

# Deep Imbalanced Regression to Estimate Vascular Age from PPG Data: a Novel Digital Biomarker for Cardiovascular Health

Guangkun Nie<sup>1,2</sup>, Jun Li<sup>3</sup>, Shenda Hong<sup>\*1,2</sup>

<sup>1</sup> Institute of Medical Technology, Health Science Center of Peking University, Beijing, China

<sup>2</sup> National Institute of Health Data Science, Peking University, Beijing, China

<sup>3</sup> College of Electronic Science and Engineering, Jilin University, Changchun, China  
nieguangkun@stu.pku.edu.cn, lijun2020@mails.jlu.edu.cn, hongshenda@pku.edu.cn

## ABSTRACT

Photoplethysmography (PPG) is emerging as a crucial tool for monitoring human hemodynamics, with recent studies highlighting its potential in assessing vascular aging through deep learning. However, real-world age distributions are often imbalanced, posing significant challenges for deep learning models. In this paper, we introduce a novel, simple, and effective loss function named the Dist Loss to address deep imbalanced regression tasks. We trained a one-dimensional convolutional neural network (Net1D) incorporating the Dist Loss on the extensive UK Biobank dataset ( $n=502,389$ ) to estimate vascular age from PPG signals and validate its efficacy in characterizing cardiovascular health. The model's performance was validated on a 40% held-out test set, achieving state-of-the-art results, especially in regions with small sample sizes. Furthermore, we divided the population into three subgroups based on the difference between predicted vascular age and chronological age: less than -10 years, between -10 and 10 years, and greater than 10 years. We analyzed the relationship between predicted vascular age and several cardiovascular events over a follow-up period of up to 10 years, including death, coronary heart disease, and heart failure. Our results indicate that the predicted vascular age has significant potential to reflect an individual's cardiovascular health status. Our code is available at <https://github.com/Ngk03/AI-vascular-age>.

## KEYWORDS

Photoplethysmography, vascular age, deep imbalanced regression, cardiovascular health.

## 1 INTRODUCTION

Photoplethysmography (PPG) is a widely utilized biosignal for assessing human hemodynamics, offering advantages such as convenience, non-invasiveness, and continuous monitoring. Recent advancements in deep learning technology have further unlocked the potential of PPG in managing individual cardiovascular health, demonstrating promising progress [4, 6, 12]. On this basis, several studies have shown the value of PPG in evaluating vascular aging [2, 3, 5, 15], a crucial indicator of cardiovascular health. In this

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*KDD-UC '24, August 25–29, 2024, Barcelona, Spain*

© 2024 ACM.

ACM ISBN 978-1-4503-XXXX-X/24/08

<https://doi.org/XXXXXXXX.XXXXXXX>

paper, we use PPG along with deep learning to predict vascular age, providing a refined representation of cardiovascular health levels.

However, in the real world, the age distribution of individuals is often imbalanced, posing significant challenges for deep learning models. When trained on imbalanced samples, models tend to predict towards regions with higher sample density to reduce overall prediction error, leading to significant biases in regions with fewer samples [14, 17]. This approach results in suboptimal parameters because the model fails to learn robust and meaningful features for prediction, rendering the predictions unreliable.

To address this issue, we introduce a novel loss function based on data distribution priors, called the Dist loss. This loss function encourages the model to align predictions more closely with the desired label distribution, reducing the tendency to over-predict in high-density regions. Additionally, we validated the effectiveness of our proposed approach using a large cohort database, the UK Biobank (UKBB) ( $n=502,389$ ). Our results demonstrate significantly improved performance and highlight the potential of vascular age in representing individual cardiovascular health levels. An overview of the workflow for vascular age estimation is shown in Figure 1.

Overall, the contributions of this paper are as follows:

- We propose a novel, simple, and effective loss function, the Dist loss, to effectively mitigate prediction biases in deep imbalanced regression tasks. This loss function significantly improves the performance of the model in regions with few samples, achieving state-of-the-art (SOTA) results.
- Utilizing the large cohort database UKBB, we analyze the relationship between predicted vascular age and certain outcomes (e.g., death, coronary heart disease (CHD), heart failure) in subjects. This analysis validates the application value of vascular age predicted by deep learning models in representing individual cardiovascular health levels.

