This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Parameter-Efficient 12-Lead ECG Reconstruction from a Single Lead

Junseok Lee[1 [0009−0003−6922−5929]], Yeonho Yoo[1 [0000−0002−2636−633X]], Jinkyu Kim[1 [0000−0001−6520−2074]], Dosun Lim[2 [0000−0001−5751−5177]], Gyeongsik Yang[1 [0000−0003−4560−2972]], and Chuck Yoo[1 [0000−0002−1115−1862]]

[1] Korea University, Seoul 02841, Republic of Korea
{jslee,yhyoo}@os.korea.ac.kr,
{jinkyukim,g_yang}@korea.ac.kr, chuckyoo@os.korea.ac.kr
[2] Korea University Anam Hospital, Seoul 02841, Republic of Korea
dslmd@kumc.or.kr

Abstract. With the rise of wearable IoT devices such as smartwatches and smart rings, ECG signals have become more accessible and made cardiovascular monitoring a reality. However, analyzing the ECG signals for complex conditions, such as bundle branch blocks and myocardial infarction, requires multi-lead ECG data. Although various deep learning models for ECG reconstruction have been proposed, they are computationally expensive and unsuitable on resource-constrained wearable IoT devices. To address this challenge, we propose mEcgNet, a parameter-efficient model for reconstructing 12-lead ECG signals from a single lead. mEcgNet introduces a modular deep learning architecture for parameter efficiency and separates the single lead-I signal into multiple frequency segments to improve accuracy. Our experiments demonstrate that mEcgNet significantly reduces the number of parameters and inference time by ∼23.1× and ∼5.4×, respectively, compared to existing state-of-the-art models. Furthermore, it reduces the reconstruction error by ∼22.1%, demonstrating its high accuracy and efficiency.

Keywords: ECG reconstruction · mEcgNet · Frequency-based segment partitioning · Parameter-efficient model · Wearable IoT device

1 Introduction

The growing use of wearable IoT devices, such as smartwatches and smart rings, enables the personalized monitoring of digital health data [4]. In particular, these devices typically have two electrodes to measure voltage; as a result, they can record a lead-I ECG signal that is the voltage difference between two distinct points on the body (e.g., between two fingers). It has been reported that this lead-I signal enables the early detection of cardiovascular diseases (e.g., atrial flutter and atrial fibrillation) [14, 30].

While such symptoms can be detected by analyzing lead-I data, more complex conditions and diseases require a full set of ECG signals which consist of several

key waveforms, such as P-wave, QRS complex, and T-wave, and each represents a different cardiac cycle phase [7]. For example, diseases like bundle branch blocks and myocardial infarction are diagnosed by QRS complex abnormalities and T-wave changes [6, 23], which are not observable with a single lead-I signal; they can only be accurately detected and localized using a 12-lead ECG [25, 16]. However, wearable IoT devices can only measure the lead-I signal, and there are a couple of studies that attempt to reconstruct 12-lead ECG signals from lead-I inputs using deep learning models. For example, EKGAN [8] used a generative adversarial network (GAN) with two generators—a 5-layer U-Net and a 5-layer autoencoder—while ECGrecover [13] used a 5-layer U-Net.

However, these studies face a major challenge: they rely on large-scale deep learning models that demand significant computational resources. In contrast, wearable IoT devices typically operate under limited power, memory, and computational constraints [19], making it impractical to deploy such models on these devices. One possible solution is to run the model in the cloud [18]. In this approach, ECG signals are transmitted from the devices to remote servers, where they are then reconstructed. However, this approach raises significant privacy concerns because ECG data itself is highly sensitive personal information and could potentially be used to identify individuals [20, 2].

To address the challenge, we propose mEcgNet, an efficient and accurate ECG reconstruction model. Instead of relying on a single large-scale deep learning model, mEcgNet decomposes the reconstruction process into multiple stages, each handled by a small and modular model. Our goal is to perform local reconstruction on the device without compromising accuracy. Specifically, it partitions ECG signals into distinct frequency segments with each segment processed by a dedicated model, termed EcgModule. This modular approach uses far fewer parameters than a single large model, while the frequency segments preserve the accuracy by capturing key ECG characteristics. We evaluate mEcgNet on two public datasets, PTB-XL [29] and Chapman-Shaoxing [31], and demonstrate that mEcgNet reduces the number of parameters and computational complexity (GFLOPs) by $\sim$23.1$\times$ and $\sim$5.5$\times$, respectively, compared to state-of-the-art (SOTA) models, EKGAN and ECGrecover. Also, our results show that the inference on wearable IoT devices becomes $\sim$5.4$\times$ faster. Furthermore, mEcgNet achieves $\sim$22.1% lower MSE in ECG reconstruction.

