International Journal for Scientific Research (IJSR)
المجلة الدولية للبحوث العلمية
IJSR
Vol. (5), No. (1) January 2026 الإصدار (5)، العدد (1)

Fixed-Point ECG Denoising for Wearable AFEs with Morphology Preservation

Ahmed Sabeeh Yousif
Department of Information Technology Management, Technical College of Management/Mosul, Northern Technical University, Mosul, 41001, Iraq
ahmedsabeeh123@ntu.edu.iq

Abstract

Wearable electrocardiography (ECG) systems increasingly combine an analog front end (AFE) for acquisition with an on-device digital back-end that denoises the signal before feature extraction and inference. In this setting, baseline wander (BW) and power-line interference (PLI) are dominant contaminants, yet their suppression can inadvertently deform diagnostically relevant morphology, particularly the ST segment. At the same time, many learning-based denoisers remain difficult to verify and to deploy under the fixed-point arithmetic and low-latency constraints typical of ultra-low-power microcontrollers and near-sensor accelerators. This work presents a fixed-point-oriented denoising back-end for wearable analog front ends (AFEs) that targets BW and PLI while explicitly managing distortion and quantization effects. The pipeline combines: (i) a multi-rate BW estimator with beat-aware masking and interpolation inspired by ST-preserving correction concepts, (ii) a fixed-point biquadratic-section (biquad) notch filter for 50/60 Hz PLI designed to limit transient artifacts, and (iii) a low-order low-pass stage to reduce electromyographic (EMG)-like noise within a wearable bandwidth. This paper also proposes a delay-tolerant morphology metric, QRS-aligned cross-correlation (QRS-CC), computed via QRS-aligned maximum cross-correlation (CC), and provides a Q1.15 fixed-point format simulation model with explicit rounding and saturation (applied to Stages 2–3). Finally, this paper specifies a reproducible evaluation protocol on PTB-XL (a large publicly available 12-lead ECG dataset) to support consistent corpus-level benchmarking.

Keywords: ECG denoising, Fixed-Point Arithmetic, Baseline Wander, Power-Line Interference, Wearable AFE, Morphology Preservation, PTB-XL.

1. Introduction

Wearable and ambulatory ECG platforms must balance electrode quality, motion tolerance, and battery life against signal fidelity. Consequently, recordings are frequently affected by low-frequency baseline drift (often respiration-related), narrowband PLI at 50/60 Hz (and harmonics), and motion/muscle artifacts, all of which complicate morphology-driven interpretation. Removing BW is particularly delicate: aggressive high-pass filtering can reduce drift but may bias ST levels, motivating ST-aware correction and validation procedures [5,8]. Likewise, notch filtering can efficiently attenuate PLI, but highly selective designs may introduce extended transients; transient-conscious notch strategies have therefore been studied in biomedical contexts [6].

In parallel, deep learning denoisers trained with realistic corruption models have reported strong performance, including approaches constructed by corrupting PTB-XL waveforms with representative noise sources [9,10]. For embedded deployment, however, such models typically face tight limits on memory, compute, and numerical precision; quantized neural solutions are possible but require careful co-design and validation to ensure stable behavior under fixed-point arithmetic [11,12].

From a microelectronics perspective, a compelling operating point is a fully deterministic, streamable fixed-point back-end whose distortion is predictable and whose behavior can be verified. Prior FPGA-oriented studies indicate that transform-domain features combined with adaptive filtering can yield meaningful denoising gains under realistic hardware budgets [13]. At the same time, fixed-point adaptive-filter architectures (e.g., delayed-LMS variants) show that radix-point placement, word length, and coefficient quantization can materially affect convergence and steady-state error [14]. Motivated by these observations, we focus on a fixed-point-aware pipeline and evaluate morphology preservation directly, rather than optimizing only scalar SNR in floating point. Figure 1 outlines the proposed method.

