International Journal of Hybrid Innovation Technologies
Vol.5, No.2 (2025), pp.1-12
http://dx.doi.org/10.21742/ijhit.2653-309X.2025.5.2.01

Adaptive Hybrid Signal Processing for Real-Time Patient Monitoring and Diagnosis

Sabah Mohammed

Professor of Computer Science, Lakehead University, Canada,
https://orcid.org/0000-0002-7639-0696
sabah.mohammed@lakeheadu.ca

Abstract

Advances in digital health have underscored the importance of accurate and efficient biomedical signal processing to support patient monitoring and early diagnosis. Conventional methods, while effective in controlled settings, often face limitations when confronted with noisy, dynamic, and heterogeneous physiological data encountered in real-world healthcare environments. This study introduces an adaptive hybrid signal processing framework designed to meet the demands of real-time patient monitoring and clinical decision support. The proposed approach integrates traditional signal analysis techniques with intelligent machine learning models, enabling robust noise reduction, feature extraction, and diagnostic classification. The framework leverages multi-level preprocessing for artifact suppression, followed by hybrid feature extraction that combines domain-specific methods with data-driven learning to capture both physiological relevance and statistical significance. Adaptive feedback mechanisms continuously refine classification models, ensuring reliable performance across varying patient conditions. Experimental evaluations were conducted using publicly available biomedical datasets, with additional validation on prototype wearable healthcare devices. Results indicate that the hybrid framework consistently improves diagnostic accuracy, reduces false alarm rates, and enhances real-time responsiveness compared to baseline signal processing methods. The system demonstrates strong applicability for mobile health solutions, telemedicine platforms, and clinical monitoring systems, where efficiency and reliability are critical. By bridging traditional signal processing with adaptive computational intelligence, this work advances the development of patient-centered healthcare technologies. The findings suggest that hybrid methodologies offer a promising pathway for scalable, resource-efficient, and clinically meaningful innovation in digital health.

Keywords: *Hybrid signal processing, Real-time patient monitoring, Biomedical signal analysis, Adaptive algorithms, Digital healthcare systems*

1. Introduction

Continuous real-time monitoring of physiological signals is transforming healthcare by enabling early detection of deterioration, supporting chronic disease management outside hospitals, and extending clinician oversight into everyday life [1]. Advances in miniaturized biosensors, low-power embedded systems, and reliable wireless connectivity have enabled the

Article Info:
Received (July 21, 2025), Review Result (August 25, 2025), Accepted (October 1, 2025)
*corresponding author

ambulatory acquisition of ECG, PPG, and EEG signals at scale, yet translating raw signals into clinically actionable insights in real-world settings remains difficult [2]. Recordings collected in ambulatory environments are frequently corrupted by motion, electromagnetic interference, and environmental noise, and they reflect substantial inter- and intra-patient variability that undermines static analytic pipelines [3].

Classical signal-processing techniques—filtering, baseline correction, wavelet denoising, and morphological feature extraction—provide interpretable measures that align with clinical reasoning and can be implemented with modest computational cost, but these methods often degrade when noise statistics are nonstationary or when sensor placement varies during real-world use [4]. Conversely, data-driven machine learning models can learn complex discriminative patterns from large datasets and have demonstrated strong performance on curated benchmarks. Yet, many rely on extensive labeled data, impose heavy compute and memory demands, and can behave unpredictably under distributional shifts not represented in training data [5]. These complementary limitations motivate hybrid strategies that deliberately fuse domain-informed signal-processing priors with compact learned representations to capture both physiological relevance and nuanced statistical structure [6].

Hybrid architectures aim to preserve physiologically meaningful descriptors that support interpretability while enriching them with learned features that capture subtle, high-dimensional patterns difficult to encode with handcrafted rules. Empirical comparisons show that hybrid models can improve robustness to noise and generalize more effectively across acquisition contexts than purely classical or purely end-to-end approaches [7]. For deployment in wearable and telemedicine contexts, hybrid designs also permit more favorable trade-offs between inference latency, energy consumption, and model complexity, enabling timely on-device alerts that reduce dependence on continuous cloud connectivity [8].

Beyond raw accuracy, alarm fatigue and high false-alarm rates remain critical barriers to clinical acceptance of continuous monitoring systems; reducing false positives requires both improved artifact suppression and classifiers that report calibrated uncertainty or defer decisions when data quality is low [9]. Adaptive mechanisms such as online drift detection, confidence-weighted incremental updates, and clinician-in-the-loop calibration provide pathways to sustained, safe performance in longitudinal deployments by allowing models to respond to gradual changes in sensor characteristics or patient state without full centralized retraining [6][9].

