

SafeECGMatch: Calibration-Aware Joint Frequency and Time Space Semi-Supervised Learning for Open-Set ECG Classification

Hongkyu Koh

202200307@hufs.ac.kr

Hankuk University of Foreign Studies
Yongin, Republic of Korea

Ikbeom Jang*

ijang@hufs.ac.kr

Hankuk University of Foreign Studies
Yongin, Republic of Korea

Abstract

Electrocardiogram (ECG) classification models often suffer from severe label scarcity, making semi-supervised learning (SSL) an attractive strategy for reducing annotation costs. In clinical settings, however, unlabeled pools frequently contain out-of-distribution (OOD) anomalies or diagnostic groups absent from the labeled set. Standard SSL forces incorrect pseudo-labels onto these unseen classes, producing overconfident predictions. To address this, we propose SafeECGMatch, a calibration-aware safe SSL framework for single-label ECG classification under label distribution mismatch. Methodologically, SafeECGMatch employs a dual-branch architecture extracting time-frequency latent representations via ECG-specific augmentations. Crucially, it dynamically aligns confidence with empirical accuracy through adaptive label smoothing and temperature scaling, calibrating both the multiclass classifier and the OOD detector across temporal and spectral domains. This joint optimization allows trustworthy OOD rejection and reliable pseudo-labeling. Evaluated on the PTB-XL and PhysioNet/CinC Challenge benchmarks, SafeECGMatch achieves state-of-the-art accuracy and calibration, advancing reliable knowledge discovery in physiological time-series. Code is available at <https://github.com/labhai/SafeECGMatch>.

CCS Concepts

• **Applied computing** → **Health informatics**; • **Computing methodologies** → **Semi-supervised learning settings**; **Neural networks**.

Keywords

Confidence Calibration, Joint Time-Frequency Representation Learning, Safe Semi-Supervised Learning, Electrocardiogram, Out-of-Distribution Detection, Label Distribution Mismatch

ACM Reference Format:

Hongkyu Koh and Ikbeom Jang. 2026. SafeECGMatch: Calibration-Aware Joint Frequency and Time Space Semi-Supervised Learning for Open-Set ECG Classification. In *Proceedings of The Undergraduate Consortium at*

*Corresponding author.

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 the author(s) 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.

KDD '26, Republic of Korea

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-2258-5/2026/08
<https://doi.org/XXXXXXX.XXXXXXX>

KDD 2026 (KDD '26). ACM, Jeju, Korea, 8 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Electrocardiography (ECG) is one of the most widely used non-invasive tools for detecting and monitoring cardiovascular disorders in clinical practice. ECG can be used to diagnose conditions such as myocardial infarction, abnormal cardiac rhythms, and angina [17]. In addition, while the diagnosis of heart failure typically requires costly echocardiography, ECG has increasingly been utilized as a powerful screening and prediction tool for identifying patients with latent asymptomatic heart failure at an early stage [1].

With the rapid progress of deep neural networks, automated ECG interpretation systems have achieved increasingly competitive diagnostic performance and have shown strong potential for supporting large-scale screening and clinical decision-making [3, 20]. Despite this progress, building reliable ECG classification models still depends on expert annotations, which are expensive, time-consuming, and often difficult to obtain at scale. This bottleneck is particularly severe in medical domains, where labeling requires specialized clinical expertise and strict quality control. Semi-supervised learning (SSL), which jointly leverages a small set of labeled samples and a large pool of unlabeled samples, has therefore emerged as a practical and effective strategy for reducing annotation cost while maintaining strong predictive performance [16].

However, the standard SSL assumption that labeled and unlabeled data share the same label space is frequently violated in real-world clinical environments. In practice, unlabeled ECG recordings collected from hospitals or monitoring systems can contain rhythms, morphologies, or patient subgroups that are not represented in the labeled training set. Such a label distribution mismatch introduces out-of-distribution (OOD) samples into the unlabeled data. Applying conventional SSL under this setting can force incorrect pseudo-labels onto unseen classes, and once these errors are fed back into training, they can severely degrade both in-distribution (ID) classification and overall model reliability. To address this issue, Safe SSL, also referred to as open-set SSL, has been proposed to combine semi-supervised classification with OOD detection so that harmful unlabeled outliers can be identified and excluded during learning [4, 6, 13, 21].

Although existing Safe SSL methods have made meaningful progress, they remain vulnerable to a persistent problem of modern deep neural networks: overconfidence. In both pseudo-labeling and OOD detection, the model may assign high confidence to erroneous predictions, causing mislabeled unlabeled samples or OOD instances to be incorporated into training with excessive certainty. This failure is particularly problematic in medical applications,

where reliability is as important as accuracy [5]. Recent work, such as CaliMatch, has shown that calibration can serve as an effective method by jointly improving the calibration of both the classifier and the OOD detector [2]. Nevertheless, CaliMatch has been developed almost exclusively for image recognition benchmarks. In the ECG domain, prior SSL studies like ECGMatch [23] have demonstrated the promise of leveraging unlabeled signals, but they do not directly address the open-set scenario or explicitly tackle overconfidence. This gap is significant because ECG signals exhibit strong temporal structure, morphology-dependent variation, and clinically meaningful inter-patient heterogeneity, making confidence estimation and distribution shift handling far more challenging than in standard image benchmarks.

To address this problem, we propose **SafeECGMatch**, a calibration-aware safe semi-supervised learning framework for single-label ECG classification under label distribution mismatch. Rather than directly adapting image-based Safe SSL to physiological signals, our method is designed around the structural characteristics of ECG time series and the reliability requirements of clinical prediction. SafeECGMatch jointly learns from time- and frequency-domain views of the same ECG sample, applying ECG-specific augmentations. It simultaneously calibrates the predictive confidence in both views, combining the calibrated time- and frequency-domain objectives to enable safer pseudo-label assignment and more reliable OOD rejection.

