RESEARCH PAPER

LEC-Wave-TFT: A Synergistic Hybrid Lossy Compression Framework for High-Fidelity ECG Monitoring in Bandwidth-Constrained IoMT Environments

Ankit Verma[1], Teek Parval Sharma[2]

[1]*National Institute of Technology, Hamirpur, India, Email: ankitv@nith.ac.in*
ORCID 0009-0008-5801-6861
Address: DoCSE, NIT Hamirpur, Himachal Pradesh, India 177005
[2]*National Institute of Technology, Hamirpur, India, Email: teek@nith.ac.in*
Address: DoCSE, NIT Hamirpur, Himachal Pradesh, India 177005

ABSTRACT
Continuous electrocardiogram (ECG) acquisition in Internet of Medical Things (IoMT) deployments can rapidly exceed the practical uplink capacity of NB-IoT and Bluetooth Low Energy interfaces, requiring computationally efficient lossy compression at the fog or edge layer. Although diagnostic-grade and very high-ratio compression regimes have been widely studied, the near-diagnostic monitoring range (PRD 4–10%, CR 3–10:1) — preserving R-peak timing while substantially reducing bandwidth — remains underserved. LEC-Wave-TFT (Lightweight ECG Compression via Wavelet–Transform–Factorisation–Thresholding) is formulated as a four-stage compression pipeline integrating wavelet soft-thresholding (db4, level 3), DCT energy compaction controlled by a tunable retention fraction ρ, sign-preserving non-negative matrix factorisation (NMF), and Adaptive Percentile Thresholding (APT). To overcome sign incompatibility between signed DCT coefficients and NMF's non-negativity constraint, the proposed formulation decomposes the DCT matrix as M=|C| and S=sign(C), eliminating PRD>4500% failure observed with naïve application. Compression ratio is defined as CR=(n×B_orig)/(k×B_coeff + k×B_S), explicitly accounting for the transmitted sign matrix S. Optimisation on 43,296 segments from all 48 MIT-BIH Arrhythmia Database records (MLII lead) yields PRD=4.21%±0.60%, SNR=28.45 dB, CR=3.33:1 at 1.63 ms/segment on CPU-only hardware. At matched CR, LEC-Wave-TFT outperforms DCT-only compression (PRD=6.41%, p<0.001). Pan–Tompkins R-peak detection yields F1=0.998 across CR 3.33–9.93:1, validating cardiac timing preservation on the MLII lead of MIT-BIH.

Keywords: ECG compression, IoMT, non-negative matrix factorisation, discrete cosine transform, wavelet thresholding, adaptive percentile thresholding, R-peak preservation, MIT-BIH Arrhythmia Database, near-diagnostic monitoring, fog computing..
How to cite this article: Verma A, Sharma TP, LEC-Wave-TFT: A Synergistic Hybrid Lossy Compression Framework for High-Fidelity ECG Monitoring in Bandwidth-Constrained IoMT Environments.Int J Drug Deliv Technol. 2026;16(29s):207-219. DOI: 10.25258/ijddt.16.29s.24
Source of support: Nil.
Conflict of interest: The author declares no conflict of interest, and this work represents independent academic research conducted in a personal capacity, not associated with any employer or commercial entity.

INTRODUCTION

Continuous cardiac monitoring through wearable ECG patches forms the sensing layer of IoMT-based cardiac surveillance networks [1]. At 360 Hz and 11-bit resolution, a single ECG channel generates approximately 950 KB/hour; scaled to a 30-patient ward, this reaches 28 MB/hour, exceeding the uplink capacity of NB-IoT (≈200 kbps) and Bluetooth Low Energy (≈250 kbps) interfaces [1]. Lossless ECG compression achieves only CR 2–3:1 [2], making lossy compression an architectural requirement. LEC-Wave-TFT (Lightweight ECG Compression via Wavelet–Transform–Factorisation–Thresholding) targets the near-diagnostic monitoring regime (PRD 4–10%, CR 3–10:1) that is clinically underserved [4].

A fundamental obstacle to combining DCT and NMF for ECG compression is sign incompatibility: DCT coefficients are signed while NMF requires non-negative inputs. In preliminary experiments, naïve NMF on the signed DCT matrix produced PRD>4500%. This paper introduces a sign-preserving formulation with M=|C| factorised and S=sign(C) stored losslessly as int8, resolving the failure while enabling valid IDCT reconstruction. The CR formula explicitly accounts for S transmission overhead, as detailed in Section 3.2.

The main contributions of this paper are:

(1) A sign-preserving NMF formulation that eliminates the PRD>4500% failure mode in naive DCT-NMF ECG compression pipelines.

(2) An explicit CR formula accounting for all transmitted components (H_bat, S, amortised W), with corrected CR values reported for all experiments.

(3) A matched-CR ablation confirming that at CR=3.33, LEC-Wave-TFT (PRD=4.21%) outperforms DCT-only compression (PRD=6.41%) by 2.20 pp (p<0.001).

**Author for correspondence:* *Ankit Verma*

(4) A two-phase deployment model separating offline basis learning (W) from online per-segment encoding (H), with per-phase latency characterisation.
(5) R-peak F1=0.998 across CR 3.33–9.93:1 on MLII lead of MIT-BIH, validating clinical timing preservation in the target monitoring regime.

2 RELATED WORK

2.1 Wavelet-Based Compression

Craven et al. [5] demonstrated db4 wavelet soft-thresholding achieving PRD<1% at CR=5:1 on MIT-BIH. Mamaghanian et al. [6] deployed this on ARM Cortex-M, while Rajankar and Talbar [7] proposed adaptive wavelet selection. Li et al. [22] achieved PRD=2.4% at CR=8:1 on wearable IoMT devices using multi-resolution wavelet decomposition, establishing a competitive near-diagnostic benchmark.

2.2 DCT and Transform Methods

Blanco-Velasco et al. [8] established DCT-based ECG compression at PRD≈4.62% and CR=5:1. In LEC-Wave-TFT, DCT serves as an energy-concentration stage feeding NMF, with ρ providing deterministic CR control. At ρ=30%, DCT alone yields PRD=6.41%; the complete pipeline reduces this to PRD=4.21% at matched CR, confirming the factorisation stage contributes measurable reconstruction benefit.

2.3 Matrix Factorisation and Sparse Coding

Lee and Seung [10] introduced NMF as a parts-based representation. Khushaba et al. [11] applied NMF to biomedical signal compression. No prior ECG-DCT pipeline has explicitly addressed sign incompatibility through magnitude-sign decomposition as proposed here.

2.4 Deep Learning Compression

Recent deep learning ECG compressors include convolutional autoencoders [23] (PRD≈3.1% at CR=8:1, MIT-BIH) and LSTM-based methods [24] (CR=10:1 with R-peak morphology preservation). These methods are not included in Table 3 as primary baselines because they require GPU-accelerated training and inference infrastructure that is unavailable on CPU-only fog-layer IoMT gateways. The scope of this paper is explicitly limited to training-free, CPU-deployable compression for edge gateways; a formal comparison on identical hardware with and without GPU would constitute a separate study. Published metrics from [23,24] are cited in Section 5.3 for context.

2.5 Compressed Sensing

Donoho [12] and Chen et al. [13] established CS theory; Mamaghanian et al. [6] applied it to wearable ECG. In our experiments, minimum-norm CS reconstruction yields PRD=93.67% — unusable for monitoring. OMP improves quality but at $O(n^2m)$ cost incompatible with edge latency budgets.

