

ACCon: Angle-Compensated Contrastive Regularizer for Deep Regression

Botao Zhao^{*}, Xiaoyang Qu^{*}, Zuheng Kang, Junqing Peng, Jing Xiao, Jianzong Wang[†]

Ping An Technology (Shenzhen) Co., Ltd.
 jzwang@188.com, zhao_bot@163.com,
 {quxiaoyang343, kangzuheng896, pengjq, xiaojing661}@pingan.com.cn

Abstract

In deep regression, capturing the relationship among continuous labels in feature space is a fundamental challenge that has attracted increasing interest. Addressing this issue can prevent models from converging to suboptimal solutions across various regression tasks, leading to improved performance, especially for imbalanced regression and under limited sample sizes. However, existing approaches often rely on order-aware representation learning or distance-based weighting. In this paper, we hypothesize a linear negative correlation between label distances and representation similarities in regression tasks. To implement this, we propose an angle-compensated contrastive regularizer for deep regression, which adjusts the cosine distance between anchor and negative samples within the contrastive learning framework. Our method offers a plug-and-play compatible solution that extends most existing contrastive learning methods for regression tasks. Extensive experiments and theoretical analysis demonstrate that our proposed angle-compensated contrastive regularizer not only achieves competitive regression performance but also excels in data efficiency and effectiveness on imbalanced datasets.

Introduction

Regression, a fundamental machine learning task, is employed when the prediction target is continuous. In the last ten years, deep regression has emerged as a more effective approach than traditional regression methods across various domains, including computer vision (Wang, Sanchez, and Li 2022), natural language process (Chandrasekaran and Mago 2021). Recent studies in deep regression have predominantly focused on developing network structures for specific application scenarios (Lee et al. 2021; Chen, Ma, and Lin 2021; Tabelini et al. 2021). However, these task-specific architectures are difficult to adapt to different tasks or domains, which remains a non-trivial research question.

A in-depth examination of deep regression by (Lathuilière et al. 2019) revealed that a well-tuned general-purpose network can achieve results close to the state-of-the-art (SOTA) models, potentially obviating the need for more intricate and specialized regression models. In a standard deep regression

learning (SDL) task, the representation is first extracted by the backbone network and subsequently processed by passing through a fully connected layer to generate the prediction for the target. The mean squared error (MSE) loss (Lathuilière et al. 2019) is broadly acknowledged as the most frequently utilized loss function in deep regression. Additionally, researchers have explored alternative loss functions, including L1 loss, Huber loss (Huber 1992), Tukey loss (Belagiannis et al. 2015). In an SDL, enhancing the quality of learned representations is a key approach to improving performance.

Given the continuous nature of regression targets, recent research in representation learning has increasingly focused on capturing the nuanced relationships within the label space. *These methods can be broadly categorized into: order-aware (Gong, Mori, and Tung 2022; Zha et al. 2024), and distance-aware approaches (Keramati, Meng, and Evans 2024; Dai et al. 2021).* Order-aware methods typically rely on ranking-based techniques to constrain models to obtain order-aware representations. However, these methods have a major flaw (Fig. 1 (A), first row), as regression targets involve not only order information but also distance information (Fig. 1 (A), second row). For example, consider face images with ages 2, 21, 46, and 80; ranking-based methods would generally assign these labels' representations to a scale with equal intervals, which might incorrectly reflect the true relationships among them. Distance-aware methods, on the other hand, aim to maintain the representation similarity among samples proportional to their corresponding label distances. Current distance-aware approaches commonly apply weighted cosine similarity between the anchor and negative pairs within the contrastive learning framework (Keramati, Meng, and Evans 2024; Dai et al. 2021). However, we posit that direct weighting may not be optimal due to the non-linearity of the cosine function. This nonlinearity can lead to unevenly distributed representations on the hypersphere, potentially distorting the relationship between labels and features. *In this paper, we introduce angle compensation as a refinement to direct weighting, aiming to advance distance-aware representation learning in regression tasks through a more nuanced and effective approach.*

To this end, we considered building upon supervised contrastive learning (SupCon), an approach that has seen rapid development with numerous notable contributions (Ermolov et al. 2021; Chen et al. 2020; Jiang et al. 2020). As shown in

^{*}These authors contributed equally to this work.

[†]Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

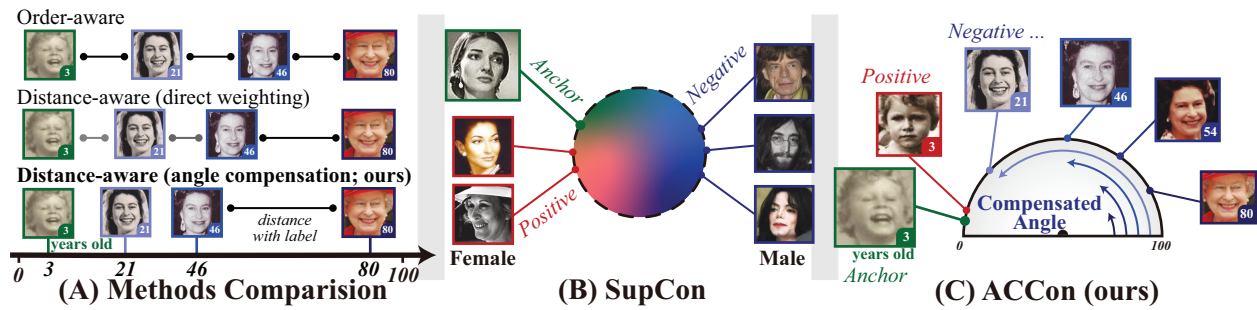


Figure 1: The subfigure (A) illustrates two types of representation learning for regression: order-aware and distance-aware. The subfigure (B) depicts the supervised contrastive learning for gender classification, where all samples from the same class are treated as positives and contrasted against negatives. The subfigure (C) presents the proposed angle-compensated supervised contrastive learning approach for age regression, which projects the input onto a semi-hypersphere while preserving the label relationship information.

