Review

Not peer-reviewed version

Advances in Interpretation of Electrocardiogram by Artificial Intelligence

S. Suave Lobodzinski * and Ryszard Piotrowicz

Posted Date: 20 May 2026

doi: 10.20944/preprints202605.1353.v1

Keywords: artificial intelligence; electrocardiography; deep learning; arrhythmia detection; digital biomarkers; clinical decision support; cardiovascular diagnostics; AI-ECG

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 May 2026 doi:10.20944/preprints202605.1353.v1

Review

Advances in Interpretation of Electrocardiogram by Artificial Intelligence

S. Suave Lobodzinski [1],* and Ryszard Piotrowicz [2]

[1] University of California Los Angeles and California State University Long Beach, USA
[2] National Institute of Cardiology, Warsaw, Poland
* Correspondence: mail@lobodzinski.com

Abstract

The 12-lead electrocardiogram is essential for cardiovascular diagnosis but limited by inter-observer variability, low sensitivity for subclinical disease, and labor-intensive telemonitoring analysis. Artificial intelligence (AI), particularly deep learning, addresses these constraints by extracting high-dimensional patterns that correlate with arrhythmias, structural abnormalities, and systemic conditions. This review synthesizes recent AI-enabled ECG advances, covering technical foundations—including foundation models and validation strategies—and clinical applications such as arrhythmia detection, structural heart disease identification, and digital biomarker derivation. We discuss emerging trends like self-supervised learning, multimodal integration, generative models, and explainability techniques. Furthermore, we address critical challenges regarding generalizability, algorithmic bias, privacy, and regulatory frameworks. Finally, we outline research priorities, including curated open datasets, personalized continuous-learning systems, and deployment in resource-limited settings. With rigorous validation, transparent governance, and human-centered design, AI-ECG has the potential to democratize cardiovascular diagnostics and improve clinical outcomes across diverse environments.

Keywords: artificial intelligence; electrocardiography; deep learning; arrhythmia detection; digital biomarkers; clinical decision support; cardiovascular diagnostics; AI-ECG

1. Introduction

The 12-lead electrocardiogram (ECG) remains the cornerstone of cardiovascular diagnosis, a century after its inception. Its ubiquity, cost-effectiveness, portability, and ease of acquisition make it indispensable for acute chest pain triage, arrhythmia evaluation, and the longitudinal management of structural heart disease. Beyond these conventional roles, the ECG encodes high-dimensional information regarding cardiac structure, conduction dynamics, and autonomic tone that can precede overt clinical disease [1]. Harnessing this complex, rich signal is now central to early detection and precision risk stratification in contemporary cardiology.

Despite its clinical importance, conventional ECG interpretation faces significant limitations. Visual analysis is inherently subjective, prone to inter- and intra-observer variability, and often lacks the sensitivity to detect subtle findings, such as minor ST-T changes or early conduction disturbances [2]. Diagnostic accuracy is further modulated by clinician training, fatigue, and clinical context, leading to diagnostic inconsistencies. Moreover, human readers are generally ill-equipped to identify the subtle, high-dimensional waveform patterns that correlate with microstructural remodeling or future adverse events. Consequently, the sensitivity of standard interpretation for subclinical cardiomyopathies, incipient left ventricular (LV) dysfunction, or future atrial fibrillation (AF) remains modest [3,4].

While long-term monitoring—particularly ECG telemonitoring—has enhanced diagnostic yield, these limitations persist, compounded by the labor-intensive analysis of the massive data volumes generated by continuous recording. Recent advances in artificial intelligence (AI), specifically

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 May 2026 doi:10.20944/preprints202605.1353.v1

machine learning and deep learning, offer a paradigm shift [1,5]. Unlike traditional methods, these algorithms are trained on vast, annotated datasets to identify patterns that extend well beyond manual features such as intervals, amplitudes, or axes [5,9]. Deep neural networks, including convolutional and transformer-based architectures, can process raw waveforms to learn hierarchical representations, demonstrating cardiologist-level performance in arrhythmia detection, identification of asymptomatic LV systolic dysfunction, and prediction of biological age or future AF [3,4,6,7].

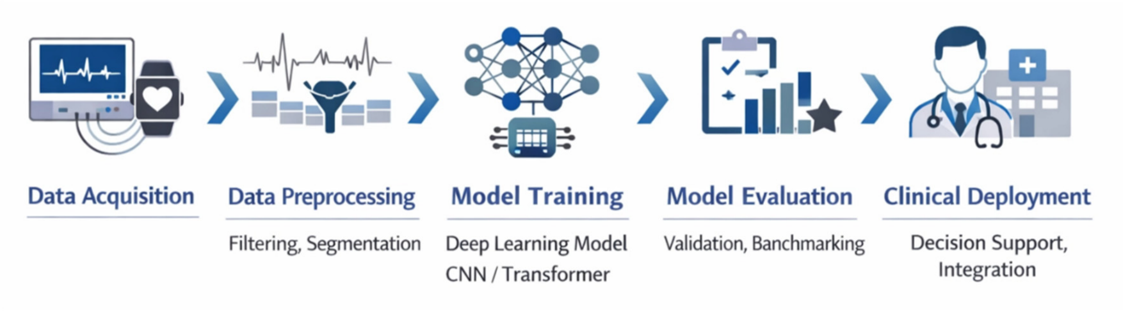

Figure 1.1. End-to-End AI-ECG Pipeline.

This review synthesizes the latest advances in automated 12-lead ECG interpretation enabled by AI. We first summarize the technical foundations of contemporary machine learning approaches, including data requirements, model architectures, and validation strategies. We then discuss established and emerging clinical applications, spanning rhythm disturbances, structural heart disease, and systemic digital biomarkers. Finally, we address critical challenges—including generalizability, algorithmic bias, regulatory governance, clinical integration and future applications—to provide a concise framework for the responsible translation of AI-ECG into routine cardiovascular care [11].

2. Technical Foundations of AI in ECG Analysis

Artificial intelligence methods for electrocardiogram interpretation rest on a set of technical foundations that span data types, model architectures, and end-to-end analytic pipelines. Understanding these elements is essential for evaluating current systems and designing new ones that are robust, generalizable, and clinically meaningful.

2.1. Types of ECG Data for AI

Digital ECG signals arise from multiple sources that differ in duration, sampling characteristics, noise profile, and clinical context, all of which shape the design and performance of AI models. Standard 10-second, 12-lead recordings obtained in resting conditions remain the most common input for supervised learning, as they are widely available in hospital archives and accompanied by diagnostic reports that can be mined as labels at scale. Longer-term Holter and patch recordings provide continuous multi-hour to multi-day monitoring, enabling detection of infrequent arrhythmias and dynamic changes in rate or conduction but introducing substantial class imbalance and higher artifact burden.

Data from implantable devices—such as pacemakers, implantable cardioverter-defibrillators, and implantable loop recorders—typically consist of single- or dual-lead intracardiac electrograms and event-triggered segments, requiring models that can handle sparse, device-specific signals while exploiting precise temporal annotations of arrhythmic episodes. Wearable devices, including smartwatches, fitness bands, and consumer ECG patches, contribute vast volumes of single-lead or reduced-lead signals sampled at varying frequencies and often acquired in nonclinical environments. These data streams broaden population coverage and capture physiology during daily activities but introduce challenges related to motion artifact, variable electrode placement, and uncertain ground truth.

Recent foundation-model efforts have begun to jointly pretrain on heterogeneous sources—resting 12-lead ECGs, Holter segments, and wearable signals—to learn representations that transfer across devices and clinical tasks [8–10]. Additionally, thermal and graphics paper ECGs remain widely used in everyday clinical practice for interpretation and insertion into patient records. Printed or scanned paper ECG records may also used as source data, providing they are of high quality comparable to original digital recordings.

2.2. Model Architectures

Early AI-ECG work relied on classical machine-learning algorithms that operate on handcrafted features derived from the signal, such as intervals, amplitudes, heart-rate variability indices, and wavelet coefficients. Common approaches included logistic regression, support vector machines, random forests, and gradient-boosted trees, which remain attractive in settings with limited data or strong prior domain knowledge because they are relatively sample-efficient and easier to interpret. However, their performance is constrained by the quality and completeness of manually engineered features and their limited capacity to model complex temporal dependencies.

Supervised and Deep learning techniques have largely supplanted these methods for high-volume ECG datasets by enabling end-to-end learning from raw or minimally processed waveforms. Convolutional neural networks (CNNs) are widely used for arrhythmia classification and detection of structural heart disease because they can capture local morphological patterns such as QRS shape, ST-segment deviations, and T-wave abnormalities [5,9,11–13]. Recurrent architectures, including long short-term memory (LSTM) and gated recurrent units, can model longer-range temporal dependencies and have been applied to rhythm analysis in Holter recordings, though they are gradually being replaced by more scalable architectures. Hybrid CNN-recurrent models further combine local feature extraction with sequence modeling and remain competitive in many tasks.

Figure 2.1. Overview of Supervised (a) and Deep learning (b) techniques for AI ECG applications.

Transformer architectures, originally developed for natural-language processing, have recently demonstrated strong performance in ECG analysis by leveraging self-attention mechanisms to model global temporal relationships. Variants such as ECG-specific transformers convert one-dimensional waveforms into tokenized or image-like representations, enabling efficient capture of both long-range context and detailed morphology. Building on these advances, several groups have introduced ECG foundation models—large transformer- or CNN-based encoders trained on hundreds of thousands to millions of ECGs to learn general representations that can be fine-tuned for diverse downstream tasks [11–15].

Self-supervised learning strategies, including contrastive learning, masked-signal modeling, and multi-segment prediction, further reduce reliance on manual labels by pretraining on unlabeled ECG corpora before task-specific fine-tuning [15–17]. These foundation and self-supervised models exhibit improved label efficiency, cross-dataset generalization, and adaptability compared with task-specific networks.

