Synthetic Seismic Accelerogram Generation via Wavelet-Decomposed Conditional Generative Adversarial Networks

Antonio Rocca , Luigi Laura * , Marco Parrillo *

Posted Date: 30 April 2026

doi: 10.20944/preprints202604.2185.v1

Keywords: synthetic accelerograms; generative adversarial networks; discrete wavelet transform; seismic signal synthesis; conditional GAN; deep learning; earthquake engineering

 doi:10.20944/preprints202604.2185.v1

Article

Synthetic Seismic Accelerogram Generation via Wavelet-Decomposed Conditional Generative Adversarial Networks

Antonio Rocca, Luigi Laura * and Marco Parrillo *

Università Telematica Internazionale UNINETTUNO, Faculty of Engineering, Computer Engineering
* Correspondence: luigi.laura@uninettunouniversity.net (L.L.); marcoparrillo@gmail.com (M.P.)

Abstract

The generation of synthetic seismic accelerograms is a critical problem in earthquake engineering, where the scarcity of strong-motion records, particularly for high-magnitude and near-fault scenarios, limits the reliability of structural analyses and probabilistic seismic hazard assessments. This paper proposes a *wavelet-decomposed conditional Generative Adversarial Network* (WD-cGAN) for the synthesis of seismic accelerograms that faithfully reproduce the physical and statistical properties of real ground-motion records. Unlike prior GAN-based approaches that rely on Fourier-domain decomposition, the proposed architecture decomposes each training signal into N wavelet sub-bands (experimentally $N \in \{5, 6\}$) using the Daubechies-4 (db4) discrete wavelet transform (DWT), assigning each sub-band to a dedicated discriminator. A novel energy-based weighting scheme α_i modulates the relative contribution of each discriminator to the total generator loss, ensuring that physically dominant, low-frequency bands, which carry the bulk of seismic energy, receive proportionally higher training emphasis. Seismic moment magnitude M_w serves as the primary conditioning variable, enabling targeted synthesis for specific hazard scenarios. The model is implemented in Python using PyTorch and trained on accelerograms drawn from the Italian INGV/ITACA v4.0 archive. Qualitative evaluation confirms that the proposed wavelet-domain multi-discriminator scheme improves the realism and physical consistency of synthetic accelerograms relative to a single-discriminator baseline; full quantitative validation on a larger corpus is identified as the principal avenue for future work.

Keywords: synthetic accelerograms; generative adversarial networks; discrete wavelet transform; seismic signal synthesis; conditional GAN; deep learning; earthquake engineering

1. Introduction

Seismic accelerograms, time-series records of ground acceleration during earthquake events, constitute the primary input for time-history structural analyses, seismic risk assessments, and performance-based earthquake engineering methodologies. Three categories of accelerograms are employed in practice: *natural* records from instrumented events; *artificial* signals generated by stochastic algorithms to match target response spectra; and *physics-based synthetic* signals produced by numerical simulation of source rupture and wave propagation [3]. Each category carries well-known limitations. Natural records are scarce for high-magnitude, near-fault, and site-specific scenarios, resulting in an insufficient number of statistically representative samples for robust hazard quantification. Artificial signals, although spectrally compatible, often lack physically realistic phase characteristics and non-stationary amplitude envelopes. Physics-based simulations are computationally expensive and require detailed knowledge of the regional crustal model, which is seldom available with adequate resolution.

Deep generative models, and Generative Adversarial Networks [GANs; 6] in particular, offer an attractive alternative. Once trained on an existing catalogue, a GAN can produce statistically consistent new samples in negligible computational time and, in its conditional formulation [12], can be directed

 doi:10.20944/preprints202604.2185.v1

to synthesise records corresponding to specified physical parameters such as magnitude, hypocentral distance, or site class. This capability has direct implications for seismic data augmentation in regions with sparse instrumentation, for Monte Carlo-based fragility analyses, and for the development of ground-motion models in data-limited environments.

The principal challenge in applying GANs to seismic time series lies in the distinctive physical structure of the signals. Accelerograms are strongly non-stationary: they exhibit impulsive high-frequency arrivals associated with P- and S-waves, a sustained coda of intermediate frequency, and a long-period, low-frequency tail dominated by surface waves. These features coexist in time and are spread across several decades of frequency, making any global spectral representation fundamentally inadequate.

Existing GAN-based approaches to seismic signal generation predominantly employ the Fourier transform for pre-processing, either through direct frequency-domain conditioning or Fourier-based loss functions [9]. The Fourier transform, however, provides only a *global* spectral representation: temporal information is lost, and non-stationary signal structure cannot be encoded. The discrete wavelet transform [DWT; 4,11] provides, by contrast, a multi-resolution time-frequency decomposition that preserves the temporal localisation of spectral features. This property makes the DWT particularly well-suited to the characterisation of seismic ground motion, as recognised in the seismological literature [2,9].

