

SemiSegECG: A Multi-Dataset Benchmark for Semi-Supervised Semantic Segmentation in ECG Delineation

Minje Park*

VUNO Inc.

Seoul, Republic of Korea

minje.park@vuno.co

Jeonghwa Lim*

VUNO Inc.

Seoul, Republic of Korea

jeonghwa.lim@vuno.co

Taehyung Yu

VUNO Inc.

Seoul, Republic of Korea

taehyung.yu@vuno.co

Sunghoon Joo[†]

VUNO Inc.

Seoul, Republic of Korea

sunghoon.joo@vuno.co

Abstract

Electrocardiogram (ECG) delineation, the segmentation of meaningful waveform features, is critical for clinical diagnosis. Despite recent advances using deep learning, progress has been limited by the scarcity of publicly available annotated datasets. Semi-supervised learning presents a promising solution by leveraging abundant unlabeled ECG data. In this study, we present **SemiSegECG**, the first systematic benchmark for semi-supervised semantic segmentation (SemiSeg) in ECG delineation. We curated and unified multiple public datasets, including previously underused sources, to support robust and diverse evaluation. We adopted five representative SemiSeg algorithms from computer vision, implemented them on two different architectures: the convolutional network and the transformer, and evaluated them in two different settings: in-domain and cross-domain. Additionally, we propose ECG-specific training configurations and augmentation strategies and introduce a standardized evaluation framework. Our results show that the transformer outperforms the convolutional network in semi-supervised ECG delineation. We anticipate that **SemiSegECG** will serve as a foundation for advancing semi-supervised ECG delineation methods and will facilitate further research in this domain. The code repository is available at <https://github.com/bakqui/semi-seg-ecg>.

CCS Concepts

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

Keywords

Semi-supervised learning; Semantic segmentation; Electrocardiogram; Delineation; LUDB; QTDB; PTB-XL; ViT

ACM Reference Format:

Minje Park, Jeonghwa Lim, Taehyung Yu, and Sunghoon Joo. 2025. SemiSegECG: A Multi-Dataset Benchmark for Semi-Supervised Semantic Segmentation in ECG Delineation. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25)*, November 10–14, 2025, Seoul, Republic of Korea. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3746252.3760790>

*Both authors contributed equally to this research.

[†]Corresponding author.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

CIKM '25, Seoul, Republic of Korea

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2040-6/2025/11

<https://doi.org/10.1145/3746252.3760790>

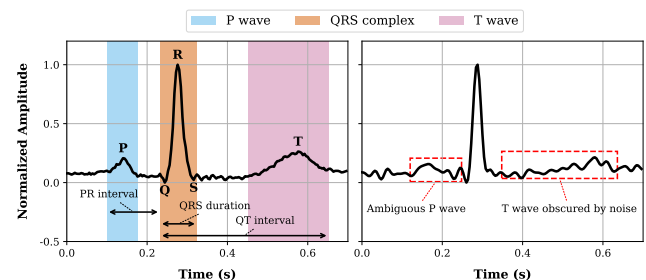


Figure 1: Examples of clean (left panel) and noisy (right panel) ECG recordings. In the left panel, standard waveform components (P wave, QRS complex, and T wave) are annotated, together with the PR interval, QRS duration, and QT interval. In the right panel, an ambiguous P wave and a noise-masked T wave make their boundaries difficult to identify.

1 Introduction

The electrocardiogram (ECG) is a non-invasive, widely used tool for monitoring cardiac electrical activity. A typical ECG waveform includes the P wave, QRS complex, and T wave, representing atrial depolarization, ventricular depolarization, and ventricular repolarization, respectively. Accurate segmentation of these components—termed ECG delineation—is critical for clinical diagnosis [9].

Traditional methods using signal processing (e.g., wavelet transforms [27]) often struggle with signal variability and noise [8, 22], as shown in Figure 1. Deep learning has recently shown promise, treating delineation as a semantic segmentation task [2, 5, 12, 13, 15, 20], but existing models rely on small-scale datasets due to costly expert labeling.

