Prediction of Atrial Fibrillation Onset from Continuous Intensive Care Unit ECG Monitoring: A Deep Learning Approach

Maria Galanty[1,2][0009−0001−2574−1775]⋆, Björn van der Ster[3][0000−0001−5190−5721], Merijn C. Reuland[4][0009−0009−3738−0397], Michiel Hulleman[5,6][0000−0002−8254−214X], Alexander P.J. Vlaar[4][0000−0002−3453−7186], and Clara I. Sánchez[1,2][0000−0001−9787−8319]

[1] Informatics Institute, University of Amsterdam (UvA), the Netherlands (NL)
[2] Department of Biomedical Engineering and Physics, Amsterdam University Medical Center (Amsterdam UMC), NL
[3] Department of Anaesthesiology, Amsterdam UMC, NL
[4] Department of Intensive Care, Amsterdam UMC, NL
[5] Department of Clinical and Experimental Cardiology, Amsterdam UMC, NL
[6] Amsterdam Cardiovascular Sciences, Amsterdam UMC, NL

Abstract. Atrial fibrillation (AF) in the intensive care unit (ICU) is associated with stroke, heart failure, and increased mortality. Although continuous ECG monitoring is routine, this data is not used to predict AF, which could enable preventive treatment. We propose a deep learning framework that combines a Residual Network (ResNet) with Mamba, a state-space model for long-range sequence modelling, to predict AF within a two-hour horizon from raw single-lead ICU ECG. We compare the ResNet–Mamba hybrid with a ResNet-only baseline using 10- and 30-minute input windows on two ICU cohorts (MIMIC-III and Amsterdam UMC). With 30-minute windows, ResNet–Mamba reached a ROC-AUC of 0.846 and an average precision of 0.519, outperforming the ResNet-only model (0.830 and 0.446). These results suggest that selective long-range modelling of raw single-lead ICU ECG is a promising approach for early AF warning within a clinically actionable horizon.

Keywords: Atrial fibrillation · Intensive care unit · ECG

1 Introduction

Atrial fibrillation (AF) is a common cardiac arrhythmia characterised by disorganised atrial electrical activity and an irregular ventricular response [18]. In the intensive care unit (ICU), where patients have complex, rapidly evolving conditions, AF is particularly prevalent and is associated with adverse outcomes including stroke, heart failure, and increased mortality [2, 10, 17]. Continuous electrocardiogram (ECG) monitoring is a standard practice in the ICU, generating longitudinal waveform data, commonly from a limited set of leads. In this

⋆ Corresponding author: m.galanty@uva.nl

clinically complex population, leveraging these data for early prediction of AF could support anticipatory interventions such as preventive treatment to reduce complications. Yet this potential remains largely unrealised.

A substantial body of literature addresses AF detection, predominantly in outpatient and Holter-monitored populations [15, 19]. More recently, attention has shifted toward AF prediction, yet most predictive models continue to target outpatient cohorts [5, 6]. ICU-focused approaches have largely relied on handcrafted clinical or signal features rather than end-to-end learning from raw ECG [1, 8]. This leaves an unexplored gap: a deep learning framework capable of predicting AF onset directly from continuous ICU ECG data.

In this work, we address this gap by proposing a hybrid architecture that combines Residual Networks (ResNet) [15], widely validated for ECG classification, with Mamba [7], a state-space model designed for efficient long-range sequence modelling. This architecture receives as input continuous ICU ECG, particularly lead II, which is commonly available across heterogeneous ICU monitoring configurations. Focusing on clinically stable ICU patients in sinus rhythm (SR), we aim to predict whether they will remain in SR or transition to AF within a predefined horizon. We hypothesise that the ResNet–Mamba combination is particularly well suited to AF prediction, where subtle precursor patterns may span extended temporal windows. Specifically, we investigate whether the ResNet–Mamba hybrid outperforms a ResNet-only baseline and whether longer input windows (10 vs 30 minutes) provide richer temporal context for predicting AF onset within a two-hour horizon for ICU patients.

