

Time-Domain GAN Compression of Intracranial EEG with Latent Quantization

Original

Time-Domain GAN Compression of Intracranial EEG with Latent Quantization / Depaoli, F., Macii, E., Patti, E., Ponzio, F., Aliberti, A.. - ELETTRONICO. - (2026), pp. 910-915. (2026 IEEE 39th International Symposium on Computer-Based Medical Systems (CBMS) Limassol (CYP) June 03-05, 2026) [10.1109/CBMS69103.2026.00187].

Availability:

This version is available at: 11583/3015880 since: 2026-10-12T07:43:31Z

Publisher:

IEEE

Published

DOI:10.1109/CBMS69103.2026.00187

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Time-Domain GAN Compression of Intracranial EEG with Latent Quantization

Fabio Depaoli
DAUIN

Enrico Macii
DIST

Edoardo Patti
DAUIN

Francesco Ponzio
DAUIN

Alessandro Aliberti
DIST

Politecnico di Torino
Turin, Italy

Politecnico di Torino
Turin, Italy

Politecnico di Torino
Turin, Italy

Politecnico di Torino
Turin, Italy

Politecnico di Torino
Turin, Italy

fabio.depaoli@polito.it enrico.macii@polito.it edoardo.patti@polito.it francesco.ponzio@polito.it alessandro.aliberti@polito.it

Abstract—Intracranial electroencephalographic (iEEG) signals are typically captured using matrices with a high density of electrodes. This makes compressing these signals a valid objective for reducing memory usage in storage and transmission. Among the available compression methods, generative adversarial networks (GANs) have not been used. GANs are primarily associated with image generation and image-related applications. They can capture the probability distribution of variables in a dataset. Compression algorithms have minimally explored this characteristic, especially when applied to time series.

We present an algorithm for multichannel signal compression based on a GAN, and we prove its efficacy on a dataset of intracranial electroencephalography (iEEG) signals. The compression method modifies the Backpropagation GAN (BPGAN) algorithm to work with signals instead of images. To get around the difficulties introduced by the peculiarities of the specific application, we propose an algorithm that separates channels of the iEEG signal into groups with high correlation and synchronization. The algorithm compresses groups separately by selecting a latent representation for each signal window among a finite set of possible representations. To validate the results, we decided to compare the results of an epilepsy seizure recognition experiment on the original and reconstructed signal.

This approach presents encouraging results regarding compression ratio despite the introduction of some artifacts in the reconstruction. This could open for future improvements in applying GANs in compressing signals by increasing the reconstruction quality while keeping the high compression ratio.

Index Terms—Generative Adversarial Networks, Compression, Intracranial EEG, Alternative direction method of multipliers, Mean phase coherence, Hjorth’s parameters

I. INTRODUCTION

Intracranial EEG (iEEG) is routinely used to localize epileptogenic tissue and brain tumors by recording activity directly from cortical or deep structures [1]. It is also valuable for BCIs and neuroscience [2]. Modern systems can record up to 256 channels at high sampling rates [3], producing data volumes that make long-term storage, online monitoring, and wireless transmission difficult.

Compression is particularly challenging because iEEG is high-resolution in time and space and has much larger amplitudes than scalp EEG. With limited skull filtering, it preserves clinically important high-frequency components linked to neuronal firing and hemodynamics [1]. Signals are also

more localized [4] and have far higher SNR [5], so aggressive compression risks removing diagnostically relevant structure.

While deep learning has shown promise for biosignal and scalp EEG compression [6], iEEG-specific methods remain limited, especially approaches that work directly in the time domain and explicitly account for inter-channel dependencies—both critical for tasks such as seizure detection and localization. GANs are a plausible alternative because adversarial training can model complex distributions and preserve fine detail [7], although their known instability and difficulty with rare events make their use in compression exploratory.