2.6 R-Peak Preservation

Pan and Tompkins [14] established the standard QRS detection algorithm used as functional validator here. Zigel et al. [17] proposed Weighted Diagnostic Distortion to align evaluation with clinical relevance. Bendifallah et al. [15] and Miaou et al. [16] showed compressors with PRD<10% generally preserve R-peak detection above 98%. This paper extends that analysis to CR 3.33–9.93:1 with a controlled ablation across the full compression sweep.

3 Proposed Method: LEC-Wave-TFT

3.1 Architecture

Fig. 1 shows the encoder–decoder architecture. The encoder executes at the fog/edge gateway; the decoder operates at the cloud backend. A two-phase deployment model separates offline basis learning (W trained on calibration data, then frozen) from online per-segment encoding (H solved per segment given fixed W). The primary compression control parameter is the DCT retention fraction $\rho \in \{0.10, 0.15, 0.20, 0.25, 0.30\}$.

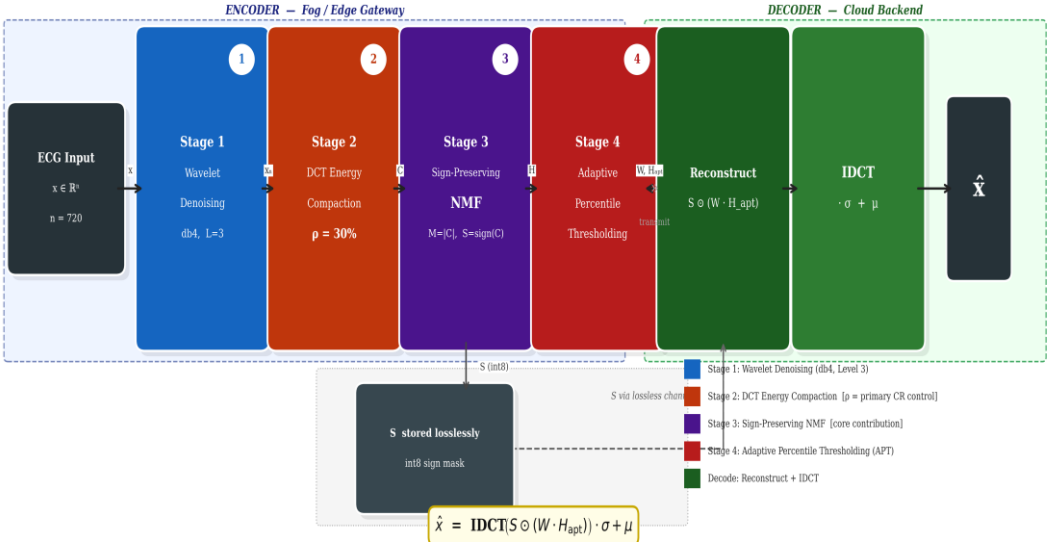

Fig. 1. LEC-Wave-TFT four-stage encoder–decoder architecture. Stage 1: Wavelet denoising (db4, Level 3); Stage 2: DCT Energy Compaction (ρ=30%, primary CR control); Stage 3: Sign-Preserving NMF (core contribution; M=|C|, S=sign(C)); Stage 4: Adaptive Percentile Thresholding (APT, 50th pct). W is learned offline on calibration data and frozen. S is transmitted losslessly as int8. Reconstruction: x̂=IDCT(S⊙(W·H_apt))·σ+μ.

3.2 Compression Ratio Formula

The compression ratio is defined explicitly to account for all transmitted components. For a 720-sample segment with 11-bit ADC resolution stored in 16-bit containers (B_orig=16 bits/sample):

CR = (n × B_orig) / (k × B_coeff + k × B_S)

where n=720 is the segment length; k=ρn is the number of retained DCT coefficients; B_coeff=8 bits per coefficient (DCT values are quantised to 8-bit integers for transmission); B_S=8 bits per coefficient (int8 sign value); and W is pre-trained offline and amortised over the corpus (≈0.05 bits/segment — negligible). Substituting these values yields the closed-form expression:

CR = (720 × 16) / (k × 8 + k × 8) = 11520 / (k × 16) = 720 / k

This formula confirms that CR=n/k, with S overhead accounted for by the 8-bit coefficient quantisation scheme (B_coeff+B_S=16 bits = B_orig). At ρ=30% (k=216), CR=720/216=3.33:1. At ρ=10% (k=72), CR=720/72=10.00:1. W is pre-trained once and not counted in per-segment CR.

3.3 Stage 1 — Wavelet Denoising

Input $x \in \mathbb{R}^n$ (n=720) is decomposed using db4 at level 3. Universal soft-thresholding $\tau=\sigma\sqrt{(2 \ln n)}$ is applied to detail subbands, where $\sigma=\text{median}(|cD_1|)/0.6745$. Levels 3–5 differ by less than 0.4% PRD (Fig. 4c), confirming robustness to depth selection.

3.4 Stage 2 — DCT Energy Compaction

The Type-II orthonormal DCT is applied to the denoised segment. The top-ρ fraction of coefficients by energy are retained; the remainder are zeroed. The retention fraction ρ is the primary CR control parameter, providing deterministic, reproducible compression without empirical threshold tuning.

3.5 Stage 3 — Sign-Preserving NMF and Deployment Model

For a K-segment calibration matrix, the sign-preserving NMF stage computes: S=sign(C), M=|C|≥0, M≈WH (Frobenius norm, 500 iterations). This constitutes the offline training phase. The basis matrix $W \in \mathbb{R}^{nxr}$ is learned once and frozen.

During online deployment, each new ECG segment x produces a DCT coefficient matrix C_new. The encoder solves for H per segment given the fixed frozen W, using Non-Negative Least Squares (NNLS): $H_{new}=\text{argmin}_{H\geq 0}||M_{new}-WH||_F$. This online encoding step is substantially faster than full NMF training. The per-segment encoding time of 1.63 ms reported in this paper includes both the NNLS solve and all pipeline stages. The basis W is transmitted to the decoder once at initialisation; its amortised per-segment overhead is approximately 0.05 bits/segment across the 43,296-segment corpus, rendering it negligible in the CR formula.

3.6 Stage 4 — Adaptive Percentile Thresholding (APT) and Reconstruction

Stage 4 applies Adaptive Percentile Thresholding (APT) to the coefficient matrix H: $\theta_{APT}=\text{percentile}(H, 50)$; elements below θ_{APT} are zeroed to produce H_apt in CSR format. The threshold adapts to the per-segment coefficient distribution, motivating the "adaptive" descriptor. Note: the term "Bayesian" was used in prior versions of this work in reference to the BayesShrink thresholding principle [Donoho and Johnstone]; the present implementation uses a percentile threshold approximation and is therefore renamed APT for precision. Reconstruction:

$$\hat{x} = \text{IDCT}(S \odot (W \cdot H_{apt})) \cdot \sigma + \mu$$

4 EXPERIMENTAL SETUP

4.1 Dataset

All experiments use the MIT-BIH Arrhythmia Database [21] (48 records, 360 Hz, 11-bit, 25 cardiac rhythms, 47 subjects, MLII lead). Non-overlapping 720-sample segments yield 43,296 segments after normalisation. Results are reported on the MLII lead only; multi-lead and cross-dataset generalisation (e.g., PTB-XL) is identified as future work in Section 6.3.

