_{i,j}$ with $2 \times |\phi(y_i) - \phi(y_j)|$ at the decision boundary to give us:

$$\cos(\theta_{i,p}) - \cos(\theta_{i,j}) = 2 \times |\phi(y_i) - \phi(y_j)|. \quad (7)$$

The value $|\phi(y_i) - \phi(y_j)|$ is equal to the sample estimate of $Pr(y_j < Y < y_i)$. Therefore, at the decision boundary, the margin between the cosine similarity values is proportional to the sample probability of observing a value between y_j and y_i . This means if there are a larger number of observed samples with labels between y_j and y_i , we will impose a larger margin between $\cos(\theta_{i,p})$ and $\cos(\theta_{i,j})$. If there are fewer samples, the margin will be smaller.

B. Comparison With Adaptive Triplet Loss

Adaptive methods for distance metric learning have been used in [17], [52] for regression tasks by introducing variable separation margins based on label differences into the standard triplet loss function. The major shortfall of triplet loss however is that it is highly dependant on the triplets sampled, but the number of possible triplets grows at a cubic rate relative to batch size, making it difficult to do so effectively [49]–[51]. Margins are also calculated and applied only within each triplet group. Unlike adaptive triplet loss, AdaCon simultaneously applies constraints to all samples within a batch, optimizing their features simultaneously during the training process. Also, the contrastive task identifies positive pairs and therefore only grows at a quadratic rate relative to batch size. It is generalizable to multiple positive and negative pairs.

We analytically demonstrate in Appendix B that AdaCon can be regarded as an improved, general form of adaptive distance metric learning by showing that the adaptive-margin loss function reduces to an approximation of adaptive triplet loss in the case of only one positive and one negative pair. We also support this empirically by demonstrating in our experiments that AdaCon outperforms adaptive triplet loss, as can be seen in Tables II and IX.

C. Full Framework

The network architecture in our proposed framework consists of a feature extraction backbone and two prediction branches: one for regression and one for contrastive learning.

The contrastive learning branch outputs feature projections and is composed of a dense layer, a ReLU activation layer, a final dense layer with an output dimension of 128, and a L2 normalization operation. For each input batch $\{x_i; i \in \mathcal{I}\}$, input samples undergo augmentation to form an augmented batch $\{x'_i; i \in \mathcal{I}\}$. The two batches are used together in the AdaCon loss function to calculate contrastive loss, \mathcal{L}_{con} . Positive pairs are defined as inputs with the same ground truth labels, which in most cases will be augmented images from the same sample source, i.e., $\{x_i, x'_i\}$.

The regression branch is used to generate the prediction output for the target value from the input batch $\{x_i; i \in \mathcal{I}\}$, and can be supervised with either MSE or robust alternatives like L1 and Huber loss. The regression loss, \mathcal{L}_{reg} , is optimized together with the contrastive loss in a multi-task training scheme. The total loss is:

$$\mathcal{L}_{total} = \gamma_{reg} \mathcal{L}_{reg} + \gamma_{con} \mathcal{L}_{con}, \quad (8)$$

where γ_{reg} and γ_{con} are weights for regression and adaptive contrastive loss respectively and control the relative importance of the two tasks. We fix $\gamma_{reg} = 1$ and set γ_{con} such that the weighted loss values for both tasks are within similar magnitudes after the first training epoch. We show in




Fig. 2. Under the AdaCon framework, a contrastive prediction branch and a regression branch are added after the feature layer for feature representation learning and regression prediction respectively. The contrastive learning branch is supervised with AdaCon loss which is trained together with a regression loss in a multi-task scheme.

Tables IV and XI that our results are robust to slight differences in selected values for the weight parameter. A visual illustration of our framework is shown in Fig. 2

IV. EXPERIMENTS

We demonstrate the effectiveness of our method on two medical image regression tasks: BMD estimation from X-ray images and LVEF prediction from echocardiogram videos.

A. BMD Estimation From X-Ray Images

1) Experimental Setting:

a) *Dataset*: Dual-energy X-ray absorptiometry (DXA) is the gold standard for diagnosing osteoporosis and monitoring changes in BMD over time. We use the BMD readings from DXA scans as ground truth labels and collect corresponding X-ray images taken within 6 months of the DXA scans as our model input, following similar settings in [17], [20]. In total, we collect 475 valid hip X-ray images from 317 patients who have had a DXA scan performed between 2016-2020. The X-ray images were collected using Carestream Health - DRX-EVOLUTION and Canon Inc. - CXDI Control Software NE machines. DXA scans were performed using Hologic - Horizon A machines. We present some summary statistics of the dataset in Tables XIV and XV of Appendix C. The anonymized data was provided by Queen Mary Hospital in Hong Kong. The study was approved by the Institutional Review Board of the University of Hong Kong/Hospital Authority Hong Kong West Cluster (IRB no: UW19-732) and requirement for informed consent was waived.

b) *Preprocessing*: The 415 X-ray images were pre-filtered by a radiologist to contain only anteroposterior views of the hip such that the femur-head and femur-neck are clearly visible. Those containing displaced fractures in the femur-neck, where the head has been entirely detached, or containing a prosthesis were also excluded. We then take square crops of the X-ray images around the femur-head and neck region, which are the region of interests corresponding to hip DXA scans. We first use the tool LabelImg² to find the square

²<https://github.com/tzutalin/labelImg>

TABLE I
COMPARISON WITH THE STATE-OF-THE-ART
METHODS FOR BMD ESTIMATION

Method	MAE↓	RMSE↓	R ² ↑
Hsieh <i>et al.</i> [20]	0.0625	0.0789	61.38%
Zheng <i>et al.</i> [17]	0.0612	0.0773	62.04%
Ours	0.0592	0.0759	64.49%

region that tightly encloses the femur head and neck. We then expand the bounds by 10% in all directions to incorporate neighbourhood regions within our crop. This yielded a total of 667 square crops centered on the femur with sizes between 800-1200 pixels. Some examples of valid and invalid crops are shown in Fig. 5 of Appendix C.

We apply random data augmentations including horizontal flips, rotations, and frame jitters during sampling. The crops are resized to 128×128 pixels before being fed into our network. We perform four-fold cross-validation for all the experiments and report the averaged performance across folds.

c) *Training details*: We use EfficientNet-B1 [70] as our backbone for BMD estimation and initialize with weights pretrained on ImageNet. We use SGD as our optimization algorithm with learning rate of 10^{-2} , momentum of 0.9, and weight decay of 10^{-4} . The network is trained for 6,000 iterations with learning rate decay of factor 0.1 at 3,000 and 4,500 iterations. For the regression prediction branch, we use L1 loss for supervision, which provides the highest regression baseline from our experiments (see Table III).

Contrastive learning benefits from a large batch size since more negative sample pairs can be formed to improve results. To obtain more samples per batch, we enlarge the batch size by following the batch augmentation methodology in [71]. We use an original batch size of 8 and an augmentation multiple of 8. Experiments were conducted with a V100 GPU.