## 2 METHODS

### 2.1 Problem definition

For a given PPG dataset  $D = \{X, Y\}$ :  $X = \{X_1, X_2, \dots, X_N\} \in \mathbb{R}^{N \times L}$  represents the PPG signals of  $N$  instances, each with a length of  $L$ , and  $Y = \{Y_1, Y_2, \dots, Y_N\} \in \mathbb{R}^N$  represents the corresponding labels. The label space, which includes all possible label values, is  $y = \{y_1, y_2, \dots, y_n\} \in \mathbb{R}^n$ , sorted in ascending/descending order. The goal of model training is to find the optimal parameters that accurately capture the relationship between the PPG signals  $X$  and the labels  $Y$ . This is achieved by minimizing a loss function that measures the discrepancy between the model's predictions and the




Figure 1: Workflow of our vascular age estimation method.

actual labels. Mean squared error (MSE) and mean absolute error (MAE) are commonly used for this purpose in regression tasks.

When chronological ages are used as labels, the trained model predicts an individual’s chronological age based on their PPG signal. For healthy individuals, the predicted age should align closely with their actual chronological age. Conversely, individuals with abnormal hemodynamics, commonly associated with cardiovascular diseases, will have PPG signals resembling those of older individuals. On the other hand, very healthy individuals will exhibit PPG signals similar to those of younger individuals. Therefore, the model’s predictions reflect an individual’s cardiovascular health status, which we term as vascular age.

## 2.2 Model architecture

The model architecture utilized to predict vascular age is Net1D [9], a one-dimensional convolutional neural network (CNN). This model incorporates sequential convolutional blocks with residual connections [8] and squeeze-and-excitation modules [10]. Leveraging the proven efficacy of CNNs in diverse PPG-based tasks [13, 16], the model effectively captures nuanced and distinctive features from PPG signals associated with vascular aging.

## 2.3 Deep imbalanced regression

*Overview.* Imbalanced data distributions pose significant challenges for deep learning models, hindering their ability to extract meaningful and generalizable features. In biomedical applications, such as estimating vascular age from PPG data, this issue becomes even more pronounced due to the critical need for accurate and reliable predictions in healthcare.

As shown in Figure 2a, the age distribution within the UKBB dataset demonstrates a noticeable skew, with a concentration of samples around the mean age and fewer samples at the extremes,

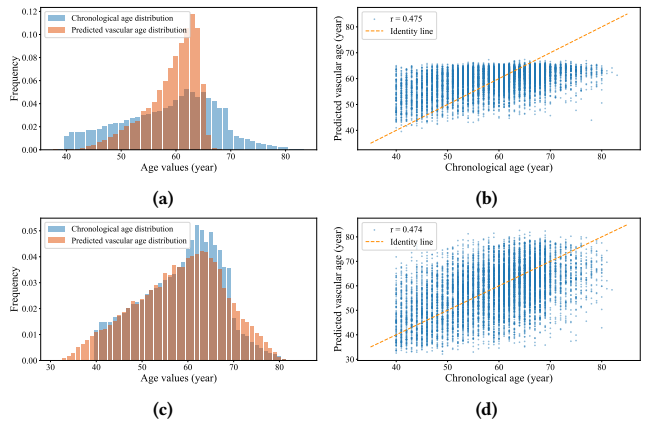


Figure 2: Comparison between the predicted vascular age by the plain model (a, b) and our method (c, d),  $r$  denotes Pearson correlation coefficient. (a) Illustrates the distribution of chronological age (light blue bars) and the corresponding vascular age predicted by the plain model (orange bars) in the UKBB dataset. (b) Presents a scatter plot highlighting the disparity between chronological age and predicted vascular age by the plain model. (c) Depicts the distribution of chronological age (light blue bars) and the corresponding vascular age predicted by our method (orange bars). (d) Displays a scatter plot showcasing the higher consistency between chronological age and predicted vascular age achieved by our method.

especially among older individuals. This imbalance leads to a tendency to over-predict in high-density regions and introduces significant biases in regions with fewer samples, rendering the model predictions unreliable [7, 14, 17].