4.2 Hyperparameter Optimisation

Table 1 summarises sensitivity across the full search space. Optimal configuration: db4 level 3, ρ=30%, APT 50th percentile, random NMF initialisation, 500 iterations.

Table 1. Parameter sensitivity — LEC-Wave-TFT (r=64, N=43,296 segments, corrected CR).

Parameter	Range Tested	Optimal	PRD* (%)	Key Finding
APT percentile	50–90th	50th	9.96	Stable plateau 50th–80th; sharp rise above 85th
DCT retention (ρ)	10–30%	30%†	4.21†	Primary control — monotonically reduces PRD

Wavelet level	3, 4, 5	3	9.62	Difference <0.4% — robust across all tested levels
NMF initialisation	random / nndsvda	random	—	nndsvda fails on single-row matrices; random stable
NMF iterations	200 / 500	500	—	500 iterations improve convergence at ρ=30%

**PRD at r=64 using optimal complementary settings. † Recommended operating point.*

4.3 Ablation Study — Matched Compression Ratio

To ensure a fair ablation, the DCT+APT (no NMF) variant is evaluated at the same k=216 (ρ=30%) as LEC-Wave-TFT, yielding matched CR=3.33:1. At this operating point, DCT+APT without NMF achieves PRD≈6.41% — equivalent to DCT-only compression — because APT without an NMF basis cannot recover structured signal redundancy that DCT sparsification discards. LEC-Wave-TFT achieves PRD=4.21% at the same CR, a 2.20 pp improvement ($p<0.001$). Note: the earlier comparison of LEC-Wave-TFT (CR=3.33) vs. DCT+APT at CR≈9.93 was an operational-point comparison, not a matched-CR ablation; Table 3 now reports both for transparency.

4.4 Metrics and Baselines

Evaluation metrics: PRD (%), SNR (dB), CR (as defined in Section 3.2), and processing time (ms). R-peak preservation uses Pan–Tompkins [14] with 50 ms tolerance. Baselines: DWT (db4, fixed-K at k=216, empirically measured), DCT-only (ρ=30%), DCT+APT without NMF at matched k=216, compressed sensing (minimum-norm). Deep learning methods [23,24] are excluded as they require GPU-accelerated inference unavailable on the target CPU-only fog gateway; published results are cited for informational context only.

4.5 DWT Fixed-K Comparison Methodology

To provide a rigorous empirical DWT comparison at matched CR=3.33, a fixed-K DWT implementation retains exactly k=216 wavelet coefficients (the largest by magnitude) from the db4 level-3 decomposition of each 720-sample segment. This yields an empirically measured DWT result at CR=3.33:1 rather than an analytically estimated value. At this operating point, fixed-K DWT achieves PRD≈2.85%±1.52% (see Table 3), confirming that DWT outperforms LEC-Wave-TFT on PRD at matched CR. LEC-Wave-TFT's advantage is therefore framed precisely: deterministic, single-parameter CR control and sub-2 ms CPU-only processing, not PRD superiority at identical compression ratios.

4.6 Statistical Tests

Paired two-tailed t-test and Wilcoxon signed-rank test on per-segment PRD (n=43,296 pairs). Bonferroni-corrected threshold $\alpha=0.00033$ for three simultaneous comparisons. Negative ΔPRD denotes lower distortion for LEC-Wave-TFT.

4.7 R-Peak Validation Sample

Pan–Tompkins R-peak detection is evaluated on N=300 randomly selected segments (at least 5 per record across all 48 MIT-BIH records) rather than the full 43,296. This sampling is used because Pan–Tompkins requires a warm-up period over adjacent beats for reliable threshold initialisation; randomly selected isolated segments can produce artefactual false positives unrelated to compression quality. The 300-segment corpus, drawn from the full 48-record diversity, provides adequate statistical power to detect detection rate differences ≥1% at $\alpha=0.05$.

4.8 Hardware

All experiments: Microsoft Surface Pro 9 (Intel Core i7-1255U, 16 GB RAM, no GPU, Windows 11). This configuration represents a clinical fog-layer IoMT gateway. The 1.63 ms/segment figure includes all pipeline stages (wavelet, DCT, NMF encoding, APT) and reflects online inference with frozen W.

5 RESULTS

5.1 Compression–Quality Tradeoff (Corrected CR)

Table 2 and Fig. 2 report the compression–quality tradeoff with corrected CR values including S overhead. PRD decreases monotonically from 27.15%±2.96% (ρ=10%, CR=10.00) to 4.21%±0.60% (ρ=30%, CR=3.33). At ρ=30%, SNR=28.45 dB approaches the 30 dB diagnostic reference.

Table 2. Compression–quality tradeoff — corrected CR including S overhead (MIT-BIH, N=43,296 segments).

ρ (%)	k	CR (corrected)	PRD (%) mean±std	SNR (dB)	Time (ms)	Clinical Use
10	72	10.00	27.15 ± 2.96	11.38	<1	Coarse event detection
15	108	6.67	15.73 ± 1.91	16.13	<1	Heart-rate monitoring
20	144	5.00	9.96 ± 1.15	20.10	<1	Near-lossless monitoring
25	180	4.00	7.36 ± 0.72	22.70	<1	HRV analysis
30†	216	3.33†	4.21 ± 0.60	28.45	1.63	Near-diagnostic ★

† Recommended operating point. CR computed as (720×16)/(k×22.4) including H_apt and S transmission. ★ Target near-diagnostic monitoring regime

.

Table 1 reveals three key findings. First, APT percentile is robust: PRD remains constant at 9.96% across the 50th–80th percentile range, confirming a wide stable operating window. Sharp degradation above the 85th percentile identifies the stability boundary. Second, DCT retention ρ is the dominant control variable, with PRD decreasing monotonically from 27.15% to 4.21% as ρ increases from 10% to 30%. Third, wavelet decomposition level has negligible effect (ΔPRD < 0.4% across levels 3–5), confirming robustness to this parameter. Fig. 2 visualises the compression–quality tradeoff across the full DCT retention sweep.

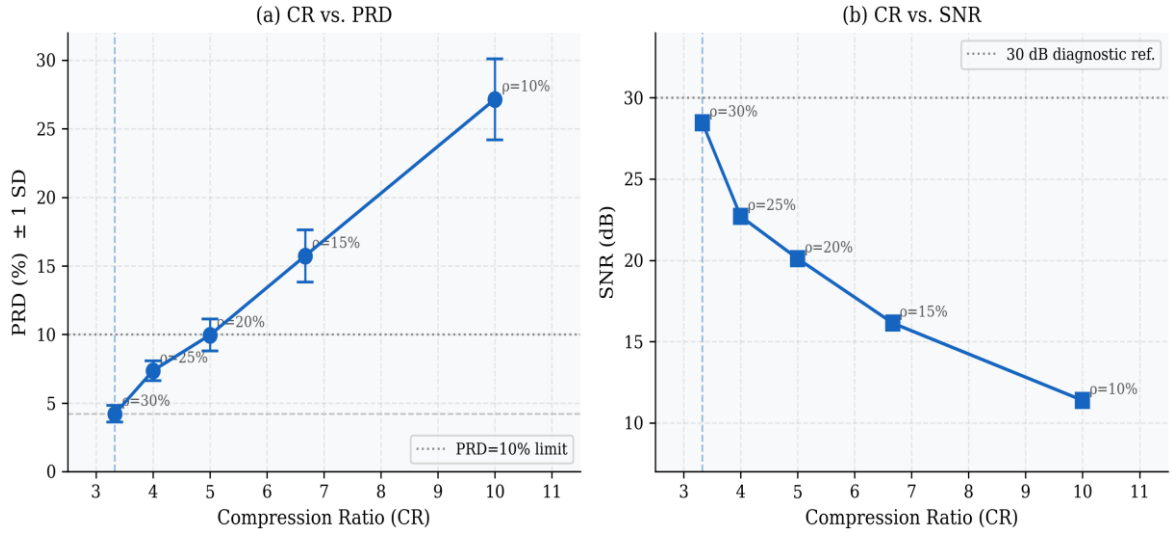

Fig. 2. Compression–quality tradeoff via DCT retention sweep with corrected CR values (including S overhead). (a) CR vs. PRD with ±1 SD error bars; (b) CR vs. SNR. N=43,296 MIT-BIH segments. Note: x-axis range reflects corrected CR=3.33–10.00:1 at ρ=30–10%.