Semi-supervised semantic segmentation (SemiSeg) can bridge this gap by utilizing abundant unlabeled ECG data [10]. In computer vision, SemiSeg methods such as consistency regularization and self-training have proven effective [30]. Nevertheless, their application to ECG delineation faces two critical challenges: (1) the absence of standardized benchmarks, and (2) insufficient evaluation in realistic ECG scenarios.

To address these challenges, we:

- Introduce **SemiSegECG**, the first standardized benchmark for semi-supervised ECG delineation.
- Curate and integrate multiple ECG datasets, including previously underutilized public resources, and develop ECG-specific augmentation and training strategies.
- Evaluate five representative SemiSeg algorithms across varying label availability and distribution shift, considering both

Table 1: ECG database characteristics. Lead types are defined as: 12-lead=standard limb (I, II, III, aVR, aVL, aVF) + precordial (V1–V6) leads; 6-lead=limb-only; 2-lead=two selected leads (e.g., MLI). Label types are lead-specific (separate onset-offset per lead), integrated (single annotation across all leads), and interval-only (PR, QRS, QT intervals without delineation annotation).

Source	# Subjects	# ECGs	Duration (labeled)	Sample rate	Lead type	Label type	# Samples			
							All	Train	Validation	Test
<i>LUDB</i>	200	200	10 s	500 Hz	12-lead	Lead-specific	2,369	1,427	468	474
<i>QTDB</i>	105	105	5.9–253.6 s	250 Hz	2-lead	Integrated	718	422	148	148
<i>ISP</i>	499	499	10 s	1000 Hz	12-lead	Integrated	5,988	3,792	1,272	924
<i>Zhejiang</i>	334	334	1.3–7.1 s	2000 Hz	12-lead	Integrated	4,008	2,400	804	804
<i>PTB-XL</i>	18,885	21,837	10 s	500 Hz	12-lead	-	270,085	270,085	-	-
<i>mECGDB</i>	205	205	2–10 s	250 Hz	6-lead	Interval only	205	-	-	205

the segmentation accuracy and clinically relevant interval error metrics.

2 Benchmark Design

2.1 ECG Databases

Five public ECG databases were curated for benchmarking, with one additional private database. The characteristics of the databases are summarized in Table 1. *LUDB* [16] and *QTDB* [18] have served as standard resources in previous ECG delineation studies [5, 14, 20]. In contrast, *ISP* [1] and *Zhejiang* [43] are relatively new and remain underutilized in delineation research. These four databases, which provide ground-truth delineation annotations, served as the core training resources. Each database (*LUDB*, *QTDB*, *ISP*, *Zhejiang*) was split subject-wise randomly into training, validation, and test sets (6:2:2 ratio except for *ISP*, which provides an official split). *PTB-XL* [40], a large-scale benchmark database for ECG classification, was used as an out-domain unlabeled dataset. We also utilized a private mobile ECG database (*mECGDB*), which consists of 6-lead ECGs measured by portable devices in a non-clinical setting unlike the others, to evaluate model generalization under distribution shift.

2.2 Data Preprocessing

Each ECG lead was treated as an independent training instance, considering its unique spatial orientation and waveform morphology [9], while also enlarging the training set. Labeled ECGs were cropped or zero-padded to a fixed 10-second length, reflecting the standard recording window for resting ECGs in routine clinical practice [33, 38]. All signals were resampled to 250 Hz, which is the lowest native rate among the benchmark databases, to ensure uniform temporal resolution and prevent artifacts introduced by upsampling. A bandpass filter (0.67–40 Hz) was applied to remove baseline drift and high-frequency noise. Z-score normalization was applied to signals before model input to improve training stability.