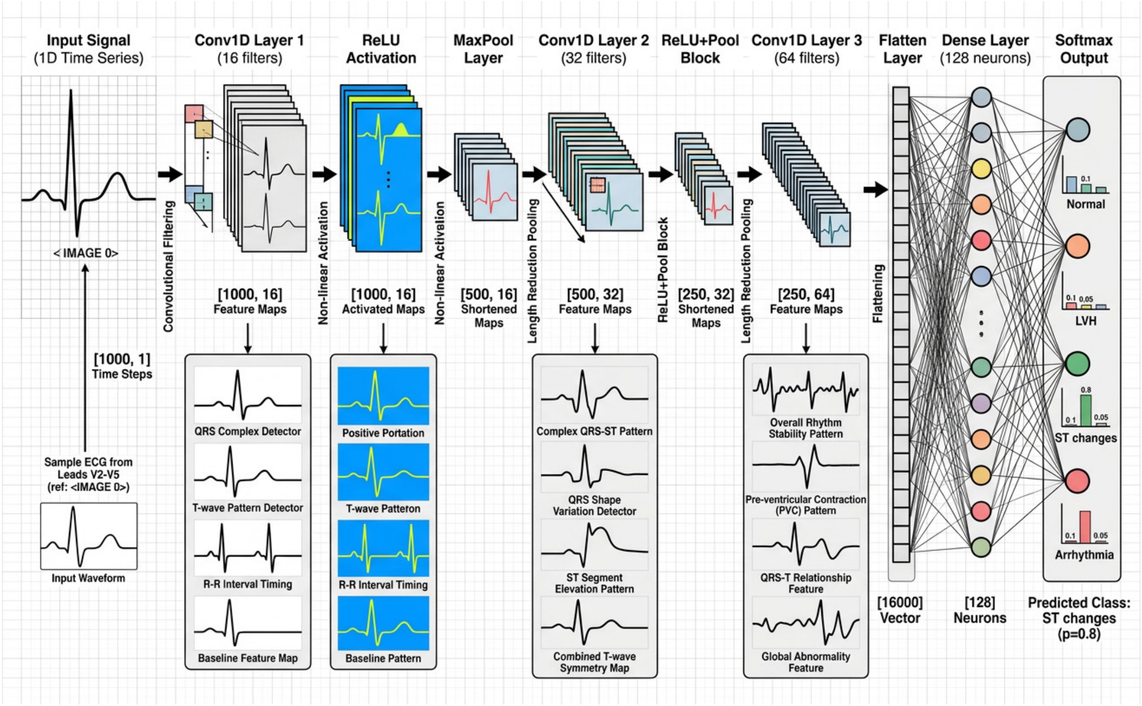

Figure 2.2. AI ECG CNN-based encoding process.

2.3. Data and Model Pipeline

Robust AI-ECG systems depend not only on architecture choice but also on a rigorous pipeline for data curation, preprocessing, labeling, and validation. Data acquisition typically begins with extraction of raw digital ECG files and associated metadata from hospital ECG management systems, Holter databases, device interrogations, or wearable platforms, followed by de-identification procedures that remove or encrypt protected health information while preserving clinically relevant attributes such as age and sex. Large multicenter projects often harmonize sampling rates, lead configurations, and file formats across vendors to facilitate joint training and reduce domain shift [13,18]

1. Digital ECG files Preprocessing Goals: Standardization and Artifact Mitigation

- **Purpose:** Preprocessing aims to enhance signal-to-noise ratio (SNR), remove interference, and enable accurate feature extraction for machine learning models.
- **Generalization:** Cleaning signals is critical for ensuring that models do not overfit to noise, thereby enhancing generalization to new data.

2. Band-pass Filtering (Baseline Wander & High-Frequency Noise)

- **Baseline Wander Removal:** Low-frequency components (typically < 0.5 Hz) caused by respiration are removed using high-pass or band-pass filters.
- **High-Frequency Noise Removal:** Powerline interference (50/60 Hz) and muscle activity (EMG) are reduced using low-pass, notch or adaptive filters.
- **Combined Filter:** A 0.05–100 Hz band-pass filter is frequently employed as an essential preprocessing step.

3. Resampling to Consistent Frequency

- **Purpose:** Because ECG signals are often acquired from various sources (e.g., MIT-BIH sampled at 360 Hz, PTB at 1000 Hz or 10 kHz), they are resampled to a consistent frequency (commonly 200 Hz or 250 Hz) to standardize input for models.
- **Impact:** Resampling lowers computational complexity, and is crucial for training deep learning models efficiently.

4. Lead Selection or Reordering

- **Standardization:** In 12-lead ECG setups, specific leads (e.g., Lead II or V1) are selected or reordered to normalize the input for diagnostic models.

5. Segmentation (Fixed-length Windows or Median Beats)

- **Windowing:** Signals are segmented into shorter fixed-length windows (e.g., 10 seconds).
- **Beat Segmentation:** Individual heartbeats are segmented around R-peaks (typically using windows like 200 ms before and 400 ms after the R-peak) to perform beat-level classification.
- **Median Beats:** Using a median filter (e.g., with 200 ms and 600 ms windows) helps estimate and subtract the baseline, effectively producing a clean "median beat".

Some pipelines transform signals into frequency or time-frequency representations using Fourier or wavelet transforms, or into image-like formats via oscillographic rendering for CNN or vision-transformer models. Importantly, recent foundation-model work has shown that standardized, automated preprocessing is critical for avoiding catastrophic performance drops when deploying models to external sites that use different ECG machines.

Label generation [19] is a central bottleneck because manual expert annotation is costly and difficult to scale. Many large-scale studies therefore rely on weak supervision, using natural-language processing to extract diagnostic codes from ECG reports, linkage to administrative or echocardiographic databases, or heuristic rules based on device logs and rhythm flags. While these approaches enable training on hundreds of thousands of records, they introduce label noise that must be mitigated through cleaning rules, probabilistic labeling, or robust loss functions. Self-supervised pretraining partially sidesteps this challenge by learning from unlabeled ECGs and requiring high-quality labels only for smaller fine-tuning datasets [13].

Model evaluation proceeds through a combination of internal and external validation designed to quantify performance, generalizability, and fairness. Internally, data are typically split into training, validation, and test sets at the patient level, with cross-validation or temporal splits used to reduce optimistic bias and mimic prospective deployment [14]. External validation on geographically and demographically distinct cohorts is now considered essential, as models frequently show significant performance degradation when applied to different healthcare systems, devices, or prevalence patterns. Recent foundation-model studies have included up to 10-11 external datasets to assess robustness across multiple tasks and institutions [14] Beyond traditional discrimination metrics such as area under the receiver-operating characteristic curve, investigators increasingly report calibration, decision-curve analysis, and subgroup performance by sex, age, and ethnicity to evaluate clinical utility and algorithmic fairness.

It is important to understand that the entire process of AI ECG system comprises two parts: 1) Model Training, 2) Inference (users interacting with trained models) [20,21].

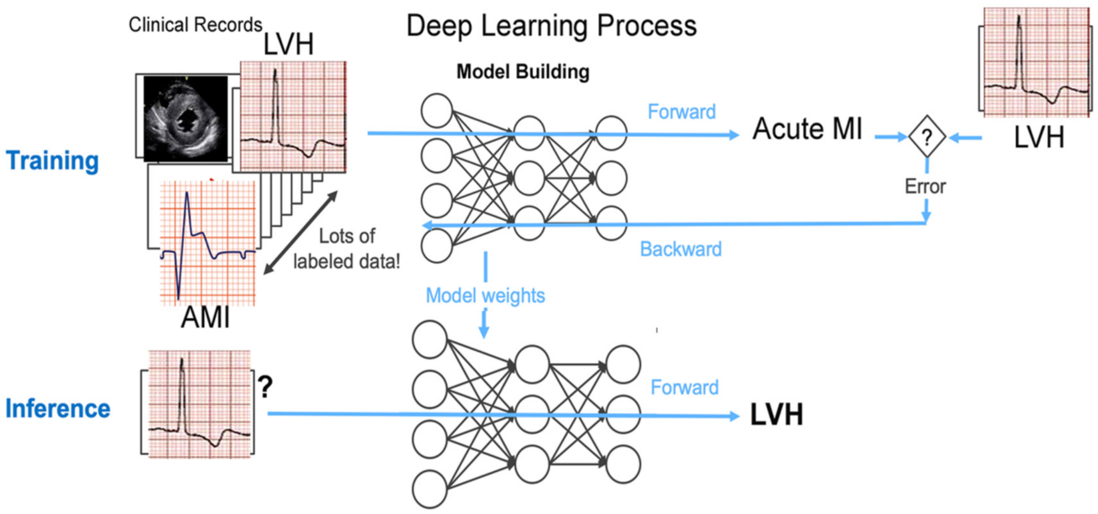

Figure 2.3. AI ECG Infrastructure. The deep learning process links the features of certain labeled ECGs through the forward/backward propagation process. The ensuing model can now be used using the inference process. An unknown ECG submission is forwarded to the model for identification resulting in LVH finding..

3. Current Clinical Applications of AI-Enhanced Electrocardiography

Artificial intelligence applied to the electrocardiogram has transitioned from experimental research to clinically actionable tools in several areas of cardiovascular medicine. By leveraging deep learning techniques, these models extract latent physiological features imperceptible to human readers, expanding the ECG's diagnostic and prognostic potential far beyond traditional rhythm and interval analysis.

3.1. Detection and Classification of Arrhythmias

Arrhythmia recognition has been the earliest and most rigorously validated domain of AI-ECG. Deep convolutional and recurrent neural networks trained on millions of labeled ECG tracings now achieve expert-level performance in detecting atrial fibrillation (AF), atrial flutter, supraventricular tachycardia (SVT), and ventricular tachycardia (VT) [5,8,22]. These models learn temporal and morphological patterns across leads, enabling robust discrimination between sinus rhythm and a wide array of rhythm disturbances.

Single-lead devices and consumer wearables, such as smartwatches and patches, have further extended arrhythmia screening to ambulatory populations. AI algorithms embedded within these platforms can identify paroxysmal AF or conduction abnormalities from brief, noisier segments with high specificity, facilitating early detection and longitudinal monitoring [8]. Validation studies have demonstrated that AI-based classification maintains accuracy across different hardware systems, promoting decentralized rhythm surveillance at scale [23].

3.2. Detection of Structural Heart Disease

AI-enabled ECG interpretation has also shown strong potential in detecting underlying structural and functional cardiac disease. Models trained to predict left ventricular (LV) systolic dysfunction from standard 12-lead ECGs have reproduced or exceeded the performance of biomarker-based screening, providing a rapid, noninvasive method to triage patients for echocardiography [24–28]. Parallel developments include algorithms detecting LV hypertrophy, hypertrophic cardiomyopathy, dilated cardiomyopathy, and pulmonary hypertension by capturing subtle repolarization and voltage abnormalities that correlate with chamber remodeling and altered hemodynamics [24,29–34].

Importantly, such models offer value in primary care and community settings, where imaging availability is limited. Prospective validation in diverse populations has emphasized both diagnostic precision and prognostic relevance, as AI-predicted LV dysfunction correlates with future heart failure events even in asymptomatic individuals [6,24,29].