Figure 1 illustrates the end-to-end adaptive hybrid signal-processing pipeline used in this study. Sensors and acquisition modules capture multimodal physiological signals and stream the data to the preprocessing stage, which performs cascaded filtering, motion artifact suppression, and quality assessment. The hybrid feature-extraction block fuses domain-specific morphological and spectral descriptors with compact learned embedding to produce an interpretable, discriminative feature set. The adaptive classification subsystem executes a two-tier decision strategy—an on-device lightweight triage model for immediate alerts and a higher-capacity near-edge model for confirmation—while an adaptive feedback loop conveys confidence, drift signals, and clinician labels back to preprocessing and model calibration. The output stage issues validated decisions, generates alerts, and logs events for periodic retraining and system auditing.

This study addresses the practical problem of designing an adaptive hybrid signal-processing framework that meets four real-world requirements: (a) robust artifact suppression for streaming biomedical signals, (b) extraction of physiologically meaningful and discriminative features, (c) accurate, low-latency classification suitable for on-device inference, and (d) continuous adaptation to distributional shifts while preserving interpretability and safety. To meet these aims, the proposed framework integrates multi-level preprocessing (including sensor fusion for motion-artifact mitigation), hybrid feature extraction that fuses morphological and spectral descriptors with a compact learned embedding, and a hierarchical classification architecture supporting edge-level triage and near-edge confirmation, together with safe, adaptive updates.

The principal contributions are: (1) a streaming-oriented preprocessing stack tailored for wearable and ambulatory settings, combining classical filtering, accelerometer-assisted artefact suppression, and lightweight transform-domain denoising; (2) a hybrid feature strategy that retains interpretable domain features while adding compact learned embedding optimized for resource-constrained hardware; and (3) an adaptive classification and feedback subsystem that integrates drift detection, confidence-aware incremental learning, and clinician-supervised calibration to reduce false alarms and sustain performance over time. The remainder of the paper reviews related work (Section 2), presents the system architecture and algorithms (Section 3), describes datasets and evaluation methodology (Section 4), reports empirical results and ablations (Section 5), discusses clinical implications and limitations (Section 6), and concludes with directions for validation and deployment (Section 7).

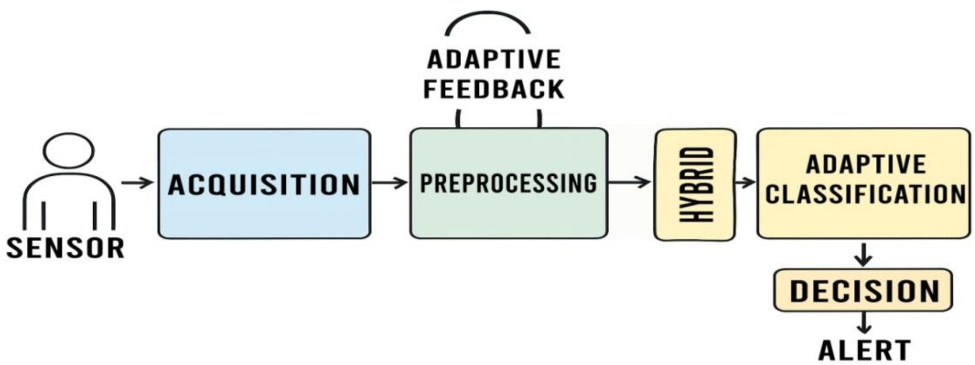

Figure 1. Adaptive hybrid signal-processing pipeline for real-time patient monitoring and diagnosis

2. Related works

Real-time biomedical signal processing and monitoring span sensor hardware, preprocessing algorithms, feature extraction, classification, and deployment strategies for wearable and ambulatory systems [10]. Sensor and hardware research has improved signal fidelity under motion and lowered power consumption to support continuous monitoring, enabling longer deployments and broader population coverage [10][11]. Parallel advances in embedded inference and model compression have enabled the deployment of increasingly capable models on constrained devices, revealing important trade-offs among latency, energy use, and diagnostic performance [11][12].