To the best of our knowledge, this is the first study investigating time-domain GAN-based compression for intracranial EEG with task-level clinical validation by adapting BPGAN [8] to iEEG. We introduce a correlation/synchrony-based channel grouping strategy to compress strongly interacting channels together, and evaluate reconstruction using standard signal measures (including Hjorth parameters [9]) and seizure-detection performance on reconstructed data. Experiments are conducted on a single long-term configuration from a public dataset, and should be viewed as an initial assessment.

We address three questions: the capability of GAN-based compression to operate in the time domain on high-resolution iEEG; which is the effect of grouping channels; and how does compression affect seizure detection. Our contributions are: (i) a GAN-based time-domain iEEG compressor, (ii) a grouping method to preserve dependencies, (iii) an RNN-based BPGAN adaptation for sequential signals, and (iv) an evaluation combining distortion metrics with a clinically relevant task. The remainder of the paper details the method (Section II), results (Section III), discussion (Section IV), and conclusions (Section V).

II. METHODS

The main motivation behind our methodology is that GANs can learn the data distribution and encode samples into compact latent codes that act as informative representations [10]–[12]. Here, we optimize a latent code for each iEEG segment and use it as a compressed, distribution-aware description of the signal.

This is challenging because iEEG has high resolution, large amplitudes, and correlations that decrease with electrode

distance, which can make joint multi-channel training prone to overfitting [1].

Figure 1 summarizes our pipeline, with training on the top and inference on the bottom. During training, we (i) group channels to reduce distance-related decorrelation, (ii) train a time-domain BPGAN for each group, and (iii) learn a discrete latent dictionary used in an optimization–quantization step to select compact codes for reconstruction and compression (Sections II-A–II-C).

At inference, we apply the same grouping, reuse the trained BPGANs and dictionaries, and run the same optimization–quantization procedure to choose the codes for compressing new recordings.




Fig. 1. Overview of the proposed GAN-based iEEG compression framework. During training (left), the channel-grouping algorithm assigns each signal channel to one of several groups. Each group is then used to train a dedicated BPGAN model, and a latent vector dictionary is built for that group from the encoder outputs. During testing (right), the incoming iEEG is partitioned into the same channel groups and passed through the corresponding trained BPGAN compressors to obtain compact latent codes (compression phase). These codes are then mapped back to latent vectors via the dictionaries and decoded by the trained generators to reconstruct the multichannel signal (reconstruction phase).

A. Channel grouping

Our algorithm does not compress the whole multichannel signal at once. Instead, it operates on groups of channels that are processed independently. The goal of this grouping step is to collect channels that exhibit similar dynamics, so that each group is as internally coherent as possible. To this end, we maximize two complementary measures: the inter-channel correlation coefficient and a synchronization index called mean phase coherence [13].

Formally, let the original recording be an N -dimensional signal, where N is the total number of channels. We partition these channels into n groups, each containing C channels,

with C treated as a hyperparameter, as later detailed. The i -th group is represented by a set S_i that lists the channel indices assigned to that group. The grouping is performed once, as a preprocessing step, using the first hour of recordings, and is then kept fixed: the same channel–group association is used during both training and testing (see the group-creation block in Figure 1). After grouping, a separate BPGAN model is trained for each group.

To quantify similarity between channels a and b over a given segment of length T_c samples, we first compute the (Pearson) cross-correlation coefficient ρ_{ab} ,

$$\rho_{ab} = \frac{\frac{1}{T_c} \sum_{t=0}^{T_c-1} [x_a(t) - \mu_a][x_b(t) - \mu_b]}{\sigma_a \sigma_b}, \quad (1)$$

where μ_a and μ_b are the sample means of the two segments and σ_a and σ_b their sample standard deviations, respectively.

