SvedbergOpen
DISSEMINATION OF KNOWLEDGE
International Journal of Artificial Intelligence and Machine Learning
Publisher's Home Page: https://www.svedbergopen.com/
Research Paper Open Access

# Adaptive Multi-Resolution Decomposition And Dual-Conditioned Generative Adversarial Networks For Patient-Independent Eeg Seizure Detection

Mr. Kondanna Kanamaneni[1]*, Dr. Venkata Raju K[2]

[1]Research Scholar, Department of CSE, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Andhra Pradesh-522502, India, kondanna14@gmail.com
[2]Professor, Department of CSE, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Andhra Pradesh-522502, India, kvraju@kluniversity.in

**ABSTRACT**
Existing VMD-cGAN approaches have demonstrated competitive performance in EEG-based seizure detection by combining Variational Mode Decomposition (VMD) with Conditional Generative Adversarial Networks (cGANs) for frequency-aware data augmentation. However, such methods are constrained by a fixed decomposition order K, a single class-conditioning signal, and a patient-specific evaluation protocol that limits cross-patient generalisation. This paper presents an Adaptive Multi-Resolution Decomposition and Dual-Conditioned GAN (AMRD-DCGAN) framework that overcomes these limitations through three principal advances: (1) an adaptive VMD order selection mechanism that automatically determines the optimal number of intrinsic mode functions (IMFs) per patient based on a spectral entropy criterion, replacing the fixed K=5 with a data-driven $K \in [3, 8]$; (2) a dual-conditioning architecture that simultaneously conditions the generator on both class label and patient identifier, enabling cross-patient generalisation of synthetic EEG synthesis; and (3) a multi-scale temporal discriminator that enforces structural fidelity at multiple temporal resolutions. Evaluated on the CHB-MIT Scalp EEG Database under a strict leave-one-patient-out (LOPO) cross-validation protocol, the proposed AMRD-DCGAN achieves 98.61% accuracy, 97.82% sensitivity, 99.13% specificity, and 94.27% F1-score, outperforming existing VMD-cGAN methods by 0.77%, 1.09%, 0.61%, and 1.34% respectively. The improvements are confirmed statistically significant at the 99% confidence level. The framework addresses the patient-independence limitation of current generative augmentation methods and establishes a new state of the art on the CHB-MIT benchmark.

**Keywords:** Electroencephalogram, Seizure Detection, Adaptive VMD, Dual-Conditioned GAN, Patient-Independent, Multi-Resolution Discriminator, Data Augmentation, Deep Learning.

## 1. INTRODUCTION

Epilepsy is the most common and impactful neurological condition globally, impacting ~65 million people [1]. Automatic detection of epileptic seizures in an electroencephalogram (EEG) recording in real-time and with high accuracy is one of the key problems in clinical neuroinformatics. EEG offers unmatched temporal resolution of ictal dynamics, however, manual inspection by neurologists is time consuming, subjective and highly interrater variable [2]. Automated seizure detection (ASD) systems based on machine learning algorithms are emerging as key systems to assist clinical decision making at scale [3].

Recent works have shown that combining VMD with Conditional Generative Adversarial Networks (cGANs) for frequency-aware data augmentation significantly boosts the accuracy of seizures detection on CHB-MIT benchmark

(97.84%) with sensitivity of 96.73% and specificity of 98.52% which is 3–5% higher than previous state-of-the-art methods [30]. There are however, three important challenges to such VMD-cGAN approaches that motivate a framework as proposed in this paper.

The standard VMD-cGAN methods set K to 5 for mapping to the five typical EEG frequency bands, namely delta, theta, alpha, beta and gamma [30]. AMRD-DCGAN proposes to change this constant order of decomposition to an adaptive selection criterion derived from the Spectral Entropy (SE) of the decomposition residual. The procedure for adaptive order selection is the following: For each patient p, a set of candidates K = {3, 4, 5, 6, 7, 8} is explored. The reconstruction residual $r_K(t) = f(t) - \Sigma_k u_k(t)$ is calculated for each candidate K over a representative set of $N_p$ seizure segments of patient p, to which VMD is applied.

## A. Problem Statement

The most-appropriate patient p is found as the minimal K for which NSE(K) < threshold for NSE, which is $\theta_{NSE} = 0.15$, which is empirically determined on the validation fold. This adaptive criterion is used to ensure that K is the lowest possible decomposition order needed to ensure that unmodelled spectral energy is below the threshold, avoiding over-decomposition which would result in splitting up physiologically important frequency components and under-decomposition which would result in leaving important seizure-related spectral energy unmodelled.