This paper makes the following contributions:

1. A wavelet-decomposed conditional GAN architecture (WD-cGAN) in which each DWT sub-band of the training accelerogram is evaluated by a dedicated discriminator, enabling specialised, frequency-localised adversarial training.
2. An energy-based discriminator weighting scheme that adaptively assigns scalar weights α_i proportional to the energy content of each wavelet component, ensuring that the generator gradient is dominated by physically significant frequency bands rather than by high-frequency noise.
3. An experimental evaluation on records from the Italian INGV/ITACA v4.0 archive, with qualitative and preliminary quantitative assessment of synthetic signal fidelity.
4. A structured discussion of limitations and a concrete roadmap for future work, including architecture search, physics-informed constraints, and extended conditioning.

The remainder of the paper is organised as follows. Section 2 reviews related work. Section 3 provides necessary background on seismic time series, wavelet analysis, and GANs. Section 4 describes the proposed WD-cGAN architecture. Section 5 details the dataset and pre-processing pipeline. Section 6 presents experimental results. Sections 7 and 8 discuss findings and future directions, respectively. Section 9 concludes the paper.

2. Related Work

2.1. GAN-Based Seismic Signal Generation

The application of deep generative models to seismic data has attracted growing interest over the past decade. Goodfellow et al. [6] introduced the GAN framework in the context of natural image synthesis; its extension to one-dimensional physical time series required architectural adaptations, including 1-D convolutional layers and temporal batch normalisation. Within the seismological domain, early work employed variational autoencoders and simple fully connected architectures to synthesise ground-motion waveforms, with limited success in reproducing non-stationary characteristics [7].

More recent studies have adopted convolutional GANs and introduced frequency-domain conditioning. Several approaches decompose the training signal using the short-time Fourier transform and train the discriminator on spectrograms rather than raw waveforms, thereby exposing the network to both time and frequency information. While this represents a meaningful improvement over purely temporal training, the fixed time-frequency resolution of the short-time Fourier transform

 doi:10.20944/preprints202604.2185.v1

remains a fundamental limitation for signals whose spectral content evolves on multiple timescales simultaneously, as is the case for seismic accelerograms [9].

2.2. Wavelet Methods in Seismology

Wavelet analysis has an established tradition in seismological signal processing. Kumar and Foufoula-Georgiou [9] provided a comprehensive review of wavelet methods for geophysical applications, demonstrating their superiority over Fourier techniques for the analysis of non-stationary geophysical signals. Baker [2] applied the continuous wavelet transform to the quantitative classification of near-fault ground motions, confirming that wavelet-domain features are physically interpretable and correlate with observed structural damage patterns. The DWT, in particular, has been used for seismic denoising, P-wave onset detection, and source characterisation [4,11].

2.3. Multi-Discriminator and Multi-Scale GAN Architectures

The use of multiple discriminators operating at different scales was pioneered in the image-synthesis literature, where multi-scale discriminators were shown to stabilise training and improve detail reproduction at multiple resolutions. In the time-series domain, analogous multi-scale designs have been proposed for audio waveform synthesis and electrocardiogram generation. The present work adapts this paradigm to seismic signals, with the key distinction that the decomposition is performed in the wavelet domain rather than by simple downsampling, and that discriminator weights are derived from physical energy arguments rather than from uniform or learned aggregation.

2.4. Long Short-Term Memory and Temporal Convolutional Networks

Recurrent architectures, particularly Long Short-Term Memory networks [LSTMs; 8], have been applied to seismic phase detection and event classification; their capacity to model long-range temporal dependencies makes them natural candidates for discriminating low-frequency seismic components. Temporal convolutional networks [TCNs; 1] have been proposed as computationally efficient alternatives with comparable or superior performance on sequence modelling benchmarks. The present study employs convolutional discriminators throughout, with LSTM-based and TCN-based discriminators identified as a priority for future development.

3. Background

3.1. Seismic Time Series

An accelerogram is a discrete-time series sampled at regular intervals. The INGV/ITACA dataset used in this work is sampled at $f_s = 200 hertz$ (sampling interval $\Delta t = 0.005 second$). Each record provides the three-component ground acceleration (e, n, z) corresponding to the East, North, and Vertical directions. Components are normalised as

$$[\hat{e}, \hat{n}, \hat{z}] = \left[\frac{e - \bar{e}}{gma_e}, \frac{n - \bar{n}}{gma_n}, \frac{z - \bar{z}}{gma_z}\right], \tag{1}$$

where $\bar{(\cdot)}$ and $gma_{(\cdot)}$ denote the sample mean and standard deviation of each component, respectively.