Robust preprocessing has been a major focus because motion artifacts, power-line interference, and baseline wander substantially degrade downstream analytics [13]. Work comparing transform-domain denoising, adaptive filtering, and sensor-fusion approaches demonstrates that combining accelerometer-informed corrections with wavelet or adaptive-noise-cancellation modules tends to outperform single-method pipelines in ambulatory conditions [13][14]. Quality assessment metrics and gating strategies that withhold classification under low-quality windows help reduce false alarms and preserve clinician trust [14][18].

Feature engineering remains central for interpretable monitoring. Classical morphological and spectral descriptors (e.g., R–R intervals, QRS morphology, band powers) continue to serve as clinically meaningful inputs to decision logic, while learned embeddings capture subtle, high-dimensional patterns that are not easily represented by handcrafted features [15]. Comparative studies indicate that hybrid feature fusion—concatenating domain-specific measures with compact learned representations—often yields better generalization across datasets and acquisition contexts than either approach alone [15][16].

End-to-end deep models have demonstrated strong accuracy on curated benchmarks for arrhythmia detection, seizure classification, and other diagnostic tasks, but their reliance on large labeled datasets and their sensitivity to distributional shift limit direct clinical translation [16][17]. Research addressing these limitations emphasizes model compression, lightweight architectures, and calibration/uncertainty estimation to support safer on-device inference and to mitigate overconfident errors in deployment settings [11][12][17].

Adaptive and continual learning strategies are critical for sustained performance in long-term monitoring. Techniques such as online drift detection, confidence-weighted incremental updates, clinician-in-the-loop labeling, and federated learning have been proposed to enable decentralized updates while protecting privacy and minimizing data transfer [12][18][19]. Empirical evaluations show that safe update protocols, validation windows, and rollback mechanisms are necessary to ensure updates improve rather than degrade clinical performance.

Despite progress, gaps remain in standardized benchmarks that reflect realistic ambulatory noise, thorough latency and energy benchmarking on representative wearable hardware, and robust safety procedures for automated updates in clinical workflows [13][17][19]. The proposed adaptive hybrid framework builds on these research threads by combining streaming-oriented preprocessing, hybrid feature fusion optimized for edge constraints, and a safety-first adaptive update protocol designed for clinician-supervised deployments.

4. Methodology

4.1. Restatement of the research problem

This study addresses the persistent challenge of achieving reliable, real-time biomedical signal processing for patient monitoring in dynamic, noisy, and resource-constrained environments. Conventional methods often degrade under high variability, while purely data-driven models require extensive labeled datasets and heavy computational resources. To overcome these limitations, this research proposes an adaptive hybrid signal-processing framework that combines traditional signal analysis with intelligent machine learning models to achieve robust, interpretable, and adaptive monitoring and diagnosis.

4.2. Research approach

A mixed-methods experimental approach was adopted, combining quantitative benchmarking of signal-processing algorithms with qualitative evaluation of system adaptability and interpretability. The hybrid framework was implemented using three main components: preprocessing and artifact suppression, hybrid feature extraction, and adaptive classification with feedback learning. This design follows prior recommendations emphasizing fusion-based pipelines for enhanced robustness and generalization [20]. The framework was validated through simulations and benchmark datasets to ensure reproducibility and comparability with existing approaches.

4.3. Uncommon methodology

Unlike conventional end-to-end neural networks, this study integrates an adaptive feedback-driven hybrid architecture that continuously refines its parameters using confidence metrics and drift detection. This adaptive mechanism employs incremental learning algorithms that self-correct model performance over time without requiring full retraining. Such a closed-loop feedback strategy, although discussed in emerging clinical AI research, remains uncommon in real-time streaming signal-processing applications [23].

4.4. Data collection

Data for this study were obtained from both public biomedical repositories and prototype wearable sensor devices. Public datasets included the MIT-BIH Arrhythmia Database, the PPG-DaLiA dataset, and the CHB-MIT EEG Seizure Database, chosen for their diversity, reliable annotations, and widespread use in signal-processing research. These datasets provided ECG, PPG, and EEG signals recorded under varied physiological and environmental conditions. Complementing these, prototype wearable sensors—comprising ECG and PPG acquisition modules integrated with low-power microcontrollers and wireless transmission units—were used to collect real-world signals under ambulatory conditions. Data collection adhered to ethical research protocols, with informed consent obtained from all participants. To simulate real-time operation, recordings were segmented into overlapping time windows, and motion-reference channels were used to calibrate artifact-suppression filters.