Official delineation labels were used for *LUDB* and *ISP*. For *QTDB* and *Zhejiang*, we used the label set released by a previous study [12]. This label set compensates for missing beats and incomplete onset-offset annotations in *QTDB* and includes new delineation labels for *Zhejiang*. *mECGDB* contains only clinically relevant interval labels, without explicit delineation labels. All the labels were annotated by at least one expert cardiologist.

2.3 Benchmarking Protocol

SemiSegECG evaluates algorithm performance under two distinct conditions: in-domain and cross-domain settings.

In-domain setting. The in-domain setting reflects a typical use case where both labeled and unlabeled data come from the same source. We simulated semi-supervised conditions by randomly selecting labeled training subsets (1/16, 1/8, 1/4, 1/2) and using the entire training set as unlabeled data. Models were independently trained and evaluated per dataset and label proportion.

Cross-domain setting. The cross-domain setting reflects a practical scenario involving heterogeneous sources across labeled, unlabeled, and potentially test data. The four labeled databases were merged into a unified dataset, preserving the original splits. *PTB-XL* served as an out-domain unlabeled dataset. Models were evaluated on both the merged in-domain test set and *mECGDB*, testing generalization across distribution shifts originating from different measuring environments (e.g., device types).

We compared convolutional [17] and Transformer-based [39] encoders, ResNet [11] and Vision Transformer (ViT) [7], selected for their proven reliability in image and ECG tasks [3, 28, 32, 44]. They were paired with a lightweight fully convolutional network (FCN) decoder [24]. Performance metrics included mean intersection over union (mIoU) for segmentation accuracy and mean absolute error (MAE) of ECG intervals (PR, QRS, QT) for clinical validity. Results were obtained from the model checkpoints that achieved the highest mIoU on the validation set.

3 Experiments

3.1 SemiSeg Algorithms

We benchmarked five SemiSeg algorithms originally developed for computer vision, each representing a distinct learning paradigm: **Mean Teacher (MT)** [36], **FixMatch** [34], **Cross Pseudo Supervision (CPS)** [4], **Regional Contrast (ReCo)** [23], and **Self-Training++ (ST++)** [41].

MT. A student model learns to match predictions from a teacher model whose weights are updated as an exponential moving average (EMA) of the student’s, providing stable pseudo-labels to guide the student’s predictions.

FixMatch. High-confidence predictions from weak augmentations of unlabeled data are used as pseudo-labels to supervise the same inputs with strong augmentations. This encourages robustness across perturbations.

CPS. Two models generate pseudo-labels for each other and are trained mutually, promoting consistency and reducing overfitting through regularized disagreement.

ReCo. Region-level contrastive regularization is applied to pixel embeddings obtained via an additional projection head from the encoder. Pixels with confidence scores between the easy and hard thresholds are selected as queries. These queries are pulled toward class prototypes of their predicted label and pushed away from those of others, sharpening segmentation boundaries.

ST++. With multi-step training, pseudo-labeled data is gradually introduced into training based on confidence and a predefined schedule. Early training focuses on labeled data, reducing noise from uncertain pseudo-labels.

A supervised baseline (**Scratch**) trained only with labeled data was included for comparison.

3.2 Augmentation Strategy Exploration

Effective augmentation is crucial for semi-supervised learning, enabling better utilization of labeled and unlabeled data [42]. However, conventional image-based strategies may not be well-suited for ECG signals [21, 31], distorting ECG-specific characteristics. To address this, we explored augmentation policies tailored to ECG segmentation. Augmentations were grouped as weak (minor global changes used to create pseudo-labels) or strong (larger perturbations applied to training inputs, yet preserving signal structure). Guided by prior studies [19, 29], we chose:

- **Weak:** random resized cropping and horizontal flipping.
- **Strong:** baseline shift, powerline noise, amplitude scaling, sine-wave noise, and white noise.

Optimal weak and strong augmentation strategies were identified on *LUDB* at the 1/16 label ratio with a ResNet backbone (**Scratch** for weak, **FixMatch** for strong) and then applied to all the algorithms.