We make the primary contributions as follows: 1) We introduce the novel Safe SSL framework for ECG classification that explicitly handles label distribution mismatch, to our knowledge; 2) We propose a dual-view calibration mechanism that jointly calibrates signal in both time and frequency domains by integrating them into a single objective function; 3) The proposed approach achieves state-of-the-art performance in both accuracy and calibration, suggesting robust domain adaptation in ECG classification.

2 Related Works

2.1 ECG Classification and Semi-Supervised Learning

Recent deep learning studies have substantially improved ECG classification in supervised settings. Medformer employs an attention-based Transformer with multi-granularity patching to capture ECG patterns across different temporal scales [20]. SuPreME couples ECG signals with large-scale physician-written reports and LLM-driven supervision to enable strong multimodal diagnosis [3]; however, this advantage relies on abundant text annotations that are often not available. Trustworthy ECG diagnosis based on OOD detection uses energy scores and ReAct-style activation truncation to filter anomalous inputs [22], yet this remains a supervised test-time safeguard that rejects unknown samples without exploiting unlabeled data to improve representation learning.

Semi-supervised learning has also been explored to alleviate the scarcity of expert-labeled recordings. ECGMatch demonstrates that SSL can improve cardiovascular disease prediction [23]; however, its emphasis on multi-label prediction is not fully aligned with practical single-label settings, which are common in benchmark data (e.g., single-label records constitute the vast majority

of PhysioNet/CinC Challenge benchmarks as evaluated in ECG-Match). Other frameworks tackle label mismatch from an anomaly-detection angle. AnoGAN learns only from normal ECGs and detects deviations as anomalies [15], while ECG-ADGAN introduces a semi-supervised one-class framework with sequence modeling [11]. These approaches reduce the task to heuristic normal-versus-abnormal detection rather than clinically useful multi-class diagnosis, leaving safe SSL for single-label ECG classification largely unexplored.

2.2 Safe Semi-Supervised Learning and Calibration

Because safe semi-supervised ECG classification is largely absent in the literature, the closest precedents come from general safe SSL. OpenMatch extends FixMatch with one-vs-all detector, providing a threshold to detect outliers with consistency regularization [13, 16]. IOMatch unifies inlier and outlier utilization by combining a closed-set and multi-binary classifier to optimize open-set targets jointly [6]. SCOMatch mitigates data overtrusting by scoring unlabeled examples via MSP, maintaining an OOD memory queue for $(K + 1)$ -class self-training [21]. SafeStudent follows a teacher-student route, using energy-discrepancy for safer teacher predictions and label distribution learning to reduce student overconfidence [4]. Adello explicitly targets proper calibration through flexible distribution alignment, distilling soft pseudo-labels from underconfident samples to manage severe distribution shifts [14].

However, these methods often suffer from overconfidence, which can lead to the generation of incorrect pseudo-labels. CaliMatch addresses this issue by incorporating calibration directly into safe SSL using learnable temperature scaling and adaptive label smoothing derived from validation reliability statistics [2]. While highly effective in the image domain, this formulation does not exploit the temporal and spectral structures central to ECG data. SafeECGMatch builds on this calibration-aware formulation but extends it to joint time-domain and frequency-domain learning.

2.3 Temporal-Frequency Co-Training

Temporal-frequency co-training for time series semi-supervised learning (TS-TFC) is the closest prior work to our dual-view motivation [8]. It treats time-domain and frequency-domain signals as complementary views, learns separate encoders through label propagation supported by a temporal-frequency supervised contrastive objective. More recent approaches such as CompleMatch further improved temporal-frequency SSL via complementary co-training and pseudo-label matching [9]. However, these methods do not address label distribution mismatch. SafeECGMatch adopts this dual-view intuition but comprehensively extends it to safe ECG SSL with explicit OOD rejection and calibration.

3 Methods

This section presents **SafeECGMatch**, our calibration-aware safe semi-supervised learning framework for single-label ECG classification under label distribution mismatch. The method is designed for the practical setting in which the unlabeled pool contains both in-distribution ECGs from the target label space and out-of-distribution recordings that should not be converted into pseudo-label supervision. SafeECGMatch introduces a dual-domain




Figure 1: SafeECGMatch Flowchart

design that jointly optimizes a time-domain branch and a frequency-domain branch, together with ECG-aware augmentations and simultaneous calibration in both domains.

3.1 Problem Setting

Given a set of ECG $^{tem}\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$, and we denote $\mathcal{D}_L = \{(\mathbf{x}_i^t, y_i)\}_{i=1}^{N_L}$ as a labeled ECG set, where $\mathbf{x}_i^t \in \mathbb{R}^{C \times T}$ is a C -lead segment of length T and $y_i \in \{1, \dots, K\}$ is its single class label. Let $\mathcal{D}_U = \{\mathbf{x}_j^u\}_{j=1}^{N_U}$ denote an unlabeled set whose hidden labels may either belong to the known target set or fall outside it, where $N_L + N_U = N$. This is the safe SSL setting with label distribution mismatch.

SafeECGMatch maintains two parallel predictors, one for the time domain(t) and one for the frequency domain(f). For branch $d \in \{t, f\}$, let $h_d(\cdot)$ be a feature encoder, $c_d(\cdot)$ be a multiclass classifier, and $r_d(\cdot)$ be a OOD detector. The frequency input is obtained by applying a magnitude transformation using the real Fast Fourier

Transform(FFT) to the original signal, i.e.,

$$freq\mathcal{X} = |\text{rFFT}(\text{tem}\mathcal{X})| = \{\mathbf{x}_i^f\}_{i=1}^N. \quad (1)$$

For any input $^d\mathbf{x}$ in branch d , where k is the class of labeled ECG set, the posterior of multiclass classifier is defined as

$$p_{d,k}(^d\mathbf{x}_i) = [\text{softmax}(c_d(h_d(^d\mathbf{x}_i)))]_k. \quad (2)$$

OOD detector produces probabilities to distinguish OOD:

$$q_{d,k}(^d\mathbf{x}_i) = \text{softmax}(r_{d,k}(h_d(^d\mathbf{x}_i))). \quad (3)$$

3.2 ECG-Specific Dual-Domain Augmentation

Applying weak and strong augmentations to unlabeled data is crucial for enhancing classifier consistency. Therefore, we propose a specialized augmentation strategy tailored for this purpose.