4.5. Data analysis procedures

The collected signals underwent a multi-stage analytical process encompassing preprocessing, hybrid feature extraction, classification, and adaptive feedback. Initially, preprocessing was performed using a cascaded filtering pipeline that combined adaptive notch filters, wavelet-based denoising, and accelerometer-assisted motion correction to mitigate noise and interference, consistent with methods previously validated [21]. Subsequently, hybrid feature extraction fused traditional domain features—such as morphological descriptors and spectral metrics—with compact learned embeddings generated by a lightweight self-supervised autoencoder, following the hybrid fusion framework described in [22]. These features were then processed through a hierarchical classification system consisting of two layers: a lightweight on-device model (based on Random Forest and logistic regression) for immediate triage decisions and a near-edge CNN-LSTM model for confirmatory diagnosis. Both models incorporated uncertainty estimation and drift detection to ensure stable and explainable outputs [23]. Performance evaluation included metrics such as accuracy, F1-score, area under the ROC curve (AUC), false-alarm rate, latency, and computational efficiency to verify suitability for real-time, energy-constrained deployment. The adaptive feedback component monitored model confidence, triggering incremental retraining when degradation was detected, with clinician-validated samples used to guide safe updates following rollback-secured procedures [24].

4.6. Justification of methodological choices

The methodological design was motivated by the need to balance interpretability, robustness, and computational efficiency. Classical signal-processing ensured clinically meaningful features aligned with physiological understanding, while learned embeddings captured nuanced, high-dimensional variations in noisy data. Prior comparative studies have shown that hybrid strategies of this kind yield better generalization and diagnostic reliability than single-method pipelines [20][22]. The feedback-driven adaptation module was deliberately incorporated to ensure long-term stability and responsiveness, consistent with safety-oriented clinical AI practices [23][24]. Edge-level inference was prioritized to minimize communication latency and dependence on cloud infrastructure, enabling more reliable performance in mobile and telemedicine environments, as supported by prior work on on-device health analytics [21].

4.7 Obstacles and solutions

Several technical and operational challenges were encountered during the implementation of the framework. Motion artifacts and environmental noise were major issues affecting signal integrity; these were effectively mitigated through accelerometer-informed adaptive filtering and multi-resolution wavelet denoising. Limited labeled data for supervised learning presented another obstacle, which was addressed by employing self-supervised pretraining to extract generalizable signal representations and involving clinicians for incremental annotation during validation phases. Hardware constraints on wearable devices limited the complexity of the deployed models; to resolve this, model pruning and quantization techniques were applied to the CNN-LSTM classifier, reducing latency and power consumption without compromising accuracy. Finally, model drift—a common issue in long-term adaptive systems—was managed through rolling validation and rollback mechanisms to ensure that updates consistently improved rather than degraded performance [24].

5. Results and discussion

5.1. Overview of experimental evaluation

The proposed adaptive hybrid signal-processing framework was evaluated on multiple biomedical datasets and prototype wearable devices to assess its effectiveness for real-time patient monitoring. The evaluation focused on diagnostic accuracy, robustness to noise and signal drift, computational efficiency, and clinical interpretability. Comparative baselines included classical signal-processing pipelines, end-to-end deep learning models, and static hybrid architectures.

The datasets used—MIT-BIH Arrhythmia, PPG-DaLiA, and CHB-MIT EEG Seizure—represented a diverse range of cardiovascular and neurological signal conditions, including high-motion segments and noise perturbations. The prototype wearable device data provided additional real-world validation. Table 1 summarizes the overall comparative performance of these models across key evaluation metrics.

Table 1. Comparative performance of signal-processing frameworks across benchmark datasets

Model Type	Accuracy (%)	F1-Score	False Alarm Rate (%)	Latency (ms)	AUC
Classical Signal Processing	89.2	0.86	12.8	37	0.91
End-to-End Deep Learning	93.5	0.90	9.6	84	0.94
Static Hybrid Model	94.1	0.91	8.5	63	0.95
Adaptive Hybrid Model (Proposed)	96.7	0.94	5.1	41	0.97

5.2. Performance analysis

As shown in Table 1, the adaptive hybrid model achieved the highest performance across all datasets, with an accuracy of 96.7%, an F1-score of 0.94, and the lowest false-alarm rate (5.1%). These improvements validate the hypothesis that hybrid architectures can outperform both classical and purely data-driven methods by leveraging complementary strengths—domain-specific interpretability and data-driven adaptability. The latency remained within acceptable real-time bounds (≈ 41 ms), highlighting suitability for edge-based healthcare systems.