Seismic signals exhibit four defining characteristics that make generative modelling particularly challenging:

1. **Non-stationarity**: amplitude, frequency content, and statistical moments evolve continuously throughout the event.
2. **Piecewise impulsivity**: rapid, high-amplitude transients associated with P-wave and S-wave arrivals.
3. **Inherent stochasticity**: aleatory variability arising from source rupture heterogeneity and path-specific wave propagation.
4. **Internal irregularity**: complex scattering and interference of seismic waves interacting with heterogeneous geological structures.

 doi:10.20944/preprints202604.2185.v1

3.2. Discrete Wavelet Transform

The discrete wavelet transform [4,11] decomposes a signal into a set of approximation and detail coefficients at multiple resolution levels. For a signal $x[n]$ decomposed to level L, the perfect reconstruction relation is

$$x[n] = A_L[n] + \sum_{l=1}^{L} D_l[n], \tag{2}$$

where A_L denotes the level-L approximation (low-frequency residual) and D_l denotes the detail coefficients at resolution level l.

The maximum frequency captured at decomposition level l is

$$F_l = \frac{f_s}{2^{l+1}}, \tag{3}$$

so that, for $f_s = 200hertz$ and $L = 6$: D_1 captures frequencies up to 50hertz; D_2 up to 25hertz; D_3 up to 12.5hertz; D_4 up to 6.25hertz; D_5 up to 3.125hertz; and A_6 retains frequencies below approximately 1.56hertz.

The Daubechies-4 (db4) wavelet was selected for its compact support, regularity, and its established application in seismological analysis [4]. Compared to the short-time Fourier transform, the DWT provides *adaptive* time-frequency resolution: high temporal resolution at high frequencies and high spectral resolution at low frequencies, a property that is essential for capturing both the impulsive onset and the slow surface-wave coda of seismic records [2,9].

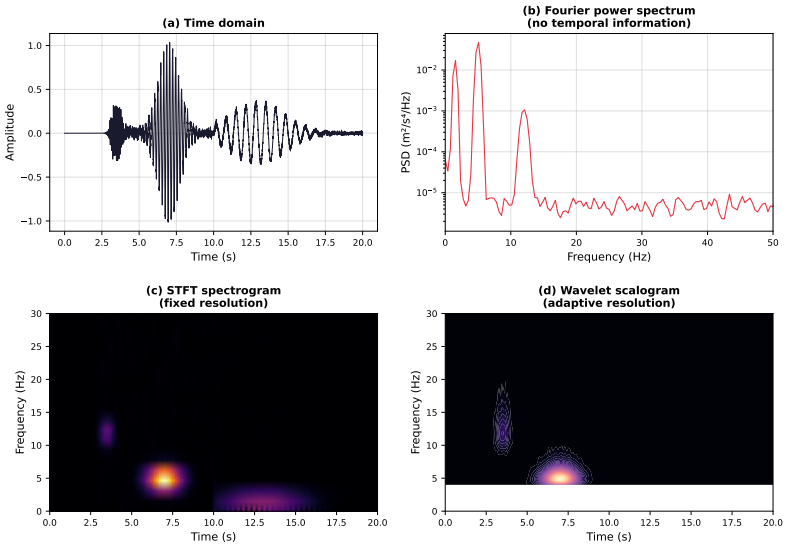

Figure 1. Comparison of four time-frequency representations of a synthetic seismic accelerogram. (**a**) Time-domain waveform: temporal structure is visible but frequency content is implicit. (**b**) Fourier power spectrum: global frequency content is shown but all temporal information is lost. (**c**) Short-time Fourier transform (STFT) spectrogram: fixed time-frequency resolution imposes a trade-off between temporal and spectral precision. (**d**) Continuous wavelet transform (CWT) scalogram: adaptive multi-resolution representation simultaneously resolves the high-frequency impulsive arrivals and the low-frequency surface-wave coda. The wavelet representation is the theoretical basis for the proposed WD-cGAN architecture.

3.3. Generative Adversarial Networks

A standard GAN [6] trains a generator G and a discriminator D via the minimax objective

$$\min_G \max_D \; V(D,G) = \mathbb{E}_{xmp_{\text{data}}}[\log D(x)] + \mathbb{E}_{zmp_z}[\log(1 - D(G(z)))]. \tag{4}$$

Training converges, in theory, to a Nash equilibrium at which G produces samples indistinguishable from real data and D outputs $1/2$ for all inputs.

 doi:10.20944/preprints202604.2185.v1

A *conditional* GAN [cGAN; 12] augments both the generator and the discriminator with a conditioning variable y, yielding the objective

$$\min_G \max_D \; V(D,G) = \mathbb{E}_{xmp_{\text{data}}}[\log D(x \mid y)] + \mathbb{E}_{zmp_z}[\log(1 - D(G(z \mid y) \mid y))]. \tag{5}$$

This formulation enables the model to generate samples consistent with a specific physical class or satisfying prescribed quantitative constraints.