3.3 Implementation Details

To ensure a fair benchmark, the following configurations were applied commonly across the experiments. The details are available in the public code repository¹.

Model architecture We adopted compact encoders: ResNet-18 [11] and ViT-Tiny [37], chosen to prevent overfitting on our modest-sized datasets while keeping parameter counts comparable. The decoder was implemented as a two-layer FCN consisting of a hidden dimension of 128 and a dropout layer [35] with $p = 0.1$.

Training schedule. All models were trained for 100 epochs with a batch size of 16. We used the AdamW optimizer [26] with a weight decay of 0.05. The learning rate followed a cosine annealing schedule [25], warming up from 0 to 0.001 over the first 10 epochs and gradually decreasing to 0.0001.

Augmentation. All augmentations were applied with a selection probability of 0.5, except for random resized cropping. For strong augmentation, we adopted the RandAugment policy [6] with the searched pool of augmentations.

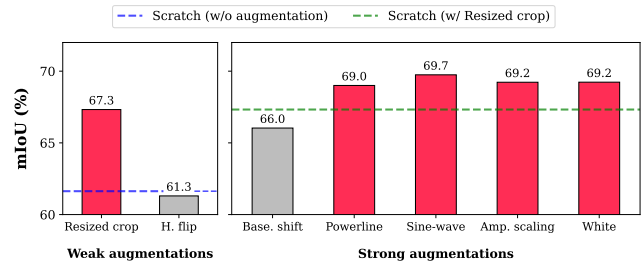


Figure 2: Comparison of augmentation strategies. Left panel: mIoU of supervised models (**Scratch**) using weak augmentations. Right panel: **FixMatch** models with fixed random resized cropping and one of the strong augmentations. Dashed lines indicate **Scratch** performance without any augmentation (blue) and with random resized cropping (green).

SemiSeg-specific hyperparameters. Hyperparameter tuning was conducted on *LUDB* at the 1/16 label ratio using a ResNet-18. For EMA-based models, the decay rate was set to 0.99. The confidence threshold was 0.8 for **FixMatch**, and 0.65 (easy) and 0.8 (hard) for **ReCo**. The projection head used in **ReCo** consisted of two convolutional layers with 128 channels each.

4 Results

4.1 Optimal Augmentation Strategy

We empirically determined the optimal augmentation strategies for all subsequent experiments. Figure 2 illustrates their impact on delineation performance. For weak augmentation, random resized cropping significantly improved **Scratch** performance, while horizontal flipping degraded performance, likely because reversing the signal order (P-QRS-T becomes T-QRS-P) confused the model’s temporal cues. For strong augmentation, baseline shift reduced the performance of **FixMatch**, but powerline noise, sine-wave noise, amplitude scaling, and white noise all led to improvements. The best results arose when RandAugment combined any three of these four strong augmentations per sample (mIoU = 70.9%). Accordingly, we adopted random resized cropping as the weak augmentation and the four-operation RandAugment policy as the strong configuration for all later experiments.

4.2 In-Domain Benchmarking Results

The in-domain benchmarking performances are summarized in Table 2. The results confirm the effectiveness of SemiSeg algorithms on ECG delineation. The gap between SemiSeg algorithms and **Scratch** widened as the label proportion decreased (e.g., 1.7%p → 2.8%p → 5.4%p → 7.6%p in **MT** with ViT on *LUDB*). At the smallest label proportion (1/16), most of the algorithms demonstrated their effectiveness, indicating successful use of unlabeled data when labels were scarce.

From a database perspective, *LUDB* had the fewest cases where SemiSeg algorithms underperformed **Scratch**, suggesting that their effectiveness was more consistently observable, likely due to the presence of accurate, lead-specific labels. From a model perspective, ViT-Tiny outperformed ResNet-18 overall. The detailed trends

¹<https://github.com/bakqui/semi-seg-ecg>

Table 2: In-domain benchmarking results (mIoU, %) on the test sets under varying label ratios (1/16, 1/8, 1/4, and 1/2). Best values are bolded.