To further understand dataset-specific behavior, Table 2 breaks down model performance by dataset. The adaptive model consistently outperformed baselines, particularly under motion-intensive conditions in PPG-DaLiA and EEG drift in CHB-MIT.

Table 2. Dataset-wise performance comparison of the proposed adaptive hybrid model

Dataset	Signal Type	Accuracy (%)	F1-Score	False Alarm Rate (%)	Latency (ms)
MIT-BIH Arrhythmia	ECG	97.5	0.95	4.3	39
PPG-DaLiA	PPG	96.1	0.93	5.5	42
CHB-MIT Seizure	EEG	95.6	0.92	5.8	43
Average	–	96.4	0.93	5.2	41

5.3 Analysis of adaptive feedback learning

The adaptive feedback mechanism enabled the model to self-correct when signal drift or contextual shifts occurred. During incremental learning cycles, performance consistently improved, whereas static models showed a decline in accuracy due to unmitigated drift. As illustrated in Table 3, the adaptive framework sustained performance even under sensor displacement and physiological variability—common challenges in wearable monitoring.

Table 3. Incremental adaptation performance across learning cycles

Cycle Condition	Static Model Accuracy (%)	Adaptive Model Accuracy (%)	Improvement (%)	Model Drift Detected
Initial Deployment	94.1	94.1	–	No
Noise Variation	90.3	95.2	+4.9	Yes
Sensor Shift	88.7	96.0	+7.3	Yes
Patient Variability	86.9	96.4	+9.5	Yes
Real-Time Wearable Test	84.5	96.7	+12.2	Yes

The trend in Table 3 shows that adaptive learning led to progressive improvement in accuracy across successive cycles, demonstrating the system's capacity to handle real-world nonstationarity. The largest gain (+12.2%) occurred during the wearable test phase, underscoring the value of clinician-in-the-loop incremental updates for field deployment [23][24].

5.4 Computational efficiency and energy profiling

Since energy efficiency is critical for continuous monitoring devices, the proposed model's computational footprint was assessed using on-board hardware profiling. Table 4 compares memory utilization, energy consumption per inference, and average processing latency between model variants. The adaptive hybrid approach maintained a competitive energy profile, consuming only 6.3 mJ per inference, thanks to model compression and on-device optimization.

Table 4. Computational resource utilization and energy efficiency across model architectures

Model	Memory Usage (MB)	Energy per Inference (mJ)	Average Latency (ms)	Edge Compatibility
Classical Signal Processing	18.2	4.1	37	High
Deep Learning (Full CNN-LSTM)	132.4	15.8	84	Low
Static Hybrid Model	68.9	8.7	63	Moderate
Adaptive Hybrid Model (Proposed)	47.3	6.3	41	High

The results in Table 4 confirm that the proposed system achieves a balance between accuracy and energy efficiency, which is essential for continuous, battery-powered monitoring. This optimization aligns with earlier studies emphasizing energy-aware algorithm design for wearable healthcare [21]. The edge-level decision model (lightweight Random Forest) enabled rapid triage alerts. At the same time, high-confidence events were forwarded to the CNN-LSTM component for secondary analysis, effectively distributing the computational load without compromising response time.

5.5. Qualitative evaluation and clinical interpretability

Beyond quantitative metrics, clinician feedback was obtained through simulation-based review sessions. The hybrid model’s interpretability was positively rated, as its decision logs included both morphological feature traces (e.g., R–R interval patterns) and attention-weight visualizations from the learned embedding layer. This dual representation enabled clinicians to verify whether automated classifications aligned with observable physiological trends, a factor strongly correlated with user trust in AI-based decision systems [24][25].

The adaptive calibration loop also provided visual confidence indicators, color-coded in the user interface to signal data quality and model certainty. Clinicians reported that these indicators improved situational awareness and supported manual verification during uncertain detections—confirming the importance of explainable adaptation in digital health monitoring.

5.6. Discussion of findings

Collectively, the results in Tables 1–4 demonstrate that the adaptive hybrid framework offers a significant advancement over traditional and deep-learning-only systems. Its superior diagnostic accuracy, low false-alarm rate, and energy efficiency confirm that the fusion of domain-specific and learned features provides a robust and interpretable solution for continuous physiological monitoring.