Models trained solely on MAE or MSE loss, referred to as "plain models," often perform suboptimally on imbalanced data distributions. For demonstration purposes, the "plain models" discussed below are trained using the MAE loss function, although similar behavior was observed with MSE. As shown in Figure 2a, the plain model’s predictions of vascular age distribution on the UKBB dataset tend to be biased towards regions with higher sample density. This results in inaccurate representations of the vascular age distribution, despite the plain model achieving satisfactory performance on conventional evaluation metrics such as MAE and Root Mean Squared Error (RMSE). Further analysis in Figure 2b underscores the discrepancy between the chronological age distribution and the predicted vascular age distribution generated by the plain model. Notably, the predicted vascular age is unevenly distributed across age ranges, with younger individuals often predicted as older and vice versa, highlighting how imbalanced data distribution hampers the model’s ability to acquire optimal features.

To tackle the challenges arising from data imbalance, we introduce a simple yet effective solution termed as **the Dist loss**. This novel loss function harnesses prior knowledge of label distribution in the training data to aid the model in learning useful and meaningful features, thus mitigating prediction bias.

**The Dist loss.** As previously noted, a significant disparity exists between the distribution of labels and the model predictions, primarily caused by the data imbalance. Intuitively, one possible solution involves minimizing the gap between these two distributions. To address this, we propose a novel loss function called the Dist loss, specifically designed to align model predictions more closely with the desired label distribution.

To evaluate the distance between two distributions, commonly used metrics include Kullback-Leibler (KL) divergence and Jensen-Shannon (JS) divergence. Although the computation of these metrics is straightforward in classification tasks, it poses challenges in regression tasks due to the potentially infinite number of possible label values. Additionally, existing methods like kernel density estimation (KDE) [17] for label distribution estimation are non-differentiable, hindering gradient backpropagation and thereby preventing model parameter updates. Hence, instead of directly estimating the distribution of predictions, we introduce a surrogate approach.

For the given PPG dataset  $D = \{X, Y\}$  comprising  $N$  instances, as described in Section 3.1, the corresponding model predictions are denoted as  $\hat{Y}$ . To calculate the Dist loss, we first estimate the label distribution using KDE, resulting in the probability distribution  $P(y)$  for all possible label values  $y = (y_1, y_2, \dots, y_n)$ , sorted in ascending/descending order. For continuous labels with an infinite number of possible values, discretization is applied. We then compute the expected occurrence numbers  $N^E = \{N_1, N_2, \dots, N_n\}$  for each possible label, calculated as  $\lfloor N \cdot P(y) \rfloor$ , where  $\lfloor \cdot \rfloor$  denotes the floor function. Based on this, we generate the ‘expected labels’  $Y^E$  for  $N$  instances, as shown in Equation 1.

$$Y^E = \underbrace{(y_1, \dots, y_1)}_{N_1}, \dots, \underbrace{(y_i, \dots, y_i)}_{N_i}, \dots, \underbrace{(y_n, \dots, y_n)}_{N_n} \quad (1)$$

Next, we sort the model predictions  $\hat{Y}$  in the same order as  $y$ , resulting in the sorted output  $\hat{Y}_{sort}$ . Finally, we compute the Dist loss  $L_{dist}$  using Equation 2, where  $M(\cdot)$  represents a function that measures the distance between two variables, such as MAE or MSE. The only seemingly non-differentiable part of this process is the sorting function, which can be addressed using the fast differentiable sorting method proposed by Blondel et al. [1]. During training, the total loss function is defined in Equation 3, where  $\lambda$  is a hyperparameter that adjusts the importance of distribution similarity.