Methods	ResNet-18 + FCN				ViT-Tiny + FCN			
	1/16	1/8	1/4	1/2	1/16	1/8	1/4	1/2
<i>LUDB</i>								
Scratch	67.3	71.3	72.9	74.1	66.0	71.3	75.9	78.5
MT	70.8	72.3	73.6	74.3	73.6	76.7	78.7	80.2
FixMatch	70.9	72.2	72.9	74.1	72.2	76.6	78.1	80.1
CPS	68.6	71.6	73.1	74.3	70.9	75.4	78.3	80.0
ReCo	71.5	72.5	73.1	73.9	72.0	74.2	74.9	75.0
ST++	69.2	71.6	73.7	74.5	71.3	76.0	78.0	80.1
<i>QTDB</i>								
Scratch	47.5	56.2	60.5	64.9	38.8	52.5	62.0	67.1
MT	47.8	44.8	63.0	66.7	55.2	58.0	64.9	69.2
FixMatch	46.7	53.3	58.2	66.3	46.7	53.4	64.2	69.3
CPS	53.2	57.1	64.7	68.0	48.9	55.3	58.9	67.4
ReCo	53.4	53.4	58.7	64.5	45.8	47.6	59.4	62.9
ST++	52.9	57.8	62.5	68.0	52.1	55.3	65.7	68.8
<i>ISP</i>								
Scratch	62.3	64.6	68.1	69.3	69.4	73.0	76.3	78.5
MT	64.2	65.9	68.0	69.1	74.8	75.7	78.3	78.6
FixMatch	63.1	64.4	67.9	68.9	74.7	75.7	78.0	79.4
CPS	64.4	66.1	68.5	69.4	72.9	73.7	78.4	79.9
ReCo	62.3	64.7	67.4	68.1	71.3	73.2	76.3	77.2
ST++	63.7	65.7	68.5	69.5	72.8	75.4	78.3	79.4
<i>Zhejiang</i>								
Scratch	76.7	78.9	80.7	82.3	69.5	73.3	76.2	80.2
MT	79.2	80.1	81.5	82.9	79.1	81.4	81.6	83.6
FixMatch	79.0	80.4	81.2	82.7	76.1	80.2	82.0	83.5
CPS	77.6	79.5	81.3	82.7	74.2	80.2	81.7	82.8
ReCo	77.6	78.9	79.7	80.6	68.6	68.3	68.7	63.3
ST++	78.8	79.8	81.9	82.9	73.9	79.0	81.3	83.5

differ: ViT-Tiny consistently performed well with MT, while ST++ remained competitive in ResNet-18. In general, ReCo showed limited benefit, notably on *QTDB*, which exhibited performance degradation at all label ratios except for 1/16.

4.3 Cross-Domain Benchmarking Results

The benchmarking performance patterns in the cross-domain setting deviated from those in the in-domain setting. SemiSeg algorithms provided limited benefit when using ResNet-18 in the cross-domain setting, as shown in Table 3 and Table 4. All the SemiSeg algorithms failed to outperform the **Scratch**, indicating poor generalization. In contrast, ViT-Tiny consistently benefited from SemiSeg algorithms. **MT**, **FixMatch**, and **ST++** achieved remarkable gains, with **ST++** yielding the highest mIoU (84.7%) and **MT/FixMatch** showing the lowest average ECG interval MAE (14.9 ms).