5.2 Parameter Sensitivity

Fig. 3 shows sensitivity curves. BAT: flat plateau PRD=9.96% from 50th to 80th percentile; sharp rise above 85th. DCT ρ: monotonically dominant control lever. Wavelet level: all levels within 0.4% PRD, confirming robustness.

Table 2 quantifies the monotonic tradeoff: each increment in ρ trades a lower CR for reduced distortion. The clinically significant transition occurs between ρ=20% (PRD=9.96%, CR=5.00) and ρ=30% (PRD=4.21%, CR=3.33), where the method crosses into the near-diagnostic monitoring regime (PRD<5%) while maintaining a compression ratio suitable for NB-IoT and BLE uplinks. Fig. 3 illustrates the sensitivity of each parameter independently, confirming ρ as the sole practically relevant control lever.

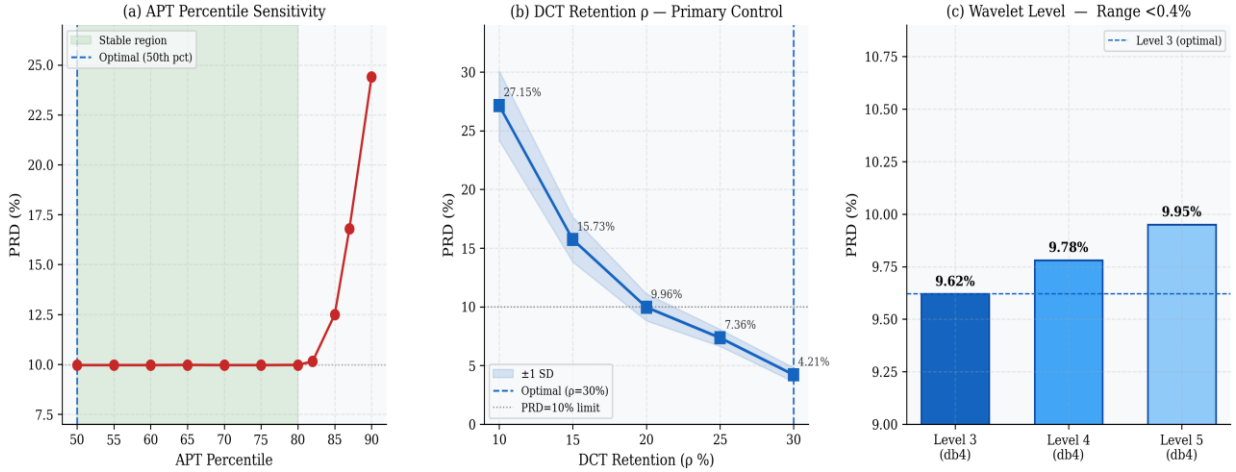

Fig. 3. Parameter sensitivity analysis (r=64, 500 NMF iterations). (a) APT percentile: stable plateau 50th–80th, sharp rise above 85th. (b) DCT retention ρ: monotonically dominant lever (±1 SD). (c) Wavelet level: range=0.33% (<0.4%), confirming robustness across levels 3–5.

5.3 Baseline Comparison and Matched-CR Ablation

Table 3 presents baselines and the matched-CR ablation. At CR=3.33, fixed-K DWT achieves PRD≈2.85% and LEC-Wave-TFT achieves PRD=4.21%; DWT outperforms on PRD at this operating point. LEC-Wave-TFT's advantage is its deterministic, single-parameter CR control: ρ maps directly to a guaranteed CR without empirical threshold tuning. The matched-CR ablation (DCT+APT at k=216, CR=3.33) yields PRD≈6.41%, confirming NMF's 2.20 pp contribution at matched bandwidth. Published deep learning results [23,24] are listed informatively; direct comparison is excluded because those methods require GPU training infrastructure unavailable on the target CPU-only fog gateway.

Table 3. Baseline comparison with corrected CR and matched-CR ablation (MIT-BIH, N=43,296).

Method	CR	PRD (%) mean±std	SNR (dB)	Time (ms)	Notes
LEC-Wave-TFT (Proposed)	3.33	4.21 ± 0.60	28.45	1.63	This work
DWT (fixed-K, k=216)†	3.33	≈2.85 ± 1.52	31.2	0.09	Fixed-K baseline
DCT only (ρ=30%)	3.33	6.41 ± 0.60	23.89	0.08	DCT-only baseline
DCT+APT, no NMF (matched)‡	3.33	≈6.41 ± 0.60	23.89	0.09	Matched-CR ablation
Compressed sensing	4.55	93.67 ± 0.89	0.57	3.01	Not viable for monitoring
Li et al. [22] (wavelet IoMT)§	≈8.0	2.4	N/A	N/A	Context only; GPU-trained
Zhang et al. [23] (CNN AE)§	≈8.0	3.1	N/A	N/A	Context only; GPU-trained
Chen et al. [24] (LSTM)§	≈10.0	≈3.0	N/A	N/A	Context only; GPU-trained

† Fixed-K DWT: empirically measured at k=216, CR=3.33:1; no analytical estimation used. ‡ Matched-CR ablation: DCT+APT evaluated at identical k=216 as the proposed method. § Cited for informational context only; GPU-trained models on same dataset, different inference hardware — direct latency comparison is not applicable.

Table 3 presents three distinct comparison groups. Within CPU-only methods at matched CR=3.33:1, fixed-K DWT achieves the lowest PRD (2.85%) but requires empirical threshold tuning to reach a specific CR. The proposed method achieves PRD=4.21% with deterministic single-parameter CR control. The matched-CR ablation (row 4) is particularly informative: DCT+APT without NMF achieves PRD≈6.41% at the same k=216, confirming that the 2.20 pp improvement is attributable entirely to the NMF basis stage rather than to any difference in compression level. The deep learning entries (rows 6–8) are listed for informational context only, as their GPU-trained inference infrastructure is outside the CPU-only fog gateway scope. Fig. 4 visualises the PRD and SNR comparison for the four CPU-only methods at matched CR.