3.3. The ECG as a Digital Biomarker for Noncardiac Conditions

Beyond cardiovascular disorders, AI-ECG has emerged as a window into systemic physiology. Algorithms have been developed to infer conditions such as obstructive sleep apnea, chronic obstructive pulmonary disease (COPD), and biological aging ("ECG-age") directly from sinus rhythm recordings. These models utilize micro-variations in beat morphology, interval variability, and repolarization heterogeneity that reflect autonomic tone and cardiopulmonary load.

In predictive cardiology, AI-ECG has been used to estimate future risk of AF, even in patients with apparently normal rhythms [3,23,35]. By identifying subclinical electrical remodeling, these approaches may enable targeted preventive strategies. Additionally, AI-ECG models can estimate biological age and predict future cardiometabolic disease risk, demonstrating utility beyond traditional cardiovascular endpoints [4,35]. Collectively, they represent the evolution of ECG interpretation from categorical diagnosis to quantitative phenotyping, bridging cardiovascular and systemic health assessment.

3.4. Clinical Decision Support and Workflow Integration

AI-assisted ECG interpretation is increasingly integrated into decision pathways across care settings. In emergency departments, rapid AI-ECG triage tools assist clinicians in identifying acute coronary syndromes, arrhythmias, and heart failure presentations, reducing time to intervention. In outpatient and oncology contexts, models forecasting drug-induced QT prolongation or cardiotoxicity enable early therapy modification and individualized monitoring [36]..

Furthermore, AI outputs can prioritize patients for advanced imaging based on latent risk signatures invisible to traditional metrics. Such "ECG-first" triage has been particularly valuable in resource-limited environments. Importantly, clinical adoption emphasizes a collaborative model—AI as an augmentative tool that enhances, rather than replaces, physician judgment. Well-designed interfaces providing explainable features, confidence estimates, and workflow compatibility remain essential for sustained clinical trust.

4. Recent Research Directions and Innovations

Recent progress in artificial intelligence-enhanced electrocardiography is increasingly shaped by four intertwined research areas: scalable pretraining via foundation and self-supervised models, multimodal and integrated architectures, generative and simulation frameworks, and advances in explainability [. Together, these developments aim to move beyond narrow, task-specific algorithms toward flexible, trustworthy systems that can be adapted across indications, populations, and healthcare environments.

4.1. Foundation Models and Self-Supervised Learning

Foundation models for ECG seek to learn generic, reusable representations from very large corpora of predominantly unlabeled recordings. In contrast to traditional supervised approaches, which require labor-intensive expert annotation for each target task, self-supervised objectives exploit the intrinsic structure of the ECG signal. Typical strategies include masked signal modeling (predicting masked segments from surrounding context), contrastive learning (bringing augmented views of the same beat or lead closer in representation space while pushing different signals apart), and sequence-to-sequence prediction across time scales [32–35]. These objectives encourage the model to encode morphology, rhythm, and temporal dependencies in a way that is not restricted to a single diagnosis.

Once pretrained, the same backbone can be fine-tuned with relatively modest labeled datasets for diverse downstream applications such as arrhythmia classification, detection of structural heart disease, or risk prediction in specific cohorts. This paradigm improves sample efficiency, allowing high performance even when only a few hundred or thousand labeled examples are available for a new indication. It also supports domain adaptation: by performing a brief fine-tuning step on data from a new institution, vendor, or ethnicity, the model can be recalibrated to local distributions without retraining from scratch. Conceptually, this shifts AI-ECG from single-use algorithms toward a platform model that can be continuously extended to emerging tasks, including rare diseases and niche clinical questions.

4.2. Multimodal and Integrated Systems

A second major trajectory involves integration of ECG with complementary data sources in multimodal architectures. The ECG captures high-resolution information about cardiac electrophysiology, but its diagnostic yield and prognostic value can be enhanced when contextualized with clinical covariates. Contemporary models therefore seek to jointly encode 12-lead (or single-lead) waveforms with elements of the electronic health record (EHR), such as demographics, comorbidities, medications, laboratory values, and vital signs, as well as imaging data including echocardiography and cardiac or thoracic computed tomography and magnetic resonance.

In such systems, each modality typically passes through a specialized encoder (e.g., a convolutional or transformer-based time-series encoder for ECG, a text encoder for clinical notes, a vision backbone for imaging) before fusion in a shared latent space. This design supports multitask learning, whereby a single model simultaneously predicts several endpoints—for instance left ventricular function, pulmonary pressures, arrhythmic risk, and short-term clinical deterioration. Multitask setups can exploit shared pathophysiological structure between outputs, often improving calibration and robustness compared with separate single-task networks. Furthermore, cross-modal prediction (e.g., inferring echocardiographic left ventricular ejection fraction or valvular disease directly from ECG plus structured EHR features) opens the possibility of "virtual imaging triage," prioritizing patients for resource-intensive tests based on an inexpensive, widely available signal.

4.3. Generative and Simulation Models

Generative modeling has emerged as a complementary line of research addressing several limitations of purely discriminative AI-ECG systems. Models such as variational autoencoders, generative adversarial networks, diffusion models, and autoregressive sequence generators are being used to synthesize realistic ECG waveforms, either de novo or conditioned on clinical attributes (e.g., rhythm label, QRS duration, presence of hypertrophy) [45–50]. Synthetic ECGs can mitigate class imbalance by augmenting underrepresented rhythms or pathologies, and they allow exploration of rare or ethically sensitive scenarios without relying solely on actual patient recordings.

Beyond simple data augmentation, generative and simulation frameworks support systematic robustness testing. By gradually introducing noise, baseline wander, electrode misplacement, or specific morphological perturbations (for example, progressive ST-segment elevation or increasing PR interval), investigators can probe how model predictions change and identify brittle regimes or failure modes. This is particularly valuable for safety-critical tasks such as detection of life-threatening arrhythmias or ischemia, where adversarial perturbations and distribution shifts may otherwise go unnoticed. In education, curated libraries of synthetic yet physiologically plausible ECGs provide standardized material for training and assessment of clinicians, while simultaneously avoiding disclosure of identifiable patient data. As generative modeling matures, there is growing interest in using such systems to simulate longitudinal trajectories, enabling in silico trials where hypothetical interventions or disease progressions are explored before real-world deployment.

4.4. Explainability and Interpretability

As AI-ECG models become more complex, explainability has moved from a purely academic topic to a practical prerequisite for clinical adoption and regulatory approval [37,38]. Local explanation methods, such as saliency maps, Grad-CAM-like techniques adapted to one-dimensional signals, and integrated gradients, aim to highlight which time points, beats, or leads contribute most strongly to a particular prediction [37,39,40]. Lead-wise relevance scores can indicate whether an algorithm's decision is driven by plausible regions—such as a prolonged QT interval, fragmented QRS in specific leads, or localized ST-T changes—or by spurious artifacts and noise. Mapping these attention patterns onto classical electrocardiographic features creates a bridge between deep learning outputs and established human heuristics, enabling clinicians to evaluate when an AI suggestion is physiologically coherent.

In parallel, global interpretability approaches seek to elucidate what a model has learned across an entire dataset rather than explaining individual predictions in isolation. Examples include extracting approximate decision rules from trained networks, fitting simpler surrogate models in latent space, clustering latent representations into prototype examples that correspond to recognizable ECG phenotypes, and performing counterfactual analyses where small, targeted modifications to the input show how predictions might change. Importantly, interpretability is now increasingly viewed as a design constraint rather than an afterthought: architectures, loss functions, and training curricula are being tailored to encourage alignment with known electrophysiological principles. When combined with transparent reporting of uncertainty and performance stratified across subgroups, these methods can foster appropriate trust, support shared human-AI decision-making, and ultimately facilitate safe, equitable integration of AI-ECG into routine cardiovascular care.

5. Practical and Regulatory Challenges

Clinical deployment of AI-enhanced ECG (AI-ECG) confronts substantial practical and regulatory hurdles. These must be addressed to ensure safety, efficacy, and equity across diverse healthcare ecosystems. These challenges coalesce around four interconnected domains: generalization and transferability, algorithmic bias and fairness, data privacy and security, and evolving regulatory frameworks for Software as a Medical Device (SaMD) [42,43].

5.1. Generalization and Transferability

Model generalizability is a primary barrier. Algorithms optimized in a specific healthcare environment often exhibit significant performance degradation when deployed elsewhere. This is driven by systematic domain shifts, including heterogeneity in ECG acquisition hardware, filtering protocols, sampling rates, lead-placement conventions, and baseline patient distributions. Mitigation necessitates robust multi-site external validation, adoption of domain-adaptive techniques (e.g., feature alignment, test-time augmentation), and continuous post-deployment monitoring with automated drift detection to maintain performance standards.

5.2. Algorithmic Bias and Equity

Algorithmic bias remains a critical concern, with documented performance disparities across sex, age, race, ethnicity, and socioeconomic strata [35,44–46]. Training datasets skewed toward majority cohorts often yield reduced sensitivity or specificity for underrepresented populations, reflecting physiologic, anatomic, and pathologic variations. Beyond aggregate metrics like the area under the receiver-operating characteristic curve (AUROC), clinical validation must prioritize subgroup-specific reporting of calibration error, demographic parity, and equalized odds. Strategies including stratified sampling, adversarial debiasing, and mandatory subgroup-specific performance disclosures are prerequisites for equitable deployment thresholds [47,48].

5.3. Privacy, Security, and Data Ownership

ECG signals function as quasi-biometric data, elevating risks regarding re-identification from latent representations. Furthermore, adversarial perturbations—subtle, imperceptible waveform alterations—can induce catastrophic misclassifications in deployed systems. Countermeasures must leverage privacy-preserving architectures, including differential privacy during training, homomorphic encryption for inference, and federated learning paradigms that update models without exchanging raw patient data [48–51]. These frameworks must be balanced against the need for transparent governance regarding intellectual property, patient consent for secondary uses, and clinical liability in breach scenarios.

5.4. Regulatory Frameworks and Deployment

Regulatory landscapes are maturing but remain fragmented. Under the U.S. Food and Drug Administration (FDA) Software as a Medical Device (SaMD) framework and the European Medicines Agency (EMA) Medical Device Regulation, high-risk AI-ECG tools are classified as Class II or III devices. Compliance requires not only retrospective validation but also prospective, ideally randomized, evidence demonstrating positive impact on clinical workflows, patient outcomes, and health economic indices [52–54]. Lifecycle management for adaptive algorithms—necessitating pre-specified change protocols, continuous safety surveillance, and human-in-the-loop oversight—

demands unprecedented collaboration among developers, regulators, and end-users to foster clinical trust without over-promising capabilities. Responsible advancement thus necessitates interdisciplinary governance: multidisciplinary review boards, transparent performance dashboards, clinician training on limitations, and patient-centered communication to foster trust without overhyping capabilities. Only through such concerted action can AI-ECG transcend proof-of-concept to deliver broad, equitable clinical value.