Fig. 1 (B), SupCon seeks to minimize the distance between an anchor and a "positive" sample (belonging to the same class) within the embedding space, while also maximizing the distance between the anchor and several "negative" samples (from different classes). However, the application of these methods has been mainly focused on classification and segmentation tasks (Khosla et al. 2020; Sun et al. 2023; Seifi et al. 2024). Extending these approaches to regression tasks presents several challenges, a key issue being the effective handling of varying distances among negative pairs. In regression tasks, for instance, negative samples with labels such as 46 and 80 are expected to exhibit different similarities to an anchor with label 3. To address this gap, we propose an angle-compensated contrastive learning (ACCon) method to fill this gap and achieve effective distance-aware representation learning for regression. Our main contributions are as follows:

- We address the challenge of learning distance-aware representations for deep regression by proposing ACCon, a straightforward yet effective method. ACCon can be integrated into various contrastive learning-based approaches to enhance performance. Specifically, ACCon projects the input onto a semi-hypersphere, preserving the relationships among labels, as illustrated in Fig. 1 (C).
- We provide theoretical proof that optimizing \mathcal{L}_{ACCon} results in features located on the semi-hypersphere according to the distances among labels.
- We perform comprehensive experiments across three datasets from both computer vision and natural language processing fields, showcasing enhanced performance relative to current state-of-the-art techniques.

Related Work

Deep Regression

Many researchers have focused on deep regressions over the past decades. Some researchers formulate regression as a classification problem (Gao et al. 2017; Shi, Cao, and Raschka 2023; Rogez, Weinzaepfel, and Schmid 2017; Pan et al. 2018). However, this formulation fails to fully leverage the relationship among labels and introduces a trade-off between optimization complexity and method performance. Recent

advancements have proposed various methods to address imbalance issues in deep regression (Yang et al. 2021; Ren et al. 2022). Although some applications aim to enhance model performance in balanced settings, such as disease prediction, many certain real-world scenarios require models to perform effectively on imbalanced naturally sampled data, where long-tail samples may be considered outliers and potentially disregarded. Several studies have proposed regularizers to constrain the embedding space, including approaches that model uncertainty within the embedding space (Li et al. 2021) and methods that learn higher-entropy feature spaces (Zhang et al. 2022). Recently, the concept of learning ordered features for regression has gained increased attention, with notable contributions such as Rank-N-Contrast (Zha et al. 2024) and RankSIM (Gong, Mori, and Tung 2022). *In contrast to existing works, we propose a distance-aware representation learning method capable of capturing relationships among continuous labels in feature space.*

Contrastive Learning

Contrastive learning is a representation learning technique that projects features onto a hypersphere (Ermolov et al. 2021; Chen et al. 2020; Jiang et al. 2020). The fundamental principle of contrastive learning involves attracting positive sample pairs while repelling negative ones. In self-supervision, positive sample pairs are generally formed by different augmentations of the same sample, while negative pairs are made up of the anchor and other samples within the minibatch. SupCon, an extension of contrastive learning in fully-supervised settings, has recently emerged as a promising approach, demonstrating significant advancements in image recognition and various classification tasks (Chen et al. 2020; Sun et al. 2023; Khosla et al. 2020; Liu et al. 2023; Seifi et al. 2024). In the SupCon framework, positive pairs are defined as samples from the same class, while negative pairs are samples from different classes. However, the adaptation of contrastive learning to regression tasks presents several challenges, particularly in the definition of positive and negative pairs and the capture of relationships within the label space. (Zha et al. 2024) proposed a redefinition of positive and negative pairs to achieve order-aware representation learning for regression, while (Dai et al. 2021) exploited a distance-weighted

negative cosine similarity to achieve distance-aware representations. (Keramati, Meng, and Evans 2024) introduced a novel approach that defines negative pairs as samples with dissimilar labels but similar predictions to the anchors. By combining this definition with distance-weighted negative cosine similarity, their method demonstrated robust performance in imbalanced regression tasks. *In contrast to these existing approaches, we refine distance-weighted negative cosine similarity by introducing angle compensation, striving to enhance distance-aware representation learning in regression tasks.*

Methodology

Problem Setting

In this paper, we consider the input $(\mathcal{X}, \mathcal{Y}) = \{x_i, y_i\}_{i=1}^N$. Similar to contrastive learning, our objective is to develop a feature representation network that could map the input x_i to an L2-normalized d-dimensional embedding, $z_i \in \mathcal{S}^{d-1}$. Furthermore, we assume that the learned representation $z \in Z$ should be distributed on the \mathcal{S}^{d-1} such that each position corresponds to its label, as illustrated in Fig. 1 (C).

Given the continuous nature of labels, we first partition the label space \mathcal{Y} into M bins with equal intervals, denoted as $[y_0, y_1), [y_1, y_2), \dots, [y_{M-1}, y_M)$. These bins represent the precision of the labels, which can be adjusted according to practical requirements.

We categorize samples from the same bin as positive pairs, while those from different bins are considered negative pairs. Specifically, for an anchor x_i , the set of the positive pairs is denoted as $\mathcal{P}(i) := \{j \in \mathcal{B} | y_i = y_j, j \neq i\}$, and the negative pairs are defined as $\mathcal{N}(i) := \{j \in \mathcal{B} | y_i \neq y_j\}$. Subsequently, we can obtain a naïve extension of supervised contrastive learning loss function (Dai et al. 2021; Keramati, Meng, and Evans 2024; Khosla et al. 2020) for regression tasks as follows:

$$\mathcal{L}_i = -\frac{1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(\cos(\theta_{i,p})/\tau)}{\sum_{k \in \mathcal{N}(i) \cup \mathcal{P}(i)} \exp(\cos(\theta_{i,k})/\tau)}, \quad (1)$$

where, $\theta_{i,k}$ denotes the angle between embedding z_i and z_k , and $\cos(\theta_{i,k}) = z_i z_k^T$. However, this formulation is unsuitable for regression tasks as it fails to capture the inherent relationships among labels. As illustrated in Fig. 1 (C), both samples with labels 46 and 80 are treated equivalently as negatives for an anchor with label 3. Yet, the dissimilarity between the anchor 3 and label 80 is significantly greater than that between the anchor and label 46. Furthermore, due to the inherent monotonic nature of labels in regression tasks, mapping x_i to a complete hypersphere faces great challenges. Therefore, our objective is to *develop a model that maps the learned representation z distributed onto a semi-hypersphere, positioning them in accordance with their respective labels.*