As a complementary measure, we evaluate the mean phase coherence R_{ab} ,

$$R_{ab} = \left| \frac{1}{T_c} \sum_{t=0}^{T_c-1} \exp(i\phi_{ab}(t)) \right| = 1 - CV_{ab}, \quad (2)$$

where i is the imaginary unit, $\phi_{ab}(t)$ is the instantaneous relative phase between the two channels, CV_{ab} is the circular variance of its angular distribution, and

$$\begin{aligned} \phi_{ab}(t) &= \phi_a(t) - \phi_b(t) \\ &= \arctan \frac{\tilde{x}_a(t)x_b(t) - x_a(t)\tilde{x}_b(t)}{x_a(t)x_b(t) + \tilde{x}_a(t)\tilde{x}_b(t)}. \end{aligned} \quad (3)$$

Here, $\tilde{x}_i(t) = \mathcal{H}[x_i(t)] = -i\mathcal{F}^{-1}\{\mathcal{F}[x_i(t)](\omega) \text{sign}(\omega)\}$ denotes the Hilbert transform of $x_i(t)$ for $i \in \{a, b\}$, $\mathcal{F}$ is the Fourier transform, and ω is the angular frequency.

Algorithm 1 describes the exact procedure to create the groups. The resulting partition $\{S_1, \dots, S_n\}$ yields groups with high internal correlation and phase synchronization, which are then used unchanged throughout training and testing. The number of groups n (equivalently, the group size C) is a hyperparameter that controls the trade-off between modeling capacity per group and the amount of inter-channel structure captured within each BPGAN. An empirical analysis of how the system behaves as a function of this parameter is reported in Section III.

B. BPGAN structure and training

BPGAN consists of an encoder E , a generator G , and a discriminator D . At compression time (Figure 2(a)), each segment $x \in \mathbb{R}^{T \times C}$ is mapped by the 1D-CNN encoder to a latent sequence $z_0 \in \mathbb{R}^{T \times L}$. The BiLSTM-based generator takes a latent sequence z and reconstructs $\hat{x} = G(z) \in \mathbb{R}^{T \times C}$, while the discriminator receives either x or $\hat{x}$ and outputs a scalar probability of being real.

The encoder output z_0 initializes an optimization–quantization step used for compression: we seek a latent sequence z that minimizes reconstruction error while being constrained to a finite dictionary of admissible codes. We solve this constraint via ADMM [14]. Each dictionary

Algorithm 1 Channel grouping procedure

Require: Ordered list of N channels (by electrode position), number of groups n (and consequently the number C of channels in each group), sampling frequency f_s

Ensure: Groups $S_1, \dots, S_n$ forming a partition of the N channels

```

1: Initialize assigned set  $A \leftarrow \emptyset$ 
2: for  $i = 1$  to  $n$  do
3:   Initialize group  $S_i \leftarrow \emptyset$ 
4: end for
5: Seed selection: choose one channel every  $C$  entries from the ordered
   list and assign one seed to each  $S_i$ 
6: Add all selected seeds to  $A$ 
7: Group construction window: use the first hour of recordings
8: Divide each channel into 10 s segments, with  $T_c = 10f_s$ 
9: Let  $n_k$  be the total number of segments in one hour
10: while  $|A| < N$  do
11:   for all channels  $j \notin A$  do
12:     for  $i = 1$  to  $n$  do
13:       for all  $l \in S_i$  do
14:         for  $k = 1$  to  $n_k$  do
15:           Compute  $\rho(k, l, j, i)$  and  $R(k, l, j, i)$ 
16:         end for
17:          $\bar{\rho}(l, j, i) \leftarrow \frac{1}{n_k} \sum_{k=1}^{n_k} \rho(k, l, j, i)$ 
18:          $\bar{R}(l, j, i) \leftarrow \frac{1}{n_k} \sum_{k=1}^{n_k} R(k, l, j, i)$ 
19:       end for
20:        $I(j, i) \leftarrow \max_{l \in S_i} |\bar{\rho}(l, j, i)| + \max_{l \in S_i} \bar{R}(l, j, i)$ 
21:     end for
22:      $i^* \leftarrow \arg \max_{i \in \{1, \dots, n\}} I(j, i)$ 
23:      $S_{i^*} \leftarrow S_{i^*} \cup \{j\}; \quad A \leftarrow A \cup \{j\}$ 
24:   end for
25: end while

```




Fig. 2. BPGAN compressor (a) and reconstructor (b) for a segment belonging to one group of channels.

element can be represented by a binary string $2z$, defining the codebook (Figure 2(b)).