6. Integration of AI-ECG into Clinical Practice

The integration of artificial intelligence-enhanced electrocardiography into clinical practice offers improvement in algorithmic performance and new challenges since it has to be tightly integrated into existing workflows, decision structures, and professional roles as shown in Figure 6.1 [56].

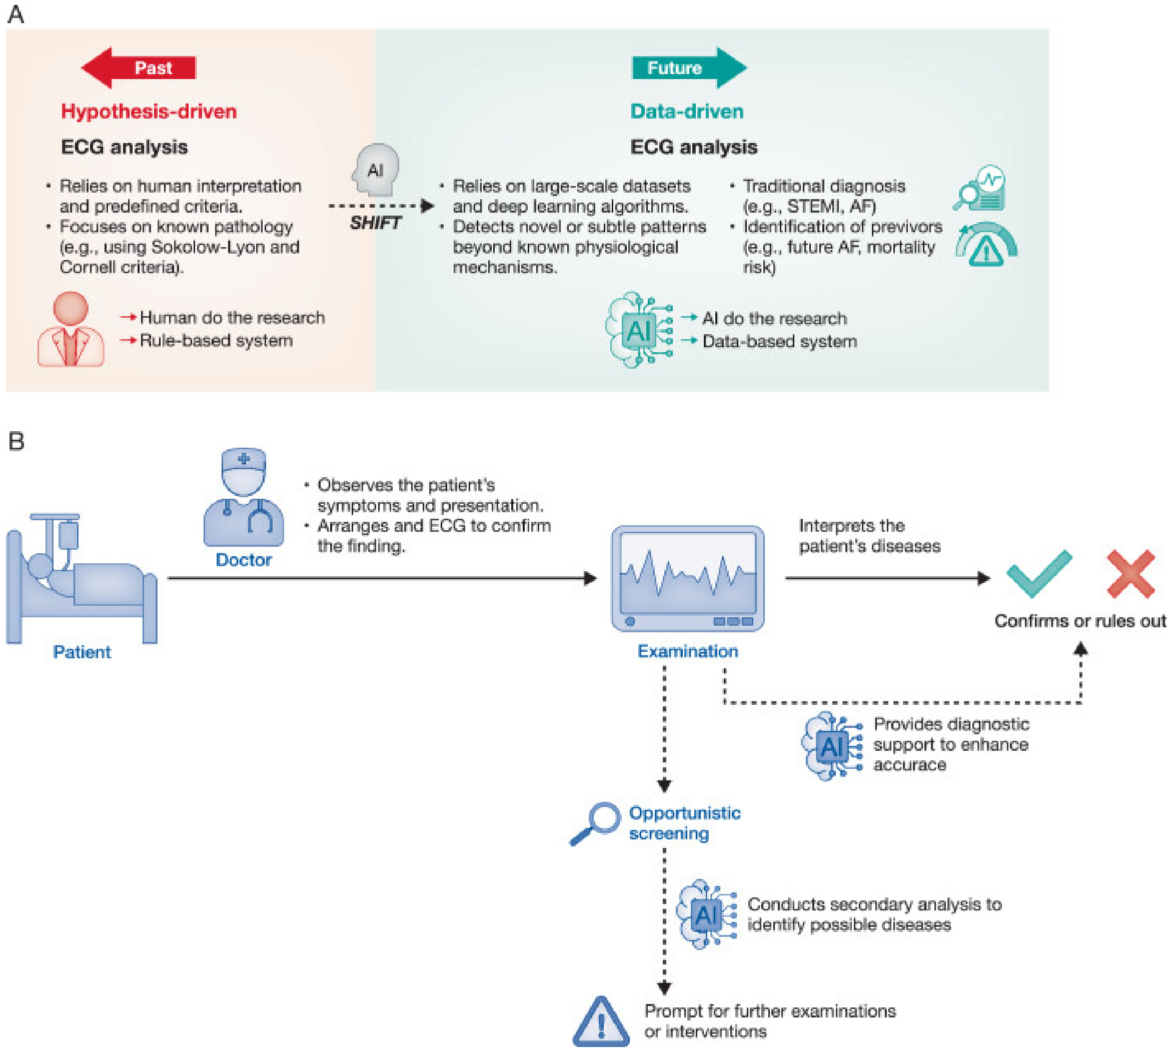

Figure 6.1. Panel (A) illustrates the paradigm shift in ECG analysis caused by the use of AI. Traditionally, a hypothesis-driven approach is used to establish criteria, and human interpretations are used to identify known pathological markers. In contrast, the data-driven approach harnesses large-scale datasets and deep learning algorithms to detect subtle or novel patterns in ECG signals—those that may not be readily apparent through conventional clinical reasoning. These signatures are used to diagnose conditions such as left ventricular dysfunction and atrial fibrillation and identify "previvors," individuals with a hidden yet elevated risk for future adverse events. This transformation enables earlier detection of abnormalities, improved risk stratification, and an expanded understanding of cardiac electrical activity beyond historically defined diagnostic criteria. Panel (B) visually depicts the shift from conventional ECG interpretation to an integrated AI-driven approach (dashed line), emphasizing how the use of AI serves as a powerful adjunct to traditional methods (solid line) in enhancing

diagnostic accuracy and patient care. From Lin et al. Artificial intelligence-enabled electrocardiography from scientific research to clinical application. EMBO Mol Med. 2026;18(1):22-40.

A central operational paradigm is the use of "dual reporting" models, in which AI-generated interpretations are presented alongside, but do not replace, physician ECG reports [57–59].

In such configurations, AI systems can pre-classify rhythm, flag potentially critical abnormalities, and generate structured draft text, thereby reducing reporting time for normal or straightforward tracings and helping to prioritize complex or high-risk studies. At the same time, the final signed report remains the responsibility of the clinician, who adjudicates discordant findings, contextualizes predictions with clinical information, and can override erroneous or implausible outputs. Early experience suggests that dual reporting can improve throughput and consistency in high-volume settings, but may transiently increase workload during the adoption phase if interfaces are poorly integrated with existing ECG management systems or if alert thresholds are not appropriately calibrated.

6.1. System Architecture and Operational Workflow

The successful implementation of AI-ECG systems requires a comprehensive understanding of the complete architecture through which signals are acquired, processed, interpreted, and ultimately delivered as clinical reports [60]. The system architecture spans multiple functional layers, from raw signal acquisition through neural network processing to final physician validation and reporting. A detailed visualization of this architecture is presented in Figure 6.2, which delineates the key components and information flow throughout the system lifecycle [61]

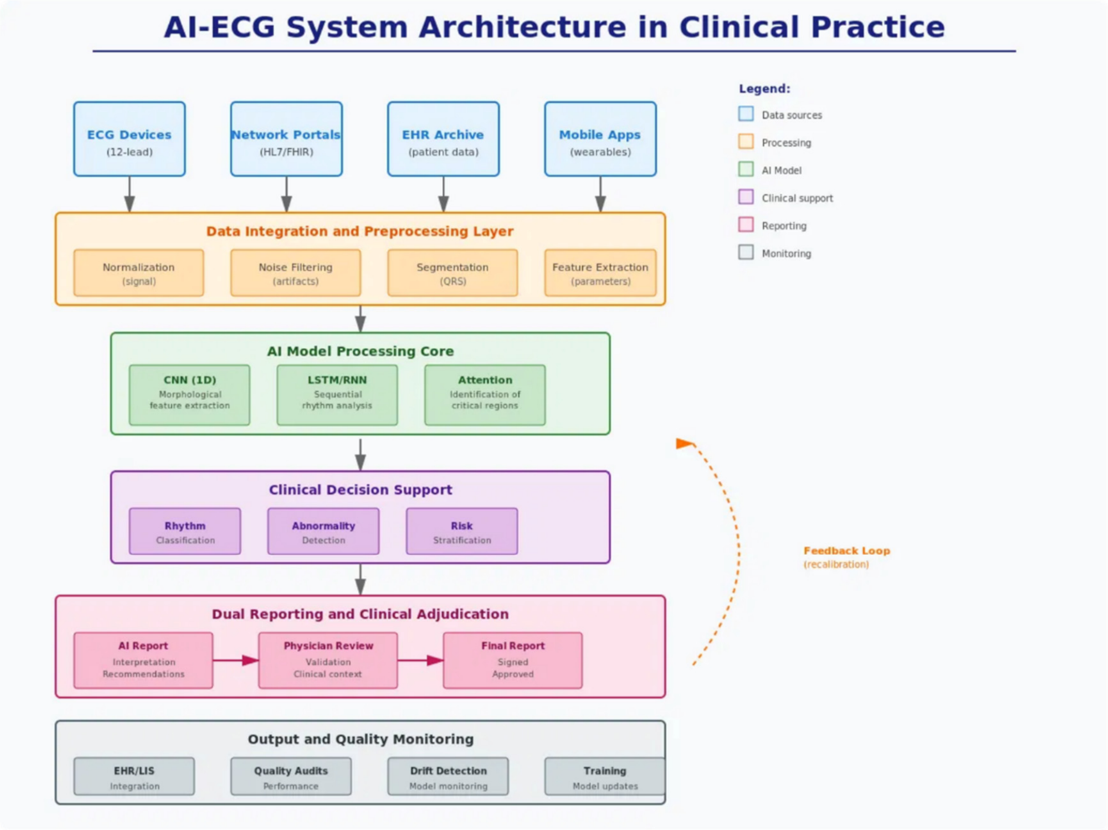

Figure 6.2. Architecture overview of AI ECG clinical workflow.

The architecture encompasses six major functional layers. The **data input layer** integrates multiple ECG sources: conventional 12-lead ECG devices, networked platforms supporting HL/FHIR standards, electronic health record (EHR) systems containing historical patient data, and mobile applications with wearable device integration for remote monitoring. The **data integration and preprocessing layer** performs signal normalization, noise filtering, segment identification (QRS complexes, intervals, and waves), and extraction of morphological and temporal features. The **AI model processing core** employs an ensemble of neural network architectures: one-dimensional convolutional neural networks (CNN- 1D) for morphological feature extraction, long short-term memory networks (LSTM) or gated recurrent units (GRU) for sequential rhythm analysis, and attention mechanisms for automatic identification of diagnostically relevant signal regions. The **clinical decision support layer** contextualizes model outputs into clinically meaningful categories: rhythm classification (normal sinus rhythm, atrial fibrillation, flutter, ectopy), structural abnormality detection (ventricular dysfunction, QRS amplitude variations, ST-segment anomalies), and clinical risk stratification. The **dual reporting and clinical adjudication layer** is where AI-generated interpretations are reviewed and validated by clinicians, maintaining physician responsibility while enabling efficient processing. Finally, the **output and quality monitoring layer** ensures integration with hospital information systems, conducts quality audits, monitors for performance degradation, and coordinates model maintenance and updates.