Notably, **FixMatch** with a ViT backbone, one of the best models on the merged in-domain test set, did not lead on *mECGDB*, surpassed by **MT** in the average MAE (20.7 ms < 21.6 ms). It likely reflected the distribution shift between clinical 12-lead ECGs and mobile ECGs collected outside clinical settings. The models may have overfitted to the large-scale unlabeled set (*PTB-XL*), limiting cross-domain robustness. Furthermore, the model with the highest mIoU did not produce the lowest average MAE. These findings

Table 3: Cross-domain benchmarking results on the merged in-domain test set (*LUDB*, *QTDB*, *ISP*, *Zhejiang*). Models were trained on the merged in-domain labeled data and *PTB-XL* unlabeled data. Best values are bolded (\uparrow : higher is better, \downarrow : lower is better).

Methods	mIoU (%) \uparrow	MAE (ms) \downarrow			
		Avg.	PR	QRS	QT
ResNet-18 + FCN					
Scratch	74.5	20.6	21.4	14.1	26.5
MT	73.9	21.5	22.7	15.6	26.3
FixMatch	73.6	21.5	24.3	13.6	26.5
CPS	74.4	21.1	21.2	15.0	27.2
ReCo	73.7	22.3	21.9	16.6	28.3
ST++	74.2	20.8	21.5	14.2	26.8
ViT-Tiny + FCN					
Scratch	81.7	18.5	21.4	10.3	23.8
MT	84.6	14.9	13.1	9.7	21.9
FixMatch	84.4	14.9	13.4	9.9	21.4
CPS	84.0	15.3	15.4	10.2	20.4
ReCo	84.1	15.9	13.8	10.4	23.5
ST++	84.7	15.3	14.4	9.7	21.9

Table 4: Out-domain generalization results on *mECGDB*. Models were trained on the merged in-domain labeled data and *PTB-XL* unlabeled data. Only MAE (ms) was reported, as *mECGDB* contained interval labels (PR, QRS, QT) without delineation annotation. Best values are bolded.

Methods	ResNet-18 + FCN				ViT-Tiny + FCN			
	Avg.	PR	QRS	QT	Avg.	PR	QRS	QT
Scratch	26.4	26.0	18.7	34.6	23.2	27.1	13.3	29.2
MT	28.2	27.1	20.6	36.9	20.7	20.8	13.0	28.2
FixMatch	27.7	28.1	19.5	35.4	21.6	21.4	13.8	29.5
CPS	28.1	28.2	20.3	35.9	21.1	23.2	13.6	26.6
ReCo	28.7	26.1	22.8	37.3	21.0	19.5	13.2	30.3
ST++	27.8	27.4	19.5	36.6	21.6	23.4	13.6	27.8

highlight the need for domain adaptation techniques and a multi-metric evaluation that considers both segmentation quality and clinical interval accuracy.

5 Conclusion

We presented **SemiSegECG**, a unified benchmark and evaluation protocol for semi-supervised ECG delineation. Across diverse public datasets, SemiSeg algorithms consistently improved delineation performances under label scarcity, with ViT backbones providing the clearest gains. Meanwhile, the inconsistency in model performance across datasets between the in-domain merged test set and *mECGDB* highlights the impact of distribution shifts and the need for domain-aware training and ECG-specific augmentation strategies. Limitations include the modest size of annotated datasets and the scope of evaluated methods, which focused on representative SemiSeg algorithms originally developed for computer vision. Future work may leverage the **SemiSegECG** to explore more advanced or ECG-specific semi-supervised learning approaches and domain adaptation techniques tailored to the challenges of physiological signals.

6 GenAI Usage Disclosure

ChatGPT was used to refine sentence structure and assist in evaluating the clarity and soundness of paragraph-level writing. All original content and ideas were created by the authors. Grammarly was also used for additional grammar and fluency checks.