Angle-Compensated Supervised Contrastive Loss

We propose an angle-compensated supervised contrastive loss to achieve our goal of preserving label relationships in feature space. Firstly, we hypothesize that a linear negative correlation exists between label distances and representation

similarities, implying that samples' representations should be positioned in alignment with their respective labels. Based on this hypothesis, if we use cosine distance to measure representation similarities, we can derive the ideal angle, $\hat{\theta}$ between anchor and negatives as follows:

$$\hat{\theta} = \frac{y_{\text{neg}} - y_{\text{anc}}}{\max(\mathcal{Y}) - \min(\mathcal{Y})} \pi, \quad (2)$$

where $y_{\text{anc}}, y_{\text{neg}}$ denote the label of anchor and negative sample, respectively. For instance, given an anchor with a label of 21 and an age range from 0 to 100, the angles between the anchor and negative samples (3, 26, 54, 80) would be $-5.4^\circ, 32.4^\circ, 41.4^\circ, 91.8^\circ$ and 138.6° , respectively.

Several studies have demonstrated that the representations of each class spontaneously collapse to the vertices of a regular simplex when the standard supervised contrastive loss reaches its minimum (Graf et al. 2021; Zhu et al. 2022). The SupCon loss effectively constrains the representations of anchors and negatives to be as far apart as possible within a minibatch, which is equivalent to constraining the included angle $\tilde{\theta} = \pi$. This observation forms the basis for the formulation of Eq. 2.

$$\cos(\tilde{\theta}) = \cos\left(\hat{\theta} + \pi - \frac{y_{\text{neg}} - y_{\text{anc}}}{\max(\mathcal{Y}) - \min(\mathcal{Y})} \pi\right). \quad (3)$$

Based on the Eq. 1, we will get the angle-compensated supervised contrastive loss for sample i as follows:

$$\mathcal{L}_i^{\text{ac}} = \frac{-1}{|\mathcal{P}(i)|} \sum_{p \in \mathcal{P}(i)} \log \frac{\exp(z_i z_p^T / \tau)}{\left(\sum_{k \in \mathcal{P}(i)} \exp(z_i z_k^T / \tau) + \sum_{m \in \mathcal{N}(i)} \exp(\cos(\tilde{\theta}_{i,m}) / \tau) \right)}, \quad (4)$$

where, $\mathcal{P}(i)$ denotes the set of positive pairs for anchor i , and $\mathcal{N}(i)$ represents the negative set. This formulation is equal to the SupCon loss, with the key distinction that it replaces the negative similarity with an angle-compensated version. Considering a minibatch, the optimization of Eq. 4 aims to drive $\tilde{\theta}_{i,m} \rightarrow \pi$ and make the $\tilde{\theta}_{i,m} \rightarrow \frac{y_{\text{neg}} - y_{\text{anc}}}{\max(\mathcal{Y}) - \min(\mathcal{Y})} \pi$, aligning with our intended goal. Then, the angle-compensated supervised contrastive loss $\mathcal{L}_{\text{ACCOn}}$ is defined as the mean value of $\{\mathcal{L}_i^{\text{ac}}\}_{i=1}^{2N}$:

$$\mathcal{L}_{\text{ACCOn}} = \frac{1}{2N} \sum_i \mathcal{L}_i^{\text{ac}} \quad (5)$$

where N is the batch size.

To compute the $\mathcal{L}_i^{\text{ac}}$, we should determine $\cos(\tilde{\theta}_{i,m})$. We firstly define the compensation angle φ based on the Eq. 3:

$$\varphi = \pi \left(1 - \frac{y_{\text{neg}} - y_{\text{anc}}}{\max(\mathcal{Y}) - \min(\mathcal{Y})} \right). \quad (6)$$

Subsequently, we derive $\cos(\tilde{\theta}_{i,m})$ as follows:

$$\cos(\tilde{\theta}_{i,m}) = z_i z_m^T \cos(\varphi) - |\sin(\varphi)| \sqrt{1 - (z_i z_m^T)^2} + \epsilon, \quad (7)$$

where ϵ is the smoothing term introduced to prevent gradient explosion. A detail derivation of $\cos(\tilde{\theta}_{i,m})$ is provided in Appendix A1.




Figure 2: The frameworks of the deep regression with angle-compensated supervised contrastive.

Theoretical Insights

In this section, we theoretically analyzed the efficiency of $\mathcal{L}_{\text{ACCOn}}$. Firstly, we have derived a lower bound of $\mathcal{L}_{\text{ACCOn}}$ in Appendix A2.

Theorem 1 (Lower bound of $\mathcal{L}_{\text{ACCOn}}$). $L^* := \frac{1}{4N^2} \sum_{i=0}^{2N} \sum_{m \in \mathcal{N}(i)} \cos(\tilde{\theta}_{i,m}) / \tau - \frac{\log N / \tau}{2N}$ is a lower bound of $\mathcal{L}_{\text{ACCOn}}$, i.e., $\mathcal{L}_{\text{ACCOn}} > L^*$.

Minimizing the $\mathcal{L}_{\text{ACCOn}}$, will lead to the minimization of L^* . Optimizing L^* results in the convergence of $\hat{\theta}$ to $\frac{y_{\text{neg}} - y_{\text{anc}}}{\max(\mathcal{Y}) - \min(\mathcal{Y})} \pi$, which aligns with our initial hypothesis.