6.2. Implementation Metrics and Performance Monitoring

Robust implementation strategies are essential to translate technical performance into clinical value. However, translating performance into practice requires systematic monitoring of multiple dimensions of system performance and clinical impact. Table 6.1 presents a comprehensive set of metrics that should be monitored throughout the implementation lifecycle to ensure both safety and clinical benefit [63].

Table 6.1. AI ECG Implementation lifecycle.

Metric Category Performance Indicator		Implementation Target	Monitoring Frequency
Operational Efficiency	ECG report turnaround time	- % reduction for normal cases	Weekly
Operational Efficiency	Reporting throughput	- % increase	Monthly
Diagnostic Accuracy	Sensitivity for LV dysfunction detection	> % (baseline diagnostics)	Monthly
Diagnostic Accuracy	Specificity for atrial fibrillation screening	> %	Monthly
Clinical Safety	Missed critical findings (false negatives)	< % for critical abnormalities	Weekly
Clinical Safety	Physician-AI concordance rate	> % for pre-normal cases	Weekly
Model Calibration	Calibration error	< % per diagnostic class	Monthly
Subgroup Performance	Performance variance (sex)	< % difference (M vs F)	Quarterly
Subgroup Performance	Performance variance (age)	< % difference between age groups	Quarterly
	Physician acceptance rate	> % positive feedback	

doi:10.20944/preprints202605.1353.v1

These metrics address multiple dimensions of system performance. Operational metrics (report turnaround time, throughput) measure the efficiency gains that justify implementation costs. Diagnostic accuracy metrics (sensitivity, specificity) ensure that the system maintains or improves upon existing diagnostic standards. Clinical safety metrics (false negatives, concordance rates) provide early warning of system failures or inappropriate reliance on AI. Calibration and subgroup performance metrics ensure equitable and reliable performance across the entire patient population, regardless of demographic characteristics or clinical device variations.

It is essential that these metrics be collected **routinely and automatically**—not merely at the conclusion of a pilot phase or during special audits, but continuously on a weekly or monthly basis depending on metric criticality. Most modern implementations utilize automated monitoring dashboards that track these metrics in real time and generate alerts when thresholds are exceeded. Without automated monitoring infrastructure, data are collected irregularly and performance degradation may go undetected for extended periods.

6.3. Continuous Quality Monitoring and Adaptive Management

Continuous quality monitoring is a critical component of this lifecycle. Routine audits should track overall model performance, calibration, and subgroup behavior (for example by sex, age, or device vendor). The transition from rigorous clinical validation studies to real-world deployment often reveals performance variations not apparent in controlled trial environments. Systematic monitoring enables early detection and correction of such problems. Furthermore, as AI-ECG systems mature and accumulate larger clinical datasets, opportunities emerge for periodic model retraining with contemporary data, which may improve performance beyond that achievable in the original development dataset [62,63].

The integration of AI-enhanced electrocardiography into clinical practice represents a substantial organizational commitment but offers demonstrable benefits in efficiency, consistency, and diagnostic accuracy when implemented thoughtfully. Success depends not primarily on algorithm performance—though that remains essential—but rather on careful attention to system architecture, operational workflow integration, comprehensive metrics monitoring, structured implementation strategies, and sustained commitment to quality assurance. Healthcare organizations that approach AI-ECG implementation with this systematic framework are well positioned to translate technological capability into genuine clinical value and improved patient outcomes.

7. Future Directions and Priorities

7.1. Personalized Continuous Learning Systems

Future AI-ECG systems will move beyond static models to personalized, continuous learning frameworks. This involves models that adapt to individual patient baselines and clinical trajectories, learning from new data streams (e.g., from wearables and EHR) to track subtle deviations signaling disease progression or therapy response, supporting truly longitudinal phenotyping. Continuous learning frameworks, drawing from continual pretraining and replay buffers, would allow models to incorporate new data without forgetting prior knowledge, while mechanisms for safe online updating (e.g., via secure cloud inference with differential privacy) address regulatory demands for lifecycle management [64–66]. These paradigms promise not just diagnostic snapshots but dynamic risk trajectories, calibrated to the individual’s baseline and contextualized by real-time inputs like activity or vital signs (Figures 7.1 and 7.2 and Table 7.1).

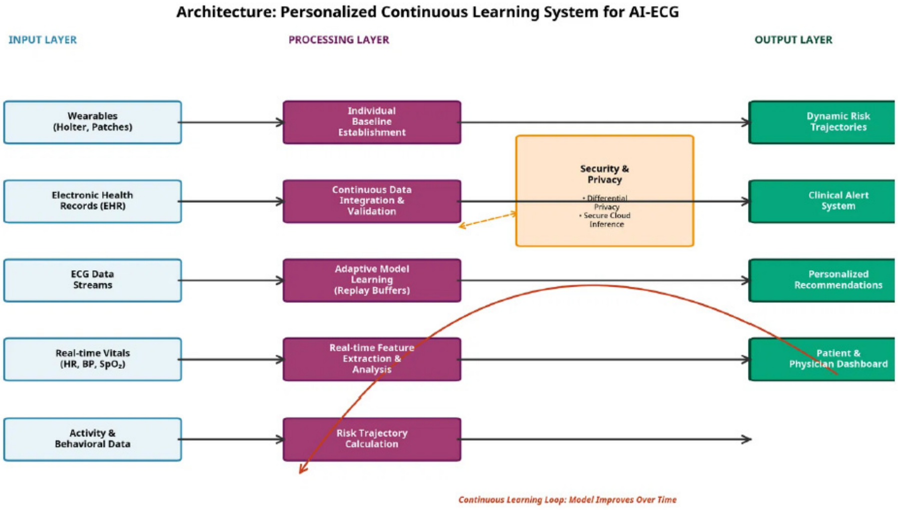

Figure 7.1. Architecture of Personalized Continuous Learning System: a multilayered AI-ECG architecture integrating data from wearables, electronic health records, vital signs, and activity monitoring through a centralized hub with continuous feature extraction, secure cloud inference with differential privacy, and replay buffers that enable safe online model updating while preventing loss of knowledge.

Table 1. Comparison of Traditional ECG Models versus Personalized Continuous Self-learning Systems. This table contrasts conventional universally-applied electrocardiographic models with next-generation personalized systems across seven dimensions, demonstrating the superior capacity of personalized approaches to generate individualized dynamic risk trajectories and detect subtle deviations from unique baseline physiology [67–69].

Aspect	Traditional ECG Models	Personalized Continuous Self-learning
Nature of Model	Static, Universal	Adaptive, Individual
Patient Adaptation	None; all patients compared to universal standard	Full; learns from each patient's baseline
Data Sources	Pre-collected ECG only	Continuous: ECG, wearables, EHR, vitals, activity
Update Frequency	New version: months/years	Real-time, continuous updates
Data Security	Standard protocols	Differential privacy, secure cloud inference
Clinical Decision Quality	Diagnostic "snapshots" only	Dynamic risk trajectories, individually calibrated
Knowledge Retention	Not applicable	Prevents catastrophic forgetting (replay buffers)

Figure 7.2. Data sources and Integration in Personalized AI-ECG Systems. Eight primary data modalities integrated into contemporary AI-ECG systems contributing distinct clinical information essential for comprehensive multimodal cardiac risk assessment [66,70].

7.2. Prevention of Sudden Cardiac Death Through AI Telemonitoring

The application of AI to continuous, long-term ECG monitoring and telemonitoring creates a previously unattainable opportunity for the prevention of sudden cardiac death (SCD) [71–74]]. Each year, more than 500,000 cases of SCD occur in the United States alone. AI-based long-term tele-ECG facilitates real-time quantitative and qualitative assessment of multiple bioelectrical parameters—including electrical instability (arrhythmias), ischemia, modulators, and triggers—that contribute to the occurrence of SCD, thereby enabling dynamic electrophysiological testing in real time. Deep learning models have demonstrated ability to predict SCD with area under the curve values ranging from 0.66 to 0.90, outperforming conventional ECG risk scores and traditional guideline criteria for implantable cardioverter-defibrillator placement [74]. When integrated with data from electronic health records, such systems could permit the identification, in accordance with catastrophe theory, of specific constellations of factors that herald the imminent onset of SCD. In turn, an appropriately configured alert mechanism could notify the monitoring center and the patient, potentially enabling life-saving intervention (Figures 7.3–7.5).

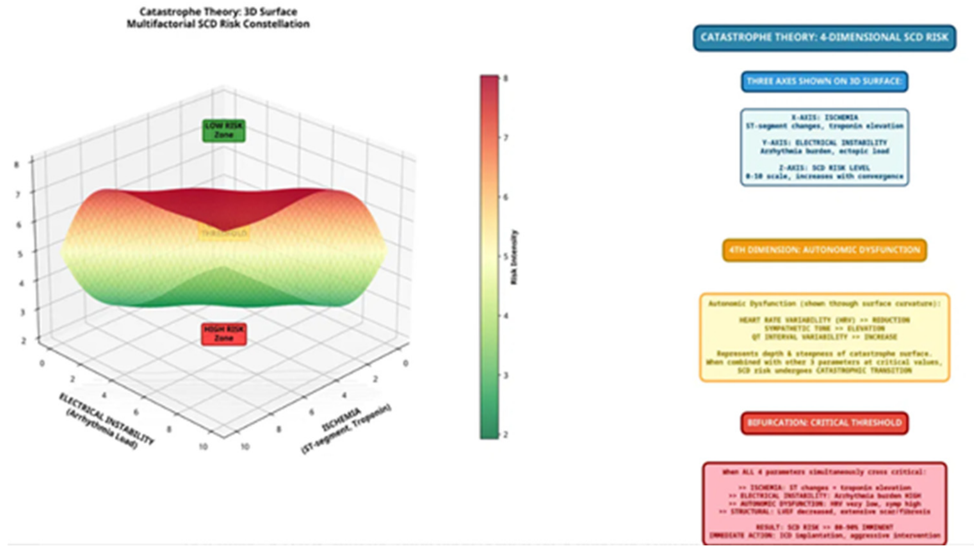

Figure 7.3. Catastrophe Theory – Multifactorial Nature of Sudden Cardiac Death Risk. This three-dimensional surface plot visualizes how four critical pathophysiological parameters monitored by AI-ECG: ischemia, electrical instability, autonomic dysfunction, and structural vulnerability interact nonlinearly to determine SCD risk, with the bifurcation point representing the catastrophic threshold where simultaneous modulation of multiple parameters is necessary for effective SCD prevention.

