Arrhythmia Classification Using Continual Normalization-Based Gated Recurrent Unit on the Chapman-Shaoxing Dataset

Athiya Jannati Zumaro[1], Najwa Muthia Aziz[1], Muhammad Nadhil Luthfirrahman[1], Dany Primanita Kartikasari[1], Rizal Setya Perdana[1], Kasyful Amron [1], and Eko Sakti Pramukantoro[1]*

Faculty of Computer Science, Universitas Brawijaya, Malang, Indonesia
ekosakti@ub.ac.id

Abstract. Cardiac arrhythmias are a leading cause of global mortality, and ECG is the primary non-invasive detection tool, yet manual interpretation is slow and prone to inter-observer variability. Deep learning models trained on large-scale ECG datasets face unstable convergence and degraded generalization due to non-stationary signal distributions and inter-patient variability. We propose GRU-CN, a Gated Recurrent Unit model with Continual Normalization, for four-class arrhythmia classification (AFIB, GSVT, SB, SR) on the 45,108-recording Chapman-Shaoxing dataset. GRU-CN uses a two-stage strategy: Stage 1 unsupervised GRU Autoencoder pre-training learns latent representations via encoder-decoder reconstruction; Stage 2 fine-tunes the pretrained encoder with a supervised classification head. Continual Normalization (CN) sequentially applies Group Normalization (GN) then Batch Normalization (BN) in the latent space, stabilizing training across intra- and inter-sample dimensions. Across five seeds (42–46), GRU-CN reaches a mean accuracy of 90.18% ± 4.17, macro F1 of 0.8909 ± 0.0463, and macro AUC of 98.17% ± 1.37, outperforming GRU-BN (88.92%), GRU-GN (88.06%), and a non-normalized baseline which was highly unstable across seeds (mean 67.47% ± 21.96, range 36.98–87.34%; McNemar's test, $p < 0.05$ for all comparisons on a representative run). Ablation confirms autoencoder pretraining yields a substantial accuracy gain. Training loss is largely reproducible across seeds (std = 0.102), though one run converged to a lower local optimum, showing CN reduces but does not fully eliminate instability. Overall, combining CN with two-stage learning stabilizes training and improves arrhythmia classification on large-scale ECG data.

Keywords: Gated Recurrent Unit · Continual Normalization · Autoencoder · Chapman-Shaoxing · ECG

1 Introduction

Cardiovascular diseases (CVDs) are the leading cause of death worldwide, accounting for approximately 39% of all non-communicable disease mortality [1].

© The Author(s) 2026
M. Murtadha et al. (eds.), *Proceedings of the 2026 11th International Conference on Sustainable Information Engineering and Technology (SIET 2026)*, Atlantis Highlights in Computer Sciences 26,
https://doi.org/10.2991/978-94-6239-828-3_26

Among these diseases, arrhythmias are particularly associated with sudden cardiac death and stroke [2]. Electrocardiography (ECG) is the primary diagnostic tool for detecting arrhythmia. However, manual interpretation is time-consuming, error-prone and highly dependent on clinical expertise [3]. In addition, the growing volume of ECG recordings in modern healthcare makes comprehensive manual review increasingly impractical, highlighting the need for reliable automated detection methods.

Recent advances in deep learning have significantly improved automated arrhythmia classification. CNNs effectively capture local spatial patterns, but they have limited ability to model long-range temporal dependencies [4]. Recurrent architectures such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) address this limitation, while GRU achieves comparable performance with lower computational cost [5]. Despite these advances, distribution shift remains a major challenge. ECG signals vary considerably between patients and exhibit non-stationary characteristics, often leading to degraded performance under inter-patient evaluation. This issue is further exacerbated by the class imbalance [6]. Batch Normalization (BN) is susceptible to momentum bias under evolving data distributions [7], whereas Group Normalization (GN) does not exploit inter-sample statistics. Continual Normalization (CN) combines the strengths of both methods through sequential GN and BN operations [8]. Nevertheless, learning robust latent representations directly from raw ECG signals remains challenging.