2 Methods

2.1 Data

Datasets We used two continuous ICU ECG datasets: the public MIMIC-III Waveform Database (folders 30-34) [9, 12] and an internal dataset from Amsterdam University Medical Centre (Amsterdam UMC) [13], both providing continuous multi-hour to multi-day recordings during ICU admission. [1] As lead configurations in the ICU are non-standardised and both datasets store only those leads selected for display on the bedside monitor, available leads vary across patients and sites. Although monitoring is continuous in principle, signal gaps due to patient transport, procedures, or electrode disconnection are common. In all experiments, models were trained on a mixed cohort formed by pooling samples from both datasets. During evaluation, performance is reported jointly as well as separately for MIMIC-III and Amsterdam UMC.

Annotation Pipeline To annotate long ICU ECG recordings, we used ECG foundation model (ECG-FM) [11], since manual labelling of such lengthy recordings would require substantial clinical expertise and time. ECG-FM was fine-tuned on public datasets and validated on a small annotated in-house subset (600

[1] This study was reviewed by the Medical Ethics Review Committee Amsterdam UMC.

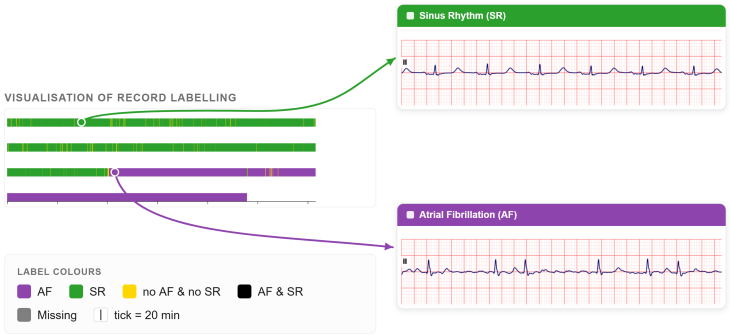

Fig. 1: Visualization of per-record rhythm annotation for a patient. Each long ECG recording is segmented into consecutive 5-second windows, annotated by ECG-FM, and colour-coded by rhythm class. The four horizontal rows show successive segments of the same recording, wrapped in time along a single x-axis; tick marks indicate 20-minute intervals. Arrows link example SR- and AF-labelled points to their corresponding lead-II ECG waveforms.

samples, 5-second duration), achieving F1 = 0.91 for SR and F1 = 0.82 for AF with stable performance across lead configurations and cross-dataset consistency confirmed on MIMIC-III. Each 5-second segment was assigned two independent binary labels, one for SR (encompassing sinus rhythm, sinus tachycardia, and sinus bradycardia) and one for AF, allowing both to be positive simultaneously. This yields four possible combinations: SR only, AF only, both, or neither (other rhythms such as ventricular tachycardia). These performance results on Amsterdam UMC and MIMIC-III data are described in detail in Galanty et al. [4]. Flat-line segments were treated as missing, with no dedicated noise detection applied. Figure 1 illustrates the annotation scheme on a representative recording of a patient transitioning from SR to AF.

Sample Construction and Data Processing We aimed to predict whether a patient in stable SR will transition to AF. Two sample types were constructed from continuous recordings, excluding the first 20 minutes to avoid early signal noise. *SR samples:* A 2-hour window was extracted from recordings containing SR for $\geq$95% of their duration, with minimal AF burden ($\leq$1%, in isolated beats only). One sample was drawn per patient. *Pre-AF samples:* A 2-hour window of predominantly SR ($\geq$90%) immediately preceding at least 10 continuous minutes of predominantly AF ($\geq$90%) was extracted. Multiple samples per patient were allowed if consecutive AF episodes were separated by at least 2 hours of SR. Table 1 summarises the class distribution across both cohorts.