References

- [1] Aram Avetisyan, Nikolas Khachaturov, Ariana Asatryan, Shahane Tigranyan, and Yury Markin. 2024. ISP ECG delineation dataset. doi:10.5281/ZENODO.11472366
- [2] Long Chen, Zheheng Jiang, Joseph Barker, Huiyu Zhou, Fernando Schindwein, Will Nicolson, G Andre Ng, and Xin Li. 2024. ECGVEDNET: A variational encoder-decoder network for ECG delineation in morphology variant ECGs. *IEEE Transactions on Biomedical Engineering* 71, 7 (2024), 2143–2153.
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2017), 834–848.
- [4] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. 2021. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2613–2622.
- [5] Zhenqin Chen, Mengying Wang, Meiyu Zhang, Wei Huang, Hanjie Gu, and Jinshan Xu. 2023. Post-processing refined ECG delineation based on 1D-UNET. *Biomedical Signal Processing and Control* 79 (2023), 104106.
- [6] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. 702–703.
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=YicbFdNTTy>
- [8] Mohamed Elgendi, Björn Eskofier, Socrates Dokos, and Derek Abbott. 2014. Revisiting QRS detection methodologies for portable, wearable, battery-operated, and wireless ECG systems. *PloS one* 9, 1 (2014), e84018.
- [9] Adam Gacek and Witold Pedrycz. 2011. *ECG signal processing, classification and interpretation: a comprehensive framework of computational intelligence*. Springer Science & Business Media.
- [10] Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. 2000. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *circulation* 101, 23 (2000), e215–e220.
- [11] 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.
- [12] Guillermo Jimenez-Perez, Juan Acosta, Alejandro Alcaine, and Oscar Camara. 2024. Generalising electrocardiogram detection and delineation: training convolutional neural networks with synthetic data augmentation. *Frontiers in Cardiovascular Medicine* 11 (2024), 1341786.
- [13] Guillermo Jimenez-Perez, Alejandro Alcaine, and Oscar Camara. 2019. U-Net architecture for the automatic detection and delineation of the electrocardiogram. In *2019 Computing in Cardiology (CinC)*. IEEE, Page–1.
- [14] Guillermo Jimenez-Perez, Alejandro Alcaine, and Oscar Camara. 2021. Delineation of the electrocardiogram with a mixed-quality-annotations dataset using convolutional neural networks. *Scientific reports* 11, 1 (2021), 863.
- [15] Chankyu Joung, Mijin Kim, Taejin Paik, Seong-Ho Kong, Seung-Young Oh, Won Kyeong Jeon, Jae-hu Jeon, Joong-Sik Hong, Wan-Joong Kim, Woong Kook, et al. 2024. Deep learning based ECG segmentation for delineation of diverse arrhythmias. *PloS one* 19, 6 (2024), e0303178.
- [16] Alena I Kalyakulina, Igor I Yusipov, Viktor A Moskalenko, Alexander V Nikolskiy, Konstantin A Kosonogov, Grigory V Osipov, Nikolai Yu Zolotykh, and Mikhail V Ivanchenko. 2020. LUBD: a new open-access validation tool for electrocardiogram delineation algorithms. *IEEE access* 8 (2020), 186181–186190.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012).
- [18] Pablo Laguna, Roger G Mark, A Goldberg, and George B Moody. 1997. A database for evaluation of algorithms for measurement of QT and other waveform intervals in the ECG. In *Computers in cardiology 1997*. IEEE, 673–676.
- [19] Byeong Tak Lee, Yong-Yeon Jo, Seon-Yu Lim, Youngjae Song, and Joon-myoung Kwon. 2022. Efficient data augmentation policy for electrocardiograms. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 4153–4157.
- [20] Xiaohong Liang, Liping Li, Yuanyuan Liu, Dan Chen, Xinpei Wang, Shunbo Hu, Jikuo Wang, Huan Zhang, Chengfa Sun, and Changchun Liu. 2022. ECG_SegNet: An ECG delineation model based on the encoder-decoder structure. *Computers in biology and medicine* 145 (2022), 105445.
- [21] Jeonghwa Lim, Yeha Lee, Wonseuk Jang, and Sunghoon Joo. 2025. Specialized ECG data augmentation method: leveraging precordial lead positional variability. *Biomedical Engineering Letters* (2025), 1–12.