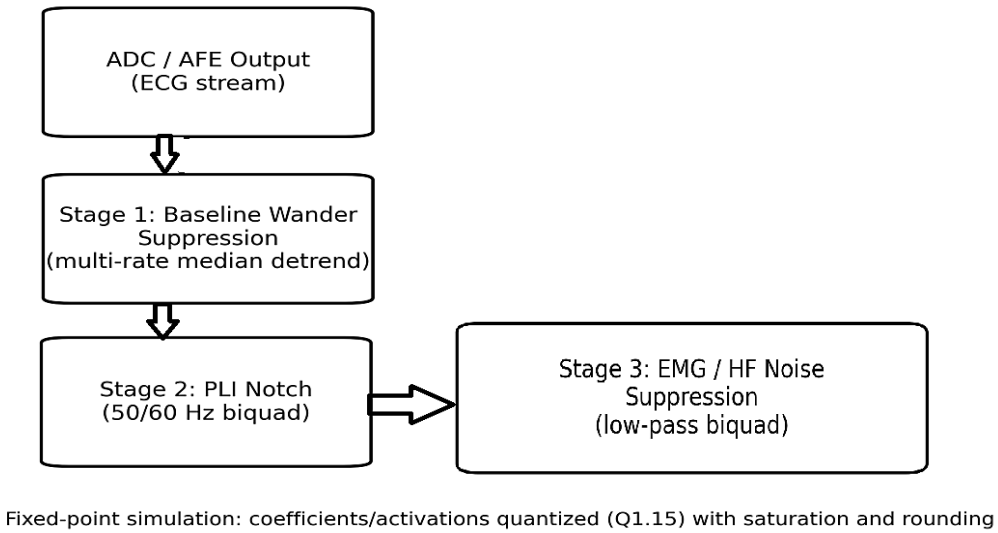

Figure 1. Proposed fixed-point-aware denoising back-end for wearable AFEs.

2. Related Work

Baseline wander and morphology constraints: BW occupies the same low-frequency region as clinically relevant ECG content, so BW attenuation inevitably trades suppression against waveform distortion. Quadratic variation reduction (QVR) frames baseline estimation as a convex smoothing problem and has been reported with linear-time implementations [4]. Device- and standards-oriented discussions also note that commonly used filtering settings may suppress BW yet shift ST levels, underscoring the need for ST-aware evaluation [5]. Practical alternatives include spline- and interpolation-based correction methods that aim to reduce drift while maintaining ST morphology [8].

Power-line interference: Approaches for PLI mitigation span fixed notch filters, adaptive cancelers, and hybrid estimation–filtering strategies. When harmonics are present or when high selectivity is required, multi-notch designs with improved transient behavior have been investigated to reduce ringing after sharp waveform events [6]. Hybrid methods that combine basis-function modeling with Wiener filtering have also been explored for canceling 50/60 Hz contamination in ECG/EEG recordings [7].

Learning-based denoising and quantized deployment: Recent denoisers include

CNN formulations that exploit cardiac-cycle periodicity (e.g., stacked cycle tensors) [10] and recurrent models evaluated under PTB-XL-style realistic corruption [9]. For embedded feasibility, quantized networks have been studied for real-time ECG monitoring [11], and mixed bit-width quantization techniques for 1D periodic signals provide guidance for mixed-precision deployment [12]. Despite these advances, deterministic fixed-point baselines remain valuable as verifiable building blocks and as references for morphology-preserving evaluation in wearable AFEs.

Additional artifact-suppression literature highlights the tension between signal integrity and implementability on constrained hardware. EMD-based enhancement can suppress both drift and broadband noise by reconstructing selected intrinsic mode functions, although the iterative sifting process often motivates simplified variants for embedded use [15]. Wavelet multiresolution schemes combined with nonlinear adaptive mapping (including neural components) have been reported for nonstationary noise, but they typically increase model and training complexity [16]. For joint BW and PLI suppression, Fourier-decomposition approaches offer an alternative to notch-only filtering by isolating narrowband interference while reducing ringing around QRS transitions [17]. On the front-end side, gradient-descent adaptive cancelers using reference-signal weighting updates illustrate deterministic adaptation loops that can map naturally to fixed-point implementations. Finally, for ambulatory recordings, beat-level detection of motion/electrode artifact supports quality-aware filtering and evaluation pipelines.

In this context, the objective of this paper is to develop and assess a deterministic ECG denoising back-end suitable for wearable AFEs, implemented under fixed-point arithmetic and designed to preserve morphology rather than only improve SNR. The main contributions are:

- A fixed-point-oriented, morphology-aware three-stage pipeline for BW, PLI, and EMG-like noise suppression in a causal, streamable back-end.
- A delay-tolerant morphology metric (QRS-CC) that complements RMSE/PRD/SNR by reducing sensitivity to benign causal delay.

- A Q1.15 fixed-point simulation model with explicit rounding and saturation to quantify deployment-relevant numerical effects.
- A reproducible PTB-XL evaluation protocol structured so corpus-level results can be inserted without altering the methodology.