The weak augmentation for time branch, is generated by per-lead amplitude scaling followed by additive Gaussian noise:

$$\mathcal{A}_w^t(\mathbf{x}_i) = \mathbf{a} \odot (\mathbf{x}_i) + \boldsymbol{\epsilon}, \quad \mathbf{a} \sim \mathcal{N}(1, 0.1^2), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, 0.01^2), \quad (4)$$

where $\mathbf{a} \in \mathbb{R}^{C \times 1}$ is broadcast along time, $\boldsymbol{\epsilon} \in \mathbb{R}^{C \times T}$, and $\odot$ denotes channel-wise multiplication. This augmentation perturbs amplitude mildly while preserving overall rhythm and morphology.

The strong augmentation for time branch, first applies harder structural corruption and then uses the weak augmentation. Let $\mathcal{P}_5(\cdot)$ denote random permutation of five contiguous temporal segments, and let $\mathcal{M}_{s,\ell}^t(\cdot)$ denote zero-masking on interval $[s, s + \ell]$ with $\ell \sim \text{Unif}\{50, \dots, 200\}$. With independent Bernoulli variables $b_p, b_m \sim \text{Bernoulli}(0.5)$,

$$\mathcal{A}_s^t(\mathbf{x}_i) = \mathcal{A}_w^t\left((\mathcal{M}_{s,\ell}^t(\mathcal{P}_5)^{b_m}(\mathcal{P}_5)^{b_p}(\mathbf{x}_i))\right), \quad (5)$$

where exponent 0 means the identity map and exponent 1 means the operation is applied. Thus the strong augmentation used for consistency learning may include segment permutation, temporal masking, or both, followed by the same amplitude-noise perturbation in Eq. (4).

The weak augmentation of frequency branch inherits the weak augmentation from the time branch and maps them into the spectral space using Eq. (1). Specifically,

$$\mathcal{A}_w^f(\mathbf{x}_i) = |\text{rFFT}(\mathcal{A}_w^t(\mathbf{x}_i))| \quad (6)$$

The strong augmentation for the frequency branch, $\mathcal{A}_s^f$, is defined by masking a contiguous frequency band in the magnitude spectrum. Let L_f denote the number of frequency bins, $m = \lfloor 0.15L_f \rfloor$, and let s_f be sampled uniformly from $\{0, \dots, L_f - m - 1\}$. Then, for each frequency-bin index $j \in \{0, \dots, L_f - 1\}$,

$$[\mathcal{A}_s^f(\mathbf{x}_i^f)]_j = \begin{cases} 0, & s_f \leq j < s_f + m, \\ [\mathcal{A}_w^f(\mathbf{x}_i^f)]_j, & \text{otherwise.} \end{cases} \quad (7)$$

3.3 Calibration-Aware Safe SSL in Each Branch

The same training routine is applied independently to the time branch and the frequency branch. For branch $d \in \{t, f\}$, labeled data minibatch ${}^d\mathcal{B}_\ell$ and unlabeled data minibatch ${}^d\mathcal{B}_u$, the supervised classification loss is

$$L_{\text{cls}}^d = - \sum_{i=1}^{|{}^d\mathcal{B}_\ell|} \sum_{k=1}^K y_{ik} \log p_{d,k}({}^d\mathbf{x}_i^l). \quad (8)$$

The OOD detector is trained on labeled data using

$$L_{\text{ood}}^d = - \sum_{i=1}^{|{}^d\mathcal{B}_\ell|} \sum_{k=1}^K [y_{ik} \log q_{d,k}({}^d\mathbf{x}_i^l) + \min_{l \neq k} \log(1 - q_{d,l}({}^d\mathbf{x}_i^l))]. \quad (9)$$

For unlabeled data, OOD detector is regularized by soft consistency regularization about the two weak views:

$$L_{\text{socr}}^d = \sum_{i=1}^{|{}^d\mathcal{B}_u|} \sum_{k=1}^K \|\mathbf{q}_{d,k}(\mathcal{A}_{w(1)}^d({}^d\mathbf{x}_i^u)) - \mathbf{q}_{d,k}(\mathcal{A}_{w(2)}^d({}^d\mathbf{x}_i^u))\|_2^2. \quad (10)$$

After warm-up, SafeECGMatch activates calibration losses for both the multiclass logits and the OOD logits in each branch. Let

$T_{\text{cls}}^d > 0$ and $T_{\text{ood}}^d > 0$ be learnable temperature parameters. The calibrated classifier and OOD probabilities are

$$\bar{p}_{d,k}({}^d\mathbf{x}_i) = \left[\text{softmax} \left(c_d(h_d({}^d\mathbf{x}_i))/T_{\text{cls}}^d \right) \right]_k, \quad (11)$$

$$\bar{q}_{d,k}({}^d\mathbf{x}_i) = \text{softmax} \left(r_{d,k}(h_d({}^d\mathbf{x}_i))/T_{\text{ood}}^d \right). \quad (12)$$

Following [2], we estimate confidence-conditioned accuracy on the validation split separately for the classifier and the OOD detector to define adaptive label smoothing. The resulting branch-specific target confidences are denoted by α_i^d for the multiclass head and β_i^d for the OOD detector. Then the calibration loss functions become

$$L_{\text{cls-cal}}^d = - \sum_{i=1}^{|{}^d\mathcal{B}_\ell|} \sum_{k=1}^K \left(\alpha_i^d \cdot y_{ik} + \frac{1 - \alpha_i^d}{K - 1} (1 - y_{ik}) \right) \cdot \log \bar{p}_{d,k}({}^d\mathbf{x}_i^l) \quad (13)$$

$$L_{\text{ood-cal}}^d = - \sum_{i=1}^{|{}^d\mathcal{B}_\ell|} \sum_{k=1}^K \left[\begin{aligned} & \left(\beta_i^d y_{ik} + (1 - \beta_i^d)(1 - y_{ik}) \right) \log \bar{q}_{d,k}({}^d\mathbf{x}_i^l) \\ & + \min_k \left(\left((1 - \beta_i^d) y_{ik} + \beta_i^d (1 - y_{ik}) \right) \right. \\ & \quad \left. \times \log(1 - \bar{q}_{d,k}({}^d\mathbf{x}_i^l)) \right) \end{aligned} \right] \quad (14)$$