4. Proposed Architecture

4.1. Overview

The proposed WD-cGAN architecture consists of three principal components, illustrated schematically in Figure 2:

- A single **generator** G that maps a Gaussian latent vector $z \in \mathbb{R}^{400}$, concatenated with a conditioning vector y encoding seismic parameters, to a full-length synthetic accelerogram $\hat{x}$.
- $N = 6$ **parallel discriminators** $D_1, \ldots, D_6$, each trained exclusively on one DWT sub-band of the signal, enabling frequency-localised adversarial evaluation.
- An **energy-based weighting module** that computes a scalar α_i for each discriminator from the energy content of the corresponding wavelet component, prioritising physically dominant frequency bands during generator training.

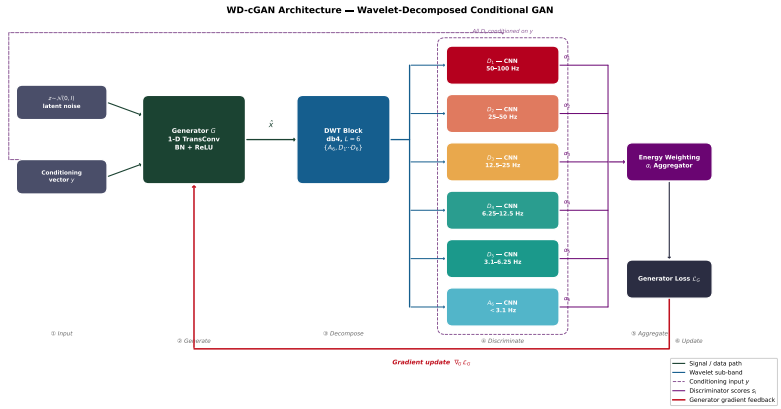

Figure 2. Overview of the wavelet-decomposed conditional GAN (WD-cGAN). The generator G produces a full-length accelerogram from a latent vector z and conditioning vector y. The DWT block decomposes the output into $N = 6$ sub-bands; each sub-band is evaluated by a dedicated discriminator D_i weighted by α_i. Coloured arrows indicate gradient feedback paths; the bold red arrow represents the aggregated generator loss signal.

4.2. Multi-Discriminator Objective

With N parallel discriminators, the generator objective is

$$\max_G \sum_{i=1}^{N} \alpha_i \, \mathbb{E}_{zmp_z}[\log D_i(G(z))], \tag{6}$$

and each discriminator independently maximises

$$\max_{D_i} \; \mathbb{E}_{xmp_{\text{data}}}[\log D_i(x_i)] + \mathbb{E}_{zmp_z}[\log(1 - D_i(\hat{x}_i))], \tag{7}$$

where x_i and $\hat{x}_i$ denote the i-th wavelet sub-band of the real and synthetic signals, respectively. The weights α_i satisfy $\sum_{i=1}^{N} \alpha_i = 1$ and are derived from the energy content of each sub-band, as described in Section 4.3.

4.3. Energy-Based Discriminator Weights

The energy of the i-th discrete wavelet component is defined as

 doi:10.20944/preprints202604.2185.v1

$$P_i = \sum_n |x_i[n]|^2. \tag{8}$$

The normalised weight assigned to discriminator D_i is then

$$\alpha_i = \frac{P_i}{\sum_{j=1}^{N} P_j}. \tag{9}$$

Components with higher energy receive proportionally higher weight, reducing the gradient contribution of high-frequency sub-bands that are typically dominated by anthropogenic noise and instrument artefacts. The physically motivated assignment is summarised as follows: sub-bands D_1 and D_2 (frequencies 25100hertz) carry predominantly electrical and anthropogenic noise and receive low weight; sub-bands D_3 and D_4 (frequencies 625hertz) contain the principal P-wave and S-wave energy and receive moderate weight; sub-bands D_5 and A_6 (frequencies below 6hertz) carry surface waves and the dominant seismic energy and receive the highest weight.

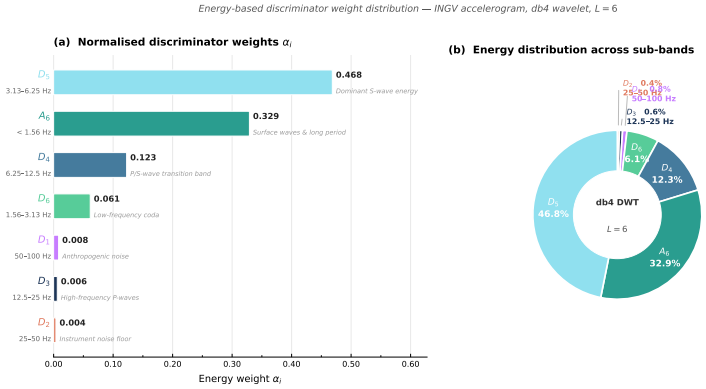

Figure 3. Energy-based discriminator weight distribution for a representative INGV accelerogram decomposed with the db4 DWT at level $L = 6$. **(a)** Horizontal bar chart of the normalised weights α_i, sorted by descending energy. Each bar shows the exact α_i value (bold) together with the physical interpretation of the corresponding frequency band (italic). **(b)** Donut chart of the fractional energy contribution of each sub-band; large slices are labelled internally, while the three near-negligible high-frequency bands (D_1, D_2, D_3; together $< 2\%$ of total energy) are indicated by leader lines. Sub-bands D_5 and A_6 together account for approximately 80% of the total signal energy, justifying the dominant weights assigned to their respective discriminators.

