
Synthetic Epileptic Brain Activities using GANs

Damián Pascual
ETHZ
dpascual@ethz.ch

Amir Aminifar
EPFL

David Atienza
EPFL

Philippe Ryvlin
CHUV

Roger Wattenhofer
ETHZ

Abstract

Epilepsy is a chronic neurological disorder affecting more than 65 million people worldwide and manifested by recurrent unprovoked seizures. Modern systems monitoring electroencephalography (EEG) signals are being currently developed with the view to detect epileptic seizures in order to alert caregivers and reduce the impact of seizures on patients' quality of life. Such systems use machine learning algorithms that require large amounts of labeled seizure data for training. However, acquiring EEG signals of seizures is a costly and time-consuming process for medical experts and patients. In this work, we generate synthetic seizure-like EEG signals, that can be used to train seizure detection algorithms, alleviating the need for recorded data. First, we train a Generative Adversarial Network (GAN) with data from 30 epilepsy patients. Then, we generate synthetic training sets for new, unseen patients, which overall yield comparable detection performance to real-data training sets. We demonstrate our results using the datasets from the EPILEPSIAE Project, one of the world's largest public databases for seizure detection.

1 Introduction

Epilepsy is the fourth most common chronic neurological disorder worldwide (Hirtz et al., 2007), affecting over 65 million people. Epilepsy manifests itself by recurrent unprovoked seizures due to abnormal activity in the brain. The length of the seizures can range from few seconds to several minutes with a large variety of symptoms, including sensory auras, loss of awareness, automatic movements and full body convulsions (Blumenfeld, 2012). These symptoms not only degrade the quality of life of the patients, but they are also associated with a mortality rate 5 times higher among patients with recurrent seizures (Sperling et al., 1999) than in the corresponding group of the general population. One third of epilepsy patients suffer from drug-resistant uncontrolled seizures, which time of occurrence is usually unpredictable. A promising solution to reduce mortality and to improve the living standard of epilepsy patients is continuous real-time monitoring using wearable technologies that collect and process EEG signals from the patient and, upon occurrence of seizures, raise alerts to caregivers or family members (Goverdovsky et al., 2017; Sopic et al., 2018; Debener et al., 2015).

However, a barrier in developing reliable epileptic seizure detection systems is a lack of sufficient training data. Indeed, modern detection systems are driven by machine-learning-based algorithms (Alotaiby et al., 2014) that require a considerable amount of recorded seizures in order to reliably detect future seizures. Collecting and labeling EEG data from epilepsy patients is a costly process that currently requires in-hospital recording in specialized units. Such recordings are performed in clinical practice in a minority of patients and over short periods of time, typically a week, enabling to only record a few seizures per patient. This is a major limitation considering the privacy concerns that exist around sharing medical data and the current trend towards personalized medicine. As a consequence, it is necessary to acquire significant amounts of new data for each patient.

The problem of scarce reliable training data is common in the field of artificial intelligence and it is particularly severe in the specific case of epilepsy monitoring. The most comprehensive solution to this problem consists of generating synthetic data that can be used to train the detection algorithms. However, generating high quality medical data is challenging, and only recently substantial progress has been made thanks to advances in deep generative models and in particular, Generative Adversarial Networks (GANs) (Goodfellow et al., 2014). Several studies have used synthetic data in areas such as medical imaging (Frid-Adar et al., 2018; Shin et al., 2018; Costa et al., 2018) and Intensive Care Unit (ICU) monitoring (Esteban et al., 2017; Che et al., 2017; Lipton et al., 2016) to augment existing training sets in order to improve detection accuracy. Although this data augmentation approach has proved effective, previous attempts to train only with synthetic data have reported such a strong degradation in performance (Esteban et al., 2017; Shin et al., 2018) that it has not been possible so far to dispense with real training data. Therefore, the scenario where no real training data can be accessed and only a purely synthetic training set may be available remains unsolved. This is, however, a common scenario in several medical applications, including epilepsy, given the difficulties and privacy concerns associated with collecting and sharing medical data (Price and Cohen, 2019).

The generation of synthetic EEG data has not been extensively studied in the existing literature and previous work on synthetic EEG generation has not tackled the generation of epileptic seizures (Hartmann et al., 2018; Aznan et al., 2019; Corley and Huang, 2018). In this work, we present the use of a GAN to produce high quality synthetic epileptic seizure electroencephalogram (EEG) signals that can be used to train detection algorithms and achieve state-of-the-art results.