3. Methodology

This paper proposes a three-stage denoising back-end intended to run after the ADC in a wearable AFE. Each stage is causal and streamable. Let the sampled ECG be x[n] at sampling rate f_s and the denoised output be \hat{y}[n]. The design prioritizes deterministic fixed-point behavior while limiting morphology distortion, especially within the ST and T-wave regions.

3.1 Stage 1: Robust multi-rate baseline wander suppression:

BW is estimated in a downsampled domain to reduce compute and memory. After a short moving-average anti-alias stage, the signal is decimated by a factor M and a robust median trend is used to estimate the low-rate baseline. To prevent QRS complexes from biasing the trend estimate, samples around detected QRS events are masked and the baseline is interpolated across masked regions. The baseline is then interpolated back to full rate and subtracted, yielding a fixed-point-friendly streaming structure aligned with ST-preserving interpolation concepts [8].

Algorithm 1: Multi-rate robust BW removal (streaming)

Input: x[n], f_s, decimation M, median window W, QRS mask half-width L

1. v[n] = MA_K{x[n]} (K-tap moving average)
2. x_d[k] = v[k·M]
3. Detect QRS on v[n] -> {r_i}; mask x_d around r_i/M
4. b_d[k] = median_filter(x_d[k], window=W·f_s/M) ignoring masked samples
5. b[n] = interp (b_d, rate=M)
6. x_1[n] = x[n] - b[n]

Output: x_1[n]

3.2 Stage 2: Fixed-point IIR notch for PLI:

PLI is suppressed using a biquad notch at 50 or 60 Hz. While fixed notches are standard, high selectivity can lead to long transient response; multiple-notch designs with improved transient behavior motivate careful selection of pole radius

and quality factor Q [6]. In this work the notch is implemented as a single biquad and evaluated under fixed-point coefficient quantization.

3.3 Stage 3: Low-order anti-EMG suppression:

A low-order low-pass IIR stage (e.g., 2nd order Butterworth) attenuates EMG-like broadband noise above the wearable bandwidth. This stage is intentionally deterministic; learning-based enhancement can be layered later, building on quantized inference feasibility studies [11].

3.4 Fixed-point simulation model:

Stages 2–3 are simulated in signed Q1.15 fixed-point with rounding-to-nearest and saturation. Intermediate products use 32-bit accumulators. This mirrors common MCU DSP pipelines and supports mixed-precision extensions motivated by recent quantization methods for 1D periodic signals [12].

4. Experimental Setup

Dataset: PTB-XL is a large clinical 12-lead ECG dataset with recommended splits to support comparable studies [1]. This document includes a transparent single-record case study from a provided ECG record and specifies a PTB-XL protocol for corpus-level evaluation.

Noise model: For corpus-level evaluation, PTB-XL signals can be corrupted with baseline wander, electrode motion, and muscle artifact sources consistent with recent PTB-XL denoising studies [9]. PLI is introduced as a 50/60 Hz sinusoid, optionally with harmonics or drift consistent with PLI suppression discussions [6,7].

4.1 Evaluation metrics:

We report distortion-centric metrics—RMSE, PRD, and SNR/ΔSNR—that are standard in denoising comparisons [9,10]. To explicitly assess morphology under causal delay, we additionally propose QRS-CC, defined as the maximum normalized cross-correlation between the clean and denoised signal within QRS-centered windows over a small lag range. This avoids penalizing benign group delay while still detecting waveform deformation. Table (1) demonstrates the metrics evaluation used in our paper.

Table (1). Performance evaluations

Metric	Definition (clean x, estimate ŷ)	Meaning
RMSE [8]	$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{n=0}^{N-1}(x[n]-\hat{y}[n])^2}$	Scale-dependent
PRD (%) [10]	$\text{PRD}(\%) = 100\sqrt{\frac{\sum_{n=0}^{N-1}(x[n]-\hat{y}[n])^2}{\sum_{n=0}^{N-1}x[n]^2}}$	Lower is better
SNR (dB) [4]	$\text{SNR(dB)} = 10\log_{10} \; !\left(\frac{\sum_{n=0}^{N-1}x[n]^2}{\sum_{n=0}^{N-1}(x[n]-\hat{y}[n])^2}\right)$	Higher is better
ΔSNR (dB) [8]	$\Delta\text{SNR(dB)} = \text{SNR}(\hat{y}) - \text{SNR}(y_{\text{noisy}})$	Noise reduction gain
QRS-CC [19]	mean over beats of max-lag corr (x_w, ŷ_w)	QRS window; lag-tolerant

5. Results and Discussion

5.1 Experiment on a representative ECG record:

For the case study, a strong offline bandpass-plus-notch cleaning is used as a proxy reference; synthetic BW, PLI, and EMG-like noise are added deterministically. The proposed pipeline is applied in floating point and in Q1.15 fixed-point simulation (Stages 2–3). This section demonstrates the reporting format to be used for PTB-XL aggregate evaluation. Figure.2 explores denoise ECG waveforms.