Deep Regression Framework

The proposed angle-compensated supervised contrastive loss integrates seamlessly into traditional supervised contrastive learning frameworks (Khosla et al. 2020), as demonstrated in Fig. 2, which depicts the two-view condition. Our method is inherently scalable, allowing for the adjustment of data augmentations according to the specific requirements of the learning task. In detail, the deep regression framework consists of four core components:

- Data augmentation module, $\text{Aug}(\cdot)$: This augments the input x to \hat{x}_1 and \hat{x}_2 , randomly, with each augmentation providing a distinct view of the input data.
- Feature encoder, $f_\theta(\cdot)$: This component maps the augmented input to an embedding, denoted as $z = f_\theta(\hat{x})$. The architectural design of the encoder may vary depending on the specific application scenario.
- Projection layer, $g_\phi(\cdot)$: This layer projects the embedding e to a representation vector $\hat{z} = g_\phi(z) \in \mathbb{R}^{d_l}$. Typically, $g_\phi(\cdot)$ is implemented by a single linear layer, followed by L_2 normalization to ensure $\hat{z} \in \mathcal{S}^{d_l-1}$.
- Predictor, $h_\psi(\cdot)$: The predictor maps the representation \hat{z} to the label \hat{y} utilizing a single linear layer without bias in our configuration.

After applying these modules, we obtain the extracted embedding \hat{z} and the predicted target \hat{y} . Subsequently, our model is trained through the simultaneous optimization of the regression loss function \mathcal{L}_{reg} and the angle-compensated supervised contrastive loss $\mathcal{L}_{\text{ACCOn}}$. The final loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{reg}} + \gamma \cdot \mathcal{L}_{\text{ACCOn}}, \quad (8)$$

where γ denotes the weight coefficient. In this study, we employ the mean square error (MSE) as \mathcal{L}_{reg} for the sentence

Algorithm 1: Training of the deep regression with ACCOn.

Input: The minibatch data $(X, Y) \subset (\mathcal{X}, \mathcal{Y})$, Feature encoder f_θ , Projection layer g_ϕ , Predictor h_ψ .

Parameter: The temperature index, τ ; Weight coefficient, γ ; Max training epoch, T ; Smoothing term, ϵ .

Output: θ^*, ϕ^*, ψ^*

- 1: Let $t = 0$.
- 2: Randomly initialize θ_0, ϕ_0, ψ_0
- 3: **while** $t < T$ **do**
- 4: draw two augmentation functions $\text{Aug}_1(\cdot), \text{Aug}_2(\cdot)$ {Forward process}
- 5: $\hat{X}^c \leftarrow \text{cat}([\text{Aug}_1(X), \text{Aug}_2(X)], \text{dim} = 0)$
- 6: $E \leftarrow f_\theta(\hat{X}^c)$
- 7: $Z \leftarrow g_\phi(E)$
- 8: $\tilde{Z} \leftarrow Z / \|Z\|$ {Calculation of loss function}
- 9: $\tilde{Y} \leftarrow \text{cat}([\hat{Y}, Y], \text{dim} = 0)$
- 10: Calculate φ by Eq. 6
- 11: Calculate $\cos(\hat{\theta})$ based on Eq. 7
- 12: Calculate $\mathcal{L}_{\text{ACCOn}}$ based on Eq. 4 and Eq. 5
- 13: Calculate the \mathcal{L}_{reg} between $h_\psi(Z)$ and Y
- 14: Calculate loss function \mathcal{L} by Eq. 8 {Optimization}
- 15: Calculate $\theta_{t+1}, \phi_{t+1}, \psi_{t+1}$ by gradient descent of \mathcal{L}
- 16: $t = t + 1$
- 17: **end while**
- 18: $\theta^*, \phi^*, \psi^* \leftarrow \theta_T, \phi_T, \psi_T$
- 19: **return** θ^*, ϕ^*, ψ^*

similarity prediction task and the mean absolute error (MAE) as \mathcal{L}_{reg} for the age estimation task. The comprehensive training process is delineated in Alg. 1.

Experiments and Results

Experiments Setup

Datasets: To rigorously evaluate our method, we selected three diverse datasets: (1) **AgeDB** (Moschoglou et al. 2017): An age estimation dataset comprising 16,488 facial images; (2) **IMDB-WIKI** (Rothe, Timofte, and Van Gool 2018): A large-scale facial age dataset containing 523,000 images with corresponding age labels; (3) **STS-B** (Cer et al. 2017; Wang et al. 2018): A natural language dataset consisting of 7,249 sentence pairs, extracted from the Semantic Textual Similarity (STS) Benchmark. To ensure a comprehensive assessment of our proposed method, we employed 2 distinct sampling strategies for partitioning the datasets into training, validation, and test sets:

(1) *Balanced Sampling*: Following the benchmark established by (Yang et al. 2021), we partitioned the AgeDB dataset into 12,208 training samples, 2,140 validation samples, and 2,140 test samples. For IMDB-WIKI, we allocated 191,500 images for training and 11,000 images each for validation and testing. From STS-B, we sampled 1,000 pairs each for validation and testing. This benchmark ensures a balanced label distribution in the validation and test sets, as illustrated in Appendix Figure 1. Similar to Yang et al.’s benchmark, we denote these balanced datasets as AgeDB-DIR, IMDB-WIKI-DIR, and