2 mEcgNet Architecture

Figure 1 illustrates the architecture of mEcgNet, which consists of two parts: 1) frequency-based segment partitioning that takes the lead-I signal as input and divides it into multiple frequency segments, and 2) EcgModules, modular deep learning models that sequentially process the partitioned segments and accumulate intermediate results to reconstruct the 12-lead ECG signals. The details are explained in the following subsections.

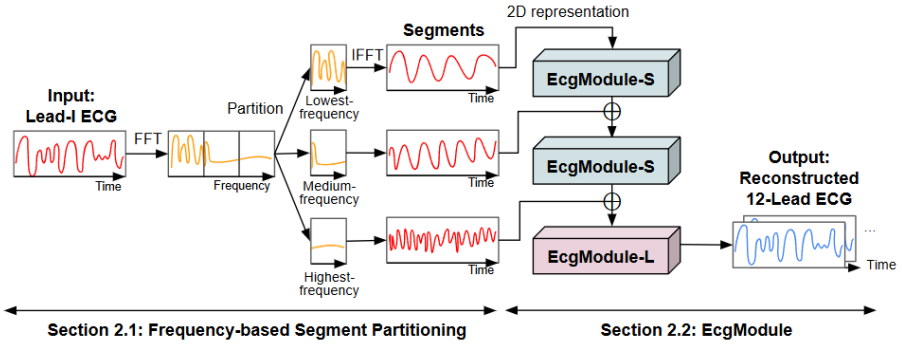

Fig. 1: mEcgNet architecture. The given lead-I ECG is first partitioned into three segments, and EcgModules reconstruct the 12-lead ECG from these segments.

2.1 Frequency-based Segment Partitioning

For parameter-efficient ECG reconstruction, we leverage the intrinsic characteristics of ECG signals. ECG signals that record the heart's electrical activity contain a wide range of frequency segments [26]. Specifically, low-frequency segments capture long-term physiological changes (e.g., circadian rhythms and respiration cycles) and the slow morphological features of heartbeats (e.g., P and T waves). High-frequency characteristics show the rapid transitions of the major heartbeat spikes (e.g., QRS complex). Thus, by analyzing the frequency segments of the lead-I signal, we can directly extract the critical diagnostic information for clinical decision-making [17].

However, existing techniques typically process ECG signals using a single large model without explicitly accounting for their distinct frequency characteristics [8, 13]. Instead, we propose partitioning the original ECG signal into multiple frequency segments in order to capture physiological changes and key ECG features specified by individual frequency segments. Specifically, a lead-I ECG signal is recorded as voltage over time, which represents the waveform of the heartbeat in sequence. mEcgNet converts this signal into a frequency-based representation (amplitude per frequency) using fast Fourier transform (FFT).

We partition the converted ECG signal into three segments: 1) the lowest-frequency segment, 2) the medium-frequency segment, and 3) the highest-frequency segment. We test different numbers of segments and find that using three segments provides the best accuracy for ECG reconstruction (details in §3.4). This choice also aligns with the physiological characteristics of ECG signals, where key waveforms are distributed across three distinct frequency segments: P/T waves lie in 0.5–10Hz, QRS complex in 8–40Hz, and anomalies above 40Hz [5]. After partitioning, each segment is converted back into the time domain using the inverse fast Fourier transform (IFFT), which restores the data back to heart electrical activity over time as lead-I ECG signals. This inverse transform

ensures that the input format for mEcgNet matches the format it is designed to reconstruct as its final output.