2 Generative model

Our model is a conditional GAN (Mirza and Osindero, 2014) that, given non-seizure (inter-ictal) EEG samples at the input, generates EEG samples of epileptic seizures (ictal). The rationale behind our design is that, while epileptic seizures are very costly to record, inter-ictal signals can be easily recorded. As a result, we condition the network on inter-ictal samples from the target patient in order to provide additional information to the generator that can be exploited to produce more realistic seizure samples. In this way, we can use an already existing database to train our GAN and then use the GAN to generate seizure samples for a new patient.

The generator of our model is a U-net (Ronneberger et al., 2015) convolutional autoencoder network with weighted skip connections where the decoder translates the latent code into an ictal sample. In order to introduce stochasticity into the model, gaussian noise with mean 0 and standard deviation 1 is concatenated to the latent code. The skip connections multiply the feature maps at each layer of the encoder with a weight which is learnt during training, and then, the result of that operation is added to the corresponding feature map of the decoder. The discriminator of our GAN has the same structure as the encoder of the generator, but it includes an additional fully connected layer at the output. The loss function of our GAN is based on the Least Squares GAN (LSGAN) (Mao et al., 2017). Consequently, the minimization objective of the discriminator is given by:

$$\min_{\theta_D} \mathcal{L}_D(\theta_D) = \mathbb{E}_{\mathbf{x} \sim p} [(D(\mathbf{x}; \theta_D) - 1)^2] + \mathbb{E}_{\mathbf{x} \sim p} [(D(G(\mathbf{x}); \theta_D))^2], \quad (1)$$

where the function D corresponds to the discriminator and G to the generator and. The input data \mathbf{x} is sampled from the input data distribution p , and θ_D are the network parameters of the discriminator. On the other hand, the loss function of the generator, with network parameters θ_G , is given by:

$$\min_{\theta_G} \mathcal{L}_G(\theta_G) = \mathbb{E}_{\mathbf{x} \sim p} [(D(G(\mathbf{x}; \theta_G)) - 1)^2] + \lambda \|G(\mathbf{x}; \theta_G) - \mathbf{y}\|_1, \quad (2)$$

The generator’s loss includes a weighted L_1 regularization term that ensures that the generated signal is similar to the reference output signal \mathbf{y} , which makes the training more stable. In Equation 2, λ is a hyperparameter that we fix to 100 in order to scale both terms of the loss function to a comparable magnitude, preventing the regularization term from dominating the optimization problem. The first term of the generator’s loss encourages the generator to produce synthetic samples that are classified as 1, i.e., real, by the discriminator, which is adversarial with respect to the discriminator’s loss function. Hence, the competing interests of the generator and the discriminator during training drive the generator to produce more and more realistic samples.

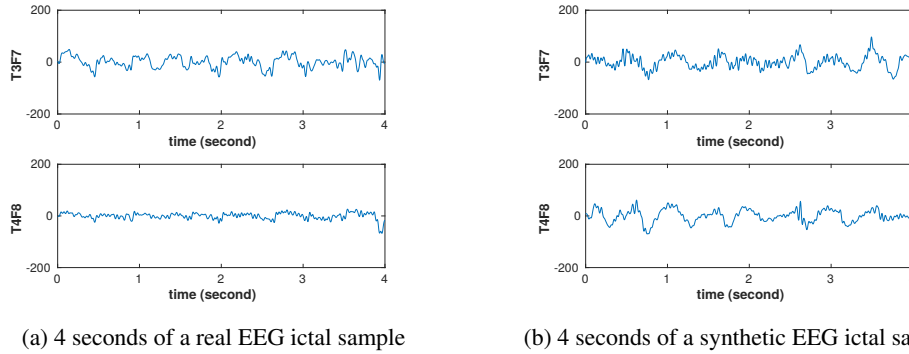


Figure 1: Comparison of a real and a synthetic ictal sample for electrodes T3F7 and T4F8. It is visible how in both cases the delta-theta rhythm are present in the signals.