7.3. AI-ECG in Preventive Medicine and Cardiac Rehabilitation

AI-ECG holds immense promise for preventive medicine and resource-limited settings, where access to specialists and advanced imaging is scarce. By distilling complex prognostic signals from ubiquitous, low-cost tracings, AI models can enable population-level screening for silent atrial fibrillation, left ventricular dysfunction, or metabolic risk in primary care and community health programs. In low- and middle-income countries, deployable on smartphones or basic ECG devices, these tools could triage patients for targeted interventions, optimizing strained healthcare pathways.

Hybrid cardiac telerehabilitation has the potential to play a central role in remote care through offering support for both patients and healthcare professionals. Real-time remote monitoring with automated analysis of long-term ECG recordings can enhance diagnostic accuracy and facilitate the early detection of clinically relevant, even asymptomatic, abnormalities. On this basis, AI systems may generate alerts for clinicians or patients, enabling prompt intervention, which is particularly important for individuals performing telerehabilitation independently, at a distance from healthcare facilities, even when they are asymptomatic.

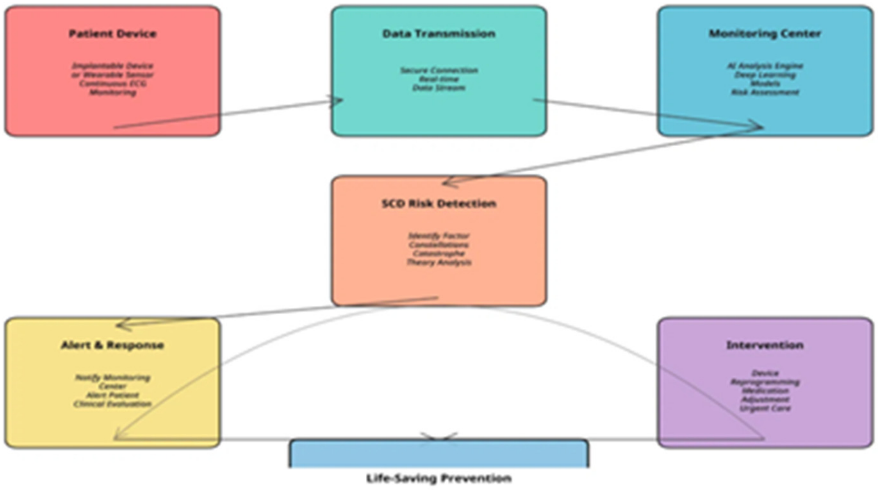

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 May 2026 doi:10.20944/preprints202605.1353.v1

Figure 7.4. An integrated AI telemonitoring platform that analyzes continuous ECG data in real time, identifies pathophysiological factor constellations predictive of imminent SCD using catastrophe theory, and triggers automated alerts to monitoring centers and patients for rapid clinical evaluation and potential life-saving intervention [67,75].

AI can integrate together ECG signals, blood pressure, and physical activity data in the context of hybrid cardiac telerehabilitation to individualize rehabilitation programs (Figure 7.5)

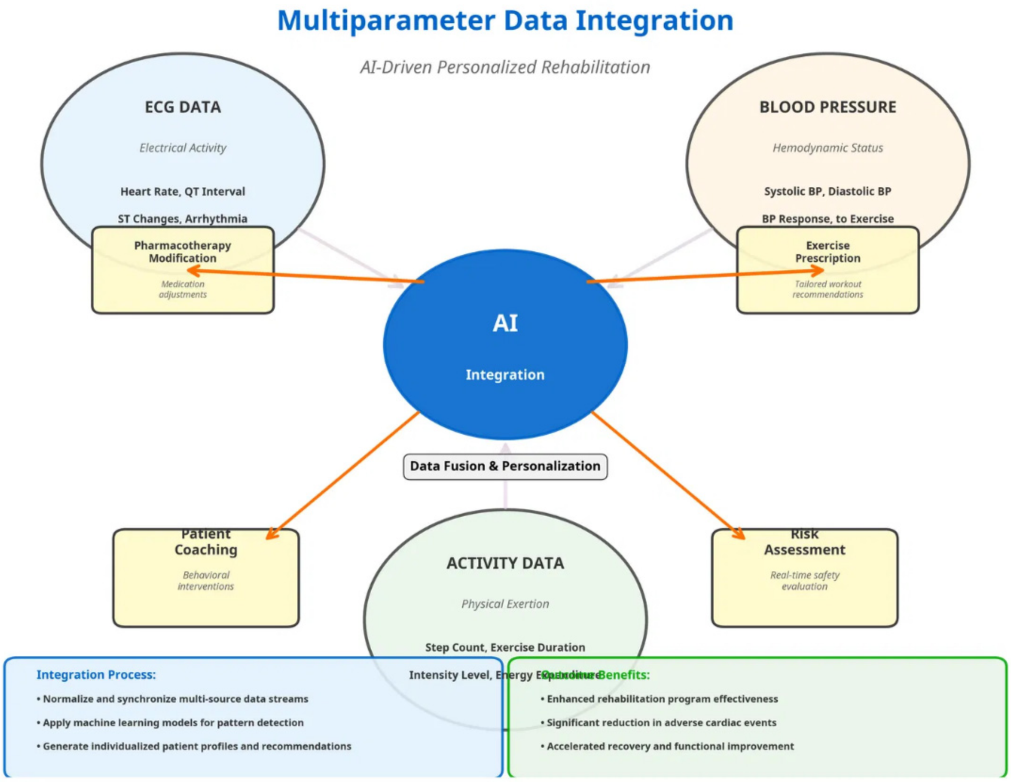

Figure 7.5. Multiparameter Data Integration for Personalized Rehabilitation: a novel approach of integrating multiple physiological data streams through artificial intelligence to create individualized cardiac rehabilitation programs. This approach emphasizes the central role of AI as an integration hub that synthesizes heterogeneous clinical data into unified, personalized therapeutic recommendations [76–78].

AI-Telerehabilitation System may also remotely adjust the course of rehabilitation by providing patients with tailored recommendations regarding exercise prescription and, when appropriate, modifications in pharmacotherapy. Such personalized management can improve the effectiveness of rehabilitation, reduce the risk of complications, and support faster recovery.

A meta-analysis shows that structured, telemedicine-based exercise cardiac rehabilitation is as effective as center-based programs, while explicitly emphasizing that telemedicine-based cardiac rehabilitation overcomes temporal and spatial constraints and expands access for patients unable to attend centers in person. In combination with telemedicine platforms, AI-assisted monitoring allows follow-up of patients irrespective of geographic location. This is especially relevant for individuals living in remote or underserved areas and may help mitigate the currently suboptimal access to cardiac rehabilitation, independent of place of residence.

Wearable technologies free cardiac telerehabilitation centers from many constraints of time and location. However, a key challenge is how to assess and properly interpret cardiac telerehabilitation progression based on the massive amounts of data collected. ECG assessment is an important, but time-consuming, element of qualifying patients, planning and monitoring the effectiveness of physical training, and making decisions in situations where patient safety is at risk during telerehabilitation. Wearable devices equipped with an accelerometer and ECG sensors can obtain heart electric parameters and estimate patients' physical exercise effort, which can serve as input data for AI systems, thereby reducing the burden on healthcare staff (Figure 7.6) .

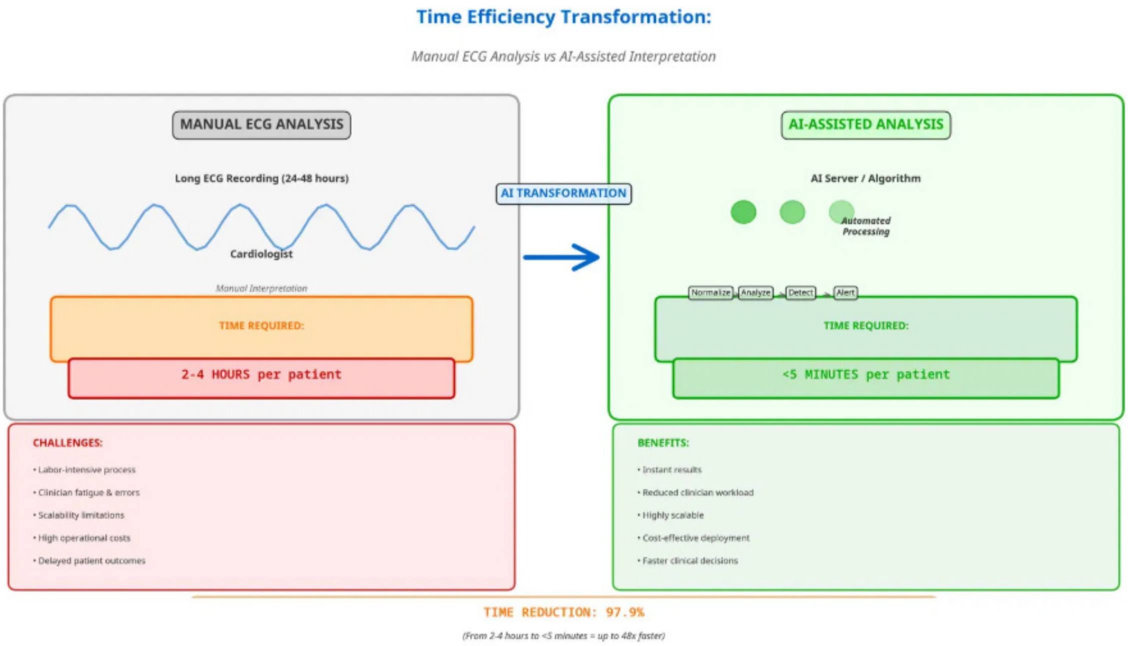

Figure 7.6. Time Efficiency Gain due to AI ECG implementation: Manual ECG Analysis vs AI-Assisted Interpretation.

Future studies should rigorously test such applications through cluster-randomized trials in diverse global contexts, measuring impacts on hard outcomes like stroke prevention or heart failure hospitalization. Ethical safeguards—ensuring models are trained on representative ancestries, validated prospectively, and paired with capacity-building for local clinicians—will be essential to avoid exacerbating disparities. Collectively, these directions position AI-ECG as a cornerstone of precision prevention, bridging high-tech innovation with universal accessibility.