Methods	AgeDB-Natural			STS-B-Natural			IMDB-WIKI-Natural		
	MAE ↓	G-means ↓	R ² ↑	MAE ↓	G-means ↓	Pearson ↑	MAE ↓	G-means ↓	R ² ↑
Vanilla	7.044	4.449	0.715	0.798	0.516	0.756	5.642	3.251	0.646
NaïveSupCon (Khosla et al. 2020)	7.070	4.549	0.718	0.788	0.518	0.764	5.542	3.221	0.642
AdaSupCon (Dai et al. 2021)	6.973	4.438	0.725	0.775	0.503	0.763	5.535	3.210	0.652
BMSE (Ren et al. 2022)	6.999	4.384	0.721	/	/	/	5.589	3.235	0.638
RankSim (Gong, Mori, and Tung 2022)	6.914	4.404	0.724	0.786	0.519	0.756	5.481	3.198	0.654
RNC (Zha et al. 2024)	6.793	4.298	0.732	0.782	0.512	0.762	5.483	3.182	0.656
ConR (Keramati, Meng, and Evans 2024)	6.808	4.324	0.733	0.795	0.517	0.760	5.492	3.199	0.655
ACCon (Ours)	6.724	4.245	0.741	0.724	0.469	0.791	5.432	3.157	0.668
Ours vs. Vanilla	↑4.54%	↑4.59%	↑3.64%	↑9.27%	↑9.11%	↑4.63%	↑3.72%	↑2.89%	↑3.41%
Ours vs. SOTA	↑1.02%	↑1.23%	↑1.08%	↑6.58%	↑6.76%	↑3.53%	↑1.91%	↑1.10%	↑1.83%

Table 1: Performance comparison with other regularizers for regression on the AgeDB-Natural, STS-Natural, and IMDB-WIKI-Natural. We trained each baseline model and the detailed information as described in Appendix B.

STS-B-DIR.

(2) *Natural Sampling*: We used the same dataset splitting ratio but randomized the division of training, validation, and test datasets, thereby preserving similar label distributions across all three sets (Appendix Figure 1). We denote the datasets obtained through natural sampling as AgeDB-Natural, IMDB-WIKI-Natural and STS-B-Natural.

Metrics: For quantitative evaluation, we utilized a diverse set of metrics including Mean Absolute Error (MAE), Geometric Mean (GM) defined as $\left(\prod_{i=1}^n e_i\right)^{\frac{1}{n}}$ where e_i represents the L_1 error for the i^{th} sample, coefficient of determination R^2 ($1 - \frac{MSE(y, \hat{y})}{VAR(y)}$), indicating the proportion of explainable variance, Mean Square Error (MSE), and Pearson correlation. The division of test set into many-shot, medium-shot, and few-shot categories follows the same protocol as outlined in (Yang et al. 2021).

Baselines

We compared our method with state-of-the-art (SOTA) approaches under both natural and balanced sampling conditions. To guarantee a fair evaluation, we utilized standard network architectures consistent with those outlined by (Yang et al. 2021). Specifically, ResNet-50 was utilized as the backbone for age estimation tasks, while BiLSTM+GloVe word embeddings were employed for sentence similarity prediction.

For STS-B-Natural, IMDB-WIKI-Natural, and AgeDB-Natural datasets, our method is compared with, **Vanilla**: using MAE or MSE loss function; **Naïve SupCon**: adopting SupCon for deep regression, as in Eq. 1; **Adaptive SupCon (AdaSupCon)**: proposing adaptive-margin contrastive loss for image regression (Dai et al. 2021); **RankSIM**: exploiting ranking similarity regularization (Gong, Mori, and Tung 2022); **BMSE**: modifying MSE loss for deep imbalanced regression (Ren et al. 2022); **RNC**: an order-aware method (Zha et al. 2024); **ConR**: exploiting contrastive learning for imbalanced regression (Keramati, Meng, and Evans 2024).

For STS-B-DIR, IMDB-WIKI-DIR, and AgeDB-DIR datasets, our approach was benchmarked against Vanilla,

NaïveSupCon, AdaSupCon, RankSIM, and ConR. Additionally, comparisons were made with re-weighting methods such as Label Distribution Smooth (LDS) and Focal-R (Yang et al. 2021), integrated with the baseline described above. Typically, Re-weighting adjusts the loss function to address label imbalance. So, we tested inverse-frequency weighting (**INV**) and its square-root variant (**SQINV**). LDS employs re-weighting based on smoothed label distribution, while Focal-R utilizes a continuous function mapping absolute error to the scaling factor. Detailed experimental setup, including training protocols, dataset specifications, and hyperparameters, are provided in Appendix B.

Basic Performance

We implement the vanilla model and five advanced regularizer methods tailored for regression tasks: NaïveSupCon (Khosla et al. 2020), AdaSupCon (Dai et al. 2021), BMSE (Ren et al. 2022), RankSim (Gong, Mori, and Tung 2022), RNC (Zha et al. 2024) and ConR (Keramati, Meng, and Evans 2024). All models were re-trained by us, with training details described in Appendix B. Notably, BMSE proved challenging to train on STS-B-Natural; consequently, its evaluation on this dataset was omitted. As evidenced in Table ??, our method outperformed all compared methods on AgeDB-Natural, STS-Natural, and IMDB-WIKI-Natural datasets. Compared to the Vanilla model, our approach achieved MAE improvements of 4.54%, 9.27%, and 3.72%, respectively. Furthermore, our method attained SOTA performance across all three datasets, with MAE improvements of 1.02%, 6.58%, and 1.91%, respectively.

Performance for Imbalanced Regression

We evaluated the efficacy of our proposed method on imbalanced regression tasks using the AgeDB-DIR, STS-B-DIR, and IMDB-WIKI-DIR datasets. Furthermore, we plugged our approach with BMSE, SqrtINV, Focal-R, and ConR. As shown in Table 2, our method demonstrated significant improvement over the vanilla model, indicating the effectiveness of ACCon. Besides, comparative analysis with LDS, NaïveSupCon, AdaSupCon, ConR, and RankSIM reveals the competitiveness of our approach in addressing deep imbalanced regression.