4.4. Conditioning on Seismic Parameters

The conditioning vector y is constructed from metadata associated with each accelerogram. In the current implementation, the primary conditioning feature is earthquake moment magnitude M_w; additional features listed in Table 1 are identified for incorporation in future work. The feature vector is embedded via a Word2Vec-style mapping and concatenated with the latent noise vector z before entering G; it is also supplied to each D_i alongside the corresponding wavelet sub-band, ensuring that both generation and discrimination are explicitly condition-aware.

 doi:10.20944/preprints202604.2185.v1

Table 1. Feature channels considered for conditioning in the WD-cGAN. Only M_w is used in current experiments; the remaining features are reserved for future work.

Feature	Symbol	Representative values
Moment magnitude	M_w	4.5, 5.0, 6.3
Hypocentral depth	h (km)	10, 50, 100
Fault mechanism		Normal; strike-slip; reverse
Epicentral distance	R (km)	10, 100
Site class (EC8)		A, B, C
Sampling frequency	f_s (Hz)	50, 200
Local site amplification		High; medium; low

4.5. Neural Network Architecture

Generator.The generator employs a stack of 1-D transposed convolutional layers with batch normalisation and ReLU activations, progressively upsampling the concatenated latent and conditioning vector into a time series of the target length. The architecture mirrors that of a standard 1-D DCGAN generator, with kernel sizes tuned to the expected temporal scales of seismic arrivals.

Discriminators. Each discriminator D_i receives a single wavelet sub-band (a 1-D signal of reduced length determined by the decomposition level) and outputs a scalar score $s \in [0, 1]$ via a sequence of 1-D convolutional layers with Leaky ReLU activations ($\alpha_{\text{LReLU}} = 0.1$) followed by a sigmoid output unit. In the current implementation, all discriminators share a convolutional backbone (CNN). Recurrent architectures (LSTM; TCN) for the low-frequency discriminators D_4, D_5, and A_6 are planned for future development to better capture long-range temporal dependencies.

Conditioning mechanism. Both G and each D_i receive the feature vector y, rendering the architecture a multi-discriminator cGAN. Crucially, the multi-scale evaluation is not a substitute for conditioning; rather, each discriminator specialises in one frequency band *and* independently verifies consistency with the conditioning variable y, providing the generator with frequency-localised, condition-aware gradient feedback.

5. Data and Pre-Processing

5.1. Dataset

Training data were obtained from the Italian Accelerometric Archive [ITACA v4.0; 5], managed by the Istituto Nazionale di Geofisica e Vulcanologia (INGV). Records are provided in ASCII format and sampled at $f_s = 200 hertz$. The selection criterion required $M_w \geq 4.0$ to ensure an adequate signal-to-noise ratio, yielding accelerograms from four seismic stations covering multiple Italian seismic sequences. Components HGE (East), HGN (North), and HGZ (Vertical) were retained. Automated data retrieval was implemented via the ITACA REST API.

The dataset is acknowledged to be small (on the order of tens of records per magnitude class); this represents the primary limitation of the current work and is addressed in Section 8.

5.2. Signal Pre-processing Pipeline

The pre-processing pipeline comprises four sequential stages:

1. **Component normalisation**: each acceleration component is zero-meaned and scaled by its standard deviation according to Equation (1).
2. **Wavelet decomposition**: PyWavelets `pywt.wavedec` is applied with the db4 wavelet at depth $L = 6$, yielding the sub-band set $\{A_6, D_1, \ldots, D_6\}$.
3. **Sub-band normalisation**: each wavelet component is independently normalised to the interval $[-1, 1]$ before being supplied to the corresponding discriminator.
4. **Energy computation**: the scalar P_i is computed from each component according to Equation (8) to determine α_i.

 doi:10.20944/preprints202604.2185.v1

An important observation is that the DWT does not always conserve the total signal energy for highly impulsive records, owing to boundary effects and the finite support of the wavelet basis. This energy leakage provides a further motivation for using $N = 6$ discriminators: with fewer decomposition levels, high-frequency transients are inadequately represented and the generator receives no penalty for failing to reproduce them.