2) Comparison With State-of-the-Art Methods:

Hsieh *et al.* [20] used only L1 loss to directly regress BMD values from femur-neck crops of X-ray images. Zheng *et al.* [17], in addition to MSE loss, applied adaptive triplet loss directly on the feature layer without using L2 normalization. Semi-supervised learning with unlabeled data was also used to boost performance. We compare AdaCon with these two methodologies to demonstrate the effectiveness of our framework. To fairly compare with their results, we use the same EfficientNet-B1 backbone.

We see from Table I that Zheng *et al.*'s method, leads to better results than Hsieh *et al.*, where only L1 loss is used for training. AdaCon gives us the best performance across all metrics, achieving 0.0592 on MAE, representing a relative improvement of 3.3% over Zheng *et al.*. In clinical applications, raw BMD values are used to track patient response to osteoporosis treatment or for disease diagnosis [9]. Our improved methodology increases accuracy of BMD estimates, which directly translates into more reliable disease tracking and classification applications for computer-aided assessment.

3) *Comparison With Other Representation Learning Losses*: To demonstrate the effectiveness of the proposed adaptive-margin contrastive loss in AdaCon, we compare its

TABLE II
COMPARISON OF DIFFERENT REPRESENTATION LEARNING
LOSS FUNCTIONS FOR BMD ESTIMATION

Method	MAE↓	RMSE↓	R ² ↑
N-Pair	0.0637	0.0816	58.70%
SupCon	0.0629	0.0794	61.15%
Regression (baseline)	0.0625	0.0789	61.38%
Adaptive Triplet	0.0601	0.0766	63.70%
Ours	0.0592	0.0759	64.49%

performance with alternative loss functions: N-pair loss [50], SupCon loss [28], and adaptive triplet loss [28]. We also compare with a regression baseline where only a single L1 loss is used. For fair comparison, we simply replace our AdaCon loss with these alternative loss functions in our framework and keep other parameters the same. For N-Pair and SupCon loss, which are typically used for classification, we treat each unique label as an individual class.

As shown in Table II, our method (AdaCon) achieves the best results whereas classification based loss functions such as N-Pair and SupCon perform worse than the baseline. This is consistent with our expectations since treating all negative samples the same without accounting for label distance can result in inadequate features being learned. Adaptive triplet loss leads to a decrease in MAE of 0.0024 compared to the regression baseline, but the improvement is still lower than ours. AdaCon achieves 0.0592 on MAE, an improvement of 0.0033 over the baseline.

We visualize the learned feature projections in Fig. 3 to better understand the performance improvement. The top row (Fig. 3a-d) shows the angular relationships of the learned feature projections from the validation set, with blue (red) lines representing low (high) BMD values. The bottom row (Fig. 3e-h) plots feature similarity against the absolute difference of $\phi(y)$ for randomly sampled pairs. We can see in Fig. 3d that by using AdaCon, the learned features are well spread out in the angular similarity space, following an orderly progression from low BMD (blue) to high BMD (red). Similarly, we see there is a clear downward trend in Fig. 3h, which demonstrates that pairwise samples with large differences in BMD values tend to have lower cosine similarity. This demonstrates that by using AdaCon, our model is able to learn separable features reflecting label distances based on feature similarities. This helps explain the stronger performance on the regression task.

On the other hand, we can see from Fig. 3b and Fig. 3c, that feature projections are not well separated using SupCon and adaptive triplet loss. There is also no distinct downward trend in their scatter plots (Fig. 3f and Fig. 3g). Although the features learned by the N-pair loss (Fig. 3a) are well separated, they do not follow the order of low BMD to high BMD. There is also no clear relationship between pair-wise feature similarity and label distance (Fig. 3e). The feature representations learned using these loss functions are therefore not as suitable, leading to worse regression performance.

4) *Ablation With Different Regression Losses*: AdaCon as a framework can be used as a plug-and-play component with existing regression losses to improve performance. We perform ablation analysis on the BMD dataset with and without

TABLE III
COMPARISON WITH DIFFERENT REGRESSION
LOSSES FOR BMD ESTIMATION

Method	MAE ↓	RMSE ↓	R ² ↑
Huber ($\delta = 0.05$)	0.0823	0.1029	34.27%
Huber ($\delta = 0.05$) + Ours	0.0732	0.0927	46.09%
MSE	0.0719	0.0894	50.29%
MSE + Ours	0.0668	0.0842	56.08%
L1	0.0625	0.0789	61.38%
L1 + Ours	0.0592	0.0759	64.49%

TABLE IV
COMPARISON WITH DIFFERENT VALUES
OF γ_{con} FOR BMD ESTIMATION

Method	MAE↓	RMSE↓	R ² ↑
$\gamma_{reg} = 1, \gamma_{con} = 0.0025$	0.0603	0.0765	63.93%
$\gamma_{reg} = 1, \gamma_{con} = 0.0050$	0.0592	0.0759	64.49%
$\gamma_{reg} = 1, \gamma_{con} = 0.0075$	0.0596	0.0763	64.16%

AdaCon using three different regression losses: L1, MSE, and Huber loss [21]. We set the delta parameter for Huber loss at 0.05, which is close to the MAE of our regression baselines.

We can see from Table III that adding AdaCon consistently boosts performance due to the additional constraints on feature representations introduced through contrastive learning. Relative improvements in MAE of 5.3%-11.1% can be seen with the loss convergence greatly benefiting from the multi-task prediction.

5) *Ablation With Different Weight Parameters*: During training, we fix $\gamma_{reg} = 1$ and set γ_{con} such that the weighted loss values for both tasks are within similar magnitudes after one epoch. We perform ablation analysis on the weight parameter γ_{con} using values of 0.0025, 0.0050, and 0.0075 to show that performance is robust to slight differences in chosen value. Results are shown in Table IV.

Using $\gamma_{con} = 0.0050$ and $\gamma_{con} = 0.0075$ for training gives us similar results. Prediction performance using $\gamma_{con} = 0.0025$, where less importance is given to contrastive learning, is slightly worse but still outperforms Zheng et al. [17] (0.0603 s 0.0612). Overall, AdaCon achieves superior results for similar magnitudes of γ_{con} .

6) *Ablation Under Different Training Schemes*: We examine the use of different training schemes on model performance to justify our choice of multi-task training. We compare with a two-stage training scheme consisting of pretraining with contrastive learning and fine-tuning with regression. Pretraining with only AdaCon loss was done using SGD with a learning rate of 10^{-2} for a total of 3,000 iterations and learning rate decay of factor 0.1 at 1,500 and 2,250 iterations. Fine-tuning using L1 loss was done with a learning rate of 10^{-3} for a total of 3,000 iterations and learning rate decay of 0.1 at 1,500 and 2,250 iterations. Results are shown in Table V.