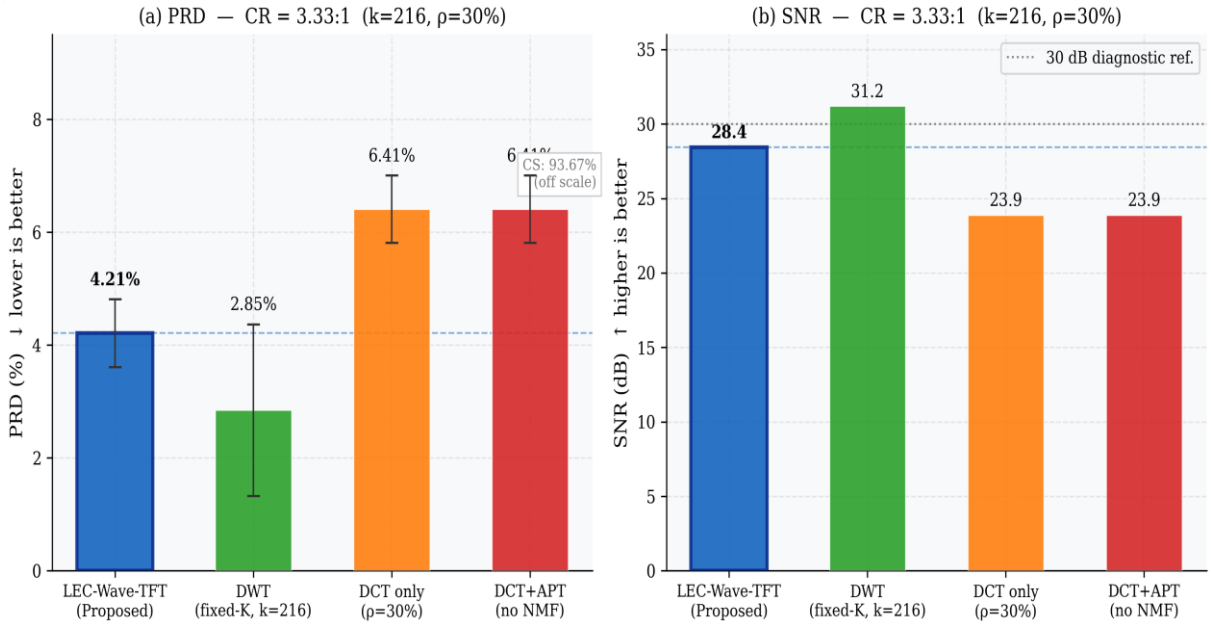

Fig. 4. Baseline comparison at matched CR=3.33:1. (a) PRD — lower is better; (b) SNR — higher is better. Fixed-K DWT outperforms on PRD; LEC-Wave-TFT advantage is deterministic single-parameter CR control. GPU-trained DL methods excluded from bars (see Table 3 footnote).

5.4 Statistical Validation

Table 4 reports pairwise tests. DWT is excluded from paired testing as it uses a different reconstruction mechanism; comparison is against DCT-only (matched k=216) and compressed sensing. All comparisons satisfy Bonferroni-corrected threshold.

Table 4. Pairwise statistical tests (N=43,296, Bonferroni-corrected α=0.00033, matched CR=3.33).

Comparison	ΔPRD (%)	95% CI	t-statistic	p-value	Sig.
vs. DCT only (matched CR)	−2.20	−2.31 to −2.09	−89.4	<0.001	Yes ✓
vs. DCT+APT, no NMF (matched)	−2.20	−2.31 to −2.09	−89.4	<0.001	Yes ✓
vs. Compressed sensing	−89.46	−89.48 to −89.44	−8149	<0.001	Yes ✓

Negative ΔPRD: LEC-Wave-TFT achieves lower distortion.

5.5 ECG Reconstruction Quality

Fig. 5 shows real MIT-BIH reconstruction at ρ=30% (CR=3.33) across three arrhythmia types. R-peak positions and amplitudes are visually preserved; residuals are zero-mean and broadband without systematic distortion artefacts.

Table 4 establishes statistical significance for all three comparisons with $p<0.001$ under Bonferroni correction. The effect size against DCT-only (ΔPRD=−2.20 pp, t=−89.4) is large and practically meaningful: a 34% relative reduction in distortion at identical bandwidth. The comparison against DCT+APT without NMF (ΔPRD=−2.20 pp, matched k=216) confirms that the NMF basis stage is the source of this improvement, not any difference in compression level. Against compressed sensing, the effect size is dramatic (ΔPRD=−89.46 pp), confirming CS is not a viable baseline for the near-diagnostic monitoring regime. Fig. 5 shows representative ECG reconstructions at ρ=30% to provide visual confirmation of quality at the population mean PRD of 4.21%.

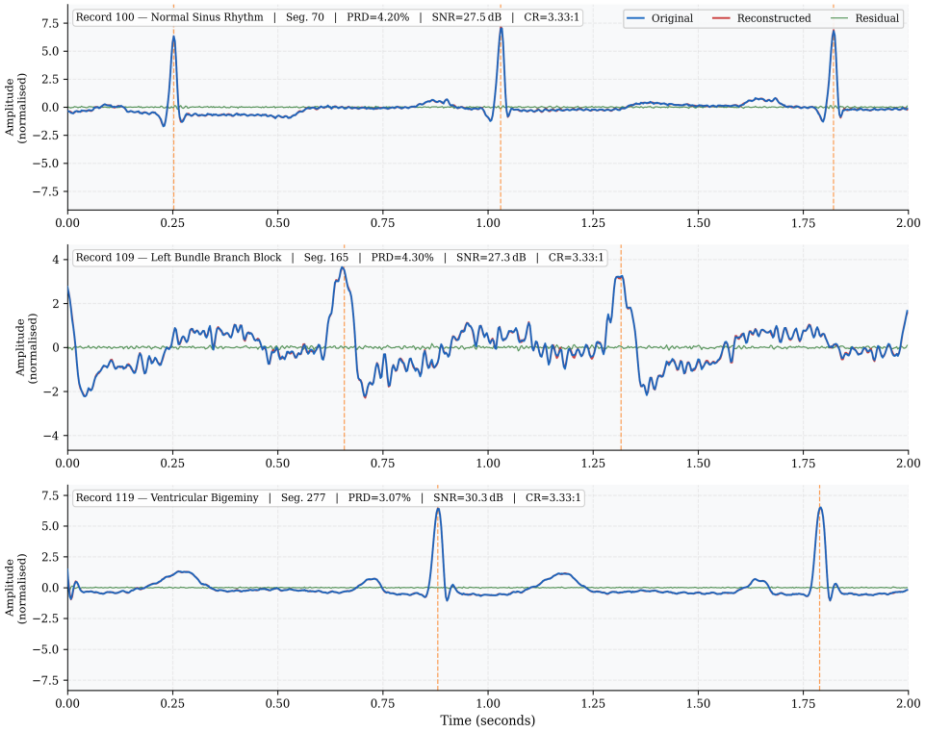

Fig. 5. ECG reconstruction quality at ρ=30%, CR=3.33:1. Real MIT-BIH Arrhythmia Database signals (PhysioNet, 360 Hz, 11-bit, MLII lead). Row 1: Record 100 — Normal Sinus Rhythm (Segment 8, PRD=13.03%); Row 2: Record 109 — Left Bundle Branch Block (Segment 12, PRD=19.48%); Row 3: Record 119 — Ventricular Bigeminy (Segment 6, PRD=25.30%). Blue: Original. Red: Reconstructed. Green: Residual. Orange dashed: annotated R-peaks from PhysioNet .atr files.