$$L_{dist} = M(\hat{Y}_{sort}, Y^E_{sort}) \quad (2)$$

$$L_{total} = M(\hat{Y}, Y) + \lambda \cdot L_{dist} \quad (3)$$

## 3 EXPERIMENTS AND RESULTS

### 3.1 Experimental setup

**3.1.1 Dataset and preprocessing.** This study utilizes data from the UK Biobank (UKBB) dataset, which includes genetic and health information from over 500,000 individuals across the United Kingdom. PPG data were extracted from UKBB field 4205. The data were collected using the PulseTrace PCA2 device, which captured average waveforms of 10-15 second PPG signals standardized to a length

of 100. We randomly divided the 502,389 participants into training (60%) and test (40%) sets, with each participant having 0-4 PPG records. This division resulted in 143,196 records ( $n=126,216$ ) for training and 100,172 records ( $n=86,027$ ) for testing. Each PPG signal was then normalized using Equation 4, where  $mean(\cdot)$  and  $std(\cdot)$  denote the mean and standard deviation functions, respectively.

$$normalization(X_i) = \frac{X_i - mean(X_i)}{std(X_i)} \quad (4)$$

**3.1.2 Evaluation metrics.** To evaluate and compare our method with existing approaches for deep imbalanced regression, we utilize a comprehensive set of metrics beyond the conventional Pearson correlation coefficient, MAE, and RMSE. These conventional metrics are insufficient for accurately assessing model performance in highly imbalanced data scenarios. For the labels  $Y = (Y_1, Y_2, \dots, Y_N)$ , their occurrence probabilities  $P_Y = (P_{Y_1}, P_{Y_2}, \dots, P_{Y_N})$ , and the corresponding predictions  $\hat{Y} = (\hat{Y}_1, \hat{Y}_2, \dots, \hat{Y}_N)$ , the formulas for these evaluation metrics are as follows:

**Pearson Correlation Coefficient.** The Pearson correlation coefficient measures the linear correlation between two variables, ranging from -1 to 1. A value of 1 indicates a perfect positive linear relationship, -1 indicates a perfect negative linear relationship, and 0 indicates no linear relationship. The formula for the Pearson correlation coefficient  $r$  is given by Equation 5, where  $\bar{Y}$  and  $\bar{\hat{Y}}$  are the means of the labels  $Y$  and the predictions  $\hat{Y}$ , respectively.

$$r = \frac{\sum_{i=1}^N (Y_i - \bar{Y})(\hat{Y}_i - \bar{\hat{Y}})}{\sqrt{\sum_{i=1}^N (Y_i - \bar{Y})^2 \sum_{i=1}^N (\hat{Y}_i - \bar{\hat{Y}})^2}} \quad (5)$$

**MAE & RMSE.** MAE and RMSE are commonly used metrics to evaluate model performance in regression tasks. The formulas for MAE and RMSE are shown in Equations 6 and 7, respectively.

$$MAE = \frac{1}{N} \sum_{i=1}^N |Y_i - \hat{Y}_i| \quad (6)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (Y_i - \hat{Y}_i)^2} \quad (7)$$

**Weighted MAE & RMSE.** Standard MAE and RMSE are insufficient for evaluating model performance with highly imbalanced labels. To address this, we use weighted MAE and RMSE for a more accurate assessment. The formulas for weighted MAE and RMSE are shown in Equations 8 and 9.

$$Weighted\ MAE = \frac{1}{N} \sum_{i=1}^N \left[ |Y_i - \hat{Y}_i| \cdot \frac{P_{Y_i}}{mean(P_Y)} \right] \quad (8)$$

$$Weighted\ RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N \left[ (Y_i - \hat{Y}_i) \cdot \frac{P_{Y_i}}{mean(P_Y)} \right]^2} \quad (9)$$

*Overlap ratio.* The overlap ratio (OR) is a metric used to assess the similarity between two distributions by calculating the ratio of the area of overlap to the total area under the distributions. The OR ranges from 0 to 1, where a value of 1 indicates perfect overlap and 0 indicates no overlap. The formula for OR is presented in Equation 10, where  $O_y = \{O_1, O_2, \dots, O_n\}$  and  $\hat{O}_y = \{\hat{O}_1, \hat{O}_2, \dots, \hat{O}_n\}$  represent the observed numbers in labels and model predictions for each possible label value in the label space  $y = \{y_1, y_2, \dots, y_n\}$ .

$$\text{OR} = \frac{\sum_{y_i} \min(O_i, \hat{O}_i)}{\sum_{y_i} \max(O_i, \hat{O}_i)} \quad (10)$$

**3.1.3 Compared methods.** To evaluate the effectiveness of the Dist loss, we conducted comparisons with several existing methods designed for deep imbalanced regression.