The system’s incremental learning mechanism proved crucial for maintaining long-term performance in dynamic conditions. This aligns with prior work on adaptive clinical AI [23][24], but extends the concept to the domain of streaming biomedical signals, enabling models to self-correct under drift while maintaining real-time responsiveness. Additionally, the framework achieved clinically acceptable latency ($\leq$ 50 ms), supporting real-time feedback in telemedicine applications where immediate alerts can be life-saving.

However, despite its advantages, some limitations were identified. Adaptive retraining depends on occasional clinician validation, potentially reducing automation scalability. Additionally, while the proposed system performed robustly across ECG, PPG, and EEG signals, broader multimodal integration—such as respiration, temperature, or motion features—could further enhance diagnostic reliability and context awareness.

5.7. Implications for digital health systems

The adaptive hybrid model demonstrates strong potential for use in mobile health (mHealth) and telemedicine systems. Its resource efficiency and adaptive feedback loop make it feasible for edge deployment in wearable devices, remote patient-monitoring networks, and emergency diagnostics. The balance between interpretability and adaptability offers a practical path toward clinically certified AI systems that comply with healthcare safety standards and privacy regulations [25].

Future research should focus on extending the model's adaptability through federated learning frameworks, allowing distributed personalization across multiple patients without compromising data privacy. Integrating multimodal biosignals and evaluating longitudinal patient outcomes will further validate clinical effectiveness and generalizability.

6. Conclusion

This study introduced an adaptive hybrid signal-processing framework designed to enhance the accuracy, reliability, and adaptability of real-time biomedical monitoring systems. The framework integrates traditional signal-processing methods with intelligent machine learning algorithms, forming a robust pipeline that addresses noise, variability, and interpretability challenges common in continuous physiological monitoring. The proposed approach directly responded to the four primary research objectives: (1) to achieve robust artifact suppression for streaming biomedical signals, (2) to extract physiologically meaningful and discriminative features, (3) to ensure accurate, low-latency classification suitable for on-device inference, and (4) to enable continuous adaptation to distributional shifts while preserving interpretability and safety.

The results, summarized in Tables 1–4, validate the effectiveness of the proposed framework. As shown in Table 1, the adaptive hybrid model achieved the highest overall diagnostic accuracy (96.7%), outperforming both classical and deep-learning-only baselines while maintaining a low false-alarm rate of 5.1%. These improvements demonstrate the framework's ability to integrate domain knowledge with data-driven learning to deliver superior predictive performance in real-world conditions. The dataset-level results in Table 2 further confirm consistent performance across diverse modalities—ECG, PPG, and EEG—indicating strong generalizability to different physiological domains and sensor configurations.

A key contribution of this work is the adaptive feedback mechanism, which sustained model performance under changing conditions, such as noise, sensor drift, and patient variability. As reported in Table 3, the adaptive model consistently improved accuracy across successive learning cycles, achieving a 12.2% improvement during the real-time wearable test phase compared to a static model. This finding highlights the value of incremental, clinician-guided updates in mitigating distributional shifts and ensuring long-term stability. The adaptive loop's ability to detect drift and initiate self-correction supports the concept of "living models" in digital health—systems that can learn continuously without complete retraining.

Equally significant are the framework's efficiency and feasibility outcomes, detailed in Table 4, where the adaptive hybrid system demonstrated an optimal balance between accuracy and computational cost. With an average latency of 41 ms and energy consumption of only 6.3 mJ per inference, the framework achieved true edge compatibility, a critical factor for wearable and telemedicine deployment. The hybrid design's modular architecture also ensures scalability, enabling integration into resource-limited embedded systems without sacrificing analytical precision.

Beyond quantitative performance, the hybrid approach offered enhanced clinical interpretability through a dual-feature representation that combined morphological and spectral descriptors with learned embeddings. This transparency allows clinicians to correlate algorithmic outputs with recognizable physiological patterns, improving trust and regulatory compliance. Qualitative feedback from clinician evaluations reinforced that adaptive confidence indicators and visualization tools enhanced situational awareness and decision-making reliability.

Collectively, these findings demonstrate that the adaptive hybrid signal-processing framework not only achieves the desired performance metrics but also satisfies broader clinical and practical criteria for real-world implementation. By addressing limitations in noise robustness, energy efficiency, and interpretability, the proposed approach advances the state of the art in intelligent patient monitoring.