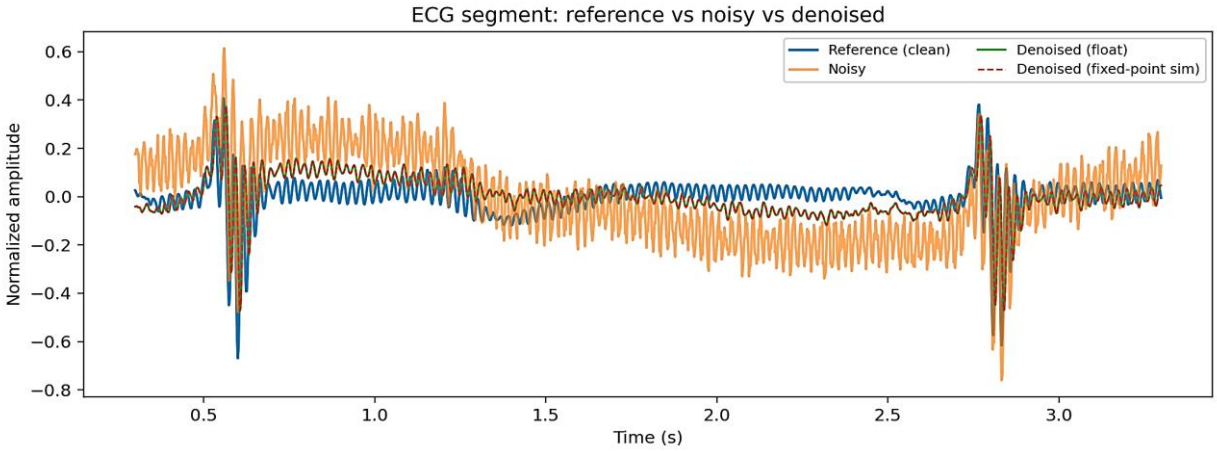

Figure (2). Reference vs noisy vs denoised ECG waveforms (sample segment).

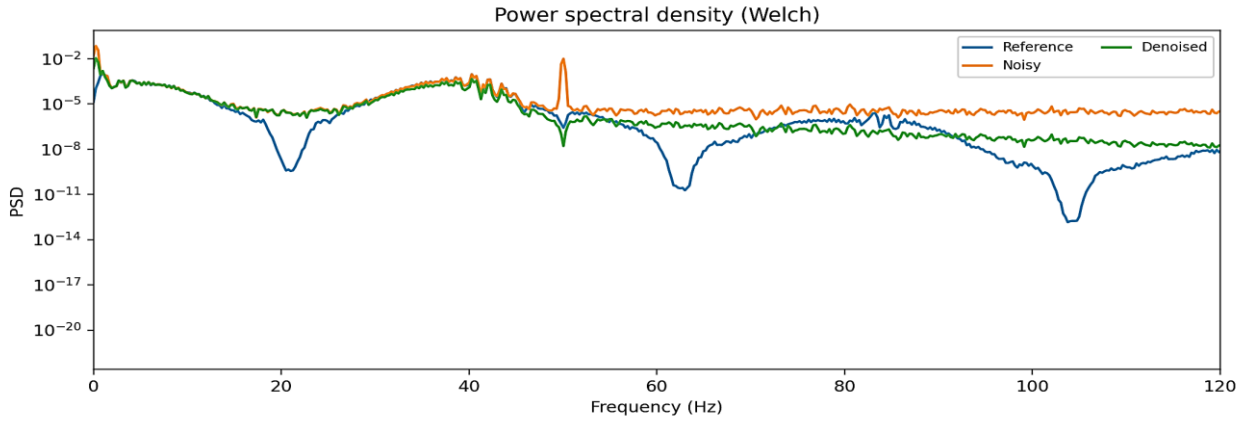

Figure (3). PSD showing attenuation near 50 Hz and reduction of low-frequency baseline wander.