Before inference, E , G , and D are trained adversarially (Figure 3). The encoder-generator pair is optimized to both reconstruct x and "fool" D , whereas D is trained to distinguish real segments from reconstructions. We follow the standard two-phase BPGAN scheme: (i) train with continuous latents (encoder output fed directly to G without optimization-quantization) using Adam; (ii) discretize the latent space by clustering encoder outputs with MiniBatchKMeans [15], using the centroids as the dictionary. Choosing n bits yields 2^n centroids and directly controls the compression ratio. In phase (ii) the encoder is frozen and G and D are fine-tuned with a linearly decreasing learning rate while latents are constrained to the dictionary.




Fig. 3. Complete BPGAN structure during training. Blue arrows represent the forward pass. Light green arrows indicate the backpropagation signals used to train the encoder and generator from the discriminator's classification. Dark green arrows denote the backpropagation used to train the discriminator as a binary classifier. The red arrow represents the optimization-quantization step used to build the latent dictionary, i.e., the association between each binary code $2z$ and its corresponding latent vector z .

The encoder/generator loss combines an adversarial term with reconstruction penalties,

$$\begin{aligned}
\mathcal{L}_{E/G}(E, G) = & -\frac{1}{B} \sum_{i=1}^B \log D(G(E(x_i))) \\
& + \alpha \frac{1}{BC} \sum_{i=1}^B \sum_{j=1}^C \|x_{ij} - G(E(x_i))_j\|_2^2 \\
& + \beta \frac{1}{B} \sum_{i=1}^B \max_{j=1, \dots, C} \|x_{ij} - G(E(x_i))_j\|_\infty,
\end{aligned} \quad (4)$$

while the discriminator is trained with the standard GAN objective,

$$\mathcal{L}_D(D) = -\frac{1}{B} \sum_{i=1}^B \log D(x_i) - \frac{1}{B} \sum_{i=1}^B \log(1 - D(G(E(x_i)))) \quad (5)$$

As usual, we alternate updates of (E, G) and D ; in second phase only G and D are updated.

Finally, we omit the additional Huffman coding used in the original BPGAN, since for iEEG the latent-code distribution is nearly uniform and Huffman coding yields negligible gains [16].

C. Optimization-quantization algorithm

Once the encoder, generator, and discriminator have been trained and the latent dictionary has been built, compressing a segment reduces to the following task: for each input x , we must choose a latent code z that (i) allows the generator to reconstruct x accurately and (ii) belongs to the finite set S of admissible latent vectors (the dictionary). Following the original BPGAN formulation, this *optimization-quantization* problem can be written as:

$$\begin{aligned}
& \text{minimize} && f(z) + I(u \in S) \\
& \text{subject to} && u = z,
\end{aligned} \quad (6)$$

where

$$f(z) = \frac{1}{C} \sum_{i=1}^C \|x_i - G(z)_i\|_2^2$$

is the reconstruction error (with C the number of channels in the group, x the segment represented by z , and x_i its i -th channel before reconstruction, and $G(z)_i$ the i -th channel

after reconstruction), and $I(u \in S)$ is the indicator function of the dictionary S , equal to 0 if $u \in S$ and $+\infty$ otherwise. Although the generator input z is a matrix in $\mathbb{R}^{T \times L}$, in this subsection we consider its flattened version in $\mathbb{R}^{TL}$.