2.2 EcgModule

We design the EcgModule, a model for ECG reconstruction aimed at improving both efficiency and accuracy. As shown in Figure 1, we use three EcgModules that process the segments sequentially. The first EcgModule processes the lowest-frequency segment and then passes its output to the next EcgModule. Afterward, the second EcgModule takes 1) the medium-frequency segment and 2) the output from the first EcgModule as input. The third EcgModule performs a similar process, taking the highest-frequency segment along with the output from the second EcgModule, and finally produces the complete 12-lead ECG signals.

The EcgModule is built upon the U-Net architecture [22] because we find that it achieves higher accuracy than other architectures (§3.4). U-Net consists of two parts: an encoder and a decoder. The encoder extracts hierarchical features by downsampling the input, and the decoder upsamples the encoded representations, which reconstructs the 12-lead ECG signals in our model. We design both the encoder and decoder using multiple convolutional blocks. Since the final output of the model consists of 12 separate ECG signals, both the input and output shapes are two-dimensional (2D) representations of 12×512, where 12 denotes 12 different ECG signals and 512 represents the length of each signal. We design input and output shapes identical to ensure the consistency for ECG reconstruction. Note that the value 512, which represents the length of each signal, is consistent with previous work [8, 13]; however, we set the signal dimension to 12 by removing redundant zero padding, thereby improving the parameter efficiency. Specifically, for the first EcgModule, we start with the lowest-frequency segment representing 1×512 and replicate it 12 times to form a 2D representation as input of 12×512.

To make mEcgNet scalable, we design two types of EcgModule: EcgModule-L (large module) and EcgModule-S (small module). As Figure 1, EcgModule-S processes the lowest- and medium-frequency segments while EcgModule-L the highest-frequency segment. We observe that this design is critical to achieving the best efficiency and accuracy together (§3.4). The details of the two types in Figure 2 are explained below.

EcgModule-L: We design EcgModule-L based on the U-Net architecture. Each convolutional block consists of convolution (Conv2D) layer, batch normalization (BatchNorm), and ReLU activation (RELU), except in the decoder's final convolutional block where we omit BatchNorm and RELU. This design makes EcgModule-L more compact than U-Net, as EcgModule-L is to handle only the highest-frequency segment, and the number of parameters of EcgModule-L is reduced by 89.6% in comparison with U-Net.

We determine hyperparameters by grid search. Through the grid search, the number of convolutional blocks in each encoder and decoder is set to three. We determine the convolutional block in the encoder to downsample the height and width of the 2D representation by two. Also, the convolutional block in the

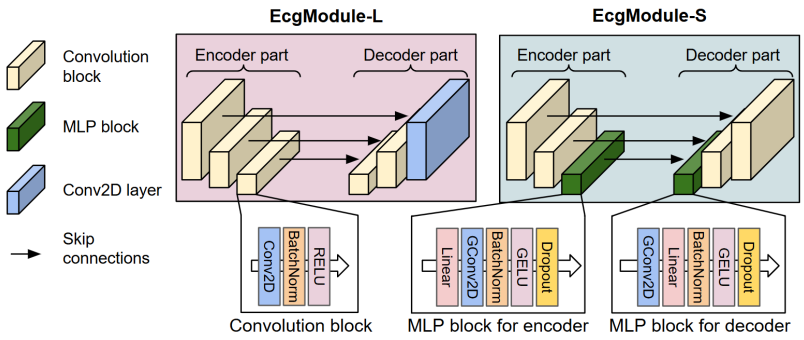

Fig. 2: Structures of EcgModule-L and EcgModule-S.

decoder upsamples the 2D representation by two. The number of channels of the Conv2D layer in the encoder is 64, 128, and 256, while in the decoder, they are 256, 128, and 64. We set each Conv2D layer to have a stride of (2,2), a kernel size of (2,4), and a padding of (0,1), except for the first convolutional block in the decoder, which uses a kernel size of (3,2) and a padding of (0,0).

EcgModule-S: We design EcgModule-S similar to EcgModule-L but with one key difference: the last convolutional block of the encoder and the first convolutional block of the decoder are replaced with MLP blocks. This design is to further reduce the number of parameters over EcgModule-L.