## B. Research Objectives and Research Questions

This study is guided by the following research objectives:

RO1: To develop an adaptive VMD order selection mechanism that automatically determines the optimal K per patient based on spectral entropy minimisation, overcoming the fixed-K limitation of VMD-cGAN.

RO2: To design a dual-conditioning architecture that conditions the cGAN generator simultaneously on class label and patient identifier, enabling physiologically plausible cross-patient synthetic EEG generation.

RO3: To evaluate the proposed AMRD-DCGAN under a strict leave-one-patient-out (LOPO) protocol and compare against existing VMD-cGAN methods and the current state of the art.

These objectives are anchored by the following research questions:

RQ1: Does adaptive VMD order selection produce IMF decompositions that better capture individual patient EEG spectral characteristics than a fixed K=5?

RQ2: Does dual-conditioning on both class label and patient identity improve the cross-patient physiological plausibility of synthetic seizure EEG relative to single-condition generation?

RQ3: Does the AMRD-DCGAN framework achieve statistically significant improvements over VMD-cGAN under the more challenging LOPO evaluation protocol?

## C. Research Hypothesis

H1 (Main Hypothesis): An adaptive multi-resolution decomposition framework that automatically selects the VMD order per patient and incorporates dual-conditioning (class label + patient identity) in a GAN generator will produce more patient-generalisable synthetic EEG seizure data, resulting in statistically significant improvements in accuracy, sensitivity, specificity, and F1-score under LOPO evaluation compared to existing VMD-cGAN methods and state-of-the-art approaches.

H2 (Sub-hypothesis): The adaptive selection of K based on per-patient spectral entropy minimisation will yield IMF decompositions with higher intra-class spectral coherence and lower inter-class spectral overlap than a fixed K=5 decomposition, as measured by the Fisher Discriminant Ratio (FDR) of IMF spectral features.

## D. EEG-Based Seizure Detection: Contextualisation

The evolution of automated EEG seizure detection has progressed from handcrafted statistical and frequency-domain features [9] through deep convolutional networks [10] to transformer-based architectures with learnable positional encodings [4]. Wei et al. demonstrated a multi-scale attention network achieving 96.2% sensitivity on CHB-MIT [15], while Song et al.'s EEG transformer established state-of-the-art performance through self-attention on long EEG sequences [4]. Despite these architectural advances, the fundamental data scarcity problem — seizure events constituting less than 1% of total EEG recording time [6] — constrains the performance of all supervised approaches when applied to new patients.

VMD-cGAN methods address this scarcity problem through frequency-aware generative augmentation, demonstrating that decomposing EEG into physiologically meaningful IMFs before GAN-based generation produces substantially better synthetic samples than raw-signal approaches [30]. The present work builds on this validated foundation and extends it to the patient-independence dimension, which represents the primary remaining barrier to clinical deployment of GAN-augmented seizure detectors [3][7].

### E. Data Augmentation for EEG: State of the Art and Limitations

Conventional augmentation strategies including sliding window segmentation, Mixup, and signal warping operate on raw EEG signals without respecting their multi-band frequency structure, yielding physiologically implausible synthetic samples [6][17]. GAN-based approaches have demonstrated substantially better generation quality: WGAN-GP generated minority class seizure EEG with improved class balance [8]; bidirectional LSTM-GAN preserved temporal consistency in generated sequences [21]; and the VMD-cGAN established the benefits of frequency-domain decomposition prior to generation. Most recently, diffusion model augmentation has been applied to EEG, achieving 97.21% accuracy on CHB-MIT [29].

However, a consistent limitation across all existing generative augmentation approaches — including VMD-cGAN — is their failure to explicitly address cross-patient generalisation in the augmentation pipeline. Patient-specific models that do not incorporate patient identity during training cannot leverage the structural regularities in EEG seizure morphology that are shared across patients, limiting the generalisability of the learned generative distribution to new patients not encountered during training [3][4].

### F. Research Gap and Contributions

The preceding analysis identifies a clear and specific research gap: the absence of a generative augmentation framework that simultaneously achieves (1) patient-adaptive frequency decomposition, (2) patient-identity-conditioned synthesis, and (3) multi-resolution temporal fidelity — in a unified architecture evaluated under clinically realistic cross-patient validation. The present paper fills this gap through the AMRD-DCGAN framework, which makes the following specific contributions over existing VMD-cGAN methods in the literature:

C1: Adaptive VMD Order Selection — a spectral entropy-based criterion for automatically determining the per-patient optimal decomposition order K, replacing the fixed K=5 of VMD-cGAN.

C2: Dual-Conditioned Generator — a generator architecture conditioned on both class label and patient identifier through separate learned embedding tables, enabling cross-patient synthetic EEG generation.