To address these challenges, we propose GRU-CN, which integrates Continual Normalization (CN) into a GRU-based classifier through a two-stage training strategy. The model is first pretrained in an unsupervised manner using a GRU Autoencoder to learn informative latent ECG representations, after which the pretrained encoder is transferred to a supervised classifier equipped with CN. Experiments on the complete Chapman-Shaoxing dataset (45,108 Lead II ECG recordings) across five independent random seeds achieved a mean accuracy of 90.18% and a mean Macro AUC of 98.17%. This study is limited to a single dataset and a single-lead ECG configuration, and the implications of these limitations are discussed in Section 5.

The main contributions of this work are:

- We propose GRU-CN, combining CN (sequential GN–BN) with a GRU classifier for arrhythmia detection on the *full-scale* Chapman-Shaoxing dataset (45,108 recordings). While CN, GRU, and autoencoder pretraining are in-dividually established, our contribution is their systematic integration and validation at full scale, with a repeated-run reliability analysis not previously reported for CN-based ECG classifiers.
- We introduce a two-stage learning strategy in which a GRU Autoencoder is pretrained on unlabeled ECG signals and its encoder is transferred to the supervised classifier, yielding a measurable accuracy gain.
- We provide a systematic ablation across non-normalized, BN, GN, and CN GRU variants, confirming improvements with paired significance testing and validating GRU-CN's reliability across five independent runs.

2 Related Work

Arrhythmia classification from ECG has used CNN-based models for spatial features [9] and recurrent architectures (RNN, LSTM, BiLSTM, GRU) for temporal modeling, with hybrid CNN-GRU architectures also proposed [10]. Abdullah et al. [10] demonstrated a CNN-GRU hybrid with LIME-based explainability on four MIT-BIH classes, but such hybrids often overlook normalization for distributional instability in large-scale multi-patient databases and do not report repeated-run reliability.

Manjunath et al. [8] first introduced GRU-CN on MIT-BIH (48 recordings, 5 classes, 99.41% accuracy) without pretraining. Sarankumar et al. [11] pro-posed BiGRU with an Autoencoder on CPSC 2018 (10,034 recordings, 99.21% accuracy) without CN. Yaser et al. [12] combined TCN-GCN on a Chapman-Shaoxing subset (10,588 recordings), reaching 95.50% accuracy on four super-classes. None report variance across repeated runs. No prior work has integrated Autoencoder-based latent learning with CN in a GRU framework on the full-scale Chapman-Shaoxing dataset, nor reported multi-seed reliability for such a model—a gap this study addresses.

3 Proposed Method

The proposed GRU-CN framework has four stages: data consolidation, preprocessing, unsupervised feature learning, and supervised classification. Raw 12-lead ECG signals are mapped into four arrhythmia superclasses; a 4th-order Butter-worth bandpass filter and Z-score normalization then remove baseline noise and standardize amplitudes. A GRU Autoencoder is trained unsupervised to extract 128-dimensional latent representations from the 5,000-timestep sequences, and the pretrained encoder weights transfer to a GRU classifier with a CN module to mitigate distribution shifts before predicting the arrhythmia class (Fig. 1).

3.1 Model Architecture: GRU-CN

Stage 1 – Unsupervised Pretraining. A GRU Autoencoder learns latent representations from unlabeled 5,000-timestep ECG signals: the encoder (hidden size 128) compresses the input into a 128-dimensional latent vector from the final hidden state h_{5000}, and a decoder GRU reconstructs the signal, minimizing MSE. Across five seeds the autoencoder consistently reaches a reconstruction R^2 of 0.99989 ± 0.00001, confirming reliable convergence regardless of initialization.

