Evaluation of an AI-EEG based Photo-Paroxysmal Response solution on healthy subjects: when false positive really matters *

Fernando Moncada Martins[1,3][0000−0002−6652−9287], Marcos Ornia[1], Víctor M. González[2,3][0000−0002−0937−1882], José R. Villar[1,3,**][0000−0001−6024−9527], María Antonia Gutiérrez[4], Pablo Calvo Calleja[4], Sara Urdiales Sánchez[4], Ricardo Díaz Pérez[4], and Alinne Dalla-Porta Acosta[4]

[1] Computer Science Department, University of Oviedo, Spain
{moncadafernando, uo277033, villarjose}@uniovi.es
[2] Electrical Engineering Department, University of Oviedo, Spain
vmsuarez@uniovi.es
[3] Biomedical Engineering Center, University of Oviedo, Spain
[4] Neurophysiology Service, Cabueñes University Hospital, Spain
mariaantoniagt@hotmail.com,pablo.calvo@sespa.es,sus.urdiales@gmail.com, cayuco4@gmail.com,hallynne73@hotmail.com

Abstract. Photosensitivity is a neurological disorder where the brain produces abnormal epileptic reactions to visual stimuli known as Photoparoxysmal Responses (PPR), which can sometimes result in epilepsy seizures. The Intermittent Photic Stimulation (IPS) protocol is used to diagnose this condition by exposing the patient to flashing lights -firstly, at increasing frequencies and, secondly, at decreasing frequencies- while recording the brain activity using an Electroencephalogram (EEG). Neurophysiologists observe the EEG signals to identify PPR, taking care to prevent triggering an epileptic seizure and halting the process if necessary. Because of the nature of the stimulation and the low prevalence of photosensitivity, automatically detecting these events is challenging because PPR activity represents minority events of unusual brain activity amidst a large volume of regular recordings. In previous research, a Variational AutoEncoder (VAE) was used to label EEG recordings from IPS sessions; with the encoder and the decoder incorporating recurrent neural networks and dense layers to deal with the EEG channels' time series. The VAE outperformed the current state of the art and a battery

* This research has been funded by the Spanish Ministry of Economics and Industry –grant PID2020-112726RB-I00–, the Spanish Research Agency –grant PID2023-146257OB-I00–, Missions Science and Innovation project MIG-20211008 (INMERBOT), and by the Ministry of Science and Innovation under CERVERA Excellence Network project CER-20211003 (IBERUS). Also, by Principado de Asturias, grant IDE/2024/000734, and by the Council of Gijón through the University Institute of Industrial Technology of Asturias grants SV-23-GIJON-1-09, SV-23-GIJÓN-1-17, and SV-24-GIJÓN-1-05. Finally, this research has also been funded by Fundación Universidad de Oviedo grants FUO-23-008 and FUO-22-450.

** Corresponding Author.

of unsupervised anomaly detection methods. However, training the VAE involved Leave-One-Subject-Out cross-validation with a short number of records (gathered from up to 9 diagnosed photosensitive patients); thus, this approach needs testing on patients who do not have photosensitivity. This study tackles this issue by gathering data from 5 patients who, although suffering from epilepsy, are not photosensitive. In accomplishing this, the VAE training uses data from 9 patients from prior research. The resulting model then evaluates each non-photosensitive patient, identifying anomalous EEG channel sequences of values. Results show the performance of the proposed VAE when analysing the data of these patients, both from the computer science or the neurophysiology point of view. This discussion leads to the design of the new generation of VAE to improve the overall performance.

Keywords: Photosensitivity · Epilepsy · Electroencephalography · EEG · Photoparoxysmal Response · PPR · Deep Learning · Anomaly Detection

1 Introduction

Photosensitivity is a neurological disorder in which the brain exhibits abnormal reactions to certain visual stimuli, such as flashing lights or rapid visual patterns, which can trigger epileptic activity. According to [2], approximately 30% of epileptic patients are photosensitive. Furthermore, about 6% of the general non-epileptic population experiences photosensitivity, as noted by [17]. The current internationally standardized clinical diagnostic procedure for this condition is known as Intermittent Photic Stimulation (IPS) [16]. It uses a flashing white light to stimulate patients at varying frequencies while an Electroencephalogram (EEG) registers their brain activity to provoke epileptic responses without causing full seizures, requiring constant attention, care and control by the clinical staff.

These reactive epileptiform discharges are called Photoparoxysmal Responses (PPR). [18] identified four different levels of PPR intensity, ranging from Type 1 to Type 4, where higher numbers represent more severe epileptic discharges and an increased likelihood of seizures. However, real PPRs often present a mixture of types, and the general morphology of brain activity can vary between patients or even between different sessions for the same patient due to factors such as treatment, sleep quality, or time of day [1]. Clinical neurophysiologists must manually inspect and mark these phenomena in EEG recordings to perform the diagnosis, making PPR detection a complex and demanding task.

The IPS procedure itself has notable limitations [2]. It requires significant human resources to track PPR activity. Due to the small prevalence of this condition [17] and the need to stop the stimulation session if epileptic activity occurs, the amount of recorded PPR data is limited. This leads to a highly imbalanced dataset that is unsuitable for Machine Learning (ML) and Deep Learning (DL) models. Consequently, another possible approach is to take advantage of the lack of PPR activity and view it as short bursts of abnormal brain behaviour

within otherwise normal patterns, making it suitable for Anomaly Detection (AD) solutions.