3.4 Reliable Unlabeled Selection and Dual-Domain Final Objective

We now aim to train the model using the reliable unlabeled data after filtering out OOD samples by OOD detector. To this end, separate thresholds are assigned to each branch, and filtering is performed independently for each branch. The procedure for obtaining reliable unlabeled data for each branch is defined as follows:

$$\begin{aligned} {}^dS_i^u &= \sum_k \bar{p}_{d,k}({}^d\mathbf{x}_i^u) \cdot \bar{q}_{d,k}({}^d\mathbf{x}_i^u) \\ {}^dC_i^u &= \max_k \bar{p}_{d,k}({}^d\mathbf{x}_i^u) \end{aligned} \quad (15)$$

$${}^d\mathcal{B}_u^I = \{ {}^d\mathbf{x}_i^u \in {}^d\mathcal{B}_u : ({}^dS_i^u > T_1) \wedge ({}^dC_i^u > T_2) \} \quad (16)$$

Following FixMatch, applying weak and strong augmentations to unlabeled data can improve classifier performance [16]. Therefore, we apply the ECG-specific augmentation strategies described in Section 3.2 and compute the following FixMatch loss function:

$$L_{\text{fix}}^d = \sum_{i=1}^{|{}^d\mathcal{B}_u^I|} \sum_{k=1}^K \mathbb{I} \left(\arg\max_\ell p_{d,\ell}(\mathcal{A}_w^d({}^d\mathbf{x}_i^u)) = k \right) \times \log p_{d,k}(\mathcal{A}_s^d({}^d\mathbf{x}_i^u)) \quad (17)$$

The complete branch loss is therefore

$$\begin{aligned} L^d &= L_{\text{cls}}^d + \lambda_{\text{ood}} L_{\text{ood}}^d + \lambda_{\text{socr}} L_{\text{socr}}^d \\ &\quad + \lambda_{\text{cls-cal}} L_{\text{cls-cal}}^d + \lambda_{\text{ood-cal}} L_{\text{ood-cal}}^d + \lambda_{\text{fix}} L_{\text{fix}}^d. \end{aligned} \quad (18)$$

Unlike other formulations that operate in a single domain, SafeECGMatch computes Eq. (18) for both the time branch and the frequency

Algorithm 1 SafeECGMatch

```

1: Input:  $\mathcal{D}_L$  (labeled ECG set),  $\mathcal{D}_U$  (unlabeled ECG set),  $\mathcal{V}$  (validation set)
2: Input: Model parameters  $\Theta$  for  $\{h_d, c_d, r_d\}_{d \in \{t, f\}}$ 
3: Input: Hyperparameters  $E_{\max}, I_{\max}, E_{\text{warm}}, \eta, T_1, T_2$ 
4: Input: Loss weights  $\lambda_{\text{ood}}, \lambda_{\text{socr}}, \lambda_{\text{cls-cal}}, \lambda_{\text{ood-cal}}, \lambda_{\text{fix}}, \lambda_{\text{sum}}$ 
5: for epoch = 1 to  $E_{\max}$  do
6:   for iteration = 1 to  $I_{\max}$  do
7:     Sample minibatch  $\mathcal{B}_\ell \subset \mathcal{D}_L$  and minibatch  $\mathcal{B}_u \subset \mathcal{D}_U$ 
8:     Generate  $\mathcal{A}_{w(1)}^t, \mathcal{A}_{w(2)}^t, \mathcal{A}_s^t$  from  $\mathcal{B}_u$  via Eqs. (4), (5)
9:     Generate  $\mathcal{A}_{w(1)}^f, \mathcal{A}_{w(2)}^f, \mathcal{A}_s^f$  from  $\mathcal{B}_u$  via Eqs. (6), (7)
10:    for each branch  $d \in \{t, f\}$  do
11:      Compute  $p_d$  and  $q_d$  via Eqs. (2), (3)
12:       $L_{\text{base}}^d \leftarrow L_{\text{cls}}^d + \lambda_{\text{ood}} L_{\text{ood}}^d + \lambda_{\text{socr}} L_{\text{socr}}^d$  (Eqs. (8)–(10))
13:    end for
14:    if epoch  $\geq E_{\text{warm}}$  then
15:      for each branch  $d \in \{t, f\}$  do
16:        Get calibrated prob.  $\tilde{p}_d, \tilde{q}_d$  via Eqs. (11), (12)
17:        Filter reliable subset  ${}^d\mathcal{B}_u^I$  via Eqs. (15)–(16)
18:         $L_{\text{warm}}^d \leftarrow \lambda_{\text{cls-cal}} L_{\text{cls-cal}}^d + \lambda_{\text{ood-cal}} L_{\text{ood-cal}}^d + \lambda_{\text{fix}} L_{\text{fix}}^d$ 
19:      end for
20:    else
21:      for each branch  $d \in \{t, f\}$  do
22:         $L_{\text{warm}}^d \leftarrow 0$ 
23:      end for
24:    end if
25:    for each branch  $d \in \{t, f\}$  do
26:       $L^d \leftarrow L_{\text{base}}^d + L_{\text{warm}}^d$  (Eq. (18))
27:    end for
28:     $L \leftarrow L^t + \lambda_{\text{sum}} L^f$  (Eq. (19))
29:    Update model parameters with stochastic gradient descent:
     $\Theta \leftarrow \Theta - \eta \nabla L$ 
30:  end for
31:  Update branch-wise statistic  $(\alpha^d, \beta^d)$  and temperatures
   $(T_{\text{cls}}^d, T_{\text{ood}}^d)$  via  $\mathcal{V}$ 
32: end for
33: return Trained SafeECGMatch model  $\{h_d, c_d, r_d\}_{d \in \{t, f\}}$ 

```

branch within every iteration. The final loss function is the weighted sum of the two branch losses,

$$L = L^t + \lambda_{\text{sum}} L^f. \quad (19)$$

This is the key design difference of SafeECGMatch: time-domain calibration and frequency-domain calibration are both optimized simultaneously, and the final training signal is obtained by adding the two calibrated safe SSL objectives. All hyperparameters, including threshold values (T_1, T_2) and balancing coefficients ($\lambda_{\text{ood}}, \lambda_{\text{socr}}, \lambda_{\text{cls-cal}}, \lambda_{\text{ood-cal}}, \lambda_{\text{fix}}$), were optimized for ECG time-series data. Sensitivity analysis of λ_{sum} is provided in Appendix B (Table 4).