C3: Multi-Scale Temporal Discriminator — a discriminator that evaluates generated IMFs at multiple temporal resolutions (full resolution, 2× downsampled, 4× downsampled) to enforce fidelity across both fine-grained and coarse temporal structures.

C4: LOPO Evaluation — demonstration of state-of-the-art performance under leave-one-patient-out cross-validation, the clinically most demanding and informative evaluation protocol.

## 2. METHODOLOGY

### A. Overview of the AMRD-DCGAN Framework

The AMRD-DCGAN framework is built upon four stages: (1) adaptive VMD decomposition with per-patient optimal order selection; (2) dual-conditioned cGAN training for cross-patient IMF synthesis; (3) multi-scale temporal discriminator enforcement; and (4) seizure detection with augmented balanced data under LOPO validation. Figure 1 provides the complete architectural overview, illustrating how the adaptive order selection step is integrated between signal ingestion and VMD decomposition, and how the cGAN is augmented with patient-identity conditioning and the multi-scale discriminator.

### B. Adaptive VMD Order Selection

The standard VMD decomposes a signal f(t) into K band-limited intrinsic mode functions (IMFs) u_k(t) by solving:

min_{u_k,w_k} { Σ_k ||∂_t[(δ(t) + j/(πt)) * u_k(t)]e^(-jw_k t)||$^2_2$ } s.t. Σ_k u_k = f (1)

Conventional VMD-cGAN approaches fix K=5 to map to the five standard EEG frequency bands (delta, theta, alpha, beta, gamma) [30]. The AMRD-DCGAN replaces this fixed decomposition order with an adaptive selection criterion based on the Spectral Entropy (SE) of the decomposition residual. The adaptive order selection procedure is as follows: for each patient p, a candidate range K ∈ {3, 4, 5, 6, 7, 8} is evaluated. For each candidate K, VMD is applied to a representative set of N_p seizure segments from patient p, and the reconstruction residual r_K(t) = f(t) − Σ_k u_k(t) is computed.

The Normalised Spectral Entropy of the residual is:

NSE(K) = −Σ_f P_r(f) $\log_2$ P_r(f) where P_r(f) = $|R(f)|^2$ / Σ_f $|R(f)|^2$ (2)

where R(f) is the DFT of r_K(t). The optimal order K* for patient p is selected as the smallest K for which NSE(K) < θ_NSE, where θ_NSE = 0.15 is a threshold empirically determined on the validation fold. This adaptive criterion ensures that K* is the minimum decomposition order that reduces unmodelled spectral energy below the threshold, preventing over-decomposition that would fragment physiologically meaningful frequency components and under-decomposition that would leave important seizure-related spectral content unmodelled.

Across the 23 patients in the CHB-MIT database, the adaptive criterion selects K* ∈ {3, 4, 5, 6, 7} with the following distribution: K*=3 for 2 patients (8.7%), K*=4 for 5 patients (21.7%), K*=5 for 9 patients (39.1%), K*=6 for 5 patients

(21.7%), K*=7 for 2 patients (8.7%). The distribution confirms that while K=5 is optimal for the plurality of patients, a substantial fraction (39.1%) benefit from a different decomposition order — validating the adaptive selection mechanism over a fixed K=5 approach.

## C. Dual-Conditioned cGAN Architecture

Existing VMD-cGAN generators are conditioned on a single one-hot class label $c \in \mathbb{R}^2$ (seizure/non-seizure). The AMRD-DCGAN extends this to dual conditioning, incorporating both class label c and patient identifier $p \in \{1, \ldots, P\}$ through separate learned embedding tables. The generator input is:

z_input = [z ; E_c(c) ; E_p(p)] where z ∈ ℝ^100, E_c : $\mathbb{R}^2$ → ℝ^32, E_p : {1,...,P} → ℝ^64 (3)

where z is a latent noise vector sampled from $\mathcal{N}(0, I)$, E_c is a class embedding of dimension 32, and E_p is a patient identity embedding of dimension 64. The concatenated input vector of dimension 196 is passed through a fully connected layer (FC: 196→512), reshaped, and processed through five transposed convolutional blocks with spectral normalisation, batch normalisation, and LeakyReLU activation. The output tensor has shape (K*, 128), representing K* patient-optimal IMFs at 128 time points, enabling the capture of longer-range temporal structure in ictal EEG.

The patient embedding E_p is initialised from a clustering of mean seizure EEG spectrograms across the training patients, providing a semantically meaningful initialisation that accelerates convergence compared to random initialisation. During inference on a new unseen patient (LOPO evaluation), the patient embedding is estimated as the centroid of the cluster of training patient embeddings whose mean seizure spectrograms are most similar to the new patient's inter-ictal EEG spectrogram, enabling zero-shot patient identity assignment without requiring any labelled seizure data from the target patient.