Problem (6) has the standard ADMM form [14] with

$$g(u) = I(u \in S), \quad A = I, \quad B = -I, \quad c = 0_{TL},$$

where I is the $TL \times TL$ identity matrix and 0_{TL} is the all-zero vector in $\mathbb{R}^{TL}$. The corresponding augmented Lagrangian is

$$L_\mu(z, u, \eta) = f(z) + I(u \in S) + \frac{\mu}{2} (\|z - u + \eta\|_2^2 - \|\eta\|_2^2), \quad (7)$$

where η is the scaled dual variable and $\mu > 0$ is the penalty parameter.

In practice, we implement the ADMM iterations as

$$z(k+1) = z(k) - \ell \nabla_z F(z, u(k), \eta(k)), \quad (8)$$

$$u(k+1) = \text{quantization}(z(k+1) + \eta(k)), \quad (9)$$

$$\eta(k+1) = \eta(k) + z(k+1) - u(k+1), \quad (10)$$

where

$$F(z, u, \eta) = f(z) + \frac{\mu}{2} \|z - u + \eta\|_2^2,$$

and $\ell > 0$ is a stepsize controlling the update of z . The z -update (8) performs a gradient step on F with respect to z ; the u -update (9) projects (quantizes) $z(k+1) + \eta(k)$ onto the discrete set S by selecting the nearest cluster centroid in the latent dictionary; finally, the η -update (10) corresponds to the usual dual ascent step in ADMM.

This scheme mirrors the structure of classical ADMM while adapting it to our non-convex setting, where f is defined by a deep generator and S is a finite (discrete) set. Although ADMM is traditionally studied in convex settings, in our framework it acts as a practical latent-projection heuristic, empirically converging in a small fixed number of iterations of (8)–(10).

III. EXPERIMENTAL RESULTS

This section reports the results obtained for the best-performing models, whose hyperparameters were selected empirically through exploratory tuning based on reconstruction quality and training stability. We considered groups of 4 channels with a latent space represented with 13 bits and groups of 8 channels with a latent space represented with 13 or 15 bits. The first subsection describes the dataset and experimental setting. The second subsection introduces the performance indexes used to evaluate the proposed methodology. The remaining subsections present and discuss the experimental results.

A. Dataset description and experimental setting

We use the public long-term SWEC-ETHZ human iEEG dataset [17], collected from 18 epilepsy-surgery patients at Inselspital Bern. Signals were recorded at 512 or 1024 Hz with a 16-bits ADC and bandpass filtered (4th-order Butterworth, 0.5–120 Hz). Recordings last 41–293 h and include 32–128 channels per subject, originally annotated for seizure identification.

Experiments were run on a PC with Intel i7-11700K, 32 GB RAM, and Nvidia RTX 3060Ti. Training is computationally expensive (about 140 h for one 8-channel group with a 15-bits code), so we focus on a representative 32-channel high-resolution configuration as a proof of concept.

The pipeline is implemented in Python: networks are trained in TensorFlow, ADMM is custom, and MiniBatchKMeans is from scikit-learn. Before training, each channel is linearly scaled to $[-1, 1]$ using bounds (m, M) estimated from the training set, and the same parameters are reused at test time, with an inverse rescaling applied after reconstruction.

B. Experimental results

We evaluate three models by varying the number of channels per group and the latent bit budget. All experiments use $T = 32$ samples (1/32 s) and the same 32-channel dataset split: 20 h training, 5 h validation, and 10 h test.

Table I reports compression and quality. The best RCR (99.68%) is obtained with 8 channels per group and 13 bits; as expected, RCR decreases with more bits and increases with fewer groups. Average reconstruction metrics are similar across models, with the best PRD (22.40%) and PSNR (34.78 dB) again achieved by the 8-channels, 13-bits configuration.

TABLE I
COMPRESSION AND RECONSTRUCTION QUALITY FOR THE THREE MODELS.