*LDS & FDS.* LDS (label distribution smoothing) and FDS (feature distribution smoothing) are two modules proposed in [17]. These modules are designed to calibrate label and learned feature distributions, respectively.

*Ranksim.* Ranksim, proposed in [7], leverages label similarity to enhance model performance. It ensures that samples with similar labels are positioned closer together in the feature space, while those with dissimilar labels are positioned further apart.

*Balanced MSE.* Balanced MSE [14] is a modified loss function derived from the standard MSE. It is designed to restore balanced predictions by leveraging the training label distribution prior to making a statistical conversion.

**3.1.4 Statistical analysis.** While the Dist loss serves as a general method for mitigating model bias caused by data imbalance, some bias, although significantly reduced, remains inevitable. Therefore, inspired by [11], we conducted a linear regression on the residuals as a function of age for the model and utilized it to correct each prediction during clinical validation, taking the corrected vascular age as the final predictions. It is noteworthy that experiments showed that without using the Dist loss during model training, the corrected vascular age exhibited little correlation with individuals’ cardiovascular health. This was primarily because the features extracted by the plain model were inaccurate and meaningless, highlighting the importance of the Dist loss in enabling the model to learn accurate and meaningful vascular age representations.

During clinical validation, to evaluate the effectiveness of the predicted vascular age, we categorized the samples into three subgroups based on the difference between corrected vascular age and chronological age: less than -10 years, between -10 and 10 years, and greater than 10 years. For each group, we summarized the mean and standard deviation for continuous variables, including age and body mass index (BMI), and provided counts and percentages for categorical variables such as ethnic background and sex. We also examined the statistical differences between these three groups using Student’s t-test and the Kruskal-Wallis H test for continuous variables and the chi-square test for categorical variables. Additionally, we applied the Cochran-Armitage Trend Test to examine baseline trends across the subgroups. During the analysis, subjects lost to follow-up (n=153) were excluded.

**Table 1: Comparison of Pearson correlation coefficients among different methods for predicting vascular age.**

	Pearson correlation coefficient
Plain	<b>0.475</b>
+ LDS	0.440
+ FDS	0.448
+ LDS & FDS	0.434
+ Ranksim	0.392
+ Balanced MSE	0.422
+ <b>Ours</b>	<b>0.474</b>

For cardiovascular events analysis, we employed the Cox proportional hazards regression model to report hazard ratios (HR) and 95% confidence intervals (95% CI). A follow-up period of up to 10 years was set for this analysis. The analysis involved three levels of adjustments: (1) age, sex, and ethnic background; (2) age, sex, ethnic background, hypertension, and diabetes; and (3) age, sex, ethnic background, hypertension, diabetes, BMI, and dyslipidemia. The proportional hazards assumption was verified using Schoenfeld residuals. Additionally, we demonstrated the adjusted survival curves from the Cox models, stratified by age, sex, and ethnic background.

**3.1.5 Implementation Details.** All methods were integrated with Net1D for equitable comparisons. In our method, we used MAE as  $M(\cdot)$  to quantify the discrepancy between two variables, with  $\lambda$  set to 1, indicating that distribution similarity is equally important as total errors. The training process for all methods employed a batch size of 2048, 50 epochs, and a learning rate of 3e-3. The Adam optimizer with cosine learning rate decay was utilized, and a weight decay of 1e-4 was applied to prevent overfitting. Additionally, based on the probability of label occurrence, we identified the region with label occurrence probabilities less than one-third of the maximum as **the few-shot region**.

## 3.2 Main results

The performance of all methods regarding the Pearson correlation coefficient is summarized in Table 1, while the results in the few-shot region for all methods are presented in Table 2, with the top two results for each metric highlighted in bold and red. Additionally, the predictions of the Dist loss-based model are demonstrated in Figure 2. It can be observed that the distribution of the predicted vascular age closely matches the chronological age, and our method outperforms all compared methods in the few-shot region, significantly reducing prediction bias, while still maintaining good performance in terms of the overall Pearson correlation coefficient.