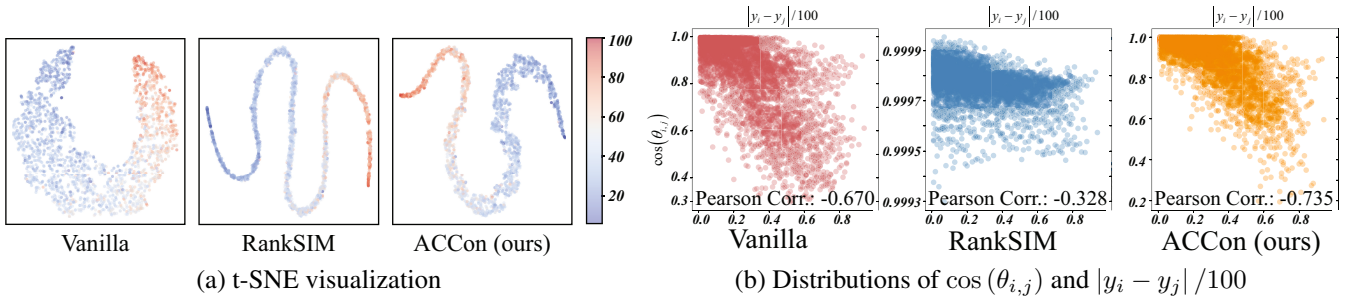


Figure 3: The visualization and quantitation analysis of feature representations. The subfigure (a) is the t-SNE visualization of feature space on the AgeDB-natural test dataset. The subfigure (b) depicts the joint distribution of $\cos(\theta_{i,j})$ and the distance of labels $|y_i - y_j|/100$ on the AgeDB-natural test dataset using kernel density estimation.

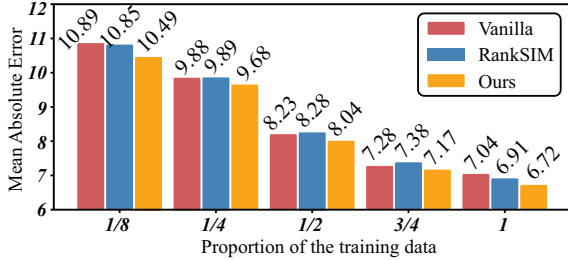


Figure 4: The performance comparison on AgeDB-Natural, when reducing the training dataset.

Our method outperformed most baseline methods across the datasets, with the exception of STS-B-DIR. We hypothesize that this exception may be attributed to the limited amount of data and constraints on data augmentation techniques for the semantic textual similarity task. Additionally, we assessed the performance of our method when plugged with other deep imbalanced regression methods. Results indicate that incorporation with ConR, SqrtINV, and Focal-R led to performance enhancements, particularly in the medium-shot and few-shot categories.

Performance on Limited Training Data

Although massive datasets have been instrumental in advancing contemporary deep learning, the labeling of extensive datasets for training purposes is often impractical due to associated costs and time constraints, particularly in scientific AI applications. Consequently, there is a growing demand for models that maintain robustness despite training data scarcity. To address this, we evaluated our method’s performance under limited training data conditions, thereby assessing its data efficiency. As shown in Fig. 4, our method demonstrates enhanced robustness to limited datasets compared to vanilla and RankSim methods.

Ablation Studies

We conducted a series of ablation studies to analyze the components of our deep regression framework. Initially, we compared the efficacy of using features extracted directly from $f_\theta(\cdot)$ versus features obtained after applying a projection layer for downstream tasks. Although it is common practice to utilize features prior to the projection layer (denoted as *Before Proj.*) for downstream tasks, our results, as presented in

Appendix Table 1, demonstrate that leveraging features post-projection layer leads to superior performance. Subsequently, we assessed the model’s performance after removing the projection layer (denoted as *w/o Proj.*). The results prove the necessity of including the projection layer in the architecture. Furthermore, we implemented a two-stage training scheme, a strategy commonly employed in self-supervised contrastive learning. However, this two-stage scheme presented challenges in model convergence and yielded inferior results compared to training with a multi-task loss function. Further analysis of feature representation, detailed in Appendix C, indicates that our framework’s design is advantageous for enhancing feature representation.

We conducted additional ablation studies to assess the impact of hyperparameters, including the weight coefficient γ and the representation dimension z . Our method demonstrated robustness to variations in these parameters. Comprehensive experimental results are presented in Appendix C.

Qualitative Visualization of Feature Space: We employed t-SNE to visualize the extracted feature representations of our method and the compared baselines on the AgeDB-Natural test dataset, as illustrated in Fig. 3 (a). Notably, the representations extracted by our method exhibit continuity, compactness, and relative symmetry in low-dimensional space, corresponding well to the label space. Specifically, RankSIM, BMSE, and our method demonstrate the capability to obtain continuous and compact features in low-dimensional space, consistent with their design to preserve label relationships.

Correlation Between Feature Representation Similarity and Label Distance To further investigate the learned representations, we present a joint distribution analysis of $\cos(\theta_{i,j})$ and the normalized label distance $|y_i - y_j|/100$ on the AgeDB-Natural test dataset, as shown in Fig. 3 (b). The extracted representations by our proposed method exhibit a negative correlation with the normalized label distance, achieving a Pearson correlation of -0.735. This indicates that our method effectively captures feature representations with inherent label ordering information, aligning with our initial motivation. The ideal cosine distance among feature representations should be distributed in the range $[-1, 1]$, whereas in our method, it mostly falls within $[0.3, 1]$. We attribute this