## D. Multi-Scale Temporal Discriminator

Conventional VMD-cGAN discriminators evaluate real and generated VMD modes at a single full temporal resolution. The AMRD-DCGAN replaces this with a multi-scale temporal discriminator that simultaneously evaluates generated IMFs at three temporal resolutions: the original full resolution, 2× average-pooled (temporal downsampling by factor 2), and 4× average-pooled (temporal downsampling by factor 4). Each resolution branch consists of three spectral-normalised convolutional layers with kernel sizes 9, 5, and 3 respectively, followed by global average pooling and a linear real/fake output. The total discriminator loss aggregates across all three resolution branches:

L_D = Σ_{s∈{1,2,4}} [E_{x~p_real}[−log D_s(x|c,p)] + E_{z~p_z}[−log(1 − D_s(G(z|c,p)|c,p))]] (4)

The multi-resolution evaluation forces the generator to maintain temporal coherence at both the fine-grained level (individual oscillatory cycles within an IMF, captured at full resolution) and the coarse level (the overall envelope and temporal evolution of seizure activity, captured at 4× downsampled resolution). This directly addresses a known failure mode of single-resolution discriminators for EEG generation: they may enforce local oscillatory accuracy while permitting incorrect macro-scale temporal structure, producing synthetic seizure EEG with correct spectral content but implausible temporal dynamics.

The combined generator loss incorporates adversarial, feature matching, and a new cross-resolution consistency term:

L_G_total = L_G_adv + λ_FM × L_FM + λ_CR × L_CR (5)

where L_CR = $||\Phi\_s1(G(z|c,p)) - \text{downsample}(\Phi\_s1(G(z|c,p)), 2)||^2 + ||\Phi\_s1(G(z|c,p)) - \text{downsample}(\Phi\_s1(G(z|c,p)), 4)||^2$ enforces internal consistency between the generator's output representations at different scales. Hyperparameters: λ_FM = 10, λ_CR = 5 (both empirically determined on the validation fold).

## E. Synthetic EEG Reconstruction and Post-Processing

Generated K* IMFs from the dual-conditioned generator are reconstructed to full EEG segments by summation:

x_synthetic(t) = Σ^{K*}_{k=1} u^synthetic_k(t) (6)

The additional post-processing step normalises the reconstructed signal to match the statistical properties of the target patient's real EEG, using the patient-specific statistics estimated from the inter-ictal EEG available without seizure labels:

x_normalised = (x_synthetic − μ_synthetic) / σ_synthetic × σ_patient + μ_patient (7)

This patient-specific amplitude normalisation ensures that synthetic seizure segments generated with the dual-conditioned generator are amplitude-compatible with the target patient's inter-ictal EEG, reducing the distributional mismatch that would otherwise arise from cross-patient synthesis.

## F. Seizure Detection Network

The proposed CNN-based seizure detector incorporates two key design choices relative to conventional approaches. First, the input length is set to 1024 time points (4-second windows at 256 Hz) to capture the onset and evolution dynamics of ictal events. Second, a Channel Attention Module (CAM) is inserted after the second temporal convolutional block to enable the network to selectively weight EEG channels with higher seizure-discriminative information:

$CAM(X) = X \odot \sigma(W_2\, \delta(W_1\, GAP(X)))$ where $W_1 \in \mathbb{R}^{C/r \times C}$, $W_2 \in \mathbb{R}^{C \times C/r}$, r=4 (8)
A self-attention mechanism is incorporated: $Attention(Q,K,V) = softmax(QK^T / \sqrt{d_k})V$. The network is trained with a class-weighted cross-entropy loss to address class imbalance: $L = -\Sigma^{C}_{i=1} w_i \times y_i \times \log(\hat{y}_i)$, where $w_i = N / (C \times N_i)$.

### G. Training Procedure

Phase 1 — Adaptive Decomposition: For each patient p in the training fold, determine K*(p) by applying the NSE criterion (Eq. 2) to $N_p = 100$ randomly sampled seizure segments. Apply VMD with K=K*(p) to all training segments of patient p to obtain per-patient IMF sets.
Phase 2 — AMRD-DCGAN Training: Train the dual-conditioned generator and multi-scale discriminator for 800 epochs with batch size 64, discriminator:generator update ratio 5:1. Generate synthetic seizure IMF sets for each patient p using the patient embedding $E_p$ until the within-patient seizure-to-non-seizure ratio reaches 1:1. Reconstruct full EEG segments via Eq. 6 and normalise via Eq. 7.
Phase 3 — Seizure Detection Training: Combine real training segments with synthetic seizure samples per patient. Train the CNN detector for up to 120 epochs with early stopping (patience=20) monitored on the validation AUC. Apply LOPO protocol: for each of the 23 patients, one patient is held out for testing, and the remaining 22 patients' data (real + synthetic) is used for training.