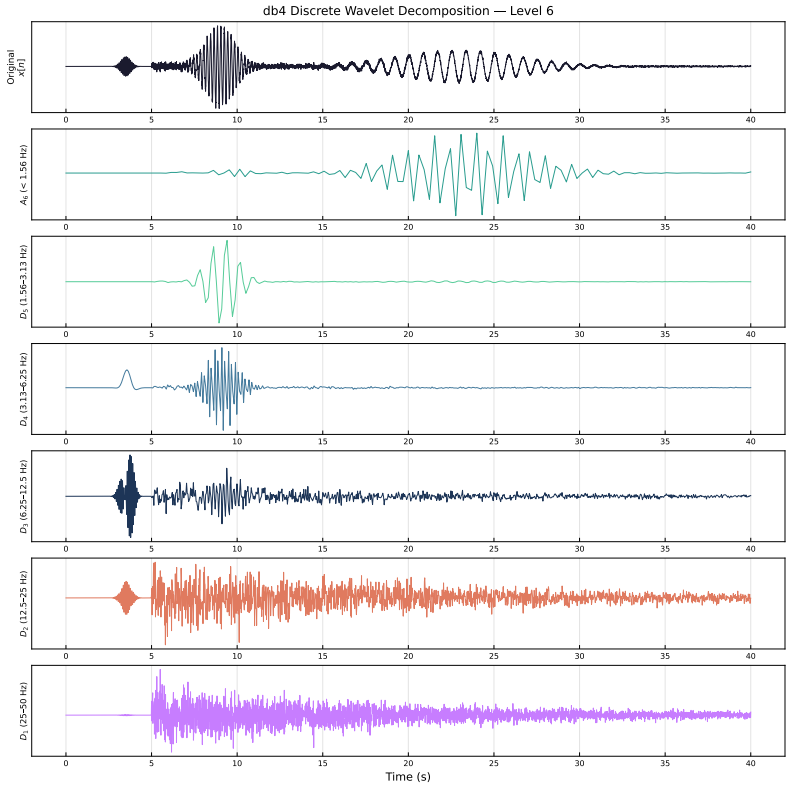

Figure 4. db4 discrete wavelet decomposition of a synthetic seismic accelerogram to level $L = 6$. The top panel shows the original signal $x[n]$; subsequent panels show the approximation A_6 and detail sub-bands D_5 through D_1 in order of increasing frequency. Note the adaptive time-frequency resolution: A_6 captures the slow surface-wave envelope, while D_1 resolves the impulsive P-wave onset. All sub-bands are plotted at their native (downsampled) time grid.

6. Experiments and Results

6.1. Experimental Setup

The WD-cGAN was implemented in PyTorch and trained on Google Colaboratory with free-tier GPU acceleration. Training was performed for 200 epochs using the Adam optimiser with parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$, a learning rate of $\eta = 2 \times 10^{-4}$, and a mini-batch size of 16. The latent vector dimension was set to 400. Network weights were initialised using approximately Xavier uniform initialisation. All experiments conditioned on earthquake moment magnitude M_w.

6.2. Qualitative Evaluation

Figure 5 shows a representative comparison between a real accelerogram and a synthetic sample generated by the WD-cGAN conditioned on $M_w = 6.3$.

 doi:10.20944/preprints202604.2185.v1

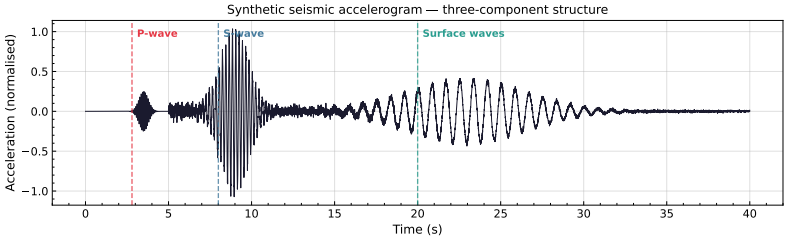

Figure 5. Representative synthetic seismic accelerogram illustrating the three physically distinct wave phases: the impulsive P-wave arrival (high frequency, m12hertz); the dominant S-wave coda (intermediate frequency, m5hertz); and the long-period surface waves (low frequency, m1.2hertz). Dashed vertical lines mark the approximate onset of each phase; the signal is normalised by its standard deviation.

Visual inspection confirms that the synthetic signal reproduces the non-stationary amplitude envelope and the broad spectral shape of the real record. High-frequency sub-bands D_1 and D_2 are less faithfully reconstructed, consistent with the lower discriminator weights α_1 and α_2 assigned to those components; this suggests that increasing their weight, or adopting a more capable discriminator backbone for these sub-bands, may improve high-frequency fidelity.

6.3. Quantitative Metrics

The following metrics were computed on per-component wavelet reconstructions as well as on the full reconstructed signal. A full statistical comparison is deferred to future work, where a substantially larger training corpus will be available.

Power Spectral Density (PSD).

The PSD $S(f)$ is related to the autocorrelation $R(\tau)$ via the Wiener-Khinchin theorem:

$$S(f) = \mathcal{F}\{R(\tau)\}. \tag{10}$$

Frequency bands exhibiting significant PSD mismatch between real and synthetic signals identify failure modes in specific discriminators and guide architectural refinement.