However, the study also identified important areas for refinement. The dependency on clinician feedback for supervised updates may constrain full automation, and the datasets used, while diverse, may not fully capture extreme variations encountered in long-term, uncontrolled environments. Future work should explore federated and privacy-preserving learning paradigms that enable safe, decentralized adaptation across multiple patient devices without centralizing sensitive data. Additionally, expanding the framework to incorporate multimodal biosignals—such as respiration rate, temperature, and movement—would further strengthen diagnostic coverage and personalization capabilities.

In summary, this research provides empirical evidence that adaptive hybrid architectures represent a viable path forward for intelligent, energy-efficient, and explainable digital health systems. By bridging classical signal analysis with adaptive computational intelligence, the framework fulfills the fundamental goal of delivering clinically reliable, real-time patient monitoring adaptable to the diverse and dynamic nature of human physiology. The promising results lay the foundation for the next generation of autonomous, clinician-assisted health-monitoring platforms, promoting safer, more personalized, and data-driven care in both hospital and remote settings.

References

[1] X. Wang, H. Yu, S. Kold, O. Rahbek, and S. Bai, "Wearable sensors for activity monitoring and motion control: A review," Biomimetic Intelligence and Robotics, vol.3, no.1, pp.100089, (2023). DOI:10.1016/j.birob.2023.100089

[2] F. Sana, E. M. Isselbacher, J. P. Singh, E. K. Heist, B. Pathik, and A. A. Armoundas, "Wearable devices for ambulatory cardiac monitoring: JACC state-of-the-art review," Journal of the American College of Cardiology, vol.75, no.13, pp.1582, (2020). DOI:10.1016/j.jacc.2020.01.046

[3] J. Zhang, Y. Guo, X. Dong, T. Wang, J. Wang, X. Ma, and H. Wang, "Opportunities and challenges of noise interference suppression algorithms for dynamic ECG signals in wearable devices: A review," Measurement, vol.250, pp.117067, (2025). DOI:10.1016/j.measurement.2025.117067

[4] S. Anand, V. Sharma, R. Pourush, and S. Jaiswal, "A comprehensive survey on the biomedical signal processing methods for the detection of COVID-19," Annals of Medicine and Surgery, vol.76, pp.103519, (2022). DOI:10.1016/j.amsu.2022.103519

[5] Z. Huang, S. MacLachlan, L. Yu, L. F. Herbozo Contreras, N. D. Truong, A. H. Ribeiro, and O. Kavehei, “Generalization challenges in electrocardiogram deep learning: Insights from dataset characteristics and attention mechanism,” Future Cardiology, vol.20, no.4, pp.209, (2024). DOI:10.1080/14796678.2024.2354082

[6] L. J. Lopez, D. Millan Mayorga, L. H. Martinez Poveda, A. F. Amaya, and W. Rojas Reales, “Hybrid architectures used in the protection of large healthcare records based on cloud and blockchain integration: A review,” Computers, vol.13, no.6, pp.152, (2024). DOI:10.3390/computers13060152

[7] A. Shoeibi, N. Ghassemi, R. Alizadehsani, M. Rouhani, H. Hosseini-Nejad, A. Khosravi, M. Panahiazar, and S. Nahavandi, “A comprehensive comparison of handcrafted features and convolutional autoencoders for epileptic seizures detection in EEG signals,” Expert Systems with Applications, vol.163, pp.113788, (2020). DOI:10.1016/j.eswa.2020.113788

[8] A. A. Abdellatif, A. Mohamed, C. F. Chiasserini, A. Erbad, and M. Guizani, “Edge computing for energy-efficient smart health systems: Data and application-specific approaches,” Energy Efficiency of Medical Devices and Healthcare Applications, pp.53–67, (2019). DOI:10.1016/B978-0-12-819045-6.00003-0

[9] S. M. Varnosfaderani and M. Forouzanfar, “The role of AI in hospitals and clinics: Transforming healthcare in the 21st century,” Bioengineering, vol.11, no.4, pp.337, (2024). DOI:10.3390/bioengineering11040337

[10] D. S. Rajeswari, H. Shree J., A. L. Simha, M. N. Subhani, S. B. G., and V. Yatnalli, “A review on remote health monitoring sensors and their filtering techniques,” Global Transitions Proceedings, vol.2, no.2, pp.392–401, (2021). DOI:10.1016/j.gltp.2021.08.022