We can see that two-stage training schemes leads to poor convergence due to inadequate network initialization for regression fine-tuning. Multi-task training on the other hand reliably leads to better performance for all cases, with best results seen from using AdaCon with multi-task training.

7) *External Validation*: We validate performance using two additional datasets: one with X-ray images obtained using

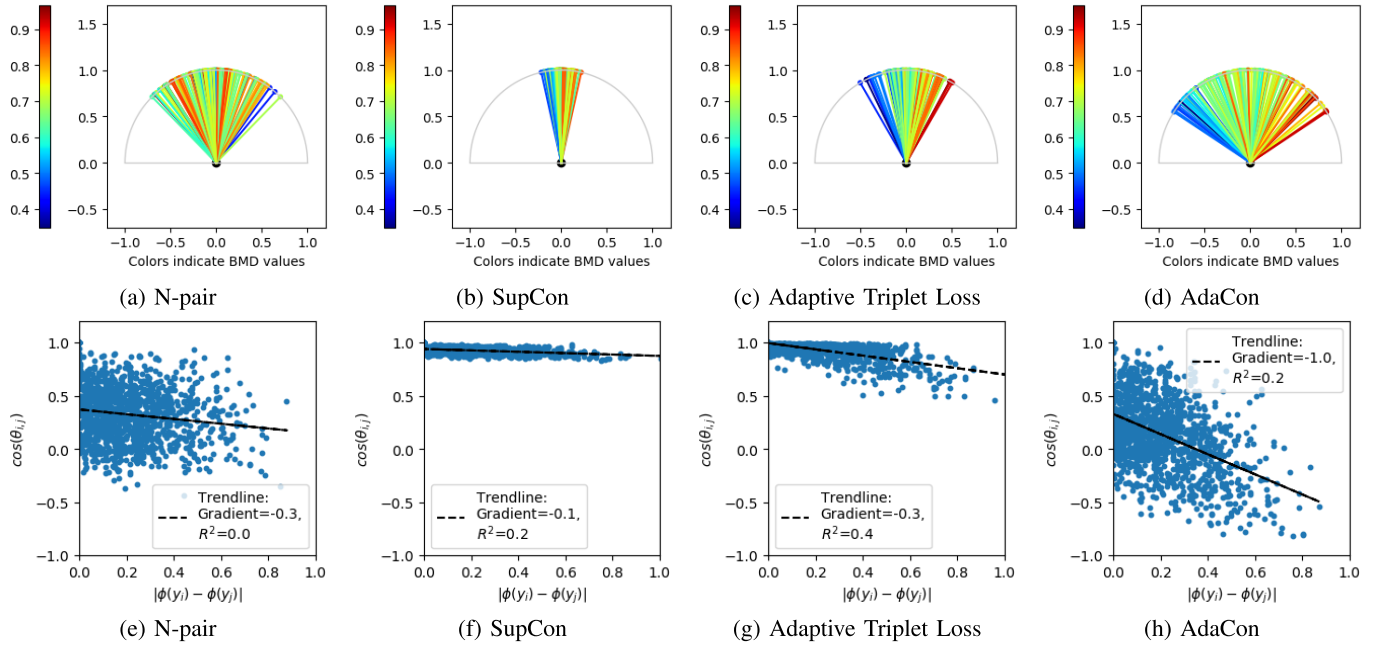


Fig. 3. Plots (a) - (d) are the angular plots of feature projections on the 2d plane. The feature representations of the validation samples are inferred, and the two samples that give the minimum cosine similarity value are used endpoints. The remaining features are distributed based on cosine similarity with the endpoints. Blue (red) color represents low (high) BMD. The wider and more orderly the spread of features from low to high BMD, the better the results. Plots (e) - (h) are scatter plots for randomly sampled pairs of feature vectors, i and j . The x axis is the label distance after applying ECDF; the y axis is feature similarity. The more distinct the downward trend in the scatter plot, the better the results. Fig. 3d and Fig. 3h clearly show AdaCon is able to learn the best feature representations for the BMD task.

TABLE V

RESULTS USING DIFFERENT TRAINING SCHEMES FOR BMD PREDICTION

Method	MAE↓	RMSE↓	R^2 ↑
N-Pair (two-stage)	0.0864	0.1140	19.76%
N-Pair (multi-task)	0.0637	0.0816	58.70%
SupCon (two-stage)	0.0927	0.1207	9.13%
SupCon (multi-task)	0.0629	0.0794	61.15%
Adaptive Triplet (two-stage)	0.0691	0.0874	52.58%
Adaptive Triplet (multi-task)	0.0601	0.0766	63.70%
Ours (two-stage)	0.0964	0.1265	1.38%
Ours (multi-task)	0.0592	0.0759	64.49%

Fujifilm machines from Queen Mary Hospital (Fujifilm dataset), and another with X-ray images obtained from Chang Gung Memorial Hospital in Taiwan (CGMH dataset). The X-ray images were paired with ground truth BMD readings taken using Hologic - Horizon A machines.

a) *Fujifilm dataset*: We perform validation using X-ray images taken by Fujifilm machines, an out of sample machine model, to demonstrate generalization across different machine types. The dataset consists of 19 X-rays from 19 patients collected from Queen Mary Hospital. We use the same pre-processing protocol as detailed in Section IV-A.1. Summary statistics are shown in Table XVI of Appendix C. We use the trained models from our four-fold cross validation experiments in Section IV-A.2 to run inference, taking the average output of the four models as our final prediction. Results are shown in Table VI.

AdaCon (ours) gives the best performance on this validation dataset, which is consistent with our results in Table I. MAE for AdaCon is 0.0071 lower than the methodology in

TABLE VI

TEST RESULTS FOR BMD ESTIMATION USING FUJIFILM X-RAY MACHINES

Method	MAE↓	RMSE↓	R^2 ↑
Hsieh <i>et al.</i> [20]	0.0684	0.0888	53.43%
Zheng <i>et al.</i> [17]	0.0651	0.0870	55.39%
Ours	0.0580	0.0814	60.98%

Zheng *et al.* [17], representing a relative decrease of 10.9%. We also see that the performance metrics are within similar range of the values in Table I (0.0580 s 0.0592). The results demonstrate that our trained model generalizes well across different X-ray machine manufacturers and that AdaCon learns more relevant features compared to alternative methods.

b) *CGMH dataset*: We perform a second external validation using data provided by Chang Gung Memorial Hospital (CGMH) to demonstrate generalization across different hospitals. The dataset consists of 100 X-ray images paired with BMD labels. A number of X-rays were excluded based on our pre-processing protocol as they contained displaced fractures or were not of anteroposterior views as determined by a radiologist. The remaining dataset consisted of 61 X-rays from 61 individual patients. Patient summary statistics are shown in Table XVII of Appendix C. Results are shown in Table VII.