The MLP blocks of both the encoder and decoder consist of grouped Conv2D (GConv2D), linear layer, BatchNorm, GELU activation (GELU), and dropout. GConv2D reduces the number of parameters and computational complexity by processing the input channels in groups rather than individually; however, this group-wise processing can limit the integration of features across channel groups [10]. We find that the linear layer can compensate for this limitation by facilitating feature integration across channel groups. As shown in Figure 2, the encoder starts with the linear layer followed by GConv2D, while the decoder begins with GConv2D followed by the linear layer. The encoder's MLP block uses GConv2D with a kernel size of (1,4), a stride of (1,4), and a group size of 64, while the decoder's MLP block uses GConv2D with a kernel size of (1,1), a stride of (1,1), and a group size of 4, both with a padding of (0,1).

Compared to conventional U-Net, EcgModule-S with our MLP blocks reduces the total number of parameters by 96.3%. Notably, when applied to the lowest- and medium-frequency segments, our evaluation results show that this modification does not degrade accuracy (§3.4).

3 Experiments

3.1 Setup

Datasets: To train and evaluate mEcgNet, we use the PTB-XL dataset and the Chapman-Shaoxing dataset. PTB-XL consists of 21,837 records of 10 s 12-

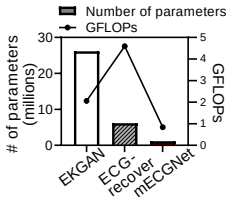

Fig. 3: Efficiency: parameter number and complexity.

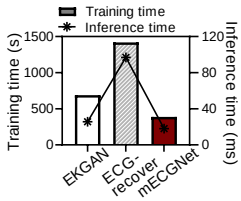

Fig. 4: Time: training and inference time.

Table 1: Accuracy: reconstruction error.

Models	PTB-XL		Chapman-Shaoxing	
	MSE	MAE	MSE	MAE
EKGAN	0.120	0.259	0.123	0.261
ECGrecover	0.145	0.310	0.120	0.268
mEcgNet	**0.113**	**0.257**	**0.109**	**0.252**

lead ECGs sampled at 500 Hz. Also, Chapman-Shaoxing dataset contains 10,646 records of 10 s ECGs sampled at 500 Hz. As these datasets originate from clinical measurements, they contain noise [1]. So we perform dataset preprocessing, following similar methods from other studies [8, 13]. Specifically, we apply min-max scaling to normalize each signal to the range of [-1,1], and we use a band-pass filter with lower and upper cutoff frequencies of [0.05, 150] Hz. We then downsample each signal to 512. This preprocessing is to preserve the ECG features while only removing high-frequency noise and baseline wandering [15, 11].

Implementation and Evaluation Details: We compare mEcgNet with two SOTA techniques: EKGAN and ECGrecover. For fair comparisons, all models are trained using the same method described in [8]. Each model is trained for 10 epochs with an initial learning rate of 1×10^{-4} for the first 5 epochs, followed by a weight decay of 0.95 applied at each subsequent epoch. We use a batch size of 512 for model training. All results are averaged across a 5-fold cross-validation to ensure robustness.

Machines: For training, we use a machine equipped with an NVIDIA A30 GPU. The inference is conducted on a Jetson Orin Nano using only a CPU. All metrics are measured on the Jetson Orin Nano, except for the training time.

3.2 Model Performance

We evaluate the model performance across three aspects: 1) model efficiency, 2) training and inference time, and 3) accuracy. For model efficiency, we report the number of parameters and computational complexity measured in floating point operations (GFLOPs). For time, we measure the average duration for 1) training until convergence and 2) inference (12-lead reconstruction) for a single ECG instance. For accuracy, we present the reconstruction error, which is the difference between the ground-truth 12-lead ECG signals from the dataset and the reconstructed signals, using mean squared error (MSE) and mean absolute error (MAE), which are standard metrics in signal reconstruction tasks [12, 27].

Model efficiency: In Figure 3, the left y-axis shows the number of parameters (in millions, bars), while the right y-axis shows GFLOPs (a line with dots). mEcgNet achieves the smallest number of parameters at 1.13M, requiring

Table 2: Disease classification performance.

Models	# of parameters (millions)	GFLOPs	F1-score
Baseline (ground-truth ECG)	-	-	0.70
EKGAN	32.54	5.71	0.53
ECGrecover	12.57	8.24	0.26
mEcgNet	**7.55**	**4.49**	**0.66**

23.1× fewer parameters than EKGAN and 5.4× fewer than ECGrecover. Also, mEcgNet exhibits the lowest complexity at 0.84 GFLOPs, achieving a 2.5× improvement over EKGAN and a 5.5× improvement over ECGrecover.