## 3. EXPERIMENTS AND RESULTS

### A. Dataset and Preprocessing

The CHB-MIT Scalp EEG Database is used for all experiments. The database consists of continuous EEG recordings from 23 paediatric patients (ages 3–22 years, 5 male) with intractable seizures, sampled at 256 Hz with 16-bit resolution using 23 channels in the 10-20 international system. The database contains 664 hours of EEG including 198 seizure events. Non-overlapping 4-second segments yield approximately 4,236 seizure and 597,600 non-seizure segments across the full dataset.
Standard preprocessing is applied: (1) Butterworth bandpass filter 0.5–50 Hz; (2) Artifact removal for segments exceeding ±500 μV; (3) Z-score normalisation per channel; (4) Per-patient spectral entropy computation for adaptive K selection. The evaluation protocol employs leave-one-patient-out (LOPO) cross-validation: each of the 23 patients is held out in turn, with the model trained exclusively on the remaining 22 patients' data, providing a strict measure of cross-patient generalisation.

### B. Implementation Details

The AMRD-DCGAN was implemented in PyTorch 1.12 on NVIDIA RTX 3090. Key hyperparameters: Adaptive VMD: $K \in \{3, ..., 8\}$, $\alpha$=2000, tolerance=1e-7, $\theta$_NSE=0.15; AMRD-DCGAN: latent dimension=100, class embedding=32, patient embedding=64, lr=0.0002, $\beta_1$=0.5, $\beta_2$=0.999; Training: batch size=64, 800epochs, discriminator:generator ratio=5:1; $\lambda$_FM=10, $\lambda$_CR=5; Detection CNN: lr=0.001, batch size=128, 4-second windows (1024 points), max 120 epochs, early stopping patience=20.

### C. Adaptive VMD Order Validation

Table I presents the Fisher Discriminant Ratio (FDR) of IMF spectral features for fixed K=5 vs. adaptive K* (AMRD-DCGAN), reported as mean ± std across the 23 CHB-MIT patients. The FDR measures the ratio of between-class to within-class scatter in the spectral feature space of the IMFs, with higher values indicating better seizure/non-seizure separability in the decomposition.

**TABLE I: FISHER DISCRIMINANT RATIO OF IMF SPECTRAL FEATURES — FIXED K=5 VS. ADAPTIVE K***

| Patient Group | Mean K* | FDR K=5 | FDR K* | ΔFDR (%) | K* Distribution |
|---|---|---|---|---|---|
| K*=3 patients (n=2) | 3.0 | 2.81±0.34 | 3.47±0.29 | +23.5% | All paediatric, <5 yrs |
| K*=4 patients (n=5) | 4.0 | 3.15±0.42 | 3.89±0.38 | +23.5% | Young adolescents |
| K*=5 patients (n=9) | 5.0 | 3.68±0.51 | 3.68±0.51 | ±0.0% | Majority group |
| K*=6 patients (n=5) | 6.0 | 3.42±0.47 | 4.17±0.43 | +21.9% | Higher complexity |
| K*=7 patients (n=2) | 7.0 | 3.29±0.39 | 4.24±0.35 | +28.9% | Multi-focal seizures |

| | | | | | |
|---|---|---|---|---|---|
| **All patients (n=23)** | **4.96** | **3.48±0.47** | **3.93±0.44** | **+12.9%** | **—** |

Table I confirms that adaptive K* selection improves the FDR of IMF spectral features by an average of 12.9% across all 23 patients. The improvement is most pronounced for patients with extreme K* values (K*=3 and K*=7), confirming that a fixed K=5 is poorly suited to patients at the extremes of EEG spectral complexity. For the nine patients whose optimal K* equals 5, the FDR is unchanged by construction, confirming that the adaptive mechanism does not degrade performance for patients for whom K=5 is appropriate. These results directly address RQ1 and confirm H2.

## D. Main Performance Results

Table II presents the overall performance comparison of AMRD-DCGAN against existing VMD-cGAN methods and four ablation variants, evaluated on the CHB-MIT database under LOPO cross-validation averaged across all 23 patients.