The MAE values based on our methodology is higher than what we obtained in Table I (0.0722 s 0.0592), but the standard deviation of BMD in the CGMH dataset is also larger, as seen in Tables XVII and XIV (0.1568 s 0.1291). Our proposed methodology performs better than both Hsieh *et al.*

TABLE VII

TEST RESULTS FOR BMD ESTIMATION USING CGMH DATASET

Method	MAE↓	RMSE↓	R ² ↑
Hsieh <i>et al.</i> [20]	0.0748	0.0988	59.10%
Zheng <i>et al.</i> [17]	0.0741	0.0980	59.77%
Ours	0.0722	0.0962	61.19%

TABLE VIII

COMPARISON WITH THE STATE-OF-THE-ART METHODS FOR LVEF PREDICTION

Method	MAE↓	RMSE↓	R ² ↑
Ouyang <i>et al.</i> [16]	4.10	5.40	80.5%
Ours	3.86	5.07	82.8%

and Zheng *et al.*, and decreases MAE by 0.0019 compared to Zheng *et al.*, a relative improvement of 3%. These results show that our method generalizes across datasets from different hospitals.

B. LVEF Prediction

1) Experimental Setting:

a) *Dataset*: We use the EchoNet-Dynamic dataset, which is a public dataset consisting of 10,036 apical-4-chamber echocardiogram videos collected from Stanford University Hospital [16]. The dataset is paired with left-ventricular ejection fraction (LVEF), end-systolic volume (ESV) and end-diastolic volume (EDV) labels. The echocardiogram videos are collected using iE33, Sonos, Acuson SC2000, Epiq 5G, and Epiq 7C ultrasound machines. Summary statistics on the dataset are shown in Table XVIII of Appendix C. Additional details of the dataset can also be found in [72]. The data has already been pre-processed by removing irrelevant text information and re-scaled to 112×112 pixels. Training, validation, and test splits have already been prepared. Frame jitter is applied as part of the augmentation process.

b) *Training details*: We follow the original model implementation and training scheme specified in [16] and add our contrastive prediction branch after the feature layer. The backbone uses a R2+1D ResNet [19] pretrained on Kinetics-400 [73] and the entire model is trained with SGD using a learning rate of 10^{-4} and momentum of 0.9 over 45 epochs, with decay of factor 0.1 every 15 epochs. MSE is used as the regression loss function. Clips of 32 frames, sampled from the video at a rate of 1 in every 2 frames, were used as the model input with a batch size of 20. Due to the large memory requirements, we do not add additional batch augmentation operations. We use the validation set to choose the best trained model over all epochs and report results on the test splits. Each experiment was run on three V100 GPUs for approximately 12 hours.

2) *Comparison With State-of-the-Art Methods*: State-of-the-art performance on the Echonet-Dynamic dataset is achieved by Ouyang *et al.* [16], where sampled frames are regressed directly on LVEF values using a R2+1D backbone and MSE loss. We compare our performance using AdaCon with their results in Table VIII.

We can see that AdaCon out-performs Ouyang *et al.*'s methodology and decreases MAE by 0.24, a relative improvement of 5.9%. LVEF is commonly used as an indicator to

TABLE IX

COMPARISON OF DIFFERENT REPRESENTATION LEARNING LOSS FUNCTIONS FOR LVEF PREDICTION

Method	MAE ↓	RMSE ↓	R ² ↑
N-Pair	4.19	5.66	78.6%
Regression (baseline)	4.10	5.40	80.5%
SupCon	4.05	5.37	80.7%
Adaptive Triplet	3.92	5.21	81.8%
Ours	3.86	5.07	82.8%

TABLE X

COMPARISON WITH DIFFERENT REGRESSION LOSSES FOR LVEF PREDICTION

Method	MAE ↓	RMSE ↓	R ² ↑
L1	4.51	6.22	74.1%
L1 + Ours	4.46	6.08	74.8%
Huber ($\delta = 4$)	4.27	5.80	77.5%
Huber ($\delta = 4$) + Ours	4.24	5.74	77.9%
MSE	4.10	5.40	80.5%
MSE + Ours	3.86	5.07	82.8%

diagnose heart conditions [10], and more accurate predictions using our methodology can lead to more reliable tracking and diagnosis of cardiac disease.

3) *Comparison With Other Representation Learning Losses*: Similar to the BMD estimation task, we compare our proposed loss function against alternatives to demonstrate its effectiveness. We compare against N-pair loss [50], SupCon loss [28] adaptive triplet loss [17], and a baseline where only regression loss is used. Results are reported in Table IX.

We see that using N-Pair loss leads to reduced performance when compared to baseline, whilst SupCon and adaptive triplet loss decreased MAE by 0.05 and 0.18 respectively. The largest performance gain occurs when using AdaCon, achieving MAE of 3.86, or an improvement of 0.24 over baseline.

We visualize the learned feature representations in Fig. 4 and see that by using AdaCon for feature representation learning, the feature vectors are able to achieve an even and orderly spread in terms of their angular similarity (Fig. 4d). A clear downward trend can also be seen between pair-wise feature cosine similarity and label distance, indicating the learned features successfully reflects label order and distance in terms of their similarity values.

On the other hand, we can see from Fig. 4a and Fig. 4e that there is little relation between feature similarity and label distance in the representations learned using N-pair loss. For SupCon, we see a downward-sloping trendline in Fig. 4f, but the relationship is not very strong. The features learned using adaptive triplet loss are well spread out in Fig. 4c, but similarity values are clustered around 1 and -1 with little observations in-between (Fig. 4g). Features learned by these alternative loss functions therefore do not reflect label relationships as well as AdaCon.

4) *Ablation Under Different Regression Losses*: We demonstrate that applying AdaCon consistently boosts performance independent of different choices for the regression prediction branch. We use three different regression loss functions: MSE, L1, and Huber Loss and present the results with and without using AdaCon in Table X. We set the delta value for Huber loss at 4, which is close to the MAE of our regression baselines.