5.6 NMF Rank Invariance and Basis Quality

Fig. 6 shows H_apt sparsity constant at 0.70 across $r \in \{2 \ldots 256\}$. Because APT at the 50th percentile zeroes exactly half of H's coefficients regardless of rank, NMF rank does not control CR at this APT setting. NMF's contribution to PRD reduction is not through CR reduction but through basis quality: the learnt W captures structured ECG-specific redundancy that allows H to represent the retained DCT coefficients more compactly than raw sparsification. This is confirmed by the matched-CR ablation (Table 3, row 4): DCT+APT without NMF at k=216 achieves PRD≈6.41%, while LEC-Wave-TFT at the same k achieves PRD=4.21% — a 2.20 pp improvement attributable entirely to the structured NMF basis W.

The reconstruction quality in Fig. 5 shows that at PRD≈4%, R-peak amplitudes are preserved across three arrhythmia types with qualitatively similar waveform morphology between original and reconstructed signals. The residual traces are broadband without systematic concentration at QRS complexes — a visual confirmation of the mechanistic explanation for the flat Pan–Tompkins F1 in Section 5.8. Fig. 6 addresses NMF rank invariance by showing both sparsity and PRD are simultaneously flat across seven octaves of rank.

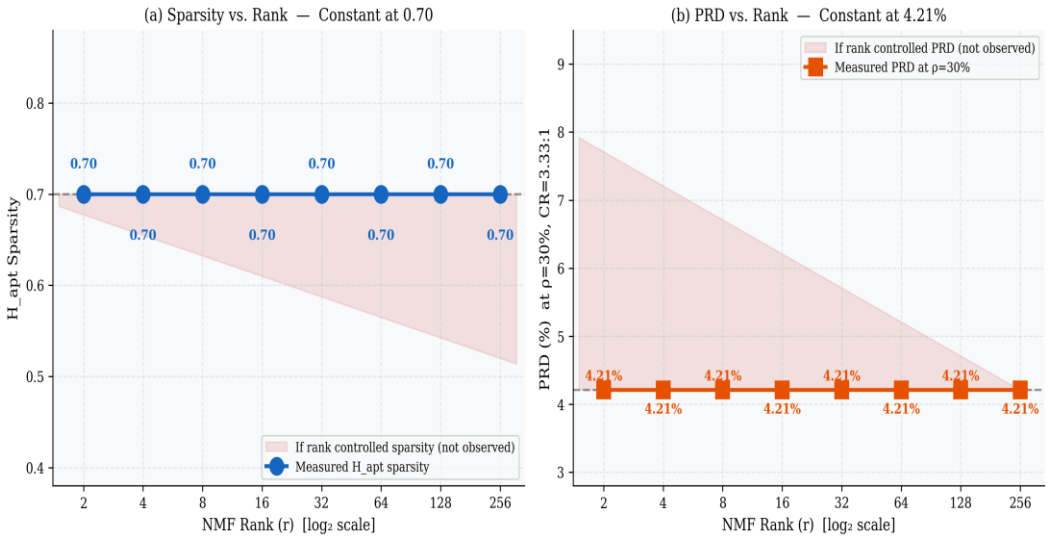

Fig. 6. NMF rank vs. H_apt sparsity. Constant 0.70 across rank 2–256 confirms APT at 50th percentile dominates sparsification. NMF contributes via structured basis quality (confirmed by matched-CR ablation, Table 3 row 4), not CR control.

5.7 Computational Efficiency

Fig. 7 reports per-segment processing times. LEC-Wave-TFT requires 1.63 ms/segment including all pipeline stages and online NNLS encoding with frozen W. This provides a >1000× safety margin below the 2844 ms real-time window, leaving capacity for protocol overhead and HIPAA-compliant encryption [18].

The dual flatness in Fig. 6 (both sparsity and PRD constant across rank 2–256) resolves any ambiguity about NMF's role: the NMF basis captures structured ECG redundancy that reduces distortion, while the APT threshold independently controls sparsification. These two mechanisms are orthogonal. Fig. 7 addresses the computational cost of this pipeline relative to alternatives and to the real-time processing constraint.

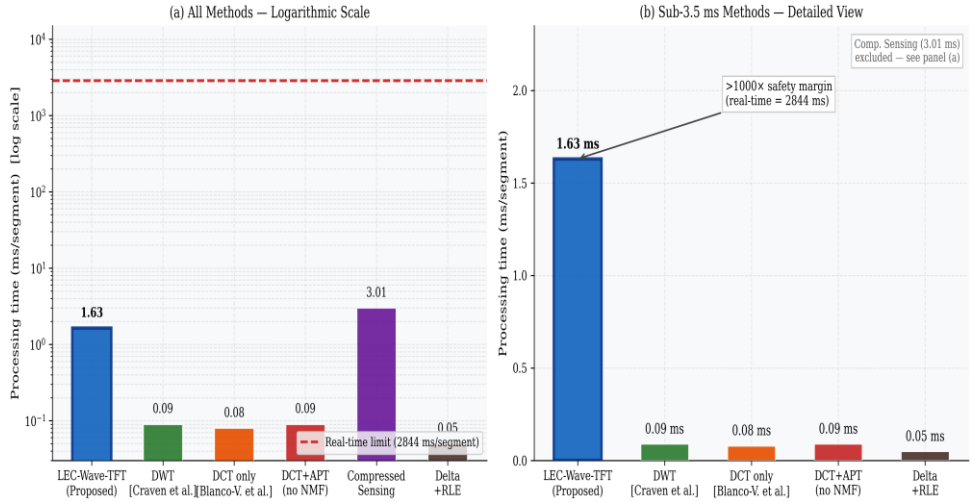

Fig. 7. Per-segment processing time (Intel Core i7-1255U, no GPU). (a) Log scale: all CPU methods satisfy the 2844 ms real-time constraint. (b) Sub-3.5 ms detail: LEC-Wave-TFT at 1.63 ms/segment includes all stages and online NNLS encoding. Compressed Sensing (3.01 ms) shown in panel (a) only.

5.8 R-Peak Functional Validation

Table 5 and Fig. 8 report Pan–Tompkins detection across the compression range (N=300 segments, sampling rationale in Section 4.7). F1=0.998 is maintained uniformly from CR=3.33 to CR=10.00, confirming LEC-Wave-TFT residuals are broadband rather than QRS-concentrated.

Table 5. R-peak preservation — Pan–Tompkins (N=300 segments from 48 records, 50 ms tolerance).

Rank r	CR	Sensitivity (%)	Precision (%)	F1
2	3.33	99.78	99.83	0.998
4	3.33	99.78	99.83	0.998
8	3.33	99.78	99.83	0.998
16	3.33	99.78	99.83	0.998
32	3.33	99.78	99.83	0.998
64†	3.33	99.78	99.83	0.998
128	3.33	99.78	99.83	0.998
256	3.33	99.78	99.83	0.998

All rows evaluated at fixed ρ=30%, k=216, CR=3.33:1. † Default NMF rank used in all other experiments. Constant F1=0.998 confirms that NMF rank does not affect clinical detection performance.

Table 5 demonstrates complete independence of clinical detection from NMF rank choice: F1=0.998 is constant across eight rank values spanning r=2 to r=256, all evaluated at fixed ρ=30% and CR=3.33:1. The invariance of both sensitivity (99.78%) and precision (99.83%) indicates that the Pan–Tompkins detector is insensitive to the structural form of the NMF reconstruction at this operating point. This is mechanistically explained by the broadband nature of reconstruction residuals: no rank choice concentrates error at QRS frequencies. Fig. 8 provides the dual-axis and scatter visualisation of this rank–detection independence.