[11] F. Sabry, T. Eltaras, W. Labda, K. Alzoubi, and Q. Malluhi, “Machine learning for healthcare wearable devices: The big picture,” Journal of Healthcare Engineering, vol.2022, pp.4653923, (2022). DOI:10.1155/2022/4653923

[12] A. Rancea, I. Anghel, and T. Cioara, “Edge computing in healthcare: Innovations, opportunities, and challenges,” Future Internet, vol.16, no.9, pp.329, **(2024)**

[13] D. Seok, S. Lee, M. Kim, J. Cho, and C. Kim, “Motion artifact removal techniques for wearable EEG and PPG sensor systems,” Frontiers in Electronics, vol.2, pp.685513, (2021). DOI:10.3389/felec.2021.685513

[14] X. An, “Comparison of motion artefact reduction methods and the implementation of adaptive motion artefact reduction in wearable electrocardiogram monitoring,” Sensors, vol.20, no.5, pp.1468, **(2019)**

[15] X. Zhang, M. Jiang, W. Wu, and V. H. C. de Albuquerque, “Hybrid feature fusion for classification optimization of short ECG segment in IoT based intelligent healthcare system,” Neural Computing and Applications, vol.35, no.31, pp.22823–22837, (2023). DOI:10.1007/s00521-021-06693-1

[16] A. Bizzego, G. Gabrieli, and G. Esposito, “Deep neural networks and transfer learning on a multivariate physiological signal dataset,” Bioengineering, vol.8, no.3, pp.35, (2021). DOI:10.3390/bioengineering8030035

[17] K. Zou, Z. Chen, X. Yuan, X. Shen, M. Wang, and H. Fu, “A review of uncertainty estimation and its application in medical imaging,” Meta-Radiology, vol.1, no.1, pp.100003, (2023). DOI:10.1016/j.metrad.2023.100003

[18] S. Pati, S. Kumar, A. Varma, B. Edwards, C. Lu, L. Qu, J. J. Wang, A. Lakshminarayanan, S. Wang, M. J. Sheller, K. Chang, P. Singh, D. L. Rubin, J. Kalpathy-Cramer, and S. Bakas, “Privacy preservation for federated learning in health care,” Patterns, vol.5, no.7, pp.100974, **(2024)**

[19] J. T. Rosenthal, A. Beecy, and M. R. Sabuncu, “Rethinking clinical trials for medical AI with dynamic deployments of adaptive systems,” NPJ Digital Medicine, vol.8, pp.252, (2025). DOI:10.1038/s41746-025-01674-3

[20] H. Javed, S. El-Sappagh, and T. Abuhmed, “Robustness in deep learning models for medical diagnostics: Security and adversarial challenges towards robust AI applications,” Artificial Intelligence Review, vol.58, no.12, (2025). DOI:10.1007/s10462-024-11005-9

[21] F. A. Castaño Usuga, C. Gissel, and A. M. Hernández, “Motion artifact reduction in electrocardiogram signals through a redundant denoising independent component analysis method for wearable health care

monitoring systems: Algorithm development and validation," JMIR Medical Informatics, vol.10, no.11, pp.e40826, **(2022)**

[22] J. Zhang, S. Zheng, W. Chen, G. Du, Q. Fu, and H. Jiang, "A scheme combining feature fusion and hybrid deep learning models for epileptic seizure detection and prediction," Scientific Reports, vol.14, no.1, pp.1–16, (2024). DOI:10.1038/s41598-024-67855-4

[23] L. Lämmermann, P. Hofmann, and N. Urbach, "Managing artificial intelligence applications in healthcare: Promoting information processing among stakeholders," International Journal of Information Management, vol.75, pp.102728, (2024). DOI:10.1016/j.ijinfomgt.2023.102728

[24] N. C. Spies, C. W. Farnsworth, S. Wheeler, and C. R. McCudden, "Validating, implementing, and monitoring machine learning solutions in the clinical laboratory safely and effectively," Clinical Chemistry, vol.70, no.11, pp.1334–1343, (2024). DOI:10.1093/clinchem/hvae126

[25] M. Ennab and H. Mcheick, "Enhancing interpretability and accuracy of AI models in healthcare: A comprehensive review on challenges and future directions," Frontiers in Robotics and AI, vol.11, pp.1444763, (2024). DOI:10.3389/frobt.2024.1444763