We used lead II, the most consistently available lead across heterogeneous ICU monitoring configurations. MIMIC-III recordings were upsampled from 125 Hz to 500 Hz; Amsterdam UMC recordings were natively acquired at 500 Hz. A

Table 1: Dataset summary. AF prevalence is reported at the sample level.

Dataset	Unique patients	Patients with ≥2 AF samples	AF samples	SR samples	AF prevalence (%)
MIMIC-III	4,995	116	563	4,637	10.8
Amsterdam UMC	117	10	40	95	29.6
Total	5,112	126	603	4,732	11.3

zero-phase third-order Butterworth bandpass filter [3] (1–47 Hz) was applied to remove baseline wander and high-frequency noise.

2.2 ResNet–Mamba Hybrid

Several ECG markers, especially P-wave characteristics (e.g. the difference between the widest and narrowest P waves), premature atrial contractions and changes in heart rate variability (HRV) are associated with future atrial fibrillation [14]. We hypothesise that a model capable of capturing evolving ECG patterns throughout a patient's stay could identify those at risk of developing AF. We propose a hybrid architecture combining a ResNet encoder with Mamba blocks for AF prediction from single-lead ICU ECG. The model was designed to capture both local morphological features and long-range temporal dependencies within extended ECG windows (10 or 30 minutes). The architecture extends the ResNet of Ribeiro et al. [15], well-established in ECG applications, by replacing its final residual stage with a stack of 4 Mamba blocks [7]. After the third residual stage, features are projected to 256 channels via a 1×1 convolution and downsampled (stride 4) before entering the Mamba stack. Each Mamba block applies layer normalisation, a state space module (d_state = 16, d_conv = 4, expand = 2), and dropout (conv dropout = 0.3, Mamba dropout = 0.2) within a residual connection. Global average pooling is used in place of a fixed dense layer to accommodate variable-duration inputs, and the output is passed through a linear binary classification head. This design leverages the ResNet convolutional stages for local morphological feature extraction and the Mamba blocks for long-range temporal modelling of cardiac rhythm.

2.3 Experimental Setup

Baseline As a baseline, we used the ResNet of Ribeiro et al. [15], with the same adaptations as the proposed model: single-lead variable-length ECG input and global average pooling in place of a fixed dense layer, followed by a linear binary classification head.

Signal Windowing and Augmentation Each two-hour recording was segmented into non-overlapping windows of fixed duration (10 or 30 minutes), yielding 12 or 4 windows per recording, respectively. Each window was independently

z-score normalised. During training, two online augmentation strategies were applied: additive Gaussian noise ($\sigma = 2$–10% of signal standard deviation, $p = 0.75$) and polarity inversion ($p = 0.10$), followed by re-normalisation. No augmentation was applied at validation or test time.

Training All models were trained using 3-fold patient-stratified cross-validation, with patients assigned entirely to a single split to prevent data leakage. Training folds were further divided 80/20 into training and validation sets. Class imbalance was addressed using a weighted random sampler enforcing equal AF and SR sampling probability per batch, combined with a positive-class weight in the binary cross-entropy loss capped at 4.0. Models were optimised using AdamW ($\mathrm{lr} = 1 \times 10^{-4}$, weight decay $= 1 \times 10^{-2}$). The learning rate was halved once validation Average Precision (AP) showed no improvement for 10 epochs. Training ran for up to 100 epochs with early stopping on validation AP and a batch size of 4.

Evaluation Performance was evaluated at the sample level. We report Receiver Operating Characteristic Area Under the Curve (ROC-AUC) and AP as means across cross-validation folds. For reference, the overall AF prevalence was 0.11, so a random classifier would achieve AP $\approx$ 0.11; any higher AP therefore reflects improvement over random. For comparisons between Amsterdam UMC and MIMIC, we additionally reported its normalised variant, which allows direct comparison across datasets regardless of class imbalance $\mathrm{AP}_{\mathrm{norm}} = (\mathrm{AP} - \pi)/(1 - \pi)$, where π is the positive-class prevalence, so that a random classifier scores 0 and a perfect classifier scores 1.