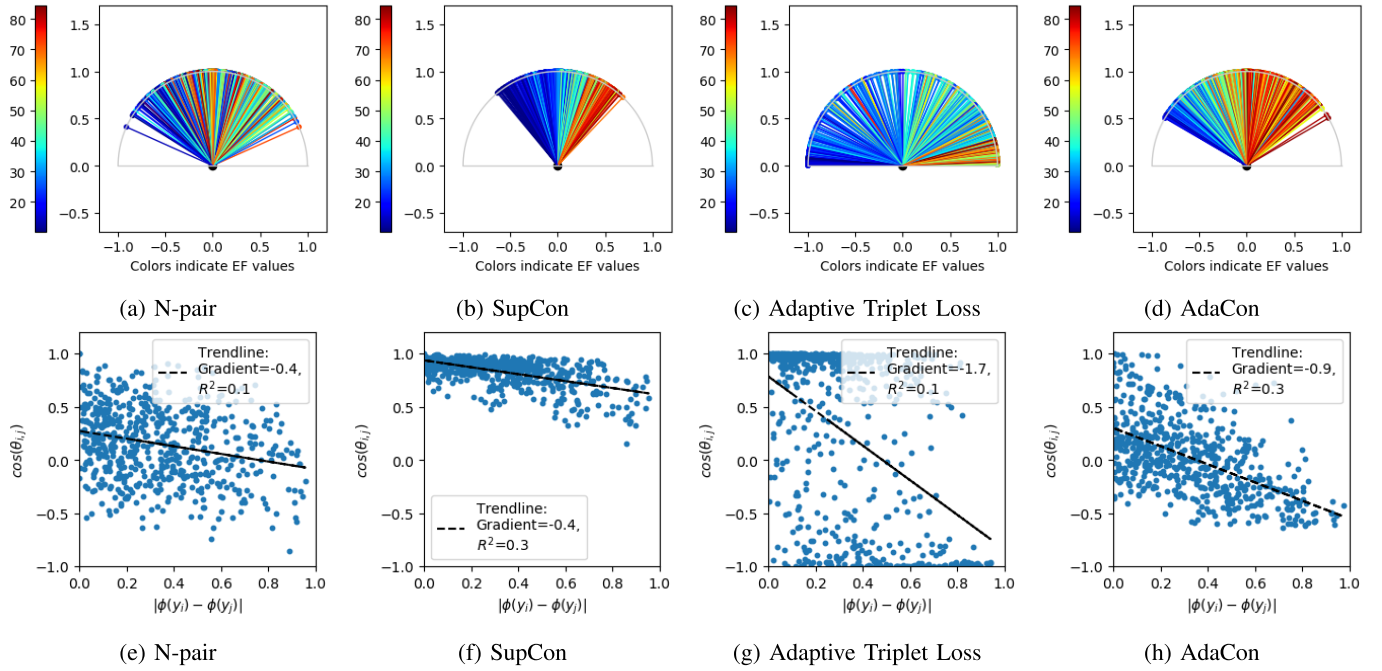


Fig. 4. Plots (a) - (d) are the angular plots of feature projections on the 2d plane. Blue (red) color represents low (high) LVEF values. Plots (e) - (h) are scatter plots for randomly sampled pairs of feature vectors, i and j . The x axis is the label distance after applying ECDF; the y axis is feature similarity. Fig. 4d and Fig. 4h clearly show AdaCon is able to learn the best feature representations for the LVEF task.

TABLE XI
COMPARISON WITH DIFFERENT VALUES OF γ_{con}
FOR LVEF PREDICTION

Method	MAE↓	RMSE↓	R^2 ↑
$\gamma_{reg} = 1, \gamma_{con} = 0.50$	3.95	5.20	82.0%
$\gamma_{reg} = 1, \gamma_{con} = 0.75$	3.86	5.07	82.8%
$\gamma_{reg} = 1, \gamma_{con} = 1.00$	3.89	5.15	82.2%

Similar to what was observed for BMD estimation, the use of our AdaCon framework consistently boosts performance. Relative improvements in MAE of between 0.7% - 5.9% can be seen independent of the regression loss used. The improvement in feature representation learning from the contrastive prediction task reliably improves performance on the regression branch.

5) *Ablation With Different Weight Parameters*: We perform ablation analysis on the weight parameter γ_{con} using values of 0.50, 0.75, and 1.00 and show results in Table XI.

The performance metrics are within the same range for all three weight values and outperform the methodology in Ouyang *et al.* [16], demonstrating robustness to values of γ_{con} within similar magnitudes.

6) *Ablation Under Different Training Schemes*: Pretraining was done by training the model with only AdaCon loss for 25 epochs, using a learning rate of 10^{-4} , momentum of 0.9, and decay of factor 0.1 at epochs 10 and 20. The model is then fine-tuned with MSE loss by training for 25 additional epochs, starting with a learning rate of 10^{-5} and decaying with factor 0.1 at epochs 10 and 20.

We see from Table XII that multi-task training is consistently more effective than two-stage training. Although the difference between the two schemes is smaller than as seen

TABLE XII
RESULTS USING DIFFERENT TRAINING SCHEMES
FOR LVEF PREDICTION

Method	MAE↓	RMSE↓	R^2 ↑
N-Pair (two-stage)	4.48	6.06	75.4%
N-Pair (multi-task)	4.19	5.66	78.6%
SupCon (two-stage)	4.46	6.06	75.5%
SupCon (multi-task)	4.05	5.37	80.7%
Adaptive Triplet (two-stage)	4.33	5.86	77.1%
Adaptive Triplet (multi-task)	3.92	5.21	81.8%
Ours (two-stage)	4.15	5.52	79.6%
Ours (multi-task)	3.86	5.07	82.8%

with the BMD dataset (Table V), the two-stage training still fails to out-perform the plain regression baseline. Overall, the best performance is observed when using AdaCon with multi-task training.

7) *External Validation*: We perform external validation on the CAMUS echocardiogram sequence dataset [74] to demonstrate generalization of LVEF prediction across hospitals. The CAMUS dataset contains a series of echocardiogram sequences taken from 500 patients at the University Hospital of St Etienne in France and is split into a training set of 450 patients and a test set of 50 patients. The dataset is acquired using GE Vivid E95 ultrasound scanners and consists of a 2-chamber sequence view and a 4-chamber sequence view for each patient. Manual segmentations and volume labels for the EDV and ESV frame are provided for the training set only. Sequences are also classified into three classes of quality: good, medium, or poor. The dataset was originally designed for a segmentation challenge using the EDV and ESV frames, but the sequence data and labels can be easily adapted to perform LVEF prediction.

TABLE XIII
TEST RESULTS FOR LVEF PREDICTION ON CAMUS
DATASET BASED ON QUALITY

Quality	Method	MAE↓	RMSE↓	R ² ↑
Good	Ouyang <i>et al.</i> [16]	6.18	7.91	49.9%
	Ours	5.95	7.41	56.1%
Medium	Ouyang <i>et al.</i> [16]	7.96	10.33	38.1%
	Ours	7.41	9.60	46.5%
Poor	Ouyang <i>et al.</i> [16]	11.20	13.09	-36.6%
	Ours	9.68	11.96	-14.2%
Overall	Ouyang <i>et al.</i> [16]	7.23	9.34	40.10%
	Ours	6.78	8.68	48.22%

We evaluate performance of our trained models using the 4-chamber training set sequences. Table XIX in Appendix C shows summary statistics of the dataset. In addition to differences in acquisition machine, the sequences in the CAMUS dataset also contain much fewer frames on average (20 frames *s* 175 frames). We resize all sequences to 112×112 pixels and use the same parameters and data processing pipeline. We show results on overall performance as well as performance categorized by sequence quality in Table XIII.