Model	CR	RCR	PRD	PSNR
4 channels-13 bits	157.54	99.36%	23.27%	34.70 dB
8 channels-13 bits	315.08	99.68%	22.40%	34.78 dB
8 channel-15 bits	273.07	99.63%	23.30%	34.41 dB

Table II shows that errors vary by frequency band, with the δ band consistently worst (up to 6.11% mean gap for the 4-channel model, versus $\approx 2.6\%$ for the 8-channel models). Channel-level PRD variability (Tables III–IV) indicates that poorly reconstructed groups also exhibit larger within-group dispersion, most notably for the 4-channel model and for the last group in the 8-channel models.

Despite good mean metrics, pointwise errors can be severe (Table V), with relative errors exceeding 50% and occasional extreme peaks; the 15-bits model improves the first groups but worsens the last one. Hjorth features (Table VI) further suggest reduced reconstructed power (lower activity), increased mobility, and a tendency for higher bit budgets to lower complexity, consistent with smoother dynamics.

C. Seizure detection validation

To evaluate whether the reconstruction preserves clinically relevant information, we apply the epileptic seizure detection algorithm presented in [18]. Following the protocol in [17], we train the detector using the first two seizures, with intervals starting 10 minutes before seizure onset and ending 30 seconds after. The same procedure is used to construct the test set. Detection results are reported in Table VII.

TABLE II
MAXIMUM MEAN GAP OF THE RELATIVE POWER DISTRIBUTION IN EEG BANDS BETWEEN RECONSTRUCTED AND ORIGINAL SIGNALS.

Model	δ	θ	α	β	γ
4 ch-13 bits	6.11%	0.88%	1.04%	2.33%	1.68%
8 ch-13 bits	2.60%	0.37%	0.49%	1.67%	0.63%
8 ch-15 bits	2.63%	0.39%	0.49%	1.53%	0.68%

TABLE III
MEAN PRD FOR EACH GROUP FOR THE MODEL WITH 4 CHANNELS PER GROUP AND CORRESPONDING MAXIMUM DEVIATION.

Group	Mean PRD	Max absolute PRD variation (w.r.t. the mean)
1	23.99%	5.33%
2	22.87%	2.37%
3	28.32%	5.44%
4	26.83%	2.98%
5	19.55%	1.93%
6	16.17%	1.38%
7	26.25%	8.99%
8	21.07%	3.59%

On the reconstructed signals, sensitivity is 90.71% and 88.21% for the 13-bit and 15-bit models, respectively, slightly higher than the 84.44% obtained on the original (uncompressed) signal. The false detection rate is slightly worse after reconstruction, with a minimum of 29.85% compared to 28.85% for the original signal. These results indicate that, in this experimental setting, the information related to epileptic activity is largely preserved by the proposed compression scheme. The slight sensitivity improvement observed on reconstructed signals may be interpreted as an implicit regularization or denoising effect introduced by latent quantization, which smooths high-frequency artifacts while preserving seizure-related dynamics.

IV. DISCUSSION

Overall, the method achieves good average reconstruction quality and high relative compression, but errors remain. Point-wise metrics show large peak errors, consistent with a common GAN issue: generators under-represent extreme values, making large iEEG spikes harder to reproduce. This is also supported by activity measures, which suggest reduced variance in the reconstructions—potentially impacting seizure recognition.

Channel grouping is crucial: the last group in the 8-channel models performs worse, likely because channels 24–26 have much larger amplitudes than the others despite apparent synchrony. Additional distortion comes from segment-wise quantization, which introduces discontinuities every $1/32$ s and visible artifacts, especially in high-amplitude regions (Figure 4).

Artifacts also affect Hjorth features: mobility and complexity increase, indicating less smooth reconstructions. The 15-bit model slightly reduces these effects, suggesting that a larger codebook can help. Frequency-wise, the larger error in the δ

TABLE IV
MEAN PRD FOR EACH GROUP FOR THE MODELS WITH 8 CHANNELS PER GROUP AND CORRESPONDING MAXIMUM DEVIATION.

Model	Group	Mean PRD	Max absolute PRD deviation (w.r.t. the mean)
8 channel-13 bits	1	20.54%	3.65%
	2	22.57%	2.40%
	3	18.55%	2.78%
	4	27.97%	7.27%
8 channel-15 bits	1	21.05%	3.35%
	2	23.57%	1.69%
	3	20.43%	2.66%
	4	28.18%	6.71%

TABLE V
MAXIMUM ABSOLUTE ERROR FOR THE MODELS WITH 8 CHANNELS PER GROUP. BOLD VALUES INDICATE THE LOWER ERROR BETWEEN THE TWO MODELS.

	8 channels-13 bits	8 channels-15 bits
Group 1	456.03 μ V	439.25 μ V
Group 2	1070.08 μ V	1062.52 μ V
Group 3	924.32 μ V	758.22 μ V
Group 4	1652.17 μ V	1876.82 μ V

band aligns with its dominant power and the observed power reduction, while β -band errors may reflect introduced artifacts and the difficulty of modelling multi-scale dynamics [19].

In summary, improvements could come from: (i) alternative quantization/optimization methods with more reliable convergence than classical ADMM, (ii) more flexible grouping (including variable group sizes), and (iii) GAN architectures better able to capture amplitude extremes.

V. CONCLUSION

We presented and empirically evaluated an iEEG compression technique achieving a relative compression ratio of 99.68%, exceeding typical scalp-EEG compression reports, with PRD 22.40% and PSNR 34.78 dB. Although lossy, seizure detection on reconstructed signals suggests that, in this setting, clinically salient information is largely retained.

Several limitations remain. First, experiments are centered on a single high-resolution configuration with fixed electrode layout; thus, results should be viewed as proof of concept rather than evidence of generalization across montages, conditions, or patients. While limited in scale, the proposed framework is architecture-agnostic and can be applied to different electrode layouts without structural modifications. Second, reconstruction is imperfect at the local level, and tuning is difficult, with optimal hyperparameters varying markedly due to channel-dependent amplitudes and spatial organization.

To improve consistency, we plan to design a more adaptive grouping strategy that allows variable group sizes, alongside a refined architecture. We also intend to run a broader evaluation on more diverse iEEG datasets to better assess robustness and clinical relevance.

TABLE VI

HJORTH'S ACTIVITY, MOBILITY, AND COMPLEXITY FOR THE ORIGINAL SIGNAL AND FOR RECONSTRUCTIONS WITH 8 CHANNELS PER GROUP. BOLD VALUES ARE CLOSER TO THE ORIGINAL SIGNAL.

Hjorth's parameter	Signal	Group 1	Group 2	Group 3	Group 4
Activity	Original	1349.97	3437.94	1343.88	6849.77
	8 channel-13 bits	1108.48	3084.30	1105.17	5180.48
	8 channel-15 bits	1164.26	3012.21	1117.35	5107.18
Mobility	Original	0.0011	0.0013	0.0011	0.0013
	8 channel-13 bits	0.0022	0.0025	0.0023	0.0030
	8 channel-15 bits	0.0023	0.0026	0.0025	0.0030
Complexity	Original	427.83	351.99	429.48	389.42
	8 channel-13 bits	553.20	476.48	541.95	480.49
	8 channel-15 bits	545.25	476.43	529.27	468.64

TABLE VII

RESULTS OF THE EPILEPTIC SEIZURE DETECTION ALGORITHM APPLIED TO THE ORIGINAL AND RECONSTRUCTED SIGNALS.

Model	Sensitivity	False detection rate
Original	84.44%	28.85%
8 channel-13 bits	90.71%	29.85%
8 channel-15 bits	88.21%	30.12%




Fig. 4. Signal segments with different levels of reconstruction quality. In (a), the reconstructed and original signals are very similar; in (b), several peaks are not correctly reconstructed and additional artifacts are present.

REFERENCES

- [1] J. Parvizi and S. Kastner, "Human intracranial eeg: Promises and limitations," *Nature neuroscience*, vol. 21, p. 474, 4 2018.
- [2] G. Schalk, C. Kapeller, C. Guger, H. Ogawa, S. Hiroshima, R. Lafer-Sousa, Z. M. Saygin, K. Kamada, and N. Kanwisher, "Facephenes and rainbows: Causal evidence for functional and anatomical specificity of face and color processing in the human brain," *Proceedings of the National Academy of Sciences*, vol. 114, pp. 12 285–12 290, 11 2017.
- [3] N. Mesgarani and E. F. Chang, "Selective cortical representation of attended speaker in multi-talker speech perception," *Nature*, vol. 485, pp. 233–236, 5 2012.
- [4] V. Menon, W. J. Freeman, B. A. Cutillo, J. E. Desmond, M. F. Ward, S. L. Bressler, K. D. Laxer, N. Barbaro, and A. S. Gevins, "Spatio-temporal correlations in human gamma band electrocorticograms," *Electroencephalography and Clinical Neurophysiology*, vol. 98, pp. 89–102, 2 1996.
- [5] T. Ball, M. Kern, I. Mutschler, A. Aertsen, and A. Schulze-Bonhage, "Signal quality of simultaneously recorded invasive and non-invasive eeg," *NeuroImage*, vol. 46, pp. 708–716, 7 2009.
- [6] X. Du, K. Liang, Y. Lv, and S. Qiu, "Fast reconstruction of eeg signal compression sensing based on deep learning," *Scientific Reports*, vol. 14, p. 5087, 3 2024.
- [7] J. Gui, Z. Sun, Y. Wen, D. Tao, and J. Ye, "A review on generative adversarial networks: Algorithms, theory, and applications," *IEEE Transactions on Knowledge and Data Engineering*, 4 2021.
- [8] B. Liu, C. Lee, A. Cao, and H.-S. Kim, "Unified signal compression using a gan with iterative latent representation optimization," 2021.
- [9] B. Hjorth, "Eeg analysis based on time domain properties," *Electroencephalography and Clinical Neurophysiology*, vol. 29, pp. 306–310, 9 1970.
- [10] B. Kazimi, K. Ruzaeva, and S. Sandfeld, "Self-supervised learning with generative adversarial networks for electron microscopy," in *IEEE/CVF*, 2024, pp. 71–81.
- [11] A. Mascolini, D. Cardamone, F. Ponzio, S. Di Cataldo, and E. Ficarra, "Exploiting generative self-supervised learning for the assessment of biological images with lack of annotations," *BMC bioinformatics*, vol. 23, no. 1, p. 295, 2022.
- [12] S. Roy, S. Dora, K. McCreadie, and G. Prasad, "Mieeg-gan: Generating artificial motor imagery electroencephalography signals," *Proceedings of the International Joint Conference on Neural Networks*, 7 2020.
- [13] F. Mormann, K. Lehnertz, P. David, and C. E. Elger, "Mean phase coherence as a measure for phase synchronization and its application to the eeg of epilepsy patients," *Physica D: Nonlinear Phenomena*, vol. 144, no. 3, pp. 358–369, 2000.
- [14] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, pp. 1–122, 2011.
- [15] S. D., "Web-scale k-means clustering," ser. WWW '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 1177–1178.
- [16] C. T. M. and T. J. A., *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. USA: Wiley-Interscience, 2006.
- [17] A. Burrello, L. Cavigelli, K. Schindler, L. Benini, and A. Rahimi, "Laelaps: An energy-efficient seizure detection algorithm from long-term human iieeg recordings without false alarms," *DATE 2019*, pp. 752–757, 5 2019. [Online]. Available: <http://ieeg-swez.ethz.ch/>
- [18] R. Hussein, H. Palangi, Z. J. Wang, and R. Ward, "Robust detection of epileptic seizures using deep neural networks," *ICASSP*, vol. 2018-April, pp. 2546–2550, 9 2018.
- [19] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville, "On the spectral bias of neural networks," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 5301–5310.