To assess how predictive performance evolves as AF onset approaches, a sliding-window evaluation was additionally performed on the test sets: windows were advanced in 10-minute increments toward the AF event (or end of a sample in the case of SR), and metrics were computed at each temporal offset. This analysis quantifies how far in advance the models can reliably distinguish pre-AF from SR.

Explainability To identify which ECG regions the model attended to the most, we construct a saliency map combining two complementary attribution signals. The first signal is obtained by computing Grad CAM [16] at the final ResNet convolutional layer to highlight locally discriminative regions. The second signal addresses the Mamba blocks; we propose a simple heuristic grounded in the architecture's own mechanics. Each Mamba block employs a residual connection [7], so its output at each position can be expressed as $\mathbf{x}_{\mathrm{out}} = \mathbf{x}_{\mathrm{in}} + f(\mathbf{x}_{\mathrm{in}})$, where $\mathbf{x}_{\mathrm{in}}$ and $\mathbf{x}_{\mathrm{out}}$ are the block's input and output representations and $f(\cdot)$ is the block's transformation. The difference $\|\mathbf{x}_{\mathrm{out}} - \mathbf{x}_{\mathrm{in}}\|_2 = \|f(\mathbf{x}_{\mathrm{in}})\|_2$ therefore directly measures the magnitude of what the block contributes to the representation at each position. We use the per-position L2 norm of the residual, averaged across all four Mamba blocks, as a proxy for where the Mamba blocks contributed the largest

Table 2: Sample-level predictive performance under 3-fold patient-stratified cross-validation for all models and input-length configurations. Metrics are reported as mean ± standard deviation across folds.

Model	Input length	ROC-AUC	AP
ResNet	10 min	0.787 ± 0.002	0.391 ± 0.031
ResNet–Mamba	10 min	0.820 ± 0.001	0.458 ± 0.027
ResNet	30 min	0.830 ± 0.015	0.446 ± 0.041
ResNet–Mamba	30 min	0.846 ± 0.009	0.519 ± 0.053

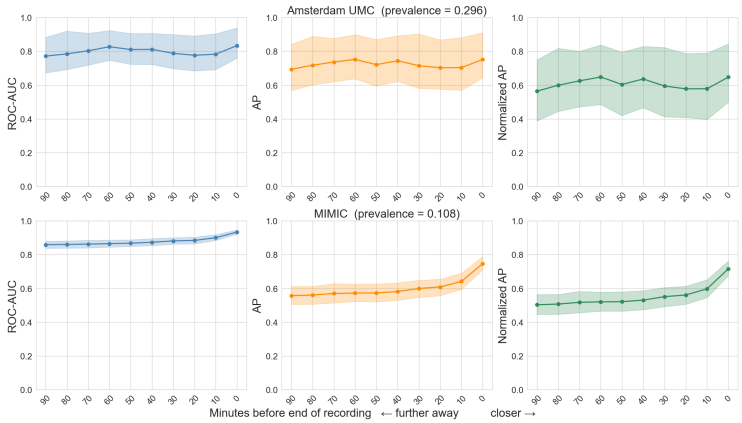

Fig. 2: Sliding-window evaluation of AF onset prediction as a function of time-to-event for the best-performing model (ResNet–Mamba, 30-minute input windows). ROC-AUC, AP, and AP_{norm} are plotted across temporal offsets in 10-minute steps before AF onset, shown separately for MIMIC-III and Amsterdam UMC test data. Shaded areas indicate 95% confidence intervals.

representational change, noting that this reflects block activity rather than prediction attribution in the strict sense. The final saliency map is the element-wise product of the two normalised attribution signals, highlighting regions that are both locally discriminative for the convolutional encoder and actively integrated into Mamba's sequential state. Explanations are visualised on a standard ECG grid and mapped back onto the raw tracing in a format familiar to clinicians.