Time: Figure 4 shows training time (bars, left y-axis) and inference time (line with dots, right y-axis). For training, mEcgNet takes 387.24 s, which is 1.8× and 3.7× faster than EKGAN and ECGrecover, respectively. For inference, mEcgNet also exhibits the shortest time at 17.99 ms, which is 1.4× and 5.4× faster than the other two, respectively.

Accuracy: Table 1 shows the ECG reconstruction errors. Note that instead of using the reported values from the original papers, we reproduce the errors for the two SOTA models because 1) the two used different datasets, and 2) EKGAN used a non-public private dataset that cannot be tested by others. For the PTB-XL dataset, mEcgNet achieves the lowest errors. Specifically, for MSE, its error is 5.8% and 22.1% lower than EKGAN and ECGrecover, respectively, and for MAE, it is 0.8% and 17.1% lower than the two models. For the Chapman-Shaoxing dataset, mEcgNet also demonstrates the lowest errors; its MSE value is 11.4% and 9.2% lower than EKGAN and ECGrecover, respectively, and its MAE value is 3.4% and 6% lower, respectively.

3.3 Disease Classification Performance

We demonstrate how to apply the reconstruction models for disease classification. We use a clinically validated disease classification model from [21] that takes 12-lead ECGs as input and detects abnormalities associated with six cardiovascular diseases. Since the classification model requires 12-lead ECGs of length 4,096, we linearly interpolate the reconstructed ECGs (length of 512 as [8, 13]) to 4,096. We evaluate model efficiency and accuracy for the combined ECG reconstruction and disease classification models. Accuracy is measured as the average F1-score for six diseases.

Table 2 shows the disease classification performance. We present the accuracy and model efficiency on the test dataset from [21]; the other datasets show similar outcomes. The baseline in the table means that the ground-truth 12-lead ECG signals are fed into the classification model without ECG reconstruction. mEcgNet demonstrates better efficiency, achieving on average 3× fewer parameters and 1.6× fewer GFLOPs than EKGAN and ECGrecover. In terms of accuracy, mEcgNet's F1-score differs from the baseline by only 5.7%, and it is 24.5% and 153.8% higher than EKGAN and ECGrecover, respectively.

Table 3: Ablation study. FSP: frequency-based segment partitioning.

Models	Params	GFLOPs	MSE
U-Net (baseline structure)	6.23	0.96	0.121
EcgModule-L + FSP	1.97	0.97	0.113
Full mEcgNet (EcgModule-L + EcgModule-S + FSP)	**1.13**	**0.84**	**0.113**

Table 4: Ablation study results for hyperparameter choices.

Models	MSE
U-Net [22]	**0.121**
Transformer [28]	0.123
LSTM autoencoder [24]	0.152
Autoencoder [3]	0.242
VAE [9]	0.279

(a) Model.

Number of segments	MSE
2	0.1179
3	**0.1132**
4	0.1133
5	0.1149

(b) Segments.

Order	MSE	Order	MSE
S, S, S	0.122	L, S, S	0.117
S, S, L	**0.113**	L, L, S	0.113
S, L, L	0.116	L, S, L	0.114
L, L, L	0.113	S, L, S	0.119

(c) Modular architecture.

3.4 Ablation Study

First, we evaluate the individual components of mEcgNet by comparing reconstruction errors, as shown in Table 3. We use the PTB-XL dataset since the results on the Chapman-Shaoxing dataset were similar. Specifically, we test: 1) U-Net, the baseline structure of mEcgNet; 2) the EcgModule-L structure with frequency-based segment partitioning; and 3) the full mEcgNet, incorporating both EcgModule-L and EcgModule-S with frequency-based segment partitioning. Table 3 shows that the full mEcgNet (third case) achieves the best efficiency and accuracy among the designs. In comparison with the U-Net baseline structure, the full mEcgNet requires 5.5× fewer parameters, has 1.1× lower GFLOPs, and achieves a 6.6% improvement in MSE.