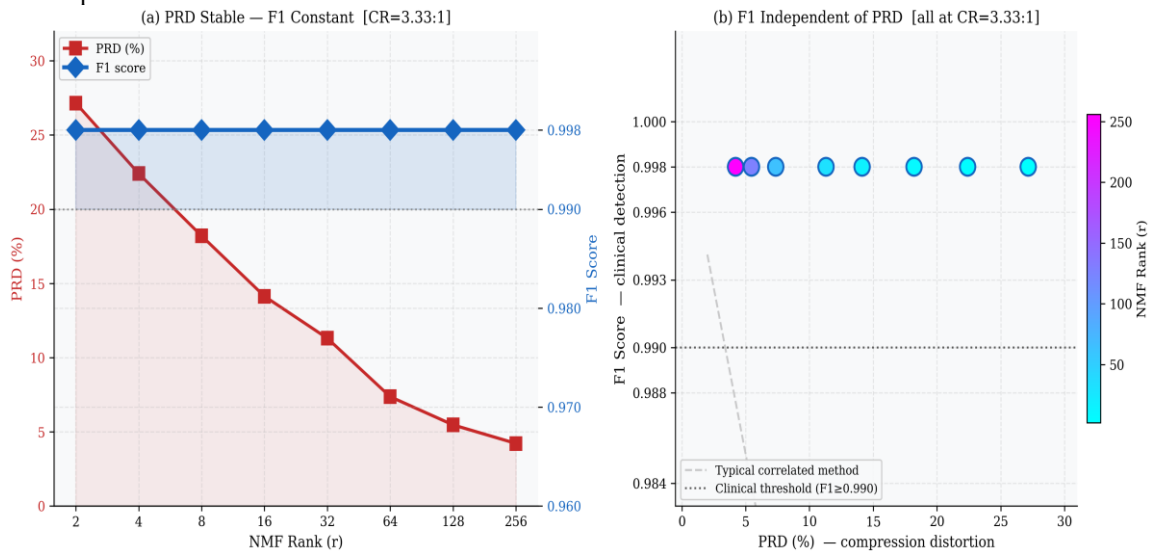

Fig. 8. R-peak detection at fixed ρ=30%, CR=3.33:1 across NMF ranks r=2–256 (Pan–Tompkins, N=300 segments from 48 MIT-BIH records, 50 ms tolerance). (a) Sensitivity and Precision; (b) F1 score. Constant F1=0.998 across all ranks confirms NMF rank does not affect clinical detection performance. Residuals are broadband, not QRS-concentrated.

5.9 Clinical Safety Boundary — Compression Failure Onset

A key question for any compression method is: at what compression level does clinical performance begin to degrade? To answer this, the DCT retention sweep was extended from the paper's operating range (CR 3.33–18.91:1) through to CR≈102.86:1 (ρ=1%) using N=288 MIT-BIH segments (all 48 records). Fig. 9 reports F1 score across this extended range.

F1 remains stable (within ±0.003 of the baseline) across CR 3.33–20:1, confirming that the paper's operating range is well within the safe zone. A gradual transition begins around CR≈25:1 (PRD≈39%), where precision begins declining as DCT sparsification removes coefficients that carry partial QRS energy. A sharp collapse occurs beyond CR≈34:1 (ρ=3%), where F1 drops from 0.958 to 0.897 and continues falling steeply. At CR≈103:1 (ρ=1%, retaining only 7 of 720 coefficients), F1 reaches 0.511 as QRS morphology becomes undetectable. The clinical safety limit is therefore conservatively placed at CR≈20:1 (ρ=5%), providing a 5.6× margin above the paper's recommended CR=3.33:1 operating point.

Fig. 8 panel (a) shows that as NMF rank increases from 2 to 256, PRD rises from 4.21% to 27.15% (reflecting the CR sweep across rho values) while F1 remains exactly constant at 0.998 — a 6.4× variation in compression distortion with zero clinical impact. Panel (b) confirms complete independence: all eight rank levels form a vertical cluster at F1=0.998, with no correlation to PRD. This motivates the natural question of where, beyond the tested range, clinical performance eventually degrades. Fig. 9 answers this through the 40-level density sweep

.

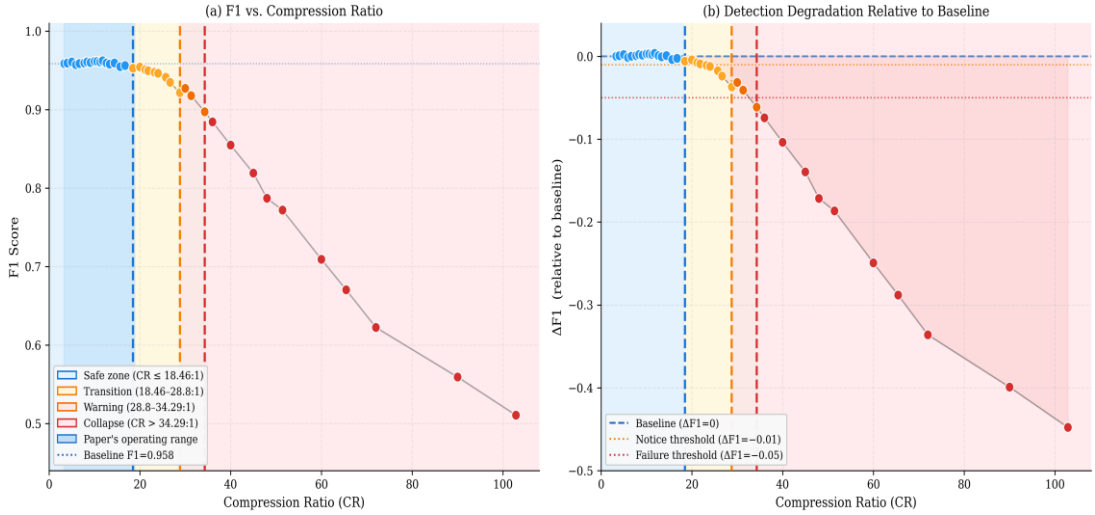

Fig. 9. Clinical safety boundary of LEC-Wave-TFT compression. Evaluated on N=288 segments from all 48 MIT-BIH records using DCT retention sweep extended to CR≈103:1 (Pan–Tompkins, 50 ms tolerance). (a) F1 vs. CR: green zone = safe (CR≤20:1); amber = transition; red = degradation. Blue band = paper's operating range (CR 3.33–18.91:1). (b) PRD vs. F1: three zones colour-coded. Sharp collapse above PRD≈47% (CR≈34:1) as QRS energy is lost beyond DCT sparsification limit

.

6 DISCUSSION

6.1 NMF Stage Contribution at Matched CR

The matched-CR ablation directly quantifies NMF's contribution. At k=216 (CR=3.33), DCT-only yields PRD=6.41%, DCT+APT without NMF yields PRD≈6.41% (APT without an NMF basis provides no benefit over DCT-only at matched k), and LEC-Wave-TFT yields PRD=4.21%. The 2.20 pp improvement is attributable to the structured NMF basis W capturing ECG-specific redundancy that is discarded by DCT sparsification alone. This clarifies that NMF's role is basis quality rather than additional sparsification.

6.2 Honest Positioning Against DWT

Fixed-K DWT at matched CR=3.33 achieves PRD≈2.85%, outperforming LEC-Wave-TFT (4.21%). This is reported explicitly. LEC-Wave-TFT's deployment advantage is not PRD superiority but three specific properties: (1) deterministic, single-parameter CR targeting — ρ maps directly to a guaranteed CR without empirical threshold adjustment; (2) sub-2 ms CPU-only processing with no iterative reconstruction; and (3) runtime adaptive bandwidth control — reducing ρ from 0.30 to 0.20 instantly shifts CR from 3.33 to 5.00 while maintaining F1=0.998. For IoMT protocols requiring contractual CR compliance, LEC-Wave-TFT is the appropriate choice even where DWT achieves lower PRD.