4 Experiments

4.1 Experimental Setup

We evaluate SafeECGMatch on two 500 Hz single-label ECG benchmarks, PTB-XL [18] and CinC21 [12], under open-set semi-supervised learning with label distribution mismatch.

Following the standard evaluation protocol of [10], all methods utilize a ResNet1D backbone. Models are optimized using Adam with a learning rate of 0.001, a total batch size of 64, and a labeled-to-unlabeled ratio of 1:99. Training spans 50,000 iterations with a 500-iteration warm-up, and checkpoints are evaluated every 1,000 iterations. Each experiment is repeated with three random seeds, and we report the mean and standard deviation. The evaluation metrics include classification accuracy (ACC) alongside calibration metrics: Expected Calibration Error (ECE), Adaptive Calibration Error (ACE), and Static Calibration Error (SCE).

For both benchmarks, we retain only single-label ECGs to construct five-class datasets, applying a random 8:1:1 train/validation/test split, with unseen class data of unlabeled data proportion set to 30% or 60%.

- **PTB-XL:** PTB-XL [18] is an open-access dataset of clinical 12-lead ECG records annotated by cardiologists. For this dataset, we define Normal ECG (NORM), Myocardial Infarction (MI), and Conduction Disturbance (CD) as seen classes, while ST/T-Change (STTC) and Hypertrophy (HYP) are treated as unseen classes. To retain the unlabeled data proportion, we utilized a subset of 15,000 ECGs. This dataset is partitioned into 150 labeled training samples (50 per seen class) and an unlabeled pool of 14,850 ECGs.
- **CinC21:** The PhysioNet/Computing in Cardiology (CinC) Challenge 2021 dataset is a comprehensive multi-center corpus used to evaluate automated cardiac abnormality identification [12]. We derive a single-label subset from this collection, preprocessed via the ECGMatch pipeline [23]. Assigning NORM, Abnormal Rhythm (RHY), and CD as seen classes, and ST/T Abnormality (ST) and Other Abnormalities (OTHER) as unseen classes results in 31,138 total ECGs. This comprises 312 labeled samples (104 per seen class) and 30,826 unlabeled samples.

We evaluate SafeECGMatch together with the supervised and latest competing methods—FixMatch [16], IOMatch [6], OpenMatch [13], SCOMatch [21], SafeStudent [4], ECGMatch [23], Adello [14], CalMatch [2], TS-TFC [8], CompleMatch [9]. Since safe ECG learning requires both correct classification and reliable confidence under label distribution mismatch [7], we evaluate accuracy and calibration jointly.

4.2 Benchmark Results

Appendix A (Table 3) summarizes the main benchmark results. SafeECGMatch achieves the best or second-best ACC in all four protocols while also obtaining the best or second-best calibration across all settings. On PTB-XL with 60% OOD, it improves ACC from 0.8360 to 0.8433 and reduces ECE from 0.054 to 0.047. On CinC21 with 60% OOD, it improves ACC from 0.7617 to 0.7803 and simultaneously lowers ECE from 0.1030 to 0.0860. Crucially, by mitigating the overconfidence typically seen when predicting unseen classes, our model successfully filters harmful out-of-distribution signals to prevent performance degradation. These results indicate that the proposed dual-domain design fundamentally improves the accuracy-calibration trade-off rather than merely trading one for the other.

Table 1: Summary of benchmark results on PTB-XL [18] and CinC21 [12]. Best results are in bold, and second-best results are underlined. See full benchmark with additional methods and metrics (SCE, ACE) in Appendix A.

Method	OOD	PTB-XL [18]		CinC21 [12]	
		ACC	ECE	ACC	ECE
Supervised	30%	0.734±0.040	0.236±0.041	0.687±0.053	0.249±0.029
OpenMatch [13]	30%	0.831±0.007	0.158±0.008	0.779±0.050	0.202±0.043
	60%	<u>0.830±0.004</u>	0.155±0.005	0.749±0.050	0.218±0.030
CompleMatch [9]	30%	0.750±0.017	0.160±0.014	<u>0.795±0.006</u>	0.135±0.006
	60%	0.734±0.032	0.183±0.018	0.786±0.011	0.124±0.038
ECGMatch [23]	30%	0.788±0.039	0.072±0.011	0.635±0.049	0.092±0.014
	60%	0.791±0.033	0.122±0.039	0.651±0.185	0.132±0.097
CaliMatch [2]	30%	<u>0.845±0.002</u>	0.039±0.010	0.819±0.010	<u>0.057±0.005</u>
	60%	0.831±0.004	<u>0.054±0.004</u>	0.758±0.006	<u>0.103±0.021</u>
SafeECGMatch	30%	0.852±0.005	0.045±0.015	0.829±0.009	0.052±0.014
	60%	0.843±0.009	0.047±0.011	<u>0.780±0.018</u>	0.086±0.027




Figure 2: Accuracy-calibration trade-off on PTB-XL with 30% OOD contamination. Since calibration-aware learning often involves an ACC-ECE trade-off, the preferred model should balance discriminative performance and confidence reliability rather than prioritizing either metric in isolation [7, 19]. SafeECGMatch provides a better compromise by achieving the highest ACC with only a slight increase in ECE.