3 Results

Table 2 reports ROC-AUC and AP for all model and input-length configurations evaluated on the test fold of a 3-fold patient-stratified cross-validation. Figure 2 shows how performance changes as AF onset approaches for the best-performing model (ResNet–Mamba, 30-minute input windows), separately for Amsterdam UMC and MIMIC. The horizontal axis represents the time remaining until AF onset, and the curves display ROC-AUC, AP, and AP_{norm} at each time point

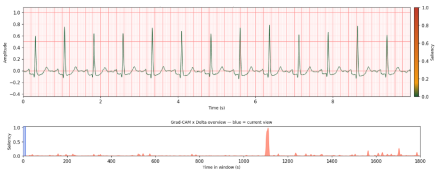

(a) Region with low predicted AF risk.

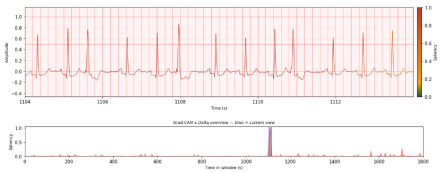

(b) Region with high predicted AF risk.

Fig. 3: Example explanations for the best-performing model on the same sample. The lower panel marks salient regions along the 30-minute ECG window, with the blue-shaded area indicating the segment zoomed in the upper panel, which shows a lead II ECG snippet with overlaid saliency values.

as the 30-minute window is advanced in 10-minute increments towards the AF event. For each time point, we report 95% confidence intervals obtained by bootstrap resampling at the patient level (1000 resamples). Specifically, we pool the out-of-fold predictions from all 3 cross-validation folds and repeatedly resample AF and SR patients with replacement within each group.

Figure 3 illustrates regions that most strongly contributed to the prediction for the best-performing model (ResNet–Mamba, 30-minute input windows) for a patient progressing from SR to AF. The bottom panel shows where salient regions are located along the ECG (peaks in the lower trace), and the top panel provides a zoomed ECG snippet with overlaid saliency values.

To compare Mamba and ResNet, we applied a paired t-test on per-fold ROC-AUC and AP scores ($n = 3$ folds). For the 10-minute configuration, the ROC-AUC difference was statistically significant (Mamba: 0.820 ± 0.000 vs. ResNet: 0.787 ± 0.002; $\Delta = +0.033$, 95%, confidence intervals $[+0.026, +0.039]$, $p = 0.002$). For the 30-minute configuration, differences in ROC-AUC ($\Delta = +0.016$, $p = 0.068$) and AP ($\Delta = +0.073$, $p = 0.059$) favoured Mamba but did not reach significance, with wide confidence intervals reflecting high metric variability under class imbalance.

4 Discussion

In this study, we proposed a hybrid architecture combining ResNet and Mamba for AF prediction from single-lead ICU ECG, with results showing consistent improvements over the ResNet baseline across both input lengths. With 30-minute windows, ResNet–Mamba achieved a ROC-AUC of 0.846 ± 0.009 and an AP of 0.519 ± 0.053 (approximately 5 times above the random baseline given an AF prevalence of 0.11) compared with 0.830 ± 0.015 (ROC-AUC gain: +0.016) and 0.446 ± 0.041 (AP gain: +0.073) for ResNet alone. Notably, improvements were also observed for 10-minute inputs (ROC-AUC gain: +0.033; AP gain: +0.067), indicating that Mamba's selective long-range modelling benefits even shorter windows, with the largest gains observed when more temporal context is available.

The sliding-window analysis showed that performance was lowest for windows far from AF onset and increased as the window moved closer to the event, with the highest scores in the final windows, which might contain short AF episodes. Intermediate windows exhibited relatively stable performance, suggesting that the model captures pre-AF signal beyond overt arrhythmia. Performance was higher on the Amsterdam UMC cohort than on MIMIC-III, although confidence intervals were wider due to the smaller sample size. Possible explanations apart from differences in the sample size include (i) native 500 Hz sampling at Amsterdam UMC versus upsampled MIMIC-III data, which may introduce artefacts; (ii) differences in ICU case-mix and clinical characteristics; and (iii) cohort-specific differences in signal quality and noise arising from distinct extraction pipelines.