3 GAN training

To train our model, we used data from the EPILEPSIAE project database (Ihle et al., 2012), which is one of the world’s largest public databases for seizure detection. The dataset contains recordings from 30 different epilepsy patients with a total of 277 epileptic seizures that sum to a duration of 21,001 seconds altogether. The EEG data is collected at a sample frequency of 256 Hz and it is divided into recordings of one hour, each one corresponding to one recording session.

In this work, we target the setup of real-world and stigma-free wearable monitoring devices (Hoppe et al., 2015) and thus we consider only the electrodes F7T3 and F8T4 in the standard 10–20 system (Klem et al., 1999), which can be easily hidden in glasses. We extract samples of four seconds of duration, since this length is effective to detect epileptic seizures. Given that the data was recorded at a frequency of 256 Hz, this results in samples of length 2048, i.e., 1024 per electrode. To construct the training set, we pair each ictal sample to an inter-ictal sample from the same patient. In this manner, the generator learns to map inter-ictal samples to ictal samples for any given patient. In order to train the GAN, the leave-one-out strategy is followed: for each target patient, the GAN is trained using the ictal and inter-ictal data coming from all other patients. The exact number of training samples depends on the number of seconds of seizure recording available in the database for all patients except for the left-out patient and, although it varies slightly, it is approximately 20,000 samples. Following this scheme, the GAN is trained independently for each patient and thus, we obtain one model per patient.

4 Evaluation of Synthetic Data

For each trained model, we generate between 2,000 and 6,000 ictal samples from inter-ictal EEG signals from the patient that was left out during training. In Figure 1, some of the generated samples are shown in the time domain. The presence of the well-known delta–theta rhythm, i.e., rhythmic slow activity with a frequency of oscillation in 0.5–4 or 4–7 hertz, is a clear indication of the correct generation of the ictal discharge and epileptic seizure segment in the synthetically-generated EEG signals (Osorio et al., 2016).

In order to evaluate the quality of the generated ictal samples beyond their visual appearance, we use them to train a state-of-the-art classifier based on the random forest algorithm (Díaz-Urriarte and De Andres, 2006). The task of the classifier is to determine whether an incoming four-second sample is an ictal or an inter-ictal sample. Here, we follow the experiments performed in Sopic et al. (2018), which are tailored to a stigma-free wearable device for epilepsy monitoring.

We target each patient independently and consider the scenario where the only data available for training the seizure detection algorithm are real inter-ictal samples and synthetic ictal samples from the target patient. As a baseline for comparison, we consider the case where real ictal samples from all other patients and inter-ictal samples from the target patient are available. Therefore, we first build a target and baseline training sets. The target training set consists of 2,000 synthetic ictal samples and 2,000 real inter-ictal samples from the target patient. The baseline training set of 2,000 samples of

real seizures randomly selected from all the patients except for the target patient and 2,000 inter-ictal samples from the target patient. In this way, the synthetic ictal samples are strictly the only aspect that differs between the target training set and the baseline training set. Then, for each patient, we construct the test set for the seizure detection task, which contains all the ictal samples of the target patient without overlap and twice as many inter-ictal samples. These test ictal samples have not been used during the training of the GAN (the target patient is left out) to ensure there is no information leakage. We build an unbalanced test set with twice as many inter-ictal as ictal samples in order to better reproduce the real-world setting where the ictal samples are under-represented in the inference phase.

Once the data is split in training and test sets, a feature extraction step is performed on the data following Sopic et al. (2018) and extract 54 features of power and non-linearity per electrode. Then, both training sets are used to train and evaluate the random forest classifier on the test set. For the sake of robustness, we repeat these experiments 15 times, splitting and shuffling the data each time.

5 Results

To assess the performance of the training sets in our experiments, we use the the geometric mean of sensitivity and specificity (Fleming and Wallace, 1986). We report the detailed results of these experiments for each patient in the Appendix A.

Our results show that training with synthetic samples not only does it not degrade the performance, but yields a 1.2% improvement overall compared to training only with real samples from a generic database. On top of that, as detailed in Figure 2, 20 out of 29 patients, i.e., 69%, improve by more than 1%, while only for four out of 29 patients the performance decreases by more than 1%. An explanation for the performance improvement when using synthetic data is that, since our GAN generates ictal samples given inter-ictal samples from the same patient, the model generates synthetic seizures that retain a number of personal features. These results demonstrate the high quality of the synthetic ictal data that we generate as well as its utility for the real-world task of seizure detection.