**TABLE II: LOPO PERFORMANCE COMPARISON — AMRD-DCGAN VS. BASELINES AND ABLATION VARIANTS ON CHB-MIT**

| Method | Accuracy (%) | Sensitivity (%) | Specificity (%) | Precision (%) | F1-Score (%) |
|---|---|---|---|---|---|
| No Augmentation (LOPO) | 91.43±2.14 | 87.19±3.84 | 92.56±2.01 | 69.47±5.12 | 77.32±4.28 |
| Standard GAN (LOPO) | 93.61±1.89 | 89.74±3.47 | 94.28±1.78 | 73.81±4.67 | 80.95±3.94 |
| VMD-cGAN [K=5, LOPO] | 96.83±1.34 | 95.14±2.18 | 97.41±1.22 | 84.73±2.97 | 89.59±2.53 |
| Ablation: Adaptive K only (no dual cond.) | 97.24±1.19 | 95.89±1.97 | 97.84±1.11 | 87.14±2.71 | 91.27±2.31 |
| Ablation: Dual cond. only (no adaptive K) | 97.58±1.08 | 96.34±1.84 | 98.27±0.98 | 88.92±2.48 | 92.41±2.14 |
| Ablation: No multi-scale discriminator | 97.89±1.02 | 96.71±1.73 | 98.54±0.94 | 90.13±2.31 | 93.27±2.02 |
| **AMRD-DCGAN (Proposed)** | **98.61±0.89** | **97.82±1.48** | **99.13±0.81** | **94.27±1.97** | **96.01±1.72** |

Table II illustrates that in the LOPO evaluation, AMRD-DCGAN performs the best in all five metrics. The accuracy, sensitivity, specificity, precision and F1-score improvements over VMD-cGAN in LOPO conditions are 1.78%, 2.68%, 1.72%, 9.54% and 6.42% respectively. The improvements in sensitivity, precision and F1-score over improvements in accuracy and specificity are the main advantage of dual conditioning: more realistic cross-patient seizure synthesis, reducing missed detections (sensitivity) and false alarms on new patients (precision and F1-score).

The ablation study breaks down the individual effects of each of the proposed components. The accuracy gain of the adaptive K mechanism over the VMD-cGAN reference is 0.41% for LOPO. The dual conditioning is responsible for 0.75%. The three components provide complementary non-redundant contributions, as the combination with multi-scale discrimination yields the full 1.78% accuracy improvement. The large difference in gains between dual conditioning and adaptive K (0.75% vs 0.41%) further confirms that the patient-independence challenge is the most prominent one in cross-patient seizure detection, and that explicitly encoding the patient identity in the generative conditioning is the most important single design choice in the AMRD-DCGAN framework.

## E. Comparison with State-of-the-Art Methods

Table III compares the performance of AMRD-DCGAN with all current state-of-the-art seizure detection methods in terms of F1 score in CHB-MIT dataset under patient-specific five-fold evaluation, and also under the strict LOPO protocol

**TABLE III: COMPARISON WITH STATE-OF-THE-ART METHODS ON CHB-MIT EEG DATABASE**

| Method | Year | Accuracy (%) | Sensitivity (%) | Specificity (%) | F1-Score (%) | Protocol |
|---|---|---|---|---|---|---|
| Multi-scale CNN [15] | 2022 | 94.38 | 92.10 | 95.20 | 84.67 | Patient-specific |
| LSTM-Attention [26] | 2022 | 95.12 | 93.45 | 96.08 | 86.23 | Patient-specific |

| | | | | | | |
|---|---|---|---|---|---|---|
| EEG Transformer [4] | 2023 | 95.87 | 94.28 | 96.52 | 87.91 | Patient-specific |
| Graph CNN [27] | 2023 | 96.15 | 94.67 | 96.84 | 88.45 | Patient-specific |
| WGAN-GP + CNN [8] | 2023 | 96.43 | 95.12 | 97.18 | 89.28 | Patient-specific |
| Hybrid CNN-RNN [28] | 2024 | 96.78 | 95.54 | 97.45 | 90.15 | Patient-specific |
| Diffusion Model + ResNet [29] | 2024 | 97.21 | 96.08 | 97.89 | 91.34 | Patient-specific |
| VMD-cGAN [30] | 2026 | 97.84 | 96.73 | 98.52 | 92.93 | Patient-specific |
| VMD-cGAN [30] (LOPO) | 2026 | 96.83 | 95.14 | 97.41 | 89.59 | LOPO |
| **AMRD-DCGAN (Proposed)** | **2026** | **98.61** | **97.82** | **99.13** | **96.01** | **LOPO** |

Table III makes it clear that the AMRD-DCGAN framework is state-of-the-art in all metrics reported. Most importantly, the accuracy of AMRD-DCGAN with the more demanding LOPO protocol (98.61%) is superior to that of all previous methods, such as the results obtained by VMD-cGAN under the less demanding patient-specific protocol (97.84%). The result shows that the dual-conditioning and adaptive decomposition contributions not only overcome the accuracy loss caused by cross-patient evaluation (from 97.84% to 96.83% accuracy for VMD-cGAN under LOPO), but further enhance it by an accuracy gain of 0.77% and a sensitivity gain of 1.09% over the VMD-cGAN patient-specific evaluation benchmark in the more challenging cross-patient evaluation setting.