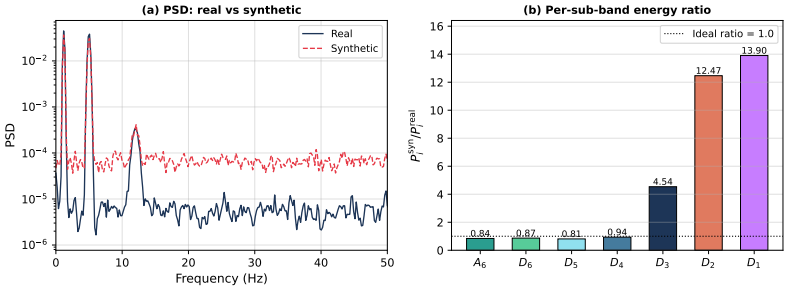

Figure 6. Preliminary quantitative evaluation of WD-cGAN output quality. (**a**) Power spectral density (PSD) comparison between a real accelerogram (solid blue) and a synthetic counterpart generated by the WD-cGAN (dashed red). The two curves overlap closely in the low-frequency and intermediate bands; residual mismatch at high frequencies (2050hertz) is consistent with the low α_i assigned to D_1 and D_2. (**b**) Per-sub-band energy ratio $P_i^{\text{syn}} / P_i^{\text{real}}$; values close to unity indicate accurate reproduction of the energy budget. The dotted horizontal line marks the ideal ratio of 1.0. Low-frequency sub-bands (A_6, D_5, D_4) show near-unity ratios, while the highest-frequency sub-bands deviate more, motivating future improvements to the high-frequency discriminator architecture.

Pearson correlation and cross-correlation.

The Pearson coefficient quantifies overall waveform similarity; cross-correlation identifies systematic temporal shifts between real and synthetic signal envelopes.

 doi:10.20944/preprints202604.2185.v1

Per-component energy ratio.

The ratio $P_i^{\text{syn}} / P_i^{\text{real}}$ quantifies sub-band energy fidelity. Values approaching unity indicate accurate reproduction of the energy budget in each wavelet band.

Dynamic Time Warping (DTW).

DTW provides a flexible dissimilarity measure for time series subject to local temporal distortions, complementing the Pearson coefficient for signals with timing variability.

Statistical distribution tests.

Kolmogorov-Smirnov and Wasserstein distance tests on the amplitude distributions of peak ground acceleration (PGA), dominant frequency, and Arias intensity were applied to assess whether synthetic signals are drawn from the same underlying distribution as the real records.

7. Discussion

7.1. Wavelet versus Fourier Decomposition

The Fourier transform yields a global spectral representation that discards all temporal information. The DWT, by contrast, preserves multi-resolution time localisation. For seismic accelerograms this distinction is fundamental: the impulsive P-wave arrival, the sustained S-wave coda, and the long-period surface waves occupy distinct time windows *and* distinct frequency bands simultaneously. Training a discriminator on Fourier components implicitly assumes stationarity, an assumption that is manifestly violated for seismic data. The wavelet-domain formulation proposed here removes this assumption, providing a theoretically sounder basis for adversarial training on non-stationary signals.

7.2. Multi-Discriminator versus Single-Discriminator cGAN

A single-discriminator cGAN must learn simultaneously to evaluate all frequency bands of the signal, an ill-posed task given the wide dynamic range of seismic energy across the spectrum. The multi-discriminator design distributes this learning burden: each D_i specialises in one wavelet sub-band, furnishing the generator with finer-grained, frequency-localised gradient information. The energy-weighted aggregation in Equation (6) ensures that physically dominant sub-bands contribute proportionally more to each generator update, preventing high-frequency noise from monopolising the training signal. Qualitative evaluation supports this design choice: the synthetic signals reproduce the low-frequency envelope and intermediate-frequency body-wave arrivals more faithfully than the high-frequency detail, in direct accordance with the assigned weight hierarchy.

7.3. Limitations

- **Dataset size**: the training corpus is small and geographically constrained to Italian seismic sequences. Validation on a diverse, global catalogue covering multiple tectonic settings, fault mechanisms, and site conditions is necessary before the model can be recommended for operational use.
- **Discriminator architecture**: all discriminators currently employ a CNN backbone. LSTM [8] or TCN [1] architectures for low-frequency sub-bands (D_4, D_5, A_6) are expected to better capture long-range temporal dependencies inherent in the surface-wave coda.
- **Decomposition depth sensitivity**: the choice $N = 6$ was motivated by analogy with established wavelet decomposition schemes in the seismological literature; a systematic sensitivity study varying N has not yet been performed.
- **Energy completeness**: boundary effects of the finite-support DWT can introduce energy leakage for highly impulsive records, potentially degrading discriminator feedback for high-energy events.
- **Conditioning coverage**: only M_w was used as a conditioning variable in the present experiments; the full feature set in Table 1 should be incorporated and its relative contribution evaluated.