4.3 Ablation Study

We conducted an ablation study under the severe 60% OOD setting. We compare: (1) **SafeECGMatch**; (2) **w/o frequency-domain calibration** (analogous to CaliMatch); and (3) **w/o frequency and time-domain calibration** (analogous to OpenMatch). As shown in Table 2, explicitly imposing calibration constraints on both domains significantly outperforms merely adding an uncalibrated auxiliary spectral branch. This confirms that joint time and frequency calibration is essential for simultaneously maximizing discriminative performance and confidence reliability.

Table 2: Ablation study at 60% OOD on PTB-XL & CinC21

Method	PTB-XL		CinC21	
	ACC	ECE	ACC	ECE
SafeECGMatch	0.843±0.007	0.047±0.009	0.780±0.014	0.086±0.022
w/o Freq. Calib.	<u>0.831±0.003</u>	<u>0.054±0.004</u>	<u>0.758±0.005</u>	<u>0.103±0.017</u>
w/o both Calib.	0.830±0.004	0.155±0.004	0.749±0.041	0.218±0.024

5 Conclusion

We investigated approaches to safely and reliably train models for ECG data to predict cardiovascular disorders—such as arrhythmias, myocardial infarction, and heart failure—in more realistic scenarios. Many cardiovascular conditions, such as heart failure, do not exhibit symptoms at rest but become apparent during physical activity. Consequently, diagnostic procedures often involve a treadmill test to measure ECG, heart rate, and blood pressure during exertion. In such dynamic settings, noise frequently occurs, and OOD instances—unseen classes presenting completely new morphologies—inevitably emerge. When training models with such contaminated data, there is a critical risk of misclassifying actual diseases as normal due to overconfident, erroneous predictions.

SafeECGMatch directly addresses this issue by reliably filtering out unseen classes and reducing model overconfidence through explicit calibration, yielding more accurate predictions for known ID diagnostic groups (e.g., NORM, CD, MI). Across PTB-XL and CinC21, our approach achieves state-of-the-art accuracy and calibration under label distribution mismatch, while using only 1% of labeled data—substantially reducing the annotation costs required from clinical experts.

More broadly, since many time-series problems admit multiple informative views (e.g., time-frequency, sensor-wise), SafeECGMatch can be seen as one instance of a calibration-aware multi-view safe SSL framework, suggesting its applicability to wider time-series settings with abundant unlabeled data and hidden distribution shift.

At the same time, this study has several limitations. Our experiments focus on single-label ECG classification, two benchmark datasets, and a fixed time-frequency two-view design. The current framework also does not yet address more realistic settings such as multi-label diagnosis. These limitations define the next step of our work. In future research, we plan to extend SafeECGMatch to multi-label classification, and evaluate the framework on broader medical and general time-series tasks.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Ministry of Science and ICT (MSIT) (RS-2024-00455720, RS-2026-25481174); the National Institute of Health (NIH) research project (2026-ER0904-00); the Korea Health Technology R&D Project through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Korea (RS-2025-02220534); the Advanced GPU Utilization Support Program and the High-Performance Computing Support project (RQT-25-070083), both funded by MSIT, Korea; and Hankuk University of Foreign Studies Research Fund of 2026.