Stage 2 – Supervised Classification. The pretrained encoder weights are transferred to a GRU classifier. The 128-dimensional latent vector is processed by a CN module that sequentially applies GN then BN. GN divides the 128 features into 8 groups of 16 channels and normalizes each group within every sample (robust to small batch sizes); BN is then applied per channel across the mini-batch:

$$\hat{x}_i = \frac{x_i - \mu_g}{\sqrt{\sigma_g^2 + \varepsilon}}, \qquad y_i = \gamma_i \frac{\hat{x}_i - \mu_{Bi}}{\sqrt{\sigma_{B,i}^2 + \varepsilon}} + \beta_i, \tag{1}$$

where γ_i, β_i are learnable affine parameters. Chaining GN and BN adjusts feature distribution at both micro (per-record) and macro (per-batch) levels. A Dropout layer (0.3) precedes a fully connected Softmax layer producing class probabilities for AFIB, GSVT, SB, and SR. Training also monitors for divergence and automatically restarts from the last stable checkpoint; for GRU-CN this triggered on average 1.6 times per run (0–2 restarts) across five seeds, reported transparently in Section 4.

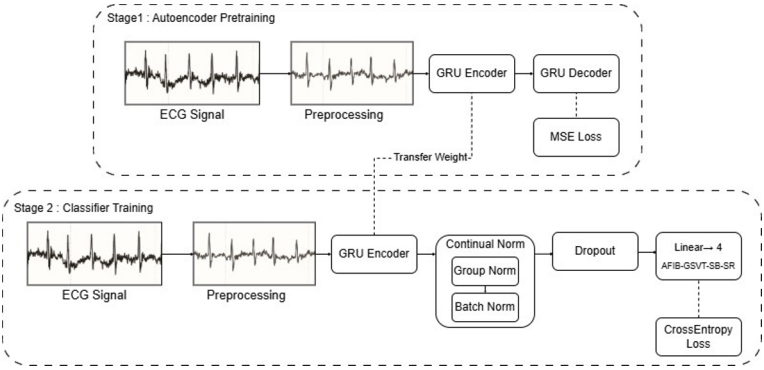


Fig. 1: Two-stage GRU-CN workflow. Stage 1: GRU Autoencoder pretraining via encoder-decoder reconstruction (MSE loss). Stage 2: pretrained encoder weights transfer to the classifier, where the latent representation passes through CN (GN→BN), Dropout, and a linear layer to predict AFIB, GSVT, SB, or SR (Cross-Entropy loss).

3.2 Dataset

The Chapman-Shaoxing dataset contains 45,152 12-lead ECG recordings sampled at 500 Hz for 10 seconds (5,000 time-steps) [13]. Following Yaser et al. [12], all 63 SNOMED-CT codes were consolidated into four superclasses; after excluding 44 unassignable recordings, the final 45,108 recordings were split 80/10/10 (train/val/test) via stratified sampling (Table 1), reused across all five seeds so only initialization and training dynamics vary. Class imbalance (SB: 36.7% vs. SR: 18.0%) was mitigated via Weighted Random Sampling with per-class weight $\omega_i = 1/N_i$.

Each arrhythmia type exhibits a distinct ECG morphology (Fig. 2): Sinus Rhythm (SR) shows a regular waveform with well-defined P, QRS, and T waves; Sinus Bradycardia (SB) shares this morphology but at a heart rate below 60 bpm; Atrial Fibrillation (AFIB) shows absent P-waves and an irregularly irregular ventricular response; and General Supraventricular Tachycardia (GSVT) shows rapid, narrow QRS complexes exceeding 100 bpm.

Table 1: Class Distribution After Superclass Mapping

Superclass	Full (*N*)	Share	Train	Val	Test
AFIB	9,802	21.7%	7,842	980	980
GSVT	10,649	23.6%	8,519	1,065	1,065
SB	16,552	36.7%	13,242	1,655	1,655
SR	8,105	18.0%	6,484	810	811
Total	45,108	100%	36,088	4,510	4,511



(a) Sinus Rhythm (SR) (b) Sinus Bradycardia (SB)

(c) Atrial Fibrillation (AFIB) (d) General Supraventricular Tachycardia (GSVT)

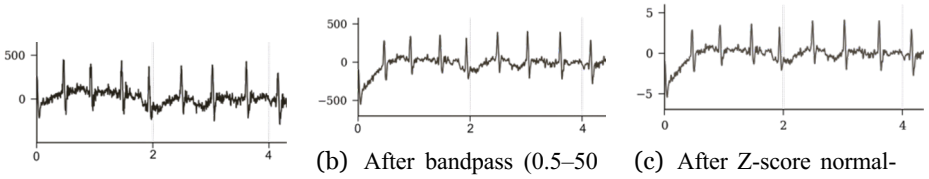
Fig. 2. Representative ECG morphologies for the four arrhythmia superclasses.