Next, Table 4 presents the ablation results for different hyperparameter choices. First, Table 4a compares various deep learning architectures commonly used in the medical domain. U-Net, our baseline structure, yields the lowest error. Second, Table 4b shows how accuracy varies when the number of partitioned segments is changed from 2 to 5. The three segments produce the lowest error, so we select three segments. Third, Table 4c reports the errors when the three segments are processed by either EcgModule-S or EcgModule-L. The order in the table indicates the EcgModule types for three segments. For example, the order "S, S, L" means that EcgModule-S, EcgModule-S, and EcgModule-L are used for the lowest-, medium-, and highest-frequency segments. The "S, S, L" order shows the lowest error; thus, we select the order for mEcgNet's structure.

4 Conclusion

In this paper, we introduce mEcgNet, a parameter-efficient model for reconstructing 12-lead ECG signals. A key idea of mEcgNet is to partition frequency

segments and to process them in a modular structure. Experiments on public ECG datasets demonstrate that it reduces the number of parameters, inference time, and reconstruction errors by ∼23.1×, ∼5.4×, and ∼22.1%, respectively, compared to SOTA models. In this work, we find that mEcgNet is lightweight enough to achieve fast yet accurate inference even on wearable IoT devices, making it suitable for real-time ECG monitoring. In future work, we will explore federated learning for ECG reconstruction to improve the training efficiency without centralizing personal ECG data.

Acknowledgments. Corresponding authors: Yeonho Yoo, Gyeongsik Yang, and Chuck Yoo. This research was partly supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (RS-2021-NR060143), by the NRF grant funded by the Korea government (MSIT) (RS-2024-00336564, RS-2023-NR077249), by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by MSIT (RS-2024-00405128), by ICT Creative Consilience Program by IITP grant funded by MSIT (IITP-2025-RS-2020-II201819), by Google Cloud Research Credits program, and by a Korea University Grant.

Disclosure of Interests. The authors declare that there are no competing interests.

References

1. Alfaouri, M., Daqrouq, K.: ECG signal denoising by wavelet transform thresholding. American Journal of applied sciences **5**(3), 276–281 (2008)
2. Arias, O., Wurm, J., Hoang, K., Jin, Y.: Privacy and security in Internet of things and wearable devices. IEEE Transactions on multi-scale computing systems **1**(2), 99–109 (2015)
3. Bank, D., Koenigstein, N., Giryes, R.: Autoencoders. Machine learning for data science handbook: data mining and knowledge discovery handbook pp. 353–374 (2023)
4. Cancela, J., Charlafti, I., Colloud, S., Wu, C.: Digital health in the era of personalized healthcare: opportunities and challenges for bringing research and patient care to a new level. Digital Health pp. 7–31 (2021)
5. Elgendi, M., Meo, M., Abbott, D.: A proof-of-concept study: Simple and effective detection of p and t waves in arrhythmic ecg signals. Bioengineering **3**(4), 26 (2016)
6. ENGEL, T.R., SHAH, R., DePODESTA, L.A., FRANKL, W.S., KRAUSE, R.L.: T-wave abnormalities of intermittent left bundle-branch block. Annals of Internal Medicine **89**(2), 204–206 (1978)
7. Hurst, J.W.: Naming of the waves in the ecg, with a brief account of their genesis. Circulation **98**(18), 1937–1942 (1998)
8. Joo, J., Joo, G., Kim, Y., Jin, M.N., Park, J., Im, H.: Twelve-lead ECG reconstruction from single-lead signals using generative adversarial networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 184–194. Springer (2023)
9. Kingma, D.P., Welling, M., et al.: Auto-encoding variational bayes (2013)

10. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems **25** (2012)
11. Kwon, O., Jeong, J., Kim, H.B., Kwon, I.H., Park, S.Y., Kim, J.E., Choi, Y.: Electrocardiogram sampling frequency range acceptable for heart rate variability analysis. Healthcare informatics research **24**(3), 198–206 (2018)
12. Lan, E.: Performer: A novel ppg-to-ecg reconstruction transformer for a digital biomarker of cardiovascular disease detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1991–1999 (2023)
13. Lence, A., Granese, F., Fall, A., Hanczar, B., Salem, J.E., Zucker, J.D., Prifti, E.: Ecgrecover: A deep learning approach for electrocardiogram signal completion. p. 2359–2370. KDD '25, ACM (2025)
14. Lin, Y.J., Chuang, C.W., Yen, C.Y., Huang, S.H., Huang, P.W., Chen, J.Y., Lee, S.Y.: Artificial intelligence of things wearable system for cardiac disease detection. In: 2019 IEEE international conference on artificial intelligence circuits and systems (AICAS). pp. 67–70. IEEE (2019)
15. Mahdiani, S., Jeyhani, V., Peltokangas, M., Vehkaoja, A.: Is 50 hz high enough ECG sampling frequency for accurate hrv analysis? In: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 5948–5951. IEEE (2015)
16. Menown, I., Mackenzie, G., Adgey, A.: Optimizing the initial 12-lead electrocardiographic diagnosis of acute myocardial infarction. European heart journal **21**(4), 275–283 (2000)
17. Mohebbi, M., Ghassemian, H.: Prediction of paroxysmal atrial fibrillation based on non-linear analysis and spectrum and bispectrum features of the heart rate variability signal. Computer methods and programs in biomedicine **105**(1), 40–49 (2012)
18. Nasser, A.R., Hasan, A.M., Humaidi, A.J., Alkhayyat, A., Alzubaidi, L., Fadhel, M.A., Santamaría, J., Duan, Y.: IoT and cloud computing in health-care: A new wearable device and cloud-based deep learning algorithm for monitoring of diabetes. Electronics **10**(21), 2719 (2021)
19. Poongodi, T., Krishnamurthi, R., Indrakumari, R., Suresh, P., Balusamy, B.: Wearable devices and IoT. A handbook of Internet of Things in biomedical and cyber physical system pp. 245–273 (2020)
20. Rabhi, E., Lachiri, Z.: Biometric personal identification system using the ECG signal. In: Computing in Cardiology 2013. pp. 507–510. IEEE (2013)
21. Ribeiro, A.H., Ribeiro, M.H., Paixão, G.M., Oliveira, D.M., Gomes, P.R., Canazart, J.A., Ferreira, M.P., Andersson, C.R., Macfarlane, P.W., Meira Jr, W., et al.: Automatic diagnosis of the 12-lead ECG using a deep neural network. Nature communications **11**(1), 1760 (2020)
22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18. pp. 234–241. Springer (2015)
23. Simson, M.B.: Use of signals in the terminal qrs complex to identify patients with ventricular tachycardia after myocardial infarction. Circulation **64**(2), 235–242 (1981)
24. Srivastava, N., Mansimov, E., Salakhudinov, R.: Unsupervised learning of video representations using lstms. In: International conference on machine learning. pp. 843–852. PMLR (2015)

25. Tan, N.Y., Witt, C.M., Oh, J.K., Cha, Y.M.: Left bundle branch block: current and future perspectives. Circulation: Arrhythmia and Electrophysiology **13**(4), e008239 (2020)
26. Tereshchenko, L.G., Josephson, M.E.: Frequency content and characteristics of ventricular conduction. Journal of electrocardiology **48**(6), 933–937 (2015)
27. Toda, D., Anzai, R., Ichige, K., Saito, R., Ueki, D.: Ecg signal reconstruction using fmcw radar and convolutional neural network. In: 2021 20th International Symposium on Communications and Information Technologies (ISCIT). pp. 176–181. IEEE (2021)
28. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
29. Wagner, P., Strodthoff, N., Bousseljot, R.D., Kreiseler, D., Lunze, F.I., Samek, W., Schaeffter, T.: PTB-XL, a large publicly available electrocardiography dataset. Scientific data **7**(1), 1–15 (2020)
30. Yun, D., Yang, H.L., Kwon, S., Lee, S.R., Kim, K., Kim, K., Lee, H.C., Jung, C.W., Kim, Y.S., Han, S.S.: Automatic segmentation of atrial fibrillation and flutter in single-lead electrocardiograms by self-supervised learning and transformer architecture. Journal of the American Medical Informatics Association **31**(1), 79–88 (2024)
31. Zheng, J., Zhang, J., Danioko, S., Yao, H., Guo, H., Rakovski, C.: A 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients. Scientific data **7**(1), 48 (2020)