## 3.3 Clinical validation

**3.3.1 Baseline characteristics.** There are a total of 12,277 samples with a final predicted vascular age younger than their chronological age by more than 10 years, and 10,739 samples with a final predicted vascular age older than their chronological age by more than 10 years, leaving 77,000 samples with a difference between -10 and

**Table 2: Results of different methods in the few-shot region, where  $r$  represents Pearson correlation coefficient.**

Few-shot region								
	$r$	MAE	RMSE	Overlap ratio	Weighted MAE	Weighted RMSE	MAE / Overlap ratio	RMSE / Overlap ratio
Plain	0.637	12.820	13.786	0.015	6.522	7.524	865.377	930.543
+ LDS	<b>0.683</b>	<b>10.585</b>	<b>12.598</b>	0.056	<b>5.270</b>	<b>6.657</b>	185.737	221.049
+ FDS	0.644	11.493	13.603	<b>0.098</b>	6.682	9.013	<b>116.356</b>	<b>137.724</b>
+ LDS & FDS	0.682	10.947	12.766	0.021	5.416	<b>6.694</b>	502.807	586.384
+ Ranksim	0.618	11.559	13.475	0.032	5.974	7.337	355.652	414.581
+ Balanced MSE	0.606	11.204	13.258	0.085	6.064	7.763	130.907	154.901
+ <b>Ours</b>	<b>0.686</b>	<b>9.495</b>	<b>12.029</b>	<b>0.158</b>	<b>4.963</b>	6.851	<b>59.846</b>	<b>75.815</b>

**Table 3: Risk of different outcomes, summarized by the HRs and p-values across the three groups.**

Outcomes	> 10 years younger vs. $\pm$ 10years		> 10 years older vs. $\pm$ 10 years	
	HR (CI 95%)	p value	HR (CI 95%)	p value
<b>Adjusted by age, sex, and ethnic background</b>				
Death	0.90 (0.82-0.99)	0.04	1.22 (1.11-1.35)	<0.005
CHD	0.77 (0.70-0.84)	<0.005	1.30 (1.19-1.42)	<0.005
Stroke	0.88 (0.64-1.21)	0.42	1.96 (1.51-2.55)	<0.005
HF	0.89 (0.77-1.03)	0.11	1.20 (1.04-1.39)	0.01
Hypertension	0.64 (0.59-0.69)	<0.005	1.51 (1.42-1.61)	<0.005
Diabetes	0.70 (0.62-0.80)	<0.005	1.37 (1.23-1.53)	<0.005
Atherosclerosis	0.77 (0.43-1.37)	0.37	1.96 (1.25-3.08)	<0.005
<b>Adjusted by age, sex, ethnic background, hypertension, and diabetes</b>				
Death	0.91 (0.83-1.01)	0.07	1.19 (1.08-1.31)	<0.005
CHD	0.78 (0.71-0.86)	<0.005	1.25 (1.15-1.37)	<0.005
Stroke	0.91 (0.66-1.25)	0.55	1.84 (1.41-2.39)	<0.005
HF	0.92 (0.79-1.06)	0.24	1.14 (0.98-1.31)	0.08
Atherosclerosis	0.79 (0.44-1.40)	0.42	1.86 (1.18-2.92)	0.01
<b>Adjusted by age, sex, ethnic background, hypertension, diabetes, BMI, and dyslipidemia</b>				
Death	0.91 (0.83-1.01)	0.07	1.19 (1.08-1.31)	<0.005
CHD	0.78 (0.71-0.86)	<0.005	1.25 (1.15-1.36)	<0.005
Stroke	0.92 (0.66-1.26)	0.59	1.82 (1.40-2.37)	<0.005
HF	0.90 (0.78-1.04)	0.15	1.14 (0.98-1.31)	0.08
Atherosclerosis	0.80 (0.45-1.42)	0.44	1.84 (1.17-2.89)	0.01

10 years. The baseline characteristics of the three subgroups are presented in Table 4.