- [22] Feifei Liu, Chengyu Liu, Xinge Jiang, Zhimin Zhang, Yatao Zhang, Jianqing Li, and Shoushui Wei. 2018. Performance analysis of ten common QRS detectors on different ECG application cases. *Journal of healthcare engineering* 2018, 1 (2018), 9050812.
- [23] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew Davison. 2022. Bootstrapping Semantic Segmentation with Regional Contrast. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=6u6N8WWwYSM>
- [24] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.
- [25] Ilya Loshchilov and Frank Hutter. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Skq89Scxx>
- [26] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg6RiCqY7>
- [27] Juan Pablo Martínez, Rute Almeida, Salvador Olmos, Ana Paula Rocha, and Pablo Laguna. 2004. A wavelet-based ECG delineator: evaluation on standard databases. *IEEE Transactions on biomedical engineering* 51, 4 (2004), 570–581.
- [28] Yeongyeon Na, Minje Park, Yunwon Tae, and Sunghoon Joo. 2024. Guiding Masked Representation Learning to Capture Spatio-Temporal Relationship of Electrocardiogram. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=WcOohbsF4H>
- [29] Naoki Nonaka and Jun Seita. 2020. Data augmentation for electrocardiogram classification with deep neural network. *arXiv preprint arXiv:2009.04398* (2020).
- [30] Adrian Peláez-Vegas, Pablo Mesejo, and Julián Luengo. 2023. A survey on semi-supervised semantic segmentation. *arXiv preprint arXiv:2302.09899* (2023).
- [31] Md Moklesur Rahman, Massimo Walter Rivolta, Fabio Badilini, and Roberto Sassi. 2023. A systematic survey of data augmentation of ECG signals for AI applications. *Sensors* 23, 11 (2023), 5237.
- [32] Antônio H Ribeiro, Manoel Horta Ribeiro, Gabriela MM Paixão, Derrick M Oliveira, Paulo R Gomes, Jéssica A Canazart, Milton PS Ferreira, Carl R Andersson, Peter W Macfarlane, Wagner Meira Jr, et al. 2020. Automatic diagnosis of the 12-lead ECG using a deep neural network. *Nature communications* 11, 1 (2020), 1760.
- [33] Hyo-Chang Seo, Gi-Won Yoon, Segyeong Joo, and Gi-Byoung Nam. 2022. Multiple electrocardiogram generator with single-lead electrocardiogram. *Computer Methods and Programs in Biomedicine* 221 (2022), 106858.
- [34] 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.
- [35] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [36] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* 30 (2017).
- [37] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. 2021. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*. PMLR, 10347–10357.
- [38] Marten E Van den Berg, Peter R Rijnbeek, Maartje N Niemeijer, Albert Hofman, Gerard Van Herpen, Michiel L Bots, Hans Hillege, Cees A Swenne, Mark Eijgelsheim, Bruno H Stricker, et al. 2018. Normal values of corrected heart-rate variability in 10-second electrocardiograms for all ages. *Frontiers in physiology* 9 (2018), 424.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [40] Patrick Wagner, Nils Strodthoff, Ralf-Dieter Boussejot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific data* 7, 1 (2020), 1–15.
- [41] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. 2022. St++: Make self-training work better for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4268–4277.
- [42] Zhen Zhao, Lihe Yang, Sifan Long, Jimin Pi, Luping Zhou, and Jingdong Wang. 2023. Augmentation Matters: A Simple-yet-Effective Approach to Semi-supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11350–11359.
- [43] Jianwei Zheng, Guohua Fu, Kyle Anderson, Huimin Chu, and Cyril Rakovsky. 2020. A 12-Lead ECG database to identify origins of idiopathic ventricular arrhythmia

- containing 334 patients. *Scientific data* 7, 1 (2020), 98.
- [44] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. 2021. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 6881–6890.