3.3 ECG Signal Preprocessing

Raw ECG signals contain baseline wander and high-frequency noise. A 4th-order Butterworth bandpass filter (0.5–50 Hz) is applied via zero-phase filtering (`filtfilt`) to remove these artifacts while preserving cardiac morphology, followed by per-recording Z-score normalization to standardize amplitude across patients (Eq. 2, Fig. 3):

$$x_{\text{norm}} = \frac{x - \mu}{\sigma} \tag{2}$$

where μ and σ are the recording's mean and standard deviation, ensuring in-variance to amplitude variation across sessions.



(a) Raw ECG signal (b) After bandpass (0.5–50 Hz) (c) After Z-score normalization

Fig. 3. ECG preprocessing pipeline applied to a sample data.

3.4 Gated Recurrent Unit

The GRU captures long-range temporal dependencies while remaining lighter than LSTM, using only two gates and no separate cell state [5]. At each time step t, hidden state h_t updates via reset gate r_t and update gate z_t:

$$r_t = \sigma(W_r x_t + U_r h_{t-1}), \tag{3}$$

$$z_t = \sigma(W_z x_t + U_z h_{t-1}), \tag{4}$$

$$\tilde{h}_t = \tanh\ W_h x_t + U_h(r_t \odot h_{t-1}), \tag{5}$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t, \tag{6}$$

where σ is the sigmoid, $\odot$ the element-wise product, and W_*, U_* the learnable weights. The reset gate controls how much past information to discard, and the update gate balances retained and new state.

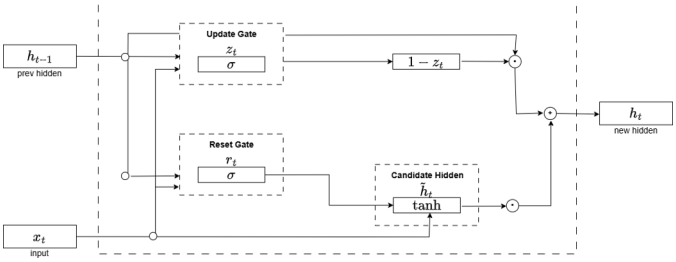


Fig. 4. Internal structure of a GRU cell: reset gate r_t and update gate z_t regulate the flow of past and candidate information into h_t.

3.5 Training Configuration

To prevent vanishing gradients over the 5,000-timestep sequence, recurrent weights used orthogonal initialization, input/classifier weights used Xavier initialization, and biases were zero-initialized. Both stages used Adam (batch size 64); the autoencoder minimized MSE, the classifier minimized Cross-Entropy with Weighted Random Sampling (Table 2). GRU-CN and the plain GRU baseline were each trained five times with different seeds (42–46); GRU-BN, GRU-GN, and the non-pretrained ablation were trained once, a single-run limitation acknowledged in Section 4.1. Experiments used an NVIDIA RTX 4070 (12 GB VRAM), CUDA 12.8, PyTorch 2.9.1, SciPy 1.15.3, and Scikit-learn 1.7.2.

Table 2. Training Hyperparameters

Hyperparameter	Stage 1 (AE)	Stage 2 (Classifier)
Optimizer	Adam	Adam
Learning rate	1×10^{-3}	5×10^{-4}
Batch size	64	64
Epochs	10	30
Loss function	MSE	Cross-Entropy
Hidden size	128	128
Weight init.	Orthogonal (rec.) / Xavier (input)	Xavier
Dropout	–	0.3
Class balancing	–	Weighted Random Sampling

4 Results and Discussion

This section evaluates GRU-CN across four aspects: overall performance, per-class metrics, autoencoder pretraining's contribution, and Continual Normalization's effectiveness, followed by comparison with prior work. GRU-CN and the plain GRU baseline are reported as mean ± std over five runs (seeds 42–46). GRU-BN and GRU-GN, trained once due to compute budget, are reported with 95% bootstrap CIs (n = 2000) on that single run; pairwise improvements against GRU-CN's representative run are validated with McNemar's test. This multi-seed-vs-single-run asymmetry is flagged as a limitation in Section 5.