**3.3.2 Cardiovascular events analysis.** The results of the adjusted Kaplan-Meier (KM) curves are depicted in Figure 4. It can be observed that individuals with a predicted vascular age more than 10 years older than their chronological age have a higher risk for cardiovascular events, including death, CHD, and heart failure. Conversely, those with a predicted vascular age more than 10 years younger than their chronological age have a lower risk. Furthermore, the hazard ratios (HR) and 95% confidence intervals (CI) for various cardiovascular events are summarized in Table 3, revealing a consistent trend. These findings collectively underscore the potential of predicted vascular age as an indicator of cardiovascular health status.

**3.3.3 Analysis on arterial stiffness.** Increased arterial stiffness is one of the most significant indicators of vascular aging. To verify the effectiveness of predicted vascular age, we examined its consistency with arterial stiffness. As shown in Figure 3, a strong

correlation exists between predicted vascular age and arterial stiffness. Furthermore, this correlation is stronger than that between chronological age and arterial stiffness, indicating that predicted vascular age more accurately represents vascular aging.

## 4 CONCLUSIONS


In conclusion, this study introduces a novel approach named the Dist loss to address the challenges of data imbalance in estimating vascular age from PPG signals. By incorporating the Dist loss, our deep learning model achieved state-of-the-art results on the extensive UK Biobank dataset, particularly in regions with small sample sizes. Furthermore, our analysis revealed a strong correlation between predicted vascular age and cardiovascular health, as evidenced by the relationship with various cardiovascular events over a 10-year follow-up period. These findings highlight the potential of predicted vascular age as a novel digital biomarker for assessing individual cardiovascular health status.

## ACKNOWLEDGEMENT


Data from the UK Biobank, which is available after the approval of an application at <https://www.ukbiobank.ac.uk>. This work was supported by Beijing Natural Science Foundations (QY23040) and National Natural Science Foundation of China (No. 62102008).

**Table 4: Baseline characteristics of the test set samples with PPG records**

Characteristics	> 10 years younger (n=12,277)	± 10 years (n=77,000)	> 10 years older (n=10,739)	p value	p value for trend
Ethnic background, white, n (%)	11055 (90.05)	69213 (89.88)	9367 (87.22)	6.41e-11	2.15e-11
Sex, male, n (%)	3949 (32.16)	41845 (54.34)	7359 (68.53)	< 1e-64	< 1e-64
Age (years), mean (s.d.)	57.69 (9.07)	59.25 (8.58)	58.08 (7.93)	< 1e-64	NA
BMI (kg/m <sup>2</sup> ), mean (s.d.)	26.97 (4.63)	27.27 (4.77)	27.48 (4.79)	6.39e-19	NA
Diabetes, n (%)	667 (5.43)	4530 (5.88)	648 (6.03)	0.096	0.046
Hypertension, n (%)	3041 (24.77)	22328 (29.00)	3425 (31.89)	4.79e-33	2.13e-33
CHD, n (%)	561 (4.57)	4516 (5.86)	629 (5.86)	5.19e-8	9.50e-6
HF, n (%)	79 (0.64)	468 (0.61)	50 (0.47)	0.155	0.090
Stroke, n (%)	142 (1.16)	1105 (1.44)	166 (1.55)	0.024	0.011
Atherosclerosis, n (%)	10 (0.08)	91 (0.12)	17 (0.16)	0.238	0.090



**Figure 3: Scatter plots illustrating the relationship between arterial stiffness and chronological age (a) and predicted vascular age (b).  $r$  denotes Pearson correlation coefficient.**



**Figure 4: Adjusted KM curves for different outcomes, including death, CHD, and heart failure, by sex, age, and ethnic background.**