References

- [1] Zachi I Attia, Suraj Kapa, Francisco Lopez-Jimenez, Paul M McKie, Dorothy J Ladewig, Gaurav Satam, Patricia A Pellikka, Maurice Enriquez-Sarano, Peter A Noseworthy, Thomas M Munger, et al. 2019. Screening for cardiac contractile dysfunction using an artificial intelligence-enabled electrocardiogram. *Nature medicine* 25, 1 (2019), 70–74.
- [2] Jinsoo Bae, Seoung Bum Kim, and Hyungrok Do. 2025. CaliMatch: Adaptive Calibration for Improving Safe Semi-supervised Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2867–2876.
- [3] Mingsheng Cai, Jiuming Jiang, Wenhao Huang, Che Liu, and Rossella Arcucci. 2025. SuPreME: A Supervised Pre-training Framework for Multimodal ECG Representation Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 11817–11844. doi:10.18653/v1/2025.findings-emnlp.633
- [4] Rundong He, Zhongyi Han, Xiankai Lu, and Yilong Yin. 2022. Safe-student for safe deep semi-supervised learning with unseen-class unlabeled data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 14585–14594.
- [5] Daehwan Kim, Haejun Chung, and Ikbeom Jang. 2026. MORCU: Margin-based ordinal classification with dynamic regularization for calibration and unimodality. *Pattern Recognition* (2026), 113820.
- [6] Zekun Li, Lei Qi, Yinghuan Shi, and Yang Gao. 2023. Iomatch: Simplifying open-set semi-supervised learning with joint inliers and outliers utilization. In *Proceedings of the IEEE/CVF international conference on computer vision*. 15870–15879.
- [7] Bingyuan Liu, Ismail Ben Ayed, Adrian Galdan, and Jose Dolz. 2022. The devil is in the margin: Margin-based label smoothing for network calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 80–88.
- [8] Zhen Liu, Qianli Ma, Peitian Ma, and Linghao Wang. 2023. Temporal-frequency co-training for time series semi-supervised learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 8923–8931.
- [9] Zhen Liu, Kun Zeng, Qianli Ma, and James T Kwok. 2025. CompeMatch: Boosting Time-Series Semi-Supervised Classification With Temporal-Frequency Complementarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025).
- [10] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. 2018. Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems* 31 (2018).
- [11] Jing Qin, Fujie Gao, Zumin Wang, David C Wong, Zhibin Zhao, Samuel D Relton, and Hui Fang. 2023. A novel temporal generative adversarial network for electrocardiography anomaly detection. *Artificial Intelligence in Medicine* 136 (2023), 102489.
- [12] Matthew A Reyna, Nadi Sadr, Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux, Ali Bahrami Rad, Andoni Elola, Salman Seyedi, Sardar Ansari, et al. 2021. Will two do? varying dimensions in electrocardiography: The physionet/computing in cardiology challenge 2021. In *2021 computing in cardiology (CinC)*, Vol. 48. IEEE, 1–4.
- [13] Kuniaki Saito, Donghyun Kim, and Kate Saenko. 2021. Openmatch: Open-set semi-supervised learning with open-set consistency regularization. *Advances in neural information processing systems* 34 (2021), 25956–25967.
- [14] Emanuel Sanchez Aimar, Nathaniel Helgesen, Yonghao Xu, Marco Kuhlmann, and Michael Felsberg. 2024. Flexible distribution alignment: Towards long-tailed semi-supervised learning with proper calibration. In *European Conference on Computer Vision*. Springer, 307–327.
- [15] Dong-Hoon Shin, Roy C Park, and Kyungyong Chung. 2020. Decision boundary-based anomaly detection model using improved AnoGAN from ECG data. *IEEE Access* 8 (2020), 108664–108674.
- [16] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems* 33 (2020), 596–608.
- [17] Christiaan Vrints, Felicita Andreotti, Konstantinos C Koskinas, Xavier Rossello, Marianna Adamo, James Ainslie, Adrian Paul Banning, Andrzej Budaj, Ronny R Buechel, Giovanni Alfonso Chiariello, et al. 2024. 2024 ESC guidelines for the management of chronic coronary syndromes: developed by the task force for the management of chronic coronary syndromes of the European Society of Cardiology (ESC) endorsed by the European Association for Cardio-Thoracic Surgery (EACTS). *European heart journal* 45, 36 (2024), 3415–3537.
- [18] Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bousselet, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific data* 7, 1 (2020), 154.
- [19] Cheng Wang, Jorge Balazs, György Szarvas, Patrick Ernst, Lahari Poddar, and Pavel Danchenko. 2022. Calibrating imbalanced classifiers with focal loss: An empirical study. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*. 145–153.
- [20] Yihe Wang, Nan Huang, Taida Li, Yujun Yan, and Xiang Zhang. 2024. Medformer: A multi-granularity patching transformer for medical time-series classification. *Advances in Neural Information Processing Systems* 37 (2024), 36314–36341.
- [21] Zerun Wang, Liuyu Xiang, Lang Huang, Jiafeng Mao, Ling Xiao, and Toshihiko Yamasaki. 2024. Scomatch: Alleviating overtrusting in open-set semi-supervised learning. In *European Conference on Computer Vision*. Springer, 217–233.
- [22] Bowen Yu, Yuhong Liu, Xin Wu, Jing Ren, and Zhibin Zhao. 2025. Trustworthy diagnosis of Electrocardiography signals based on out-of-distribution detection. *Plos one* 20, 2 (2025), e0317900.
- [23] Rushuang Zhou, Lei Lu, Zijun Liu, Ting Xiang, Zhen Liang, David A Clifton, Yining Dong, and Yuan-Ting Zhang. 2023. Semi-supervised learning for multi-label cardiovascular diseases prediction: a multi-dataset study. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 5 (2023), 3305–3320.

A Detailed Benchmark Results

Table 3: Main benchmark results on PTB-XL [18] and CinC21 [12]. Best results are in bold and second-best results are underlined for each metric within each OOD setting. Lower ECE, ACE, and SCE represent better calibration.