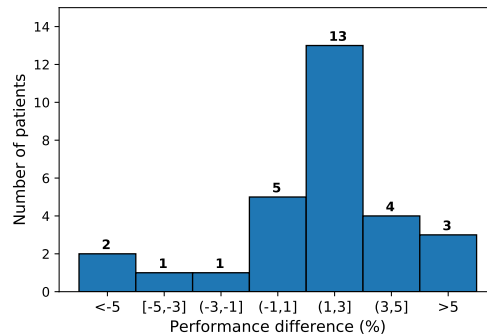


Figure 2: Performance difference in the classification task between synthetic and real training sets. The vertical axis represents the number of patients, where larger is better.

6 Conclusion

In this work, we have presented a GAN model that generates synthetic EEG signals of epileptic seizures. To the best of our knowledge, for the first time in the medical domain, we have generated synthetic data sets that can train detection algorithms achieving results comparable to training with real data. Our results underline that, in the most common scenario in which no recordings of epileptic seizures of a new patient are available, training can be done using exclusively synthetic seizures. Hence, using an existing database, deep generative models can generate data to train a system to monitor a new patient. Our work emphasizes that the application in medicine of deep generative models, such as GANs, can help solving some of the open challenges in the field and help bridging the gap on the adoption of continuous monitoring systems for patients suffering from chronic disorders. Further research into unpaired and conditional deep generative models may improve the quality and performance of synthetic training sets allowing for personalized medicine based on synthetic data.

References

- Turkey N Alotaiby, Saleh A Alshebeili, Tariq Alshawi, Ishtiaq Ahmad, and Fathi E Abd El-Samie. Eeg seizure detection and prediction algorithms: a survey. *EURASIP Journal on Advances in Signal Processing*, 2014(1):183, 2014.
- Nik Khadijah Nik Aznan, Amir Atapour-Abarghouei, Stephen Bonner, Jason Connolly, Noura Al Moubayed, and Toby Breckon. Simulating brain signals: Creating synthetic eeg data via neural-based generative models for improved ssvep classification. *arXiv preprint arXiv:1901.07429*, 2019.
- Christoph Bandt and Bernd Pompe. Permutation entropy: a natural complexity measure for time series. *Physical review letters*, 88(17):174102, 2002.
- Hal Blumenfeld. Impaired consciousness in epilepsy. *The Lancet Neurology*, 11(9):814–826, 2012.
- Zhengping Che, Yu Cheng, Shuangfei Zhai, Zhaonan Sun, and Yan Liu. Boosting deep learning risk prediction with generative adversarial networks for electronic health records. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 787–792. IEEE, 2017.
- Isaac A Corley and Yufei Huang. Deep eeg super-resolution: Upsampling eeg spatial resolution with generative adversarial networks. In *2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pages 100–103. IEEE, 2018.
- Pedro Costa, Adrian Galdran, Maria Ines Meyer, Meindert Niemeijer, Michael Abràmoff, Ana Maria Mendonça, and Aurélio Campilho. End-to-end adversarial retinal image synthesis. *IEEE transactions on medical imaging*, 37(3):781–791, 2018.
- Stefan Debener, Reiner Emkes, Maarten De Vos, and Martin Bleichner. Unobtrusive ambulatory eeg using a smartphone and flexible printed electrodes around the ear. *Scientific reports*, 5:16743, 2015.
- Ramón Díaz-Uriarte and Sara Alvarez De Andres. Gene selection and classification of microarray data using random forest. *BMC bioinformatics*, 7(1):3, 2006.
- Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.
- Philip J Fleming and John J Wallace. How not to lie with statistics: the correct way to summarize benchmark results. *Communications of the ACM*, 29(3):218–221, 1986.
- Maayan Frid-Adar, Idit Diamant, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification. *Neurocomputing*, 321:321–331, 2018.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- Valentin Goverdovsky, Wilhelm von Rosenberg, Takashi Nakamura, David Looney, David J Sharp, Christos Papavassiliou, Mary J Morrell, and Danilo P Mandic. Hearables: Multimodal physiological in-ear sensing. *Scientific reports*, 7(1):6948, 2017.
- Kay Gregor Hartmann, Robin Tibor Schirrmeyer, and Tonio Ball. Eeg-gan: Generative adversarial networks for electroencephalographic (eeg) brain signals. *arXiv preprint arXiv:1806.01875*, 2018.
- D Hirtz, DJ Thurman, K Gwinn-Hardy, M Mohamed, AR Chaudhuri, and R Zalutsky. How common are the “common” neurologic disorders? *Neurology*, 68(5):326–337, 2007.
- Christian Hoppe, Mieke Feldmann, Barbara Blachut, Rainer Surges, Christian E Elger, and Christoph Helmstaedter. Novel techniques for automated seizure registration: patients’ wants and needs. *Epilepsy & Behavior*, 52:1–7, 2015.