Accurate identification of anomalies is essential for all diagnostic processes and presents significant challenges. Therefore, AD models have gained prominence in the medical field in recent years due to their ability to detect abnormalities from patient-specific baseline data. For example, a Generative Adversarial Network (GAN) model to reconstruct adjacent brain MRI slices to identify abnormalities at various MRI stages [3] or an autoencoder to detect subtle abnormalities in chest X-rays and lymph metastase images [13] are some uses in clinical imaging. In time-series data analysis, particularly for epilepsy and seizure detection using EEG recordings, RNN-based models are promising for identifying temporal anomalies [10]. Unsupervised GAN and semi-supervised Variational Autoencoders (VAE) were explored by the same authors [12, 11] using behind-the-ear EEG recordings. Simple ML algorithms are still being tested to detect anomalies in this domain [5].

To the best of our knowledge, few studies have explored the analysis of PPR for diagnosing photosensitivity beyond our previous research. For instance, [4] proposed a PPR prediction method based on analyzing Fourier Components derived from EEG segments captured just before PPR onset during IPS stimulation. Additionally, [14] examined high-frequency brain oscillations in response to a light stimulation procedure different from IPS. Our prior research has focused primarily on PPR detection when IPS is applied. This includes, but is not limited to, the initial application of basic ML algorithms [8], followed by the development of an ad-hoc Data Augmentation technique to address dataset imbalance [6] and its effect in an Inception-based model pre-trained with an epilepsy dataset following a Transfer Learning approach [9].

Our latest study [7] focused on comparing multiple unsupervised ML and DL algorithms for detecting PPR anomalies. The results demonstrated that a VAE model with Recurrent layers [11] significantly outperformed other state-of-the-art methods. However, the performance of all models may have been impacted by the presence of diverse EEG anomalies that were not considered during training, potentially leading to a high number of False Positives when these unlabeled anomalies were mistakenly detected as PPR events –but correctly identified as anomalies–. This research aims to ascertain whether the VAE model can identify these EEG anomalies. To do so, the model from the previous study, trained exclusively on PPR anomalies, will be evaluated using EEG recordings from non-photosensitive patients. These recordings contain no PPR activity but include various correctly labeled EEG anomalies, providing a clearer understanding of the model's sensitivity to different types of anomalous brain activity.

The structure of this paper is as follows: the next Section describes the datasets and the experimentation setup, Section 3 presents the results and their analysis, and Section 4 provides the conclusions of this study and ideas for future work.

2 Materials and Methods

This section details the datasets collected at Cabueñes University Hospital used to train and evaluate the VAE model in 2.1, along with the experimental setup in 2.2. The experiment was implemented in Python.

2.1 Clinical Datasets

The clinical dataset comprises 14 anonymized EEG recordings from 9 photosensitive and 5 non-photosensitive patients. Brain activity was continuously recorded for 1.5 to 2.5 hours per session, –except for 1 session lasting approximately 1 hour and 2 sessions lasting about 30 minutes–. The recordings began long time before the IPS session and were collected with an EEG cap with 19 electrodes placed on the scalp following the 10–20 standardized system [15] as shown in Fig.1. All recordings were visually analyzed by a clinical neurophysiologist, who marked various phenomena depending on the case.

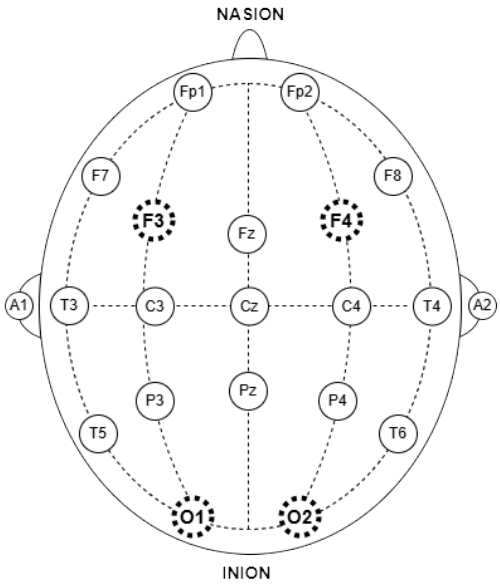

Fig. 1: Position of the 19 electrodes used according to the 10-20 international standardized system. The Nasion is located at the centre of the frontonasal area; the Inion is located at the centre of the back of the neck

On the one hand, the training dataset (*Phot_Data*) was built with the 9 photosensitive patients, captured at a sampling rate of 256Hz. After showing

PPR activity during the IPS process, these patients were diagnosed with photosensitivity. A total of 22 PPR discharges with durations in the range of 0.8–3.4 were identified and marked.

On the other hand, the testing dataset (*NONPhot_Data*) contained the data from the 5 non-photosensitive patients, sampled at 128Hz. These patients did not exhibit any PPR activity during the stimulation session. However, all other anomalies of the EEG were marked, including electrode, muscular and ocular artifacts, as well as random epileptiform discharges unrelated to PPR activity or sleep waveforms.

Following the behind-the-ear EEG approach proposed in [11], the four most relevant EEG channels for the PPR analysis based on clinical experience were selected to create two pairs of cross-head channels: F3–F4 and O1–O2 –see Fig.1–. The signals were preprocessed using a 6^{th}-order Butterworth band-pass filter (0.5–120Hz) and a Notch cut-off filter at 50Hz. The data were then segmented into 1-second length windows with 90% overlapping. For both datasets, each window was manually labeled based on the clinical marks using a binary system: "PPR" or "non-PPR" in *Phot_Data*, distinguishing between PPR activity and other brain activity that includes EEG anomalies; and "anomaly" or "not anomaly" in *NONPhot_Data*, separating all EEG anomalies from normal brain activity. The differences in sampling rate between both datasets made it necessary to apply cubic spline interpolation in *NONPhot_Data* to increase the number of samples per second up to 