4.1 Model Performance

Table 3 compares the performance of all model variants across five independent runs. GRU-CN achieved the highest performance on all evaluation metrics. In contrast, the non-normalized GRU exhibited large performance variation, with

accuracy ranging from 36.98% to 87.34% (mean 67.47% ± 21.96). This wide range indicates unstable convergence and a strong dependence on random initialization. GRU-BN markedly improved performance but remained affected by its reliance on batch statistics under high inter-patient variability. GRU-GN removed the dependency on batch statistics but did not utilize inter-sample in-formation. By combining the advantages of both normalization methods, GRU-CN obtained a mean accuracy 1.26 percentage points higher than GRU-BN while maintaining lower variability than the non-normalized model. McNemar's test further showed that the performance improvements of GRU-CN over GRU- BN (χ^2 = 5.53, p = 0.019), GRU-GN (χ^2 = 16.04, p < 0.001), and the non-normalized GRU (χ^2 = 2170.42, p < 0.001) were statistically significant.

Across the five experimental runs, GRU-CN achieved accuracies ranging from 82.78% to 92.46%. Four runs converged to a similar accuracy range of approximately 91–92%, whereas one run (seed 42) converged to a lower local optimum. Even with this variation, GRU-CN showed much lower variability (std = 4.17) than the non-normalized GRU (std = 21.96), indicating more stable convergence across different random initializations.

Table 3. Model performance comparison. GRU baseline and GRU-CN: mean ± std over 5 seeds (42–46). GRU-BN/GRU-GN: single run, 95% bootstrap CI (n = 2000).

Model	Accuracy (%)	Macro F1	Macro AUC (%)
GRU Baseline	67.47 ± 21.96	0.6355 ± 0.2456	84.15 ± 15.37
GRU-BN	88.92 ± 0.47	0.8744 ± 0.0051	98.15
GRU-GN	88.06 ± 0.48	0.8651 ± 0.0053	97.64
GRU-CN (Proposed)	**90.18 ± 4.17**	**0.8909 ± 0.0463**	**98.17 ± 1.37**

4.2 Per-Class Performance

Table 4 reports per-class metrics for GRU-CN averaged over five seeds, reflecting average rather than single-run performance. SB attains the highest F1 (0.9630), benefiting from the clear morphology of slow rhythms. AFIB also performs strongly (F1 = 0.8993), reflecting the model's ability to capture atrial fibrillation's irregular rhythm. SR follows with a balanced precision/recall profile (F1= 0.8769), indicating Weighted Random Sampling reduces majority-class bias for this underrepresented class. GSVT remains the weakest class (F1 = 0.8244), with the lowest, most variable recall (79.48% ± 8.62), suggesting its heterogeneous subtypes remain the main confusion source across seeds.

Table 4. Perc lass performance of GRU-CN, mean over 5 seeds (42–46)

Class	Precision (%)	Recall (%)	F1-Score
AFIB	88.14	91.82	0.8993
GSVT	85.82	79.48	0.8244
SB	95.30	97.34	0.9630
SR	87.72	87.82	0.8769

4.3 Impact of Autoencoder Pretraining

To isolate pretraining's contribution, GRU-CN was compared against a non-pretrained variant trained from scratch. With pretraining, GRU-CN reaches a mean accuracy of 90.18% and macro F1 of 0.8909, versus 82.13% accuracy and 0.8161 macro F1 without pretraining (single run)—a gain of about 8.05 points, consistent with McNemar's test ($\chi^2 = 213.49$, $p < 0.001$). This suggests the encoder, by learning to reconstruct 5,000-step ECG signals unsupervised, produces latent representations that stabilize downstream classification, focusing the classifier on cleaner morphological features (R-R intervals, QRS complexes) rather than raw noise. The autoencoder reconstructs signals with near-perfect fidelity across all five seeds (R^2 = 0.99989±0.00001), supporting a reliably informative latent space regardless of initialization. A limitation: the non-pretrained baseline uses a 2-layer encoder versus GRU-CN's 1-layer encoder, confounding depth with pretraining, and was trained only once, so the reported gain is an upper-bound estimate; a controlled re-run with matched depth is left for future work.