- Matthias Ihle, Hinnerk Feldwisch-Drentrup, César A Teixeira, Adrien Witon, Björn Schelter, Jens Timmer, and Andreas Schulze-Bonhage. Epilepsiae—a european epilepsy database. *Computer methods and programs in biomedicine*, 106(3):127–138, 2012.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- George H Klem, Hans Otto Lüders, HH Jasper, C Elger, et al. The ten-twenty electrode system of the international federation. *Electroencephalogr Clin Neurophysiol*, 52(3):3–6, 1999.
- Zachary C Lipton, David C Kale, and Randall Wetzel. Modeling missing data in clinical time series with rnns. *Machine Learning for Healthcare*, 2016.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3, 2013.
- Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- I. Osorio, H.P. Zaveri, M.G. Frei, and S. Arthurs. *Epilepsy: The Intersection of Neurosciences, Biology, Mathematics, Engineering, and Physics*. CRC Press, 2016. ISBN 9781439838860. URL <https://books.google.ch/books?id=097hKvyyYgsC>.
- Santiago Pascual, Antonio Bonafonte, and Joan Serrà. Segan: Speech enhancement generative adversarial network. In *INTERSPEECH*, 2017.
- W Nicholson Price and I Glenn Cohen. Privacy in the age of medical big data. *Nature medicine*, 25(1):37, 2019.
- Joshua S Richman and J Randall Moorman. Physiological time-series analysis using approximate entropy and sample entropy. *American Journal of Physiology-Heart and Circulatory Physiology*, 278(6):H2039–H2049, 2000.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- Hoo-Chang Shin, Neil A Tenenholtz, Jameson K Rogers, Christopher G Schwarz, Matthew L Senjem, Jeffrey L Gunter, Katherine P Andriole, and Mark Michalski. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 1–11. Springer, 2018.
- Dionisije Sopic, Amir Aminifar, and David Atienza. e-Glass: A wearable system for real-time detection of epileptic seizures. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2018.
- Michael R Sperling, Harold Feldman, Judith Kinman, Joyce D Liporace, and Michael J O’Connor. Seizure control and mortality in epilepsy. *Annals of neurology*, 46(1):45–50, 1999.

A Results

Table 1 contains the per patient results of our experiments. Patient 22 performs extremely poorly for both the baseline and the synthetic training sets and, therefore, it is not a relevant indicator of the classification quality. Consequently, it has been removed from the calculation of the total difference in performance. This total difference is calculated as the difference between the geometric mean of all patients in the synthetic case and the geometric mean of all patients in the baseline case. These results pass the Wilcoxon statistical significance test with a p -value of 0.0098 when patient 22 is already excluded, which indicates that the difference between the results obtained for the baseline and synthetic training sets is statistically significant.

Patient ID	Baseline (%)	Synthetic (%)	Difference (%)
1	73.49	80.06	+6.57
2	79.36	70.30	-9.06
3	77.59	82.84	+5.25
4	76.34	78.33	+1.99
5	64.86	68.19	+3.33
6	74.10	74.74	+0.64
7	68.11	68.59	+0.48
8	81.41	86.14	+4.73
9	76.74	80.67	+3.93
10	66.84	65.87	-0.97
11	81.03	83.66	+2.63
12	63.00	66.56	+3.56
13	77.20	78.54	+1.34
14	74.32	76.51	+2.19
15	74.25	74.07	-0.18
16	78.11	80.64	+2.53
17	65.27	67.84	+2.57
18	66.20	71.62	+5.42
19	76.95	78.13	+1.18
20	73.42	68.67	-4.75
21	79.18	71.61	-7.57
22	26.88	12.62	-14.26
23	77.05	78.62	+1.57
24	78.28	77.87	-0.41
25	77.02	75.65	-1.37
26	74.36	76.15	+1.79
27	76.00	78.00	+2.0
28	81.97	83.07	+1.10
29	75.73	78.41	+2.68
30	79.80	82.29	+2.49
TOTAL	74.57	75.78	+1.21

Table 1: Geometric mean of sensitivity and specificity per patient of our evaluation.