We additionally proposed an explainability approach combining Grad-CAM and a Mamba-specific residual attribution signal to highlight the ECG regions with the highest model activity, presented in a clinician-familiar format. Regular SR segments were typically assigned low importance, while regions featuring visible rhythm changes, premature atrial contractions, and occasionally noisy fragments were most frequently highlighted — consistent with known pre-AF markers and providing a basis for clinical review alongside model predictions.

These findings support the feasibility of single-lead deep learning systems for early prediction of AF onset in complex ICU populations. Computationally, Mamba's recurrent state updates also enable efficient continuous monitoring, as only the hidden state needs to be updated (if the attached residual blocks are adapted to shorter inputs) rather than recomputing predictions over full 30-minute windows from scratch.

However, this study has several limitations. First, we restricted inclusion to patients transitioning directly from SR to AF, which simplifies the rhythm landscape and may overestimate performance relative to a more heterogeneous ICU population; future work should relax this constraint to include patients with a broader range of preceding rhythms and structural abnormalities, to better reflect real-world ICU conditions. Second, ECG-FM was used for annotation, meaning that any bias or misclassification in ECG-FM bounds the predictive validity of our model. Although its performance on Amsterdam UMC and MIMIC-III ICU ECG has been evaluated in detail [4], a dedicated validation study with expert-annotated long ECG segments is still needed.Third, the overall sample size is modest, limiting statistical power and generalisability, as reflected in the wider confidence intervals for Amsterdam UMC; larger, prospective cohorts will ultimately be required to support clinical translation. Fourth, statistical comparisons across three folds have limited power; non-significant trends for the 30-minute configuration likely reflect this constraint rather than a true absence of effect. Fifth, the explainability approach for the Mamba component is approximate, as the L2 norm of each block's residual reflects activity rather than directly measuring the influence of each position on the final prediction. Moreover, the saliency maps do not capture HRV despite its relevance for AF; incorporating HRV-based descriptors alongside gradient-based attributions could enhance clinical interpretability. Finally, we did not apply explicit noise rejection, so the

model predicts on all windows regardless of signal quality, which reflects real-world ICU conditions but may affect reliability; dedicated noise detection could improve robustness in practice. Together, addressing these limitations will be essential for validating these models and supporting their translation into clinically useful early warning systems.

Prospects of Application As ICU patients are already subject to continuous ECG monitoring, an AF prediction system could be seamlessly embedded into existing infrastructure, enabling proactive intervention before AF becomes established. Given the volume of simultaneously monitored signals and gradual AF onset, manual tracking is impossible, while automated early warning could reduce AF-related complications, including stroke and increased mortality.

Acknowledgments. This work was supported by the University of Amsterdam Research Priority Area Artificial Intelligence for Health Decision-making.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Availability of data and materials The MIMIC-III Waveform dataset is publicly available [12]. The internal Amsterdam UMC dataset cannot be made publicly available to protect patient privacy. Code can be found on `https://github.com/MariaGalanty/ecg-icu-af-prediction-mamba`