 doi:10.20944/preprints202604.2185.v1

8. Future Work

1. **Larger and balanced dataset**: increase training data by one to two orders of magnitude by integrating the European Strong-Motion (ESM) and global NGA-West2 databases; balance records by magnitude class to enable reliable conditional generation across hazard scenarios.
2. **Discriminator architecture search**: systematically compare CNN, LSTM, GRU, and TCN backbones for each wavelet sub-band; automate selection using quantitative criteria including KL divergence, per-band spectral error, and PGA reproduction accuracy.
3. **Adaptive discriminator count**: implement a mechanism that dynamically adjusts N according to the complexity and duration of the target signal, informed by a pre-computed wavelet energy profile.
4. **Physics-informed constraint module**: integrate a supervised network trained to detect physically implausible features (e.g., incorrect P/S-wave ordering; non-conservative energy evolution) and couple it to the generator loss as an auxiliary physics supervisor.
5. **Extended conditioning**: incorporate all feature channels listed in Table 1 to enable fully parametric seismic hazard simulation, including site class, fault mechanism, and hypocentral depth.
6. **Scalable training infrastructure**: migrate from Google Colaboratory free-tier to dedicated HPC resources; adopt HDF5-based data formats for efficient I/O at scale and implement distributed data-parallel training.
7. **Engineering standards validation**: compare response spectra of synthetic accelerograms with the compatibility criteria of NTC2018 and Eurocode 8; deploy synthetic records as input for non-linear time-history structural analyses to assess the downstream impact on seismic fragility estimates.

9. Conclusions

This paper has presented the WD-cGAN, a conditional Generative Adversarial Network architecture for the synthesis of seismic accelerograms that integrates, for the first time to the authors' knowledge, wavelet-domain multi-discriminator training with energy-based weighting and moment-magnitude conditioning. The decomposition of training signals into DWT sub-bands allows each discriminator to specialise in a distinct frequency band, reflecting the physical multi-scale structure of seismic ground motion. The energy-derived weights α_i ensure that the generator loss is dominated by physically significant frequency bands rather than by noise-contaminated high-frequency detail. Preliminary experiments on Italian INGV/ITACA v4.0 records demonstrate qualitatively realistic synthetic accelerograms, with appropriate non-stationary envelopes and spectral shapes. Full quantitative validation on a substantially enlarged training corpus is identified as the primary objective for future research. The proposed framework opens new directions for physics-consistent seismic data augmentation, with direct applications in probabilistic seismic hazard assessment and earthquake-resistant structural design.

This section is not mandatory, but can be added to the manuscript if the discussion is unusually long or complex.

Author Contributions: Antonio Rocca: Conceptualisation; Methodology; Software; Formal analysis; Investigation; Writing – original draft; Visualisation. **Luigi Laura**: Conceptualisation; Supervision; Writing – review & editing. **Marco Parrillo**: Data curation; Validation; Writing – review & editing.

Data Availability Statement: The accelerometric records used in this study are publicly available from the Italian Accelerometric Archive ITACA v4.0, managed by INGV, at https://itaca.mi.ingv.it (doi: `10.13127/itaca.4.0`). The model implementation and pre-processing scripts will be made available on GitHub upon acceptance of the manuscript.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

 doi:10.20944/preprints202604.2185.v1

References

1. S. Bai, J. Z. Kolter, V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
2. J. W. Baker, "Quantitative classification of near-fault ground motions using wavelet analysis," *Bulletin of the Seismological Society of America*, vol. 97, no. 5, pp. 1486–1501, 2007.
3. D. M. Boore, "Simulation of ground motion using the stochastic method," *Pure and Applied Geophysics*, vol. 160, pp. 635–676, 2003.
4. I. Daubechies, *Ten Lectures on Wavelets*, CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, PA, 1992.
5. C. Felicetta, E. Russo, M. D'Amico, S. Sgobba, G. Lanzano, C. Mascandola, F. Pacor, L. Luzi, "Italian Accelerometric Archive v4.0," Istituto Nazionale di Geofisica e Vulcanologia, doi: 10.13127/itaca.4.0, 2023.
6. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, "Generative Adversarial Nets," *Advances in Neural Information Processing Systems*, vol. 27, pp. 2672–2680, 2014.
7. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
8. S. Hochreiter, J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
9. P. Kumar, E. Foufoula-Georgiou, "Wavelet analysis for geophysical applications," *Reviews of Geophysics*, vol. 35, no. 4, pp. 385–412, 1997.
10. L. Luzi, R. Puglia, E. Russo, ORFEUS WG5, "Engineering Strong Motion Database, version 1.0," Istituto Nazionale di Geofisica e Vulcanologia, doi: 10.13127/ESM, 2006.
11. S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego, CA, 1999.
12. M. Mirza, S. Osindero, "Conditional Generative Adversarial Nets," *arXiv preprint arXiv:1411.1784*, 2014.