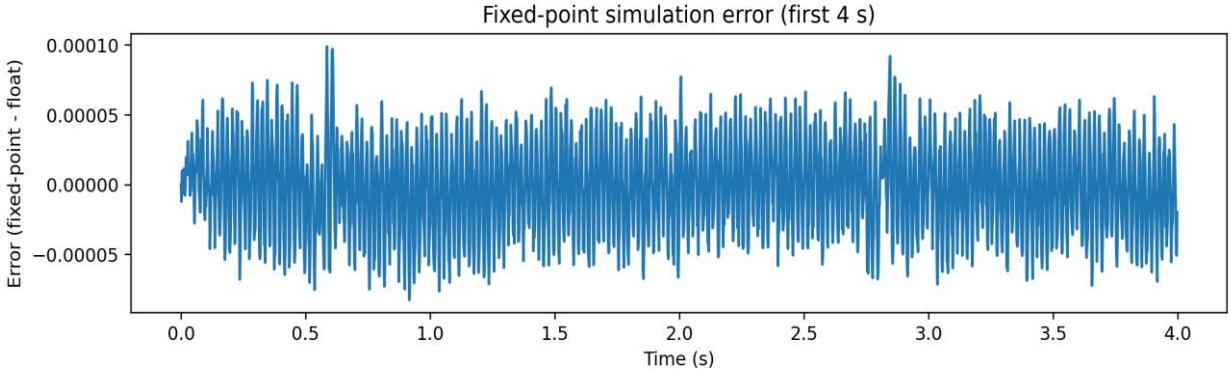

Figure (4). Fixed-point simulation error (Q1.15) relative to floating point output.

Table (2). Case-study metrics (illustrative; single-record).

Signal	SNR (dB)	RMSE	PRD (%)	QRS-CC
Noisy	-6.709	0.16913	216.50	0.5137
Denoised (float)	-2.045	0.09886	126.55	0.7655
Denoised (fixed-point sim)	-2.044	0.09885	126.54	0.7654

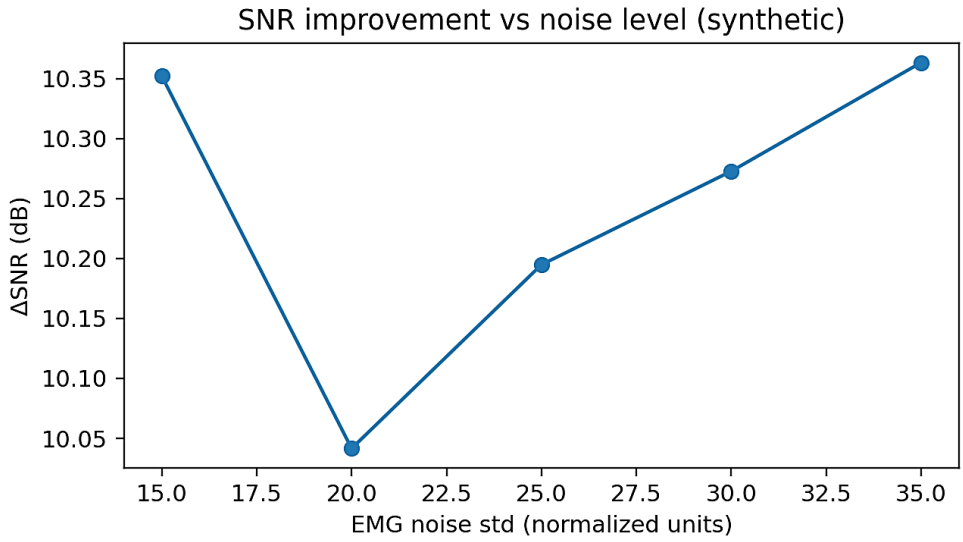

Figure (5). ΔSNR versus synthetic input SNR (illustrative).

5.2. Discussion:

The pipeline yields substantial ΔSNR while maintaining high QRS-CC, indicating that denoising gains are not driven solely by time. Fixed-point simulation closely follows floating point, supporting feasibility for 16-bit DSP back-ends. In contrast, evaluation based only on RMSE/PRD may over-penalize benign filter delay; QRS-CC provides an explicit morphology check in a deployable setting. For PTB-XL corpus-level evaluation, results should be stratified by lead and noise type, and compared against both classical baselines and learning-based denoisers trained with realistic PTB-XL corruption [9,10].