Overall performance shows that AdaCon (ours) is better than Ouyang *et al.* and decreases MAE by 0.45, a relative improvement of 6%. MAE is higher than on EchoNet-Dynamic results in Table VIII however (6.78 *s* 3.86) because CAMUS sequences have much fewer frames on average. We also see that our methodology performs consistently better in all three quality classes compared to the model trained only using regression loss. The use of our AdaCon framework for improved feature representation learning consistently leads to stronger performance, even across different data sources.

V. DISCUSSION

We demonstrated the effectiveness of our AdaCon loss function through ablation experiments in Tables II and IX. We also discussed several desirable properties of our adaptive margin function in Section III-A. Although we achieve good empirical results, there are still possible limitations to our current formulation that could potentially be improved. The current adaptive margin function imposes a prior assumption on the mapping between label space distance and feature space distance, which may not be the most optimal. It is possible that a learned mapping or one with looser constraints may perform better. Also, the margin function relies on the ECDF for transformation, which implicitly assumes that the sample label distribution of the training set is representative of true label distribution. This may be reasonable for relatively large datasets but may not be the case for smaller datasets or those that have been artificially skewed. Thus, there may exist alternative formulations of the adaptive margin function that lead to better performance. The task of finding a better adaptive margin function is a potential direction for future research.

VI. CONCLUSION

We present a novel adaptive contrastive learning framework, AdaCon, for image regression in computer-aided disease assessment. The adaptive-margin loss function used

TABLE XIV
PATIENT STATISTICS FOR BMD DATASET

Metric	Value
Total number of patients	317
Females (%)	280 (88.0%)
Mean age (sd), years	74.79 (11.95)
Mean BMD (sd), g/cm ²	0.6968 (0.1291)
Average no. of films per patient	1.31

by AdaCon allows us to capture label ordering and probability-normalized distance relationships in our learned features. The features better reflect continuous relationships between sample labels, thus improving performance when applied to regression tasks. The effectiveness of our method has been demonstrated on two medical image regression tasks: BMD estimation from X-ray images and LVEF prediction from echocardiogram videos. Results are also consistent when tested on external validation datasets for both tasks. Our method is general and can be flexibly combined with existing regression loss functions, such as MSE and L1, as well as different feature-extraction backbones. This improved methodology can help increase accuracy and reliability of deep regression tasks for various medical applications.

APPENDIX A DERIVATION OF ADACON LOSS

We wish to incorporate adaptive margins as defined in (3) within the decision boundaries of the standard SupCon loss function, (1), for positive pair classification.

We first consider the case of a batch with three samples: anchor sample i , positive sample p , and negative sample q . Under SupCon loss, the sigmoid function for identifying the positive pair is:

$$\frac{\exp(s \cos(\theta_{i,p}))}{\exp(s \cos(\theta_{i,p})) + \exp(s \cos(\theta_{i,q}))}. \quad (9)$$

The positive pair is correctly classified if:

$$\exp(s \cos(\theta_{i,p})) > \exp(s \cos(\theta_{i,q})), \quad (10)$$

or equivalently:

$$\cos(\theta_{i,p}) > \cos(\theta_{i,q}). \quad (11)$$

We incorporate our constraint in (3) into our decision boundary, such that the positive pair is correctly classified only if $\cos(\theta_{i,p}) > \cos(\theta_{i,q}) + d_{i,q}$. This gives us the modified sigmoid:

$$\frac{\exp(s \cos(\theta_{i,p}))}{\exp(s \cos(\theta_{i,p})) + \exp(s(\cos(\theta_{i,q}) + d_{i,q}))}, \quad (12)$$

and the corresponding cross-entropy loss:

$$-\log \frac{\exp(s \cos(\theta_{i,p}))}{\exp(s \cos(\theta_{i,p})) + \exp(s(\cos(\theta_{i,q}) + d_{i,q}))}, \quad (13)$$

Extending this to an arbitrary number of positive and negative pairs in the sample batch gives us:

$$\sum_{i \in \mathcal{I}} \frac{-1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(s \cos(\theta_{i,p}))}{\sum_{a \in \mathcal{I} \setminus i} \exp(s(\cos(\theta_{i,a}) + d_{i,a}))} \quad (14)$$

where by definition of our adaptive margin in (4), $d_{i,a} = 0$ if $a \in \mathcal{P}(i)$. This gives us our final AdaCon loss function in (6).



Fig. 5. Examples of valid and invalid crops.

APPENDIX B ADACon LOSS WITH ONE POSITIVE AND ONE NEGATIVE PAIR

We show that the AdaCon loss function can be seen as a general formulation for distance metric learning with variable margins and reduces to approximate adaptive triplet loss when using one positive and one negative pair. Making use of Taylor approximations for log and exp, we can see:

$$\mathcal{L}_{con} = -\log \frac{\exp(s(z_i^T z_p + d_{i,p}))}{\exp(s(z_i^T z_p + d_{i,p})) + \exp(s(z_i^T z_n + d_{i,n}))}$$

$$\begin{aligned} &= \log(1 + \exp(s(z_i^T z_n - z_i^T z_p + d_{i,n}))) \\ &\approx \exp(s(z_i^T z_n - z_i^T z_p + d_{i,n})) \\ &\approx 1 + s(z_i^T z_n - z_i^T z_p) + s d_{i,n} \\ &= 1 + s d_{i,n} - \frac{s}{2}(\|z_i - z_n\|^2 - \|z_i - z_p\|^2) \\ &= \frac{s}{2}\{\|z_i - z_p\|^2 - \|z_i - z_n\|^2 + 2d_{i,n}\} + 1 \\ &\propto \|z_i - z_p\|^2 - \|z_i - z_n\|^2 + 2d_{i,n}. \end{aligned}$$

The final line of the derivation approximates the adaptive triplet loss function [17] where the margin value $d_{i,n}$ varies depending on the labels of the sample pairs.

TABLE XV

MANUFACTURER INFORMATION FOR X-RAYS IN BMD DATASET

Manufacturer - Model	Number (% of total)
Canon Inc. - CXDI Control Software NE	176 (37%)
Carestream Health - DRX-EVOLUTION	299 (63%)
Total	475 (100%)

TABLE XVI

PATIENT STATISTICS FOR FUJIFILM DATASET

Metric	Value
Total number of patients	19
Females (%)	14 (73.6%)
Mean age (sd), years	68.6 (14.52)
Mean BMD (sd), g/cm ²	0.7149 (0.1303)

TABLE XVII

PATIENT STATISTICS FOR CGMH DATASET

Metric	Value
Total number of patients	61
Females (%)	48 (78.7%)
Mean age (sd), years	68.95 (12.14)
Mean BMD (sd), g/cm ²	0.7325 (0.1568)

TABLE XVIII

SUMMARY STATISTICS FOR LVEF DATASET

Metric	Value
Total number of patients	10,036
Average number of frames (sd)	175 (57)
Average frames per second (sd)	50.9 (6.8)
Mean LVEF (sd), %	55.75 (12.37)

TABLE XIX

SUMMARY STATISTICS FOR CAMUS DATASET

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

APPENDIX C

DETAILS ON DATASETS

See Tables XIV–XIX and Figure 5.