References

1. Bashar, S.K., Ding, E.Y., Walkey, A.J., McManus, D.D., Chon, K.H.: Atrial fibrillation prediction from critically ill sepsis patients. Biosensors **11**(8), 269 (2021)
2. Bosch, N.A., Cimini, J., Walkey, A.J.: Atrial fibrillation in the icu. Chest **154**(6), 1424–1434 (2018)
3. Butterworth, S.: On the theory of filter amplifiers. Wireless Engineer **7**(6), 536–541 (1930)
4. Galanty, M., Hulleman, M., Reuland, M.C., van der Ster, B., Vlaar, A.P., Sánchez, C.I., AF Expert Panel: Leveraging ecg foundation models in critical care for sinus rhythm and atrial fibrillation classification. Intensive Care Medicine Experimental (2026). https://doi.org/10.1186/s40635-026-00949-7, accepted, not yet published
5. Gavidia, M., Zhu, H., Montanari, A.N., Fuentes, J., Cheng, C., Dubner, S., Chames, M., Maison-Blanche, P., Rahman, M.M., Sassi, R., et al.: Early warning of atrial fibrillation using deep learning. Patterns **5**(6) (2024)
6. Grégoire, J.M., Gilon, C., Marelli, F., Bersini, H., Groben, L., Nguyen, T., Deruyter, B., Godart, P., Carlier, S.: Short-term atrial fibrillation onset prediction using machine learning. European Heart Journal-Digital Health **6**(6), 1159–1168 (2025)
7. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
8. Guan, C., Gong, A., Zhao, Y., Yin, C., Geng, L., Liu, L., Yang, X., Lu, J., Xiao, B.: Interpretable machine learning model for new-onset atrial fibrillation prediction in critically ill patients: a multi-center study. Critical Care **28**(1), 349 (2024)

9. Johnson, A.E., Pollard, T.J., Shen, L., Lehman, L.w.H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Anthony Celi, L., Mark, R.G.: Mimic-iii, a freely accessible critical care database. Scientific data **3**(1), 1–9 (2016)
10. Ling, L.H., Kistler, P.M., Kalman, J.M., Schilling, R.J., Hunter, R.J.: Comorbidity of atrial fibrillation and heart failure. Nature Reviews Cardiology **13**(3), 131–147 (2016)
11. McKeen, K., Masood, S., Toma, A., Rubin, B., Wang, B.: Ecg-fm: An open electrocardiogram foundation model. JAMIA open **8**(5), ooaf122 (2025)
12. Moody, B., Moody, G., Villarroel, M., Clifford, G.D., Silva, I.: MIMIC-III Waveform Database. PhysioNet (Apr 2020). https://doi.org/10.13026/c2607m, `https://doi.org/10.13026/c2607m`, version 1.0
13. Noteboom, S.H., Kho, E., Galanty, M., Sánchez, C.I., Ten Bookum, F.C., Veelo, D.P., Vlaar, A.P., van der Ster, B.J.: From intensive care monitors to cloud environments: a structured data pipeline for advanced clinical decision support. EBioMedicine **111** (2025)
14. Perez, M.V., Dewey, F.E., Marcus, R., Ashley, E.A., Al-Ahmad, A.A., Wang, P.J., Froelicher, V.F.: Electrocardiographic predictors of atrial fibrillation. American heart journal **158**(4), 622–628 (2009)
15. 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)
16. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
17. Sibley, S., Bedford, J., Wetterslev, M., Johnston, B., Garside, T., Kanji, S., Whitehouse, T., Welters, I., Ostermann, M., Balik, M., et al.: Atrial fibrillation in critical illness: state of the art. Intensive Care Medicine **51**(5), 904–916 (2025)
18. Van Gelder, I.C., Rienstra, M., Bunting, K.V., Casado-Arroyo, R., Caso, V., Crijns, H.J., De Potter, T.J., Dwight, J., Guasti, L., Hanke, T., et al.: 2024 esc guidelines for the management of atrial fibrillation developed in collaboration with the european association for cardio-thoracic surgery (eacts) developed by the task force for the management of atrial fibrillation of the european society of cardiology (esc), with the special contribution of the european heart rhythm association (ehra) of the esc. endorsed by the european stroke organisation (eso). European heart journal p. ehae176 (2024)
19. Zhang, P., Lin, F., Ma, F., Chen, Y., Fang, S., Zheng, H., Xiang, Z., Yang, X., Li, Q.: Automatic screening of patients with atrial fibrillation from 24-h holter recording using deep learning. European Heart Journal-Digital Health **4**(3), 216–224 (2023)