4.4 Effectiveness of Continual Normalization

The ablation in Table 3 shows CN is the primary driver of GRU-CN's gains over single-normalization variants, addressing internal covariate shift on the non-stationary Chapman-Shaoxing data. The sequential GN→BN pipeline uses GN for group-wise, batch-size-independent stability, then BN to leverage batch-level statistics, adjusting feature distribution at both micro and macro levels.

To substantiate the stability claim directly, training-loss trajectories were analyzed across the same five seeds. Final loss has a standard deviation of 0.102 (variance = 0.0103), with mean within-run volatility of 0.111. Together with the accuracy spread in Section 4.1, this indicates CN meaningfully reduces, but does not fully eliminate, training instability: most runs converge to a consistent loss and accuracy band, and an automatic restart-on-divergence safeguard (triggered 1.6 times per run on average) helped recover stable convergence in most seeds, yet one seed still settled at a lower optimum despite this safeguard—a partial, not complete, resolution of the distribution-shift problem; investigating the outlier run is left to future work.

4.5 Comparison with Prior Work

Table 5 compares GRU-CN + AE against representative prior methods using AUC. Hu et al. [14] report 99.71% AUC on Chapman using a multimodal self-supervised framework with dual 1D time-domain and 2D spectrogram pathways. Our supervised model achieves a mean AUC of 98.17% (best single run: 98.98%) on a similarly large dataset (45,108 records) with a more streamlined architecture, avoiding self-supervised pretraining overhead. Sarankumar et al. [11] and Yaser et al. [12] report a comparable 99.00% AUC, but on the 9-class CPSC 2018 database or a smaller 10,588-record Chapman subset, and neither reports variance across repeated runs. The AUC values in Table 5 are **not directly comparable head-to-head**: studies differ in dataset size, class definitions, pre-processing, and split protocol, and none besides ours report multi-run reliability. Table 5 should be read as contextual positioning, not evidence of outright superiority.

Table 5: Comparison with State-of-the-Art

Study	Method	Dataset (records)	Classes	AUC. (%)
Hu et al. [14]	SimECG	Chapman (10,628)	4	99.71
Sarankumar et al. [11]	BiGRU-AE	CPSC 2018 (10,034)	9	99.00
Yaser et al. [12]	TCN-GCN	Chapman (10,588)	4	99.00
Proposed	GRU-CN + AE	Chapman (45,108)	4	98.17 ± 1.37

5 Conclusion

This paper presents GRU-CN, a Continual Normalization-based GRU architecture for arrhythmia classification on the full-scale Chapman-Shaoxing ECG dataset, combining unsupervised GRU Autoencoder pretraining with supervised classification under sequential GN→BN normalization. Across five runs, GRU-CN achieves a mean accuracy of 90.18% ± 4.17, macro F1 of 0.8909 ± 0.0463, and Macro AUC of 98.17% ± 1.37 on 45,108 recordings, outperforming single normalization variants (BN: 88.92%, GN: 88.06%) and a non-normalized base-line that was itself highly unstable (mean 67.47% ± 21.96), with improvements confirmed significant via McNemar's test. The repeated-run analysis shows CN substantially, but not completely, stabilizes training: four of five seeds converge to a consistent accuracy band while one reaches a lower optimum despite an automatic restart safeguard, a nuance reported transparently rather than presenting only best-case numbers. This study is scoped to a single dataset (Chapman-Shaoxing, Lead II) and four superclasses; we do not claim generalizability with-out further validation. The results confirm CN meaningfully mitigates internal covariate shift in large-scale non-stationary ECG data and that two-stage trans-fer learning improves classification accuracy, positioning GRU-CN as a candidate for further study toward Clinical Decision Support Systems for arrhythmia screening. Future work includes multi-lead ECG inputs, Bidirectional GRU, validating generalization on additional datasets (e.g., PTB-XL, CPSC 2018) under a matched protocol, running GRU-BN, GRU-GN, and the non-pretrained ablation across multiple seeds, and re-running the pretraining ablation with matched encoder depth.