8. Conclusions

Artificial intelligence AI has rapidly transformed the electrocardiogram from a heart's electrical activity diagnostic recording into a rich source of quantitative information that can be mined for clinical insight beyond the limits of visual interpretation. The 12-lead ECG remains one of the most widely used and accessible tools in cardiovascular medicine, yet traditional interpretation continues to be constrained by observer variability, limited sensitivity for subtle or subclinical abnormalities, and the practical burden of reviewing large volumes of tracings generated by modern monitoring systems. AI-based methods, especially deep learning, offer a powerful approach to addressing these limitations by learning high-dimensional patterns directly from raw or minimally processed ECG signals and linking them to rhythm disturbances, structural heart disease, and broader systemic phenotypes.

The evidence reviewed in this article shows that AI-ECG has already progressed from proof-of-concept studies to clinically meaningful applications. In arrhythmia detection, AI algorithms have demonstrated strong performance in identifying atrial fibrillation, atrial flutter, supraventricular tachycardia, ventricular tachycardia, and other rhythm abnormalities, including in single-lead wearable devices and ambulatory monitoring systems. In parallel, AI ECG has shown its utility as a screening technique to detect left ventricular dysfunction, left ventricular hypertrophy, hypertrophic

doi:10.20944/preprints202605.1353.v1

cardiomyopathy, dilated cardiomyopathy, pulmonary hypertension, and other forms of structural heart disease that may be impossible to recognize by conventional ECG criteria alone. These applications are particularly important because they may allow earlier diagnosis, more efficient triage, and improved selection of patients for confirmatory imaging or specialist evaluation.

Beyond classical cardiovascular disease, AI-ECG is increasingly being used as a digital biomarker for noncardiac and systemic conditions. Models have been developed to estimate biological age, predict future atrial fibrillation, infer obstructive sleep apnea, and identify patterns associated with cardiometabolic risk. This expansion reflects a broader shift in the role of the ECG, from its traditional role toward a quantitative physiologic signal capable of capturing subtle information about autonomic balance, myocardial remodeling, and overall health status. In this respect, AI-ECG may ultimately serve not only as a diagnostic tool but also as a longitudinal marker of disease risk, treatment response, and physiologic aging.

Several methodological advances have accelerated progress in the field. Foundation models trained on large ECG datasets, self-supervised learning strategies, multimodal architectures that combine ECG data with electronic health records or imaging, and generative models for synthetic waveform creation all represent important developments that may improve performance, generalization, and adaptability. These methods reduce dependence on labor-intensive annotation and support transfer to new tasks and populations. At the same time, explainability techniques such as saliency mapping, integrated gradients, and lead-wise relevance analysis are helping to bridge the gap between black-box predictions and clinically familiar ECG concepts. Although interpretability remains imperfect, it is becoming increasingly clear that transparency will be essential for clinical acceptance, safety oversight, and regulatory approval.

Despite these advances, major challenges remain before AI-ECG can be fully integrated into routine practice. Generalizability across healthcare systems, vendors, patient populations, and acquisition settings remains a persistent concern. Models that perform well in their development environment may show reduced accuracy when deployed externally because of differences in preprocessing, sampling rate, lead placement, prevalence of disease, or demographic composition. Algorithmic bias is another important issue, since uneven representation of sex, age, race, ethnicity, and socioeconomic background can lead to disparities in model performance. For that reason, future studies must place greater emphasis on external validation, subgroup analysis, calibration, and fairness metrics rather than relying only on summary discrimination statistics.

Privacy, data governance, and security also require careful attention. ECG signals are increasingly recognized as potentially identifiable biometric data, and the use of large-scale centralized datasets raises questions about consent, ownership, and downstream reuse. Emerging approaches such as federated learning, differential privacy, and secure computation may help reduce these risks, but they add technical complexity and do not eliminate the need for robust governance frameworks. In parallel, regulatory expectations for software as a medical device are evolving, and AI-ECG systems will increasingly need to demonstrate not only analytic performance but also prospective clinical benefit, workflow compatibility, and long-term safety. This is especially true for adaptive models that change over time or are updated after deployment.

Successful implementation in clinical practice will depend on how well AI systems are integrated into existing workflows. The most realistic model is not replacement of physician interpretation but augmentation through dual reporting, triage support, and decision assistance. In this setting, AI can help classify routine studies, flag urgent abnormalities, prioritize complex tracings, and reduce reporting delays, while the final interpretation and responsibility remain with the clinician. For this approach to work, systems must provide usable outputs, clear confidence information, and interfaces that fit within current clinical operations. Without careful attention to workflow design, even highly accurate models may fail to improve care or may introduce alert fatigue, mistrust, or operational inefficiency.

Looking ahead, the most important opportunities lie in personalized and continuous monitoring, real-time risk stratification, and broader access to cardiovascular diagnostics. Future AI-

ECG systems will likely combine waveform analysis with longitudinal clinical data, wearable signals, and multimodal information to produce more individualized and dynamic assessments of risk. Such systems may be particularly valuable for early detection of silent atrial fibrillation, prevention of sudden cardiac death, monitoring of heart failure progression, and remote surveillance in resource-limited settings. If developed and validated responsibly, AI-ECG could expand access to high-quality cardiovascular interpretation in settings where expert readers are scarce and could support more proactive and preventive models of care.

In summary, AI-enhanced electrocardiography represents a major advance in cardiovascular diagnostics and risk assessment. Its promise lies in the ability to extract clinically relevant information from a signal that has long been available but only partially utilized. At the same time, its limitations highlight the need for rigorous validation, transparent reporting, bias mitigation, and thoughtful clinical integration. With continued methodological innovation, multidisciplinary collaboration, and careful governance, AI-ECG has the potential to become an important component of future cardiovascular medicine and to improve outcomes across a wide range of clinical environments.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Raghunath S, Pfeifer JM, Ulloa Cerna AE, et al. Artificial intelligence-enabled electrocardiography from scientific discovery to clinical practice. Nat Rev Cardiol. 2023;20(10):621-639.
2. Cook DA, Oh SY, Pusic MV. Accuracy of physicians' electrocardiogram interpretations: a systematic review and meta-analysis. JAMA Intern Med. 2020;180(11):1461-1471.
3. Raghunath S, Ulloa Cerna AE, Jing L, et al. Deep neural networks can predict new-onset atrial fibrillation from the 12-lead electrocardiogram and help identify those at risk of AF-related stroke. Circulation. 2021;143(13):1287-1298.
4. Lima EM, Ribeiro AH, Paixão GMM, et al. Deep neural network-estimated electrocardiographic age as a mortality predictor. Nat Commun. 2021;12(1):117.
5. Hannun AY, Rajpurkar P, Haghpanahi M, et al. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nat Med. 2019;25(1):65-69.
6. Attia ZI, Kapa S, Lopez-Jimenez F, et al. Screening for cardiac contractile dysfunction using an artificial intelligence–enabled electrocardiogram. Nat Med. 2019;25(1):70-74.
7. Attia ZI, Noseworthy PA, Lopez-Jimenez F, et al. An artificial intelligence–enabled ECG algorithm for the identification of patients with atrial fibrillation during sinus rhythm: a retrospective analysis of outcome prediction. Lancet. 2019;394(10201):861-867.
8. Steinhubl SR, Waalen J, Edwards AM, et al. Effect of a home-based wearable continuous ECG monitoring patch on detection of undiagnosed atrial fibrillation: the mSToPS randomized clinical trial. JAMA. 2018;320(2):146-155.
9. Ribeiro AH, Ribeiro MH, Paixão GMM, et al. Automatic diagnosis of the 12-lead ECG using a deep neural network. Nat Commun. 2020;11(1):1760.
10. Naqvi TZ, Marinelli T, Tadi P, et al. State-of-the-art deep learning methods on electrocardiogram data: systematic review. J Med Internet Res. 2022;24(8):e38454.
11. Liu X, Wang H, Li Z, Qin L. Deep learning in ECG diagnosis: a review. Knowl Based Syst. 2021;227:107187.
12. Hong S, Zhou Y, Shang J, Xiao C, Sun J. Opportunities and challenges of deep learning methods for electrocardiogram data: a systematic review. Comput Biol Med. 2020;122:103801.
13. Strodthoff N, Wagner P, Schaeffter T, Samek W. Deep learning for ECG analysis: benchmarks and insights from PTB-XL. IEEE J Biomed Health Inform. 2021;25(5):1519-1528.
14. Raghunath S, Pfeifer JM, Ulloa-Cerna AE, et al. Foundation models for electrocardiogram interpretation. Eur Heart J. 2024;47(3):257-268.
15. Kim YJ, Cho YS, Kim Y, et al. A foundation transformer model with self-supervised learning for electrocardiograms. Commun Med (Lond). 2025;5:223.