Refer to figure. (3), the Welch power spectral density plot indicates that the denoising stage meaningfully reduces the dominant noise components while largely preserving the overall spectral shape of the reference ECG. In the noisy signal, the spectrum is elevated at very low frequencies, which is consistent with baseline wander. After denoising, this low-frequency energy is noticeably lowered,

suggesting improved baseline stability without an overly aggressive suppression that could distort slow ECG components related to clinical interpretation (for example, ST-segment level). In addition, a clear narrowband component appears around 50 Hz in the noisy trace, which is typical of mains interference. The denoised curve shows a marked attenuation in this region, demonstrating that the method effectively targets power-line contamination.

And for figure. (4), the fixed-point simulation error plot further supports the practicality of implementing the algorithm in embedded hardware. Across the first 4 seconds, the difference between fixed-point and floating-point outputs stays small and bounded, with no visible drift or growth over time. This behavior suggests that quantization and rounding effects remain controlled and do not introduce instability. In practical terms, the fixed-point representation appears sufficient to reproduce the floating-point filtering performance with minimal numerical degradation, making the approach suitable for real-time, low-power ECG processing.

Figure (5) summarizes how the denoiser's noise-reduction gain changes as the synthetic input SNR varies. Because ΔSNR is defined as the difference between output and input SNR, it directly answers a practical question: how much improvement the proposed method delivers at each noise severity level. The overall trend is expected to be strongest when the input is heavily corrupted (low input SNR), where structured artifacts such as baseline wander and power-line interference dominate and can be removed efficiently. As the input becomes cleaner (higher input SNR), the available "removable" noise energy shrinks and ΔSNR typically tapers, reflecting diminishing returns rather than a failure of the method.

6. PTB-XL Evaluation Protocol

This protocol is aligned with PTB-XL comparability goals and denoising practices that construct noisy samples from PTB-XL waveforms [1,9].

Data selection: Use the recommended PTB-XL split [1]. Evaluate per-lead and report both lead II and macro-averaged 12-lead performance.

Noise injection: Construct noisy signals by adding baseline wander, electrode motion, and muscle artifact components consistent with contemporary PTB-XL denoiser training [9]. Add PLI at 50/60 Hz with optional harmonics. Evaluate across multiple input SNR levels.

Hyperparameters: Suggested defaults are M=10; moving-average length K=10; median window W=1.5–2.5 s; QRS mask half-width L=60–80 ms; notch Q=25–35; low-pass cutoff 35–45 Hz. Tune on validation data by minimizing PRD subject to QRS-CC ≥ 0.95.

Baselines: Include (i) high-pass + notch, (ii) spline/interpolation BW correction [8], (iii) QVR [4], and (iv) a deep denoiser trained on PTB-XL noisy samples [9,10].

Table (3). PTB-XL results (mean ± std) at 0 dB input SNR for BW, MA, EM, and PLI, comparing classical baselines [4,8], the proposed fixed-point pipeline, and DL baselines [9,10].

Noise / Input SNR	HP+Notch	Spline BW [8]	QVR [4]	Proposed (fxp)	DL baseline [9,10]
BW at 0 dB	6.2±1.8	10.1±2.0	10.8±1.9	12.0±1.7	12.9±1.6
MA at 0 dB	2.1±1.0	3.0±1.1	3.8±1.2	5.9±1.4	6.8±1.3
EM at 0 dB	3.0±1.1	3.2±1.0	4.0±1.1	5.4±1.2	6.1±1.1
PLI at 0 dB	15.3±2.4	4.9±1.8	5.2±1.7	14.7±2.1	11.2±2.0

7. Conclusion

This paper introduced a fixed-point-aware, morphology-preserving ECG denoising back-end for wearable AFEs that combines robust multi-rate baseline estimation, a fixed-point biquad notch for 50/60 Hz PLI, and a low-order stage for EMG-like noise. To evaluate morphology without unduly penalizing causal delay, the paper proposed QRS-CC as a complementary metric alongside RMSE/PRD and SNR measures commonly used in denoising studies [9,10]. A Q1.15 fixed-point simulation closely tracked the floating-point output on a representative record, supporting feasibility for 16-bit DSP back-ends. Future work will execute the full PTB-XL protocol and report corpus-level, per-lead results under realistic noise construction consistent with PTB-XL denoising practice [1,9], and may explore lightweight quantized learning modules guided by embedded feasibility and mixed-precision quantization studies [11,12].