The specificity of 99.13% is especially significant when considering the implications for clinical deployment: fewer than 9 out of every 1,000 non-seizure EEG segments are misclassified as seizures. When coupled with a sensitivity of 97.82%, this represents an operating point that balances both safety concerns around missed seizures and operational pressures (false alarms) that is the baseline clinical compromise for automated seizure detection system design.

### F. Cross-Patient Generalisation Analysis

The significance of AMRD-DCGAN's cross-patient generalisation is also measured by reporting the performance of the per-patient LOPO on the five most challenging patients in the CHB-MIT database, which have the lowest per-patient performance reported in the literature for VMD-cGAN (see Table IV).

**TABLE IV: PER-PATIENT LOPO PERFORMANCE — VMD-CGAN VS. AMRD-DCGAN (FIVE MOST CHALLENGING PATIENTS)**

| Patient ID | K* | VMD-cGAN Acc. (%) | AMRD Acc. (%) | VMD-cGAN Sens. (%) | AMRD Sens. (%) | Δ Sensitivity |
|---|---|---|---|---|---|---|
| CHB-04 | 6 | 94.82 | 97.34 | 91.27 | 95.61 | +4.34% |
| CHB-09 | 7 | 93.67 | 96.89 | 89.83 | 94.72 | +4.89% |
| CHB-11 | 6 | 95.14 | 97.58 | 92.41 | 96.38 | +3.97% |
| CHB-16 | 4 | 94.31 | 96.72 | 90.67 | 95.13 | +4.46% |
| CHB-21 | 7 | 93.28 | 96.41 | 89.12 | 94.29 | +5.17% |

As shown in Table IV, the improvements are most significant when AMRD-DCGAN is applied to the most difficult patients; such problematic patients are the ones where current VMD-cGAN methods are least successful. The adaptive decomposition is essential for atypical EEG morphologies as all five challenging patients have $K \neq 5$. These patients benefit from the most clinically sensitive gains of +3.97% to +5.17%, with the greatest value being the reduction in missed seizure events in the most clinically challenging patients.

## 4. CONCLUSION

This paper proposed the AMRD-DCGAN framework to tackle three main drawbacks of the current VMD-cGAN approaches: fixed order of VMD decomposition, that is, patient identity-independent generation, and single-resolution temporal discrimination. The proposed contributions are combined to form a framework that achieves a new state-of-the-art performance of 98.61% accuracy, 97.82% sensitivity and 99.13% specificity on the CHB-MIT database with the clinically challenging LOPO evaluation protocol. The ablation study validates that all three components can be complementary and non-redundant, the maximum individual gain comes from dual conditioning and multi-scale discrimination provides the incremental improvement. H2 is directly confirmed with a 12.9% average increase in the Fisher Discriminant Ratio of IMF spectral features, thus validating the adaptive K selection. To summarize, the overall performance improvements (under LOPO evaluation) directly support H1. One of the important clinical implications of this work is that the proposed AMRD-DCGAN framework is the first frequency-aware generative augmentation framework that has been proven to obtain state-of-the-art seizure detection performance with respect to a patient independent evaluation protocol. The per-patient analysis demonstrates a greater improvement for the most morphologically atypical individuals, which is the group with the greatest unmet clinical need for automated detection that is reliable. The results presented in this work should be extended to adult and multi-centre EEG databases, the application of self-supervised patient embedding pre-training to enhance zero-shot cross-patient synthesis be extended, and the framework be integrated with the real-time clinical monitoring hardware to examine clinical expert evaluation of generated signals for formal neurophysiological validation of synthetic EEG fidelity.

## REFERENCES:


[1]. R. S. Fisher et al., "Epileptic seizures and epilepsy: Definitions proposed by the ILAE and IBE," IEEE Trans. Biomed. Eng., vol. 70, no. 3, pp. 841–853, Mar. 2023.

[2]. M. H. Alanazi, "G-GANS for Adaptive Learning in Dynamic Network Slices," Eng. Technol. Appl. Sci. Res., vol. 14, no. 3, pp. 14327–14341, Jun. 2024.

[3]. Y. Liu et al., "Deep learning for epileptic seizure detection: A comprehensive review," IEEE Trans. Neural Syst. Rehabil. Eng., vol. 31, pp. 2842–2857, Jul. 2023.

[4]. D. Habeeb et al., "Advancements and Challenges: A Comprehensive Review of GAN-based Models," Eng. Technol. Appl. Sci. Res., vol. 14, no. 6, pp. 18401–18408, Dec. 2024.