doi:10.20944/preprints202605.1353.v1

16. Mehari T, Strodthoff N. Self-supervised representation learning from 12-lead ECG data. Comput Biol Med. 2022;141:105114.
17. Diamant N, Di Achille P, Weng LC, et al. Patient contrastive learning: a performant, expressive, and practical approach to electrocardiogram modeling. PLoS Comput Biol. 2022;18(2):e1009862.
18. Lenis G, Pilia N, Loewe A, Schulze WH, Dössel O. Comparison of baseline wander removal techniques considering the preservation of ST changes in the ischemic ECG: a simulation study. Comput Math Methods Med. 2017;2017:9295029.
19. Chang KC, Hsieh PH, Wu MY, et al. Usefulness of multi-labelling artificial intelligence in detecting rhythm disorders and acute ST-elevation myocardial infarction on 12-lead electrocardiogram. Eur Heart J Digit Health. 2021;2(2):299-310.
20. Nolin-Lapalme A, et al. Artificial intelligence-enhanced electrocardiograms: building on solid foundations. Eur Heart J Digit Health. 2026;7.
21. ExChanGeAI Investigators. End-to-end platform for electrocardiogram analysis and model fine-tuning. J Med Internet Res. 2026;28:e81116.
22. Hannun AY, Rajpurkar P, Haghpanahi M, Tison GH, Bourn C, Turakhia MP, et al. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nat Med. 2019;25:65-69.
23. Attia ZI, Noseworthy PA, Lopez-Jimenez F, et al. An artificial intelligence–enabled ECG algorithm for the identification of patients with atrial fibrillation during sinus rhythm: a retrospective analysis of outcome prediction. Lancet. 2019;394(10201):1861-1867.
24. Yao X, Rushlow DR, Inselman JW, et al. Artificial intelligence–enabled electrocardiograms for identification of patients with low ejection fraction: a pragmatic, randomized clinical trial. Nat Med. 2021;27(5):815-819.
25. Khurshid S, Friedman S, Reeder C, et al. ECG-based deep learning and clinical risk factors to predict atrial fibrillation. Circulation. 2022;145(21):122-133.
26. Cohen-Shelly M, Attia ZI, Friedman PA, et al. Electrocardiogram screening for aortic valve stenosis using artificial intelligence. Eur Heart J. 2021;42(30):2885-2896.
27. Sangha V, Nargund S, Dhingra LS, et al. Identification of hypertrophic cardiomyopathy on electrocardiographic images with deep learning. Circ Arrhythm Electrophysiol. 2024;17(1):e012284.
28. Kwon JM, Kim KH, Jeon KH, et al. Deep learning for predicting in-hospital mortality among heart disease patients based on echocardiography. Echocardiography. 2019;36(2):213-218.
29. Adedinsewo D, Carter RE, Attia Z, et al. Artificial intelligence–enabled ECG algorithm to identify patients with left ventricular systolic dysfunction presenting to the emergency department with dyspnea. Circ Arrhythm Electrophysiol. 2020;13(8):e008437.
30. Kwon JM, Jo YY, Lee SY, et al. Artificial intelligence for early prediction of pulmonary hypertension using electrocardiography. J Heart Lung Transplant. 2020;39(8):805-814.
31. Kalscheur MM, Kipp RT, Tattersall MC, et al. Machine learning algorithm predicts cardiac resynchronization therapy outcomes: lessons from the COMPANION trial. Circ Arrhythm Electrophysiol. 2018;11(1):e005499.
32. Feeny AK, Rickard J, Patel D, et al. Machine learning prediction of response to cardiac resynchronization therapy: improvement versus current guidelines. Circ Arrhythm Electrophysiol. 2019;12(1):e007316.
33. Giudicessi JR, Schram M, Bos JM, et al. Artificial intelligence-enhanced electrocardiography-derived body mass index as a predictor of future cardiometabolic disease. Mayo Clin Proc. 2022;97(11):2070-2083.
34. Khurshid S, Sangha V, Patel AP, et al. Deep learning to predict cardiac magnetic resonance-derived left ventricular mass and hypertrophy from 12-lead ECGs. Circ Cardiovasc Imaging. 2021;14(5):e012281.
35. Noseworthy PA, Attia ZI, Brewer LC, et al. Assessing and mitigating bias in medical artificial intelligence: the effects of race and ethnicity on a deep learning model for ECG analysis. Circ Arrhythm Electrophysiol. 2020;13(3):e007988.
36. Somani S, Russak AJ, Richter F, et al. Deep learning and the electrocardiogram: review of the current state-of-the-art. Europace. 2021;23(8):1179-1191.

37. van de Leur RR, Taha K, Bos MN, et al. Improving explainability of deep neural network-based electrocardiogram interpretation using variational auto-encoders. Eur Heart J Digit Health. 2022;3(3):390-404.
38. Singstad BH, Schirmer H, Ytrehus K. Explaining deep neural networks for knowledge discovery in electrocardiogram analysis. Sci Rep. 2021;11(1):10949.
39. Wickramaratne SD, Mahmud MS. Conditional-GAN based data augmentation for deep learning task classifier improvement using fNIRS data. Front Big Data. 2021;4:659146.
40. Ribeiro AH, Giatti L, Barreto SM, Duncan BB. Interpretable machine learning for 12-lead ECG classification based on ensembles of symbolic representations. Comput Methods Programs Biomed. 2021;205:106097.
41. Baek YS, Lee SC, Choi W, Kim DH. A new deep learning algorithm of 12-lead electrocardiogram for identifying atrial fibrillation during sinus rhythm. Sci Rep. 2021;11(1):12818.
42. Ribeiro AH, Oliveira GMM, Ribeiro ALP, et al. Artificial intelligence-enhanced electrocardiography for accurate diagnosis and risk prediction: opportunities, barriers, and future directions. Comput Struct Biotechnol J. 2024;22.
43. Harris S, et al. Artificial intelligence-enhanced electrocardiography for early diagnosis and management: technical, ethical, and regulatory considerations. Sci Rep. 2023;13.
44. Gichoya JW, Banerjee I, Bhimireddy AR, et al. AI recognition of patient race in medical imaging: a modelling study. Lancet Digit Health. 2022;4(6):e406-e414.
45. Pfohl SR, Foryciarz A, Shah NH. An empirical characterization of fair machine learning for clinical risk prediction. J Biomed Inform. 2021;113:103621.
46. Cirillo D, Catuara-Solarz S, Morey C, et al. Sex and gender differences and biases in artificial intelligence for biomedicine and healthcare. NPJ Digit Med. 2020;3:81.
47. Chen IY, Pierson E, Rose S, Joshi S, Ferryman K, Ghassemi M. Ethical machine learning in healthcare. Annu Rev Biomed Data Sci. 2021;4:123-144.
48. Kaissis GA, Makowski MR, Rückert D, Braren RF. Secure, privacy-preserving and federated machine learning in medical imaging. Nat Mach Intell. 2020;2(6):305-311.
49. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119.
50. Dwork C, Roth A. The algorithmic foundations of differential privacy. Found Trends Theor Comput Sci. 2014;9(3-4):211-407.
51. Warnat-Herresthal S, Schultze H, Shastry KL, et al. Swarm learning for decentralized and confidential clinical machine learning. Nature. 2021;594(7862):265-270.
52. U.S. Food and Drug Administration. Artificial intelligence and machine learning (AI/ML)-enabled medical devices [Internet]. Silver Spring, MD: FDA; 2021 [cited 2026 Mar 4]. Available from: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.
53. European Medicines Agency. Good machine learning practice for medical device development: guiding principles [Internet]. Amsterdam: EMA; 2021 [cited 2026 Mar 4]. Available from: https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/good-machine-learning-practice-medical-device-development-guiding-principles_en.pdf.
54. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25(1):44-56.
55. Peiffer-Smadja N, Rawson TM, Ahmad R, et al. Machine learning for clinical decision support in infectious diseases: a narrative review of current applications. Clin Microbiol Infect. 2020;26(5):584-595.
56. Lin CS, Liu WT, Chen YH, Lin SH, Lin C. Artificial intelligence-enabled electrocardiography from scientific research to clinical application. EMBO Mol Med. 2026;18(2):e18240.
57. Johnson LS, Zadrozniak P, Jasina G, et al. Artificial intelligence for direct-to-physician reporting of ambulatory electrocardiography. Nat Med. 2025;31(9):925-931.
58. Pors SR, Attia ZI, Kapa S, et al. A multicenter pragmatic implementation study of AI-ECG-based clinical decision support software to identify low LVEF: clinical trial design and methods. Am Heart J Plus. 2025;5:100528. doi: 10.1016/j.ahjo.2025.100528.

59. Lin CS, Al-Khatib SM, Boriani G, et al. State of the art of artificial intelligence in clinical electrophysiology in 2025: a scientific statement of the European Heart Rhythm Association. Europace. 2025;27.
60. Topol EJ, Narayan S, et al. From innovation to implementation: addressing the AI adoption gap in cardiovascular medicine. iScience. 2026;29(1):11234.
61. Kwan GF, Verbrugge FH, et al. Cognitive alignment in cardiovascular AI: designing predictive tools that fit clinician reasoning and workflows. Front Cardiovasc Med. 2025;21:651324.
62. Wiens J, et al. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019;25(11):1337-1340.
63. Andersen ES, Birk-Korch JB, Hansen RS, et al. Monitoring performance of clinical artificial intelligence in health care: a scoping review. JBI Evid Synth. 2024;22(12):42423-42444.
64. Yao X, Rushlow DR, Inselman JW, et al. Artificial intelligence-enabled electrocardiograms for identification of patients with low ejection fraction: a pragmatic, randomized clinical trial. Nat Med. 2021;27(10):1888-1894.
65. Attia ZI, Noseworthy PA, Lopez-Jimenez F, et al. Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. Nat Rev Cardiol. 2021;18(12):881-909.
66. Tasmurzayev N, Amangeldy B, Imanbek B, et al. Digital cardiovascular twins, AI agents, and sensor data: a narrative review from system architecture to proactive heart health. Sensors (Basel). 2025;25(17):5272.
67. Lin CS, Liu WT, Tsai DJ, et al. AI-enabled electrocardiography alert intervention and all-cause mortality: a pragmatic randomized clinical trial. Nat Med. 2024;30(5):1461-1470.
68. Attia ZI, Noseworthy PA, Lopez-Jimenez F, et al. Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. Nat Rev Cardiol. 2021;18(12):881-909.
69. Ding C, Yao T, Wu C, Ni J. Advances in deep learning for personalized ECG diagnostics: a systematic review addressing inter-patient variability and generalization constraints. Biosens Bioelectron. 2025;259:116423.
70. Palermi S, Vecchiato M, Ng FS, Attia Z, et al. Artificial intelligence and the electrocardiogram: a modern renaissance. Eur J Intern Med. 2025;125:55-63.
71. Thom R. Structural stability and morphogenesis: an outline of a general theory of models. Reading, MA: Benjamin/Addison-Wesley; 1975.
72. Zeeman EC. Catastrophe theory: selected papers 1972-1977. Reading, MA: Addison-Wesley; 1977.
73. Bhattacharya J, Lu DY, Kudchadkar SM, et al. Deep learning a person's risk of sudden cardiac death. Lancet. 2022;399(10342):2019-2031.
74. Sau A, Pastika L, Sieliwonczyk E, et al. Artificial intelligence-enabled electrocardiogram for mortality and cardiovascular risk estimation: a model development and validation study. Lancet Digit Health. 2024;6(11):e787-e797.
75. Silva-Cardoso J, Juanatey JRG, et al. The future of telemedicine in the management of heart failure patients. Curr Heart Fail Rep. 2021;27(5).
76. Su J, Zhang Y, Ke Q, et al. Mobilizing artificial intelligence to cardiac telerehabilitation. Rev Cardiovasc Med. 2022;23(2):45.
77. Antoniou V, Davos CH, Kapreli E, et al. Effectiveness of home-based cardiac rehabilitation, using wearable sensors, as a multicomponent, cutting-edge intervention: a systematic review and meta-analysis. J Clin Med. 2022;11(13):3772.
78. Cavalheiro AH, Silva-Cardoso J, et al. Effectiveness of tele-rehabilitation programs in heart failure: a systematic review and meta-analysis. Health Serv Res. 2021;56(1):158-173.