References

1. Hinga, I.A.T., et al.: Edukasi Kesehatan dalam Upaya Pencegahan Penyakit Jantung. Jurnal Pengabdian Masyarakat Bangsa **3**(4), 1323–1327 (2025). https://doi.org/10.59837/jpmba.v3i4.2415
2. Botto, G.L., et al.: Impact of the pattern of atrial fibrillation on stroke risk and mortality. Arrhythmia & Electrophysiology Review **10**(2), 68 (2021). https://doi.org/10.15420/aer.2021.01
3. Xiao, Q., et al.: Deep Learning-Based ECG Arrhythmia Classifica-tion: A Systematic Review. Applied Sciences **13**(8), 4964 (2023). https://doi.org/10.3390/app13084964

4. Ullah, A., et al.: A Hybrid Deep CNN Model for Abnormal Arrhyth-mia Detection Based on Cardiac ECG Signal. Sensors **21**(3), 951 (2021). https://doi.org/10.3390/s21030951
5. Tekin, H., Kaya, Y.: A novel approach for arrhythmia classification using CI-1D-LBP with LSTM, 1D-CNN and GRU Models. Bitlis Eren Üniversitesi Fen Bilimleri Dergisi (2024). https://doi.org/10.17798/bitlisfen.1542941
6. Pramukantoro, E.S., Gofuku, A.: A heartbeat classifier for continu-ous prediction using a wearable device. Sensors **22**(14), 5080 (2022). https://doi.org/10.3390/s22145080
7. Pham, Q., Liu, C., Hoi, S.: Continual normalization: Rethinking batch normal-ization for online continual learning. arXiv preprint arXiv:2203.16102 (2022). https://doi.org/10.48550/arXiv.2203.16102
8. Manjunath, C., et al.: A Gated Recurrent Unit Based Continual Nor-malization Model for Arrythmia Classification Using ECG Signals. In: 2023 International Conference on Ambient Intelligence, Knowledge Informat-ics and Industrial Electronics (AIKIIE), pp. 1–5. IEEE, Ballari (2023). https://doi.org/10.1109/AIKIIE60097.2023.10390056
9. Ardiyansyah, A., Mandala, S.: Accurate Model for Arrhythmia Classification Using Hybrid ACNN and GRU with XAI. In: ICOICT, pp. 57–64 (2024). https://doi.org/10.1109/icoict61617.2024.10698622
10. Abdullah, T.A.A., Zahid, M.S.B.M., Tang, T.B., Ali, W., Nasser, M.: Explainable deep learning model for cardiac arrhythmia classification. In: 2022 International Conference on Future Trends in Smart Communities (ICFTSC), pp. 87–92 (2022). https://doi.org/10.1109/ICFTSC57269.2022.10039860
11. Sarankumar, R., Ramkumar, M., Vijaipriya, K., Velselvi, R.: Bidirec-tional gated recurrent unit with auto encoders for detecting arrhyth-mia using ECG data. Knowledge-Based Systems **294**, 111696 (2024). https://doi.org/10.1016/j.knosys.2024.111696
12. Yaser, S., Chowdhury, M.E.H., Sarmun, R.: SER inspired deep learn-ing approach to detect cardiac arrhythmias using TCN and graph neu-ral network. Computers in Biology and Medicine **195**, 110536 (2025). https://doi.org/10.1016/j.compbiomed.2025.110536
13. Zheng, J., et al.: A 12-lead electrocardiogram database for arrhythmia re-search covering more than 10,000 patients. Scientific Data **7**(1), 48 (2020). https://doi.org/10.1038/s41597-020-0386-x
14. Hu, J., Li, C., Cao, J., Kou, B.: A novel multimodal self-supervised framework for ECG arrhythmia classification. Computers in Biology and Medicine **198**, 111137 (2025). https://doi.org/10.1016/j.compbiomed.2025.111137

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