Methods	AgeDB-DIR (MAE)				IMDB-WIKI-DIR (MAE)				STS-B-DIR (MSE)			
	All	Many	Media	Few	All	Many	Media	Few	All	Many	Media	Few
Vanilla	7.77	6.62	9.55	13.67	8.06	7.23	15.12	26.33	0.974	0.851	1.520	0.984
+LDS	7.67	6.98	8.86	10.89	7.83	7.30	12.43	22.51	0.914	0.819	1.319	0.955
+NaïveSupCon	7.51	6.50	9.23	12.36	7.96	7.27	13.82	23.34	0.962	0.893	1.187	1.147
+AdaSupCon	7.29	6.43	8.59	11.80	7.76	7.16	12.52	23.96	0.983	0.956	1.062	1.073
+RankSIM	7.13	6.51	8.17	10.12	7.72	6.93	14.48	25.38	0.873	0.908	0.767	0.705
+ConR	7.20	6.50	8.04	9.73	7.33	6.75	11.99	22.22	0.937	0.874	1.172	1.044
+ACCon	6.81	6.26	7.56	9.91	7.20	6.61	11.92	22.54	0.802	0.812	0.756	0.765
+ACCon+ConR	6.82	6.36	7.52	9.22	7.19	6.65	11.52	21.64	0.820	0.845	0.721	0.699
SqrtINV	7.81	7.16	8.80	11.20	7.87	7.24	12.44	22.76	1.005	0.894	1.482	1.046
+LDS	7.67	6.98	8.86	10.89	7.83	7.31	12.43	22.51	0.914	0.819	1.319	0.955
+LDS+NaïveSupCon	7.54	6.94	8.20	11.44	7.81	7.26	12.23	22.25	1.030	0.998	1.133	1.125
+LDS+AdaSupCon	7.43	6.82	8.24	10.93	7.75	7.18	12.36	22.06	0.994	0.974	1.098	0.972
+LDS+RankSIM	6.99	6.38	7.88	10.23	7.57	7.00	12.16	22.44	0.889	0.911	0.848	0.690
+LDS+ConR	7.16	6.61	7.97	9.62	7.43	6.84	12.38	21.98	0.927	0.901	1.021	1.087
+LDS+ACCon	6.89	6.38	7.55	9.89	7.34	6.76	11.98	22.13	0.806	0.821	0.742	0.732
+LDS+ACCon+ConR	6.84	6.41	7.42	9.26	7.34	6.79	11.78	21.91	0.820	0.846	0.715	0.692
Focal-R	7.64	6.68	9.22	13.00	7.97	7.12	15.14	26.96	0.951	0.843	1.425	0.957
+LDS	7.56	6.67	8.82	12.40	7.90	7.10	14.72	25.84	0.930	0.807	1.449	0.993
+LDS+NaïveSupCon	7.67	6.94	8.87	11.31	7.84	7.02	14.89	25.61	0.957	0.917	1.075	1.089
+LDS+AdaSupCon	7.59	6.82	8.92	11.18	7.80	7.06	14.02	25.12	1.056	0.932	1.477	1.358
+LDS+RankSIM	7.25	6.40	8.71	11.24	7.71	6.99	13.65	25.97	0.887	0.889	0.918	0.745
+LDS+ConR	7.23	6.63	8.30	11.89	7.85	7.01	14.31	25.23	0.944	0.912	1.102	0.998
+LDS+ACCon	6.90	6.43	7.48	9.74	7.38	6.68	13.12	24.43	0.801	0.822	0.717	0.702
+LDS+ACCon+ConR	6.83	6.44	7.36	9.02	7.40	6.72	13.02	24.19	0.799	0.819	0.722	0.681
Ours vs. Vanilla	↑12.4%	↑5.44%	↑22.9%	↑34.0%	↑10.8%	↑8.58%	↑23.8%	↑17.8%	↑18.0%	↑4.58%	↑53.0%	↑30.8%
Ours vs. SOTA	↑2.58%	↑1.88%	↑6.60%	↑6.24%	↑1.91%	↑2.07%	↑3.92%	↑1.55%	↑8.48%	↑0.85%	↑6.78%	↑3.40%

Table 2: Performance comparison with various imbalanced regression methods on the AgeDB-DIR, IMDB-WIKI-DIR, and STS-B-DIR. The Vanilla, LDS, SqrtINV, Focal-R, RankSim and ConR models are sourced from (Gong, Mori, and Tung 2022; Yang et al. 2021; Keramati, Meng, and Evans 2024). The optimal outcomes for each method are highlighted in **bold font**.

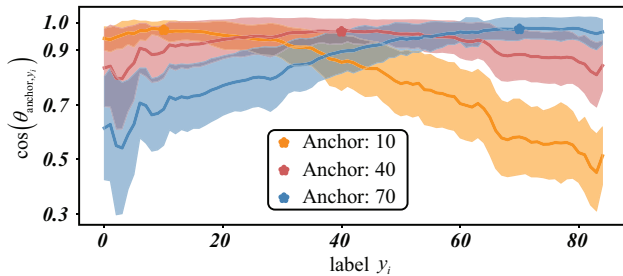


Figure 5: The cosine similarity changes as the label distance varies between the anchor and contrast samples in the AgeDB-Natural test dataset. The stars represent the anchors, while the plot depicts the mean cosine similarity between each anchor and all its contrastive samples. The shaded region indicates the standard deviation of these cosine similarities.

deviation to label imbalance, a known challenge in supervised contrastive learning for image recognition, as discussed in previous studies (Zhu et al. 2022; Li et al. 2022). Fig. 5 illustrates the changes in cosine similarity as the label distance varies between anchor and contrast samples on the AgeDB-Natural test dataset. This analysis highlights the alignment between label-space neighbors and feature-space neighbors, supporting the rationale behind our proposed approach. Notably, we observed increased standard variance as the distance grows, indicating challenges in performance improvement, particularly in the few-shot region. In summary, this visual-

ization provides evidence that our ACCon contributes to the model’s ability to learn feature representations that preserve label-space neighbor information.

Conclusion

Although continual, long-tail, imbalanced, and even missing data in regression targets complicate the regression task, the additional information embedded within these targets offers valuable insights to enhance regression performance. However, current methods fall short of capturing sufficient relationship information. For instance, order-aware approaches overlook the distance between labels, while existing distance-aware methods struggle to ensure that representation similarities accurately reflect label distances. To address this, we propose a hypothesis of a linear negative correlation between feature similarities and label distances, implementing this through an angle-compensated supervised contrastive regularizer. Through extensive experiments on age estimation and semantic textual similarity tasks, we show that our method not only surpasses several state-of-the-art approaches but also excels in terms of data efficiency and performance on imbalanced datasets. Furthermore, additional analysis and theoretical validation confirm that our method aligns well with its underlying motivation, supporting the reasonableness of our assumption and the effectiveness of our approach. Our method offers promising potential for various applications in deep regression, contributing to the advancement of this field.