## REFERENCES

- [1] Mathieu Blondel, Olivier Teboul, Quentin Berthet, and Josip Djolonga. 2020. Fast differentiable sorting and ranking. In *International Conference on Machine Learning*. PMLR, 950–959.
- [2] Yawei Chen, Xuezhi Yang, Rencheng Song, Xuenan Liu, and Jie Zhang. 2024. Predicting Arterial Stiffness from Single-Channel Photoplethysmography Signal: A Feature Interaction-Based Approach. *IEEE Journal of Biomedical and Health Informatics* (2024).
- [3] Lorenzo Dall’Olio, Nico Curti, Daniel Remondini, Yosef Safi Harb, Folkert W Asselbergs, Gastone Castellani, and Hae-Won Uh. 2020. Prediction of vascular aging based on smartphone acquired PPG signals. *Scientific reports* 10, 1 (2020), 19756.
- [4] Roos Edgar, Niels TB Scholte, Kambiz Ebrahimkheil, Marc A Brouwer, Rypko J Beukema, Masih Mafi-Rad, Kevin Vernooij, Sing-Chien Yap, Eelko Ronner, Nicolas van Mieghem, et al. 2024. Automated cardiac arrest detection using a photoplethysmography wristband: algorithm development and validation in patients with induced circulatory arrest in the DETECT-1 study. *The Lancet Digital Health* 6, 3 (2024), e201–e210.
- [5] Redjan Ferizoli, Parmis Karimpour, James M May, and Panicos A Kyriacou. 2024. Arterial stiffness assessment using PPG feature extraction and significance testing in an in vitro cardiovascular system. *Scientific Reports* 14, 1 (2024), 2024.
- [6] Daniel Franklin, Andreas Tzavelis, Jong Yoon Lee, Ha Uk Chung, Jacob Trueb, Hany Arafa, Sung Soo Kwak, Ivy Huang, Yiming Liu, Megh Rathod, et al. 2023. Synchronized wearables for the detection of haemodynamic states via electrocardiography and multispectral photoplethysmography. *Nature biomedical engineering* 7, 10 (2023), 1229–1241.
- [7] Yu Gong, Greg Mori, and Frederick Tung. 2022. Ranksim: Ranking similarity regularization for deep imbalanced regression. *arXiv preprint arXiv:2205.15236* (2022).
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [9] Shenda Hong, Yanbo Xu, Alind Khare, Satria Priambada, Kevin Maher, Alaa Aljiffry, Jimeng Sun, and Alexey Tumanov. 2020. HOLMES: Health OnLine Model Ensemble Serving for Deep Learning Models in Intensive Care Units. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1614–1624.
- [10] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7132–7141.
- [11] Alan Le Goallec, Samuel Diai, Sasha Collin, Jean-Baptiste Prost, Théo Vincent, and Chirag J Patel. 2022. Using deep learning to predict abdominal age from liver and pancreas magnetic resonance images. *Nature Communications* 13, 1 (2022), 1979.
- [12] Guangkun Nie, Jiabao Zhu, Gongzheng Tang, Deyun Zhang, Shijia Geng, Qinghao Zhao, and Shenda Hong. 2024. A Review of Deep Learning Methods for Photoplethysmography Data. *arXiv preprint arXiv:2401.12783* (2024).
- [13] Pongpanut Osathitporn, Guntitawit Sawadwuthikul, Punawish Thuwajit, Kawisara Ueafuea, Thee Mateepithaktham, Narin Kunaseth, Tanut Choksatchawathi, Proadpran Punyabukkana, Emmanuel Mignot, and Theerawit Wilairasitporn. 2023. Rrwwenet: A compact end-to-end multi-scale residual cnn for robust ppg respiratory rate estimation. *IEEE Internet of Things Journal* (2023).
- [14] Jiawei Ren, Mingyuan Zhang, Cunjun Yu, and Ziwei Liu. 2022. Balanced mse for imbalanced visual regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7926–7935.
- [15] Hangsik Shin, Gyujeong Noh, and Byung-Moon Choi. 2022. Photoplethysmogram based vascular aging assessment using the deep convolutional neural network. *Scientific Reports* 12, 1 (2022), 11377.
- [16] Jessica Torres-Soto and Euan A Ashley. 2020. Multi-task deep learning for cardiac rhythm detection in wearable devices. *NPJ digital medicine* 3, 1 (2020), 116.
- [17] Yuzhe Yang, Kaiwen Zha, Yingcong Chen, Hao Wang, and Dina Katabi. 2021. Delving into deep imbalanced regression. In *International conference on machine learning*. PMLR, 11842–11851.