ACKNOWLEDGMENT

The authors would like to thank Dr. Fang Xin Hao, Benjamin from Queen Mary Hospital, Alistair Yap, and Richard Du for their help in obtaining the BMD dataset. They also thank the Chang Gung Memorial Hospital and Dr. Kuo Chang-Fu for kindly making their data available for external validation use.

REFERENCES

- [1] X. Li, X. Hu, L. Yu, L. Zhu, C.-W. Fu, and P.-A. Heng, "CANet: Cross-disease attention network for joint diabetic retinopathy and diabetic macular edema grading," *IEEE Trans. Med. Imag.*, vol. 39, no. 5, pp. 1483–1493, Nov. 2019.
- [2] H. Xie *et al.*, "Cross-attention multi-branch network for fundus diseases classification using SLO images," *Med. Image Anal.*, vol. 71, Jul. 2021, Art. no. 102031.
- [3] H. Huang *et al.*, "UNet 3+: A full-scale connected UNet for medical image segmentation," in *Proc. ICASSP*, May 2020, pp. 1055–1059.
- [4] E. Gibson *et al.*, "Automatic multi-organ segmentation on abdominal CT with dense V-networks," *IEEE Trans. Med. Imag.*, vol. 37, no. 8, pp. 1822–1834, Aug. 2018.
- [5] K. Kamnitsas, L. Chen, C. Ledig, D. Rueckert, and B. Glocker, "Multi-scale 3D convolutional neural networks for lesion segmentation in brain MRI," *Ischemic Stroke Lesion Segmentation*, vol. 13, p. 46, Oct. 2015.
- [6] K. Kamnitsas *et al.*, "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation," *Med. Image Anal.*, vol. 36, pp. 61–78, Feb. 2017.
- [7] X. Li, H. Chen, X. Qi, Q. Dou, C.-W. Fu, and P.-A. Heng, "H-DenseUNet: Hybrid densely connected UNet for liver and tumor segmentation from CT volumes," *IEEE Trans. Med. Imag.*, vol. 37, no. 12, pp. 2663–2674, Dec. 2018.
- [8] 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 Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 2, pp. 523–534, Feb. 2021.
- [9] E. F. Kranioti, A. Bonicelli, and J. G. García-Donas, "Bone-mineral density: Clinical significance, methods of quantification and forensic applications," *Res. Rep. Forensic Med. Sci.*, vol. 9, pp. 9–21, Jul. 2019.
- [10] M. T. Maeder and D. M. Kaye, "Heart failure with normal left ventricular ejection fraction," *J. Amer. College Cardiol.*, vol. 53, no. 11, pp. 905–918, 2009.
- [11] P. Pulkkinen, T. Jämsä, E.-M. Lochmüller, V. Kuhn, M. T. Nieminen, and F. Eckstein, "Experimental hip fracture load can be predicted from plain radiography by combined analysis of trabecular bone structure and bone geometry," *Osteoporosis Int.*, vol. 19, no. 4, pp. 547–558, Apr. 2008.
- [12] M. Y. El Dib and H. M. Onsi, "Human age estimation framework using different facial parts," *Egyptian Informat. J.*, vol. 12, no. 1, pp. 53–59, Mar. 2011.
- [13] J. Kim, D. Han, S. Sohn, and J. Kim, "Facial age estimation via extended curvature Gabor filter," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2015, pp. 1165–1169.
- [14] A. Saxena *et al.*, "Learning depth from single monocular images," in *Proc. NeurIPS*, vol. 18, 2005, pp. 1–8.
- [15] G. Luo, S. Dong, K. Wang, W. Zuo, S. Cao, and H. Zhang, "Multi-views fusion CNN for left ventricular volumes estimation on cardiac MR images," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 9, pp. 1924–1934, Sep. 2018.
- [16] D. Ouyang *et al.*, "Video-based AI for beat-to-beat assessment of cardiac function," *Nature*, vol. 580, no. 7802, pp. 252–256, 2020.
- [17] K. Zheng *et al.*, "Semi-supervised learning for bone mineral density estimation in hip X-ray images," 2021, *arXiv:2103.13482*.
- [18] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.
- [19] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *Proc. CVPR*, Jun. 2018, pp. 6450–6459.
- [20] C.-I. Hsieh *et al.*, "Automated bone mineral density prediction and fracture risk assessment using plain radiographs via deep learning," *Nature Commun.*, vol. 12, no. 1, pp. 1–9, Dec. 2021.
- [21] P. J. Huber, "Robust estimation of a location parameter," in *Breakthroughs in Statistics*. Springer, 1992, pp. 492–518.
- [22] M. Carvalho, B. L. Saux, P. Trounev-Peloux, A. Almansa, and F. Champagnat, "On regression losses for deep depth estimation," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2018, pp. 2915–2919.
- [23] V. Belagiannis, C. Rupprecht, G. Carneiro, and N. Navab, "Robust optimization for deep regression," in *Proc. ICCV*, Dec. 2015, pp. 2830–2838.
- [24] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. ICML*, 2020, pp. 1597–1607.
- [25] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," 2020, *arXiv:2006.09882*.
- [26] J. Li, P. Zhou, C. Xiong, and S. C. H. Hoi, "Prototypical contrastive learning of unsupervised representations," 2020, *arXiv:2005.04966*.
- [27] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," 2020, *arXiv:2003.04297*.
- [28] P. Khosla *et al.*, "Supervised contrastive learning," 2020, *arXiv:2004.11362*.
- [29] I. Dave, R. Gupta, M. N. Rizve, and M. Shah, "TCLR: Temporal contrastive learning for video representation," 2021, *arXiv:2101.07974*.