Acknowledgments

This study was supported by the Key Research and Development Program of Guangdong Province (grant No. 2021B0101400003) and the corresponding author is Jianzong Wang (jzwang@188.com).

References

- Belagiannis, V.; Ruppert, C.; Carneiro, G.; and Navab, N. 2015. Robust optimization for deep regression. In *Proceedings of the IEEE international conference on computer vision*, 2830–2838.
- Cer, D.; Diab, M.; Agirre, E.; Lopez-Gazpio, I.; and Specia, L. 2017. SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Association for Computational Linguistics.
- Chandrasekaran, D.; and Mago, V. 2021. Evolution of semantic similarity—a survey. *ACM Computing Surveys (CSUR)*, 54(2): 1–37.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, Z.; Ma, Q.; and Lin, Z. 2021. Time-Aware Multi-Scale RNNs for Time Series Modeling. In *IJCAI*, 2285–2291.
- Dai, W.; Li, X.; Chiu, W. H. K.; Kuo, M. D.; and Cheng, K.-T. 2021. Adaptive contrast for image regression in computer-aided disease assessment. *IEEE Transactions on Medical Imaging*, 41(5): 1255–1268.
- Ermolov, A.; Siarohin, A.; Sangineto, E.; and Sebe, N. 2021. Whitening for self-supervised representation learning. In *International Conference on Machine Learning*, 3015–3024. PMLR.
- Gao, B.-B.; Xing, C.; Xie, C.-W.; Wu, J.; and Geng, X. 2017. Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing*, 26(6): 2825–2838.
- Gong, Y.; Mori, G.; and Tung, F. 2022. RankSim: Ranking Similarity Regularization for Deep Imbalanced Regression. In *International Conference on Machine Learning*, 7634–7649. PMLR.
- Graf, F.; Hofer, C.; Niethammer, M.; and Kwitt, R. 2021. Dissecting supervised contrastive learning. In *International Conference on Machine Learning*, 3821–3830. PMLR.
- Huber, P. J. 1992. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, 492–518. Springer.
- Jiang, W.; Huang, K.; Geng, J.; and Deng, X. 2020. Multi-scale metric learning for few-shot learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(3): 1091–1102.
- Keramati, M.; Meng, L.; and Evans, R. D. 2024. Conr: Contrastive regularizer for deep imbalanced regression.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- Lathuilière, S.; Mesejo, P.; Alameda-Pineda, X.; and Horaud, R. 2019. A comprehensive analysis of deep regression. *IEEE transactions on pattern analysis and machine intelligence*, 42(9): 2065–2081.
- Lee, S.; Lee, J.; Kim, B.; Yi, E.; and Kim, J. 2021. Patch-wise attention network for monocular depth estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1873–1881.
- Li, T.; Cao, P.; Yuan, Y.; Fan, L.; Yang, Y.; Feris, R. S.; Indyk, P.; and Katabi, D. 2022. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6918–6928.
- Li, W.; Huang, X.; Lu, J.; Feng, J.; and Zhou, J. 2021. Learning probabilistic ordinal embeddings for uncertainty-aware regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13896–13905.
- Liu, K.; Zhu, W.; Shen, Y.; Liu, S.; Razavian, N.; Geras, K. J.; and Fernandez-Granda, C. 2023. Multiple instance learning via iterative self-paced supervised contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3355–3365.
- Moschoglou, S.; Papaioannou, A.; Sagonas, C.; Deng, J.; Kotsia, I.; and Zafeiriou, S. 2017. Agedb: the first manually collected, in-the-wild age database. In *proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 51–59.
- Pan, H.; Han, H.; Shan, S.; and Chen, X. 2018. Mean-variance loss for deep age estimation from a face. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5285–5294.
- Ren, J.; Zhang, M.; Yu, C.; and Liu, Z. 2022. Balanced mse for imbalanced visual regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7926–7935.
- Rogez, G.; Weinzaepfel, P.; and Schmid, C. 2017. Lcr-net: Localization-classification-regression for human pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3433–3441.
- Rothe, R.; Timofte, R.; and Van Gool, L. 2018. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision*, 126(2-4): 144–157.
- Seifi, S.; Reino, D. O.; Chumerin, N.; and Aljundi, R. 2024. OOD Aware Supervised Contrastive Learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1956–1966.
- Shi, X.; Cao, W.; and Raschka, S. 2023. Deep neural networks for rank-consistent ordinal regression based on conditional probabilities. *Pattern Analysis and Applications*, 26(3): 941–955.
- Sun, W.; Shi, Z.; Gao, S.; Ren, P.; de Rijke, M.; and Ren, Z. 2023. Contrastive learning reduces hallucination in conversations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 13618–13626.

- Tabelini, L.; Berriel, R.; Paixao, T. M.; Badue, C.; De Souza, A. F.; and Oliveira-Santos, T. 2021. PolyLANE: Lane estimation via deep polynomial regression. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 6150–6156. IEEE.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *EMNLP 2018*, 353.
- Wang, H.; Sanchez, V.; and Li, C.-T. 2022. Improving face-based age estimation with attention-based dynamic patch fusion. *IEEE Transactions on Image Processing*, 31: 1084–1096.
- Yang, Y.; Zha, K.; Chen, Y.; Wang, H.; and Katabi, D. 2021. Delving into deep imbalanced regression. In *International Conference on Machine Learning*, 11842–11851. PMLR.
- Zha, K.; Cao, P.; Son, J.; Yang, Y.; and Katabi, D. 2024. Rank-N-Contrast: Learning Continuous Representations for Regression. *Advances in Neural Information Processing Systems*, 36.
- Zhang, S.; Yang, L.; Mi, M. B.; Zheng, X.; and Yao, A. 2022. Improving Deep Regression with Ordinal Entropy. In *The Eleventh International Conference on Learning Representations*.
- Zhu, J.; Wang, Z.; Chen, J.; Chen, Y.-P. P.; and Jiang, Y.-G. 2022. Balanced contrastive learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6908–6917.