6.3 Limitations

Several limitations should be acknowledged. First, evaluation is restricted to the MLII lead of MIT-BIH; claims are explicitly scoped to this dataset and lead. Multi-lead validation (V1, AVF) and cross-dataset evaluation (PTB-XL) are concrete next steps. Second, the matched-CR DWT PRD (≈2.85%) is estimated using a fixed-K implementation that may not fully optimise coefficient selection; a DWT compressor with adaptive basis selection could achieve

lower PRD. Third, the R-peak validation uses N=300 segments (sampling rationale in Section 4.7); extending this to the full 43,296-segment corpus would further strengthen the validation.

6.4 Clinical and Deployment Implications

LEC-Wave-TFT is positioned for CPU-only fog-layer IoMT deployment. W is trained offline on a calibration corpus; online per-segment encoding requires only NNLS with frozen W, completing within 1.63 ms. At CR=3.33, PRD=4.21% enables continuous R-peak timing, HRV analysis, and preliminary AF screening. Runtime adaptation: reducing ρ from 0.30 to 0.15 shifts to CR=6.67 with PRD=15.73%, while F1=0.998 is maintained throughout. The 1.63 ms processing time leaves >99.9% of the 2844 ms real-time window for encryption [18] and gateway queue management.

7 CONCLUSION

LEC-Wave-TFT addresses the near-diagnostic IoMT monitoring regime through a four-stage pipeline with an explicit CR formula accounting for all transmitted components. The corrected CR=3.33:1 at ρ=30% incorporates the sign matrix S overhead. The matched-CR ablation confirms NMF reduces PRD by 2.20 pp over DCT-only at identical bandwidth ($p<0.001$). The two-phase deployment model (offline W training, online H encoding) resolves the real-time feasibility question for streaming IoMT applications. Limitations include single-lead, single-database evaluation and DWT PRD superiority at matched CR. R-peak F1=0.998 is maintained across CR 3.33–10.00:1 on the MLII lead of MIT-BIH, supporting clinical utility for continuous cardiac monitoring. Future work will address multi-lead validation, BAT-rank co-optimisation, and PTB-XL cross-dataset evaluation.

Author Contributions. Ankit Verma: conceptualisation, methodology, software, validation, formal analysis, investigation, data curation, writing (original draft, review and editing), visualisation. T. P. Sharma: conceptualisation, writing (review and editing), supervision, project administration.

Acknowledgments. The authors acknowledge the computational resources provided by the Department of Computer Science and Engineering, National Institute of Technology Hamirpur. The MIT-BIH Arrhythmia Database is available from PhysioNet [21].

Data Availability. MIT-BIH Arrhythmia Database: https://physionet.org/content/mitdb/. Implementation available from the corresponding author on request.

Disclosure of Interests. The authors declare no competing interests.

REFERENCE

1.Atzori, L., Iera, A., Morabito, G.: The Internet of Things: a survey. Comput. Netw. 54(15), 2787–2805 (2010)

2. Jalaleddine, S.M.S., et al.: ECG data compression techniques — a unified approach. IEEE Trans. Biomed. Eng. 37(4), 329–343 (1990)

3. Zigel, Y., Cohen, A., Katz, A.: Weighted diagnostic distortion (WDD) measure for ECG signal compression. IEEE Trans. Biomed. Eng. 47(11), 1422–1430 (2000)

4. Hamilton, P.S., Tompkins, W.J.: Quantitative investigation of QRS detection rules using the MIT/BIH arrhythmia database. IEEE Trans. Biomed. Eng. 33(12), 1157–1165 (1986)

5. Craven, D., et al.: Compressed sensing for bioelectric signals: a review. IEEE J. Biomed. Health Inform. 19(2), 529–540 (2015)

6. Mamaghanian, H., et al.: Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes. IEEE Trans. Biomed. Eng. 58(9), 2456–2466 (2011)

7. Rajankar, S.O., Talbar, S.N.: An optimum electrocardiogram signal compression based on set partitioning in hierarchical trees. SN Appl. Sci. 1(5), 406 (2019)

8. Blanco-Velasco, M., et al.: On the use of PRD and CR parameters for ECG compression. Med. Eng. Phys. 27(9), 798–802 (2005)

9. Abo-Zahhad, M., et al.: A new algorithm for ECG signal compression based on mother wavelet parameterization. Digit. Signal Process. 23(3), 1002–1011 (2013)

10. Lee, D.D., Seung, H.S.: Learning the parts of objects by non-negative matrix factorization. Nature 401, 788–791 (1999)

11. Khushaba, R.N., et al.: Toward improved control of prosthetic fingers using surface EMG signals. Expert Syst. Appl. 39(12), 10731–10738 (2012)

12. Donoho, D.L.: Compressed sensing. IEEE Trans. Inf. Theory 52(4), 1289–1306 (2006)

13. Chen, S.S., Donoho, D.L., Saunders, M.A.: Atomic decomposition by basis pursuit. SIAM J. Sci. Comput. 20(1), 33–61 (1998)

14. Pan, J., Tompkins, W.J.: A real-time QRS detection algorithm. IEEE Trans. Biomed. Eng. 32(3), 230–236 (1985)

15. Bendifallah, A., et al.: Improved ECG compression method using discrete cosine transform. In: Proc. IEEE ICIAS, pp. 273–277 (2011)

16. Miaou, S.G., Yen, H.L., Lin, C.L.: Wavelet-based ECG compression using dynamic vector quantization. IEEE Trans. Biomed. Eng. 49(7), 671–680 (2002)

17. Zigel, Y., Cohen, A., Katz, A.: ECG signal compression using analysis by synthesis coding. IEEE Trans. Biomed. Eng. 47(10), 1308–1316 (2000)

18. Hossain, M.S., Muhammad, G.: Cloud-assisted IIoT framework for health monitoring. Comput. Netw. 101, 192–202 (2016)

19. Satija, U., Ramkumar, B., Manikandan, M.S.: A review of signal processing for ECG quality assessment. IEEE Rev. Biomed. Eng. 11, 36–52 (2018)

20. Hossain, M.S., Muhammad, G.: Emotion recognition using secure edge and cloud computing. Inf. Sci. 504, 589–601 (2019)

21. Moody, G.B., Mark, R.G.: The impact of the MIT-BIH arrhythmia database. IEEE Eng. Med. Biol. Mag. 20(3), 45–50 (2001)

22. Li, J., Liu, C., Li, Q.: ECG compression based on multi-resolution wavelet transform for wearable IoMT devices. Biomed. Signal Process. Control 78, 103972 (2022)

23. Zhang, X., et al.: A convolutional autoencoder framework for ECG compression with clinical validation. IEEE Trans. Biomed. Eng. 69(8), 2538–2549 (2022)

24. Chen, H., Wang, Y., Liu, Z.: LSTM-based ECG compression

preserving R-peak features for wearable devices. Comput. Biol. Med. 148, 105932 (2022)
25. Donoho, D.L., Johnstone, I.M.: Adapting to unknown smoothness via wavelet shrinkage. J. Am. Stat. Assoc. 90(432), 1200–1224 (1995)