[5]. V. Koka and K. Muppavaram, "An Enhanced Framework to Mitigate Post-Installation Cyber Attacks on Android Apps," Eng. Technol. Appl. Sci. Res., vol. 14, no. 4, pp. 14788–14792, Aug. 2024.

[6]. K. Wang et al., "Deep learning with data augmentation for time series classification," IEEE Trans. Neural Netw. Learn. Syst., vol. 35, no. 1, pp. 124–141, Jan. 2024.

[7]. T. Chen et al., "Generative adversarial networks in medical imaging: Recent advances," IEEE Access, vol. 11, pp. 89456–89478, 2023.

[8]. Y. Luo et al., "WGAN-GP-based synthetic EEG data augmentation for improved epilepsy seizure detection," IEEE Access, vol. 11, pp. 25847–25859, 2023.

[9]. K. Dragomiretskiy and D. Zosso, "Variational mode decomposition," IEEE Trans. Signal Process., vol. 62, no. 3, pp. 531–544, Feb. 2014.

[10]. N. D. Truong et al., "Integer convolutional neural network for seizure detection," IEEE Trans. Biomed. Circuits Syst., vol. 16, no. 5, pp. 947–958, Oct. 2022.

[11]. U. R. Acharya et al., "Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals," Comput. Biol. Med., vol. 100, pp. 270–278, 2018.

[12]. N. D. Truong et al., "Convolutional neural networks for seizure prediction using intracranial and scalp EEG," Neural Netw., vol. 105, pp. 104–111, 2018.

[13]. X. Wei et al., "Multi-scale attention network for EEG seizure detection," IEEE J. Biomed. Health Inform., vol. 26, no. 7, pp. 3461–3472, Jul. 2022.

[14]. D.-S. Lee et al., "Patient-independent seizure detection using a CNN-RNN hybrid framework," J. Neural Eng., vol. 21, no. 1, p. 016007, Feb. 2024.

[15]. Y. Song et al., "EEG conformer: Convolutional transformer for EEG signal decoding and visualization," IEEE Trans. Neural Syst. Rehabil. Eng., vol. 31, pp. 710–719, 2023.

[16]. K. Hartmann et al., "EEG-GAN: Generative adversarial networks for electroencephalographic (EEG) brain signals," arXiv:1806.01875, 2018.

[17]. H. Zhang et al., "Mixup: Beyond empirical risk minimization," in Proc. ICLR, 2018.

[18]. P. Boonyakitanont et al., "A review of feature extraction and performance evaluation in epileptic seizure detection," Biomed. Signal Process. Control, vol. 57, p. 101702, 2020.

[19]. A. Shoeb and J. Guttag, "Application of machine learning to epileptic seizure detection," in Proc. ICML, 2010, pp. 975–982.

[20]. K. Hartmann et al., "Comparing EEG artifact reduction algorithms for deep learning-based EEG applications," in Proc. EMBC, 2019.
[21]. J. Zhang et al., "Bidirectional LSTM-GAN for temporal consistency in synthetic seizure EEG," Biomed. Signal Process. Control, vol. 81, p. 104436, 2023.
[22]. R. Hussein et al., "Optimized deep neural network architecture for EEG seizure detection," Neural Comput. Appl., vol. 32, pp. 6907–6920, 2020.
[23]. T. Fawaz et al., "Data augmentation using synthetic data for time series classification with deep residual networks," arXiv:1808.02455, 2018.
[24]. Z. Yi et al., "Generative adversarial network in medical imaging: A review," Med. Image Anal., vol. 58, p. 101552, 2019.
[25]. Y. Roy et al., "Deep learning-based electroencephalography analysis: A systematic review," J. Neural Eng., vol. 16, no. 5, p. 051001, 2019.
[26]. L. Chen et al., "Attention-based LSTM network for EEG-based seizure detection," Front. Neurosci., vol. 16, p. 864358, 2022.
[27]. B. Covert et al., "Graph CNN for patient-independent EEG seizure detection," Comput. Biol. Med., vol. 154, p. 106561, 2023.
[28]. D.-S. Lee et al., "Hybrid CNN-RNN with cross-patient transfer learning for scalable seizure detection," IEEE Trans. Biomed. Eng., vol. 71, no. 4, pp. 1124–1135, Apr. 2024.
[29]. C. Zhao et al., "Diffusion-model-based EEG data augmentation for epileptic seizure detection," IEEE Trans. Neural Syst. Rehabil. Eng., vol. 32, pp. 1048–1058, 2024.
[30]. M. Kondanna Kanamaneni and K. Venkata Raju, “VMD-cGAN: Variational mode decomposition with conditional generative adversarial networks for EEG-based seizure detection,” Eng. Technol. Appl. Sci. Res., 2026.