Method	OOD	PTB-XL [18]				CinC21 [12]			
		ACC $\uparrow$	ECE $\downarrow$	ACE $\downarrow$	SCE $\downarrow$	ACC $\uparrow$	ECE $\downarrow$	ACE $\downarrow$	SCE $\downarrow$
Supervised	30%	0.734 $\pm$ 0.040	0.236 $\pm$ 0.041	0.160 $\pm$ 0.030	0.164 $\pm$ 0.028	0.687 $\pm$ 0.053	0.249 $\pm$ 0.029	0.169 $\pm$ 0.020	0.172 $\pm$ 0.021
	60%								
FixMatch [16]	30%	0.835 $\pm$ 0.010	0.155 $\pm$ 0.013	0.098 $\pm$ 0.006	0.105 $\pm$ 0.008	0.787 $\pm$ 0.025	0.191 $\pm$ 0.025	0.125 $\pm$ 0.016	0.129 $\pm$ 0.017
(NeurIPS'20)	60%	0.819 $\pm$ 0.005	0.166 $\pm$ 0.003	0.104 $\pm$ 0.004	0.113 $\pm$ 0.001	0.751 $\pm$ 0.018	0.210 $\pm$ 0.032	0.142 $\pm$ 0.022	0.146 $\pm$ 0.021
IOMatch [6]	30%	0.767 $\pm$ 0.036	0.209 $\pm$ 0.031	0.141 $\pm$ 0.023	0.144 $\pm$ 0.022	0.796 $\pm$ 0.010	0.179 $\pm$ 0.004	0.119 $\pm$ 0.006	0.123 $\pm$ 0.004
(ICCV'23)	60%	0.771 $\pm$ 0.024	0.208 $\pm$ 0.024	0.135 $\pm$ 0.014	0.143 $\pm$ 0.016	0.762 $\pm$ 0.002	0.211 $\pm$ 0.002	0.143 $\pm$ 0.002	0.146 $\pm$ 0.002
OpenMatch [13]	30%	0.831 $\pm$ 0.007	0.158 $\pm$ 0.008	0.098 $\pm$ 0.008	0.107 $\pm$ 0.005	0.779 $\pm$ 0.050	0.202 $\pm$ 0.043	0.133 $\pm$ 0.033	0.138 $\pm$ 0.031
(NeurIPS'21)	60%	0.830 $\pm$ 0.004	0.155 $\pm$ 0.005	0.097 $\pm$ 0.002	0.106 $\pm$ 0.004	0.749 $\pm$ 0.050	0.218 $\pm$ 0.030	0.147 $\pm$ 0.024	0.152 $\pm$ 0.023
SCOMatch [21]	30%	0.830 $\pm$ 0.010	0.139 $\pm$ 0.013	0.089 $\pm$ 0.006	0.097 $\pm$ 0.006	0.785 $\pm$ 0.010	0.175 $\pm$ 0.004	0.118 $\pm$ 0.004	0.121 $\pm$ 0.004
(ECCV'24)	60%	<u>0.836$\pm$0.010</u>	0.145 $\pm$ 0.008	0.091 $\pm$ 0.009	0.100 $\pm$ 0.006	0.737 $\pm$ 0.014	0.214 $\pm$ 0.009	0.145 $\pm$ 0.010	0.149 $\pm$ 0.009
SafeStudent [4]	30%	0.782 $\pm$ 0.019	0.147 $\pm$ 0.021	0.105 $\pm$ 0.019	0.108 $\pm$ 0.018	0.713 $\pm$ 0.051	0.181 $\pm$ 0.039	0.124 $\pm$ 0.026	0.126 $\pm$ 0.027
(CVPR'22)	60%	0.771 $\pm$ 0.032	0.164 $\pm$ 0.020	0.115 $\pm$ 0.016	0.119 $\pm$ 0.014	0.735 $\pm$ 0.003	0.169 $\pm$ 0.008	0.129 $\pm$ 0.001	0.130 $\pm$ 0.003
ECGMatch [23]	30%	0.788 $\pm$ 0.039	0.072 $\pm$ 0.011	0.065 $\pm$ 0.001	0.069 $\pm$ 0.001	0.635 $\pm$ 0.049	0.092 $\pm$ 0.014	0.143 $\pm$ 0.030	0.144 $\pm$ 0.029
(TPAMI'23)	60%	0.791 $\pm$ 0.033	0.122 $\pm$ 0.039	0.098 $\pm$ 0.042	0.103 $\pm$ 0.036	0.651 $\pm$ 0.185	0.132 $\pm$ 0.097	0.159 $\pm$ 0.125	0.159 $\pm$ 0.125
Adello [14]	30%	0.830 $\pm$ 0.008	0.152 $\pm$ 0.007	0.093 $\pm$ 0.005	0.104 $\pm$ 0.004	0.771 $\pm$ 0.043	0.204 $\pm$ 0.038	0.135 $\pm$ 0.028	0.139 $\pm$ 0.027
(ECCV'24)	60%	0.818 $\pm$ 0.006	0.165 $\pm$ 0.006	0.102 $\pm$ 0.002	0.112 $\pm$ 0.004	0.739 $\pm$ 0.028	0.225 $\pm$ 0.025	0.152 $\pm$ 0.017	0.155 $\pm$ 0.018
TS-TFC [8]	30%	0.736 $\pm$ 0.045	0.175 $\pm$ 0.042	0.139 $\pm$ 0.025	0.139 $\pm$ 0.025	0.783 $\pm$ 0.013	0.148 $\pm$ 0.011	0.103 $\pm$ 0.012	0.105 $\pm$ 0.010
(AAAI'23)	60%	0.759 $\pm$ 0.030	0.162 $\pm$ 0.022	0.122 $\pm$ 0.016	0.123 $\pm$ 0.016	0.775 $\pm$ 0.002	0.147 $\pm$ 0.004	0.109 $\pm$ 0.004	0.111 $\pm$ 0.002
CompleMatch [9]	30%	0.750 $\pm$ 0.017	0.160 $\pm$ 0.014	0.128 $\pm$ 0.013	0.129 $\pm$ 0.012	0.795 $\pm$ 0.006	0.135 $\pm$ 0.006	0.094 $\pm$ 0.002	0.097 $\pm$ 0.001
(TPAMI'25)	60%	0.734 $\pm$ 0.032	0.183 $\pm$ 0.018	0.135 $\pm$ 0.015	0.137 $\pm$ 0.013	0.786$\pm$0.011	0.124 $\pm$ 0.038	0.094 $\pm$ 0.014	0.096 $\pm$ 0.014
CaliMatch [2]	30%	<u>0.845$\pm$0.002</u>	0.039$\pm$0.010	<u>0.043$\pm$0.006</u>	<u>0.041$\pm$0.005</u>	<u>0.819$\pm$0.010</u>	<u>0.057$\pm$0.005</u>	<u>0.064$\pm$0.010</u>	<u>0.064$\pm$0.010</u>
(ICCV'25)	60%	0.831 $\pm$ 0.004	<u>0.054$\pm$0.004</u>	<u>0.046$\pm$0.003</u>	<u>0.049$\pm$0.003</u>	0.758 $\pm$ 0.006	<u>0.103$\pm$0.021</u>	<u>0.095$\pm$0.001</u>	<u>0.096$\pm$0.002</u>
SafeECGMatch	30%	0.852$\pm$0.005	<u>0.045$\pm$0.015</u>	0.036$\pm$0.003	0.037$\pm$0.007	0.829$\pm$0.009	0.052$\pm$0.014	0.057$\pm$0.003	0.057$\pm$0.003
(Ours)	60%	0.843$\pm$0.009	0.047$\pm$0.011	0.043$\pm$0.006	0.042$\pm$0.007	<u>0.780$\pm$0.018</u>	0.086$\pm$0.027	0.079$\pm$0.018	0.077$\pm$0.018

B Sensitivity Analysis

Table 4: Sensitivity analysis of λ_{sum} on PTB-XL with 60% OOD. Each experiment is repeated with three random seeds. Higher ACC is better, while lower ECE, ACE, and SCE indicate better calibration.

λ_{sum}	ACC $\uparrow$	ECE $\downarrow$	ACE $\downarrow$	SCE $\downarrow$
$\lambda_{\text{sum}} = \frac{1}{3}$	0.842 $\pm$ 0.009	0.056 $\pm$ 0.008	0.046 $\pm$ 0.008	0.046 $\pm$ 0.008
$\lambda_{\text{sum}} = 1$	0.843 $\pm$ 0.009	0.047 $\pm$ 0.011	0.043 $\pm$ 0.006	0.042 $\pm$ 0.007
$\lambda_{\text{sum}} = 3$	0.847 $\pm$ 0.010	0.048 $\pm$ 0.005	0.041 $\pm$ 0.002	0.042 $\pm$ 0.003