- [30] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [31] S. Singh, A. Gupta, and A. A. Efros, "Unsupervised discovery of mid-level discriminative patches," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2012, pp. 73–86.
- [32] C. Doersch, A. Gupta, and A. A. Efros, "Mid-level visual element discovery as discriminative mode seeking," in *Proc. NeurIPS*, 2013, vol. 1, no. 2, p. 3.
- [33] C. Huang, C. C. Loy, and X. Tang, "Unsupervised learning of discriminative attributes and visual representations," in *Proc. CVPR*, Jun. 2016, pp. 5175–5184.
- [34] G. E. Hinton, S. Osindero, and Y. Teh, "A fast learning algorithm for deep belief nets," *Neural Comput.*, vol. 18, no. 7, pp. 1527–1554, Jul. 2006.
- [35] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [36] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*.
- [37] M. Noroozi and P. Favaro, "Unsupervised learning of visual representations by solving jigsaw puzzles," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 69–84.
- [38] S. Gidaris, P. Singh, and N. Komodakis, "Unsupervised representation learning by predicting image rotations," 2018, *arXiv:1803.07728*.
- [39] L. Zhang, G.-J. Qi, L. Wang, and J. Luo, "AET vs. AED: Unsupervised representation learning by auto-encoding transformations rather than data," in *Proc. CVPR*, Jun. 2019, pp. 2547–2555.
- [40] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proc. CVPR*, Jun. 2020, pp. 9729–9738.
- [41] X. Chen, S. Xie, and K. He, "An empirical study of training self-supervised vision transformers," 2021, *arXiv:2104.02057*.
- [42] Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, and P. Isola, "What makes for good views for contrastive learning?" 2020, *arXiv:2005.10243*.
- [43] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, "Graph contrastive learning with augmentations," in *Proc. NeurIPS*, vol. 33, 2020, pp. 5812–5823.
- [44] P. Wang, K. Han, X.-S. Wei, L. Zhang, and L. Wang, "Contrastive learning based hybrid networks for long-tailed image classification," 2021, *arXiv:2103.14267*.
- [45] X. Wang, R. Zhang, C. Shen, T. Kong, and L. Li, "Dense contrastive learning for self-supervised visual pre-training," in *Proc. CVPR*, Jun. 2021, pp. 3024–3033.
- [46] X. Zhao *et al.*, "Contrastive learning for label-efficient semantic segmentation," 2020, *arXiv:2012.06985*.
- [47] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proc. CVPR*, vol. 1, 2005, pp. 539–546.
- [48] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," *J. Mach. Learn. Res.*, vol. 10, no. 2, pp. 207–244, 2009.
- [49] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proc. CVPR*, Jun. 2015, pp. 815–823.
- [50] K. Sohn, "Improved deep metric learning with multi-class n-pair loss objective," in *Proc. NeurIPS*, 2016, pp. 1857–1865.
- [51] B. Yu, T. Liu, M. Gong, C. Ding, and D. Tao, "Correcting the triplet selection bias for triplet loss," in *Proc. ECCV*, 2018, pp. 71–87.
- [52] X. Zhao, H. Qi, R. Luo, and L. Davis, "A weakly supervised adaptive triplet loss for deep metric learning," in *Proc. ICCVW*, Oct. 2019, pp. 1–4.
- [53] X. Zhang *et al.*, "A new window loss function for bone fracture detection and localization in X-ray images with point-based annotation," 2020, *arXiv:2012.04066*.
- [54] J. D. Krogue *et al.*, "Automatic hip fracture identification and functional subclassification with deep learning," *Radiol. Artif. Intell.*, vol. 2, no. 2, Mar. 2020, Art. no. e190023.
- [55] B. Guan, G. Zhang, J. Yao, X. Wang, and M. Wang, "Arm fracture detection in X-rays based on improved deep convolutional neural network," *Comput. Electr. Eng.*, vol. 81, Jan. 2020, Art. no. 106530.
- [56] Y.-J. Huang *et al.*, "Radiographic bone texture analysis using deep learning models for early rheumatoid arthritis diagnosis," Preprint (Version 1), Sep. 2020. [Online]. Available: <https://www.researchsquare.com/article/rs-76193/v1>, doi: [10.21203/rs.3.rs-76193/v1](https://doi.org/10.21203/rs.3.rs-76193/v1).
- [57] J.-T. Lu *et al.*, "Deep Spine: Automated lumbar vertebral segmentation, disc-level designation, and spinal stenosis grading using deep learning," in *Proc. MLHC*, 2018, pp. 403–419.
- [58] A. Sekuboyina, J. Kukačka, J. S. Kirschke, B. H. Menze, and A. Valentiniš, "Attention-driven deep learning for pathological spine segmentation," in *Proc. Int. Workshop Comput. Methods Clin. Appl. Musculoskeletal Imag.* Springer, 2017, pp. 108–119.
- [59] C. Chappard *et al.*, "Prediction of femoral fracture load: Cross-sectional study of texture analysis and geometric measurements on plain radiographs versus bone mineral density," *Radiology*, vol. 255, no. 2, pp. 536–543, May 2010.
- [60] P. Chu *et al.*, "Using octuplet Siamese network for osteoporosis analysis on dental panoramic radiographs," in *Proc. EMBC*, Jul. 2018, pp. 2579–2582.
- [61] F. Wang *et al.*, "Opportunistic screening of osteoporosis using plain film chest X-ray," 2021, *arXiv:2104.01734*.
- [62] L. R. Loehr, W. D. Rosamond, P. P. Chang, A. R. Folsom, and L. E. Chambless, "Heart failure incidence and survival (from the atherosclerosis risk in communities study)," *Amer. J. Cardiol.*, vol. 101, no. 7, pp. 1016–1022, Apr. 2008.
- [63] P. A. Pellikka *et al.*, "Variability in ejection fraction measured by echocardiography, gated single-photon emission computed tomography, and cardiac magnetic resonance in patients with coronary artery disease and left ventricular dysfunction," *JAMA Netw. Open*, vol. 1, no. 4, 2018, Art. no. e181456.
- [64] X. Zhen, Z. Wang, A. Islam, M. Bhaduri, I. Chan, and S. Li, "Direct estimation of cardiac bi-ventricular volumes with regression forests," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Springer, 2014, pp. 586–593.
- [65] M. H. Jafari, N. V. Woudenberg, C. Luong, P. Abolmaesumi, and T. Tsang, "Deep Bayesian image segmentation for a more robust ejection fraction estimation," in *Proc. ISBI*, Apr. 2021, pp. 1264–1268.
- [66] 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," *Med. Image Anal.*, vol. 67, Jan. 2021, Art. no. 101873.
- [67] H. Reynaud, A. Vlontzos, B. Hou, A. Beqiri, P. Leeson, and B. Kainz, "Ultrasound video transformers for cardiac ejection fraction estimation," 2021, *arXiv:2107.00977*.
- [68] A. W. Van der Vaart, *Asymptotic Statistics*, vol. 3. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [69] J. E. Angus, "The probability integral transform and related results," *SIAM Rev.*, vol. 36, no. 4, pp. 652–654, Dec. 1994. [Online]. Available: <http://www.jstor.org/stable/2132726>
- [70] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019, pp. 6105–6114.
- [71] E. Hoffer, T. Ben-Nun, I. Hubara, N. Giladi, T. Hoefler, and D. Soudry, "Augment your batch: Better training with larger batches," 2019, *arXiv:1901.09335*.
- [72] D. Ouyang *et al.*, "EchoNet-Dynamic: A large new cardiac motion video data resource for medical machine learning," in *Proc. NeurIPS*, 2019, pp. 1–11.
- [73] W. Kay *et al.*, "The kinetics human action video dataset," 2017, *arXiv:1705.06950*.
- [74] S. Leclerc *et al.*, "Deep learning for segmentation using an open large-scale dataset in 2D echocardiography," *IEEE Trans. Med. Imag.*, vol. 38, no. 9, pp. 2198–2210, Sep. 2019.