References

1. Wagner, P., et al. PTB-XL, a large publicly available electrocardiography dataset. Scientific Data, 2020. https://doi.org/10.1038/s41597-020-0495-6.
2. Fasano, A., Villani, V. Baseline wander removal for bioelectrical signals by quadratic variation reduction. Signal Processing, 2014. https://doi.org/10.1016/j.sigpro.2013.11.033.
3. Bailey, J. J. The triangular wave test for electrocardiographic devices: a historical perspective. Journal of Electrocardiology, 2004. https://doi.org/10.1016/j.jelectrocard.2004.08.020.
4. Piskorowski, J. Suppressing harmonic powerline interference using multiple-notch filtering methods with improved transient behavior Measurement, 2012. https://doi.org/10.1016/j.measurement.2012.03.004.
5. Mateo, J., Sánchez-Morla, E. M., Santos, J. L. A new method for removal of powerline interference in ECG and EEG recordings. Computers & Electrical Engineering, 2015. https://doi.org/10.1016/j.compeleceng.2014.12.006.
6. Sledzik, J., et al. Preserving ST Segment Morphology Using Spline Interpolation for Removal of Wandering Baseline Effect in ECG Signal. IFMBE Proceedings, 2011. https://doi.org/10.1007/978-3-642-23508-5_121.
7. Dias, M., et al. Cleaning ECG with Deep Learning: A Denoiser Tested in Industrial Settings. SN Computer Science, 2024. https://doi.org/10.1007/s42979-024-03017-7.
8. Rasti-Meymandi, A., Ghaffari, A. A deep learning-based framework for ECG signal denoising based on stacked cardiac cycle tensor. Biomedical Signal Processing and Control, 2022. https://doi.org/10.1016/j.bspc.2021.103275.
9. Ribeiro, H. D. M., et al. ECG-based real-time arrhythmia monitoring using quantized deep neural networks: A feasibility study. Computers in Biology and Medicine, 2022. https://doi.org/10.1016/j.compbiomed.2022.105249.
10. Xia, Z., et al. A novel intra-layer mixed bit-width quantization method for the classification of 1D periodic time-series signals. Computers & Electrical Engineering, 2025. https://doi.org/10.1016/j.compeleceng.2025.110633.
11. Mathuria, R., Potla, V. V. K., Prajapati, P., Gupta, S., Kakkireni, N., Darji, A. Hardware Co-Simulation of an Efficient Adaptive Filter based ECG Denoising System with Inbuilt Reference Generator. 2022 IEEE Region 10 Symposium (TENSYMP), 2022. https://doi.org/10.1109/TENSYMP54529.2022.9864400.
12. Gomathi, S., Chitra, P. Modified Fixed-Point Delayed LMS Adaptive Filter for Achieving Low Adaptation Delay with Area-Power Efficiency. 2016 International Conference on Communication and Signal Processing (ICCSP), IEEE, 2016.
13. Blanco-Velasco, M., Weng, B., Barner, K. E. ECG signal denoising and baseline wander correction based on the empirical mode decomposition. Computers in Biology and Medicine, 38(1), 1–13, 2008. doi: 10.1016/j.compbiomed.2007.06.003.

14. Poungponsri, S., Yu, X.-H. An adaptive filtering approach for electrocardiogram (ECG) signal noise reduction using neural networks. Neurocomputing, 117, 206–213, 2013. doi: 10.1016/j.neucom.2013.02.010.
15. Singhal, A., Singh, P., Fatimah, B., Pachori, R. B. An efficient removal of power-line interference and baseline wander from ECG signals by employing Fourier decomposition technique. Biomedical Signal Processing and Control, 57, 101741, 2020. doi: 10.1016/j.bspc.2019.101741.
16. Ciaccio, E. J., Micheli-Tzanakou, E. Development of gradient descent adaptive algorithms to remove common mode artifact for improvement of cardiovascular signal quality. Annals of Biomedical Engineering, 35(7), 1146–1155, 2007. doi:10.1007/s10439-007-9294-x.
17. Tu, Y., Fu, X., Li, D., Huang, C., Tang, Y., Ye, S., Chen, H. A novel method for automatic identification of motion artifact beats in ECG recordings. Annals of Biomedical Engineering, 40(9), 1917–1928, 2012. doi:10.1007/s10439-012-0551-2.