Regarding the patients for whom the performance degrades most significantly, i.e., Patients 2 and 21, their seizures are dominated by repetitive spiking. This pattern is relatively rare and is not well represented in the dataset (only 10.5% of the seizures). Therefore, our GAN model does not capture this behavior with as much precision as it does capture other patterns such as theta or delta rhythms. In fact, Patient 22 also suffers from seizures with repetitive spiking and our experiments show that even the state-of-the-art techniques fail in detection of such seizures. Finally, Patient 20 has only 4 seizures in this dataset, which is the lowest number of seizures in the entire dataset and hinders robust evaluation of our model.

B GAN architectural details

The architecture of our GAN is modeled after the SEGAN from (Pascual et al., 2017). Our generator is a U-net (Ronneberger et al., 2015) convolutional autoencoder with weighted skip connections in all the layers except for the input and the latent code layers. The weights of the skip connections are learnt during training. The input to the generator are samples of length 2048 points (4 seconds of signal from 2 electrodes recorded at a frequency of 256 Hz). The encoder consists of eight blocks that alternate a convolutional layer with a max pooling layer with 2x2 filters and stride 2. The feature maps extracted at each block of the encoder yield the following shapes: 2048x1, 1024x64, 512x64, 256x128, 128x128, 64x256, 32x256, 16x512, 8x1024; where 2048x1 is the shape of the input and 8x1024 that of the latent encoding. Gaussian noise with mean 0 and standard deviation 1 and shape 8x1024 is concatenated to the latent code in order to introduce stochasticity into the model. The decoder is symmetric to the encoder, but it uses deconvolutions and dilations. Thus, the shapes of the feature maps at the decoder are 16x1024, 16x512, 32x256, 64x256, 128x128, 256x128, 512x64, 1024x64, 2048x1; where 2048x1 is the final output of the generator. The skip connections multiply the feature maps output at each layer of the encoder with the learned weights and adds it to the feature map of the same shape at the decoder. The activation used is the leakyReLU function (Maas et al., 2013), except for the last block of the decoder, where we use the hyperbolic tangent function. All convolutions and deconvolutions are unbiased and spectral normalization (Miyato et al., 2018) is applied before each block.

The architecture of the discriminator has the same shape as the encoder of the generator, but it includes an additional fully connected layer at the output. In this way, the discriminator outputs a single value between 1 and 0, where 1 represents the real class and 0 the synthetic class. On top of that, in the discriminator we apply virtual batch normalization (Salimans et al., 2016), as well as spectral normalization.

To train the model we use the Adam (Kingma and Ba, 2014) optimizer with 0 and 0.9 for the values of β_1 and β_2 , respectively, and learning rates 0.0001 for the generator and 0.0004 for the discriminator. The size of the minibatches of data employed during training is 100 samples. All the hyperparameters employed are summarized in Table 2.

Parameter	Value
β_1	0
β_2	0.9
G learning rate	0.0001
D learning rate	0.0004
λ	100
Mini-batch size	100

Table 2: Model hyperparameters.

C Evaluation details

In the feature extraction stage we follow Sopic et al. (2018) and extract 54 features of power and non-linearity per electrode, i.e., a total of 108 features. To calculate the non-linear features, the signal is decomposed using the discrete wavelet transform down to level seven. The nonlinear features extracted are: sixth and seventh level sample entropy (Richman and Moorman, 2000) for $k = 0.2$ and $k = 0.35$; third, fourth, fifth, sixth and seventh level permutation entropy (Bandt and Pompe, 2002) for $n = 3$, $n = 5$ and $n = 7$; third, fourth, fifth, sixth and seventh level, as well as raw signal, Shannon, Renyi and Tsallis entropies. The power features are: total power, total and relative band power in the bands delta [0.5,4] Hz, theta [4,8] Hz, alpha [8,12] Hz, beta [13,30] Hz, gamma [30,45] Hz as well as in the bands [0,0.1] Hz, [0.1,0.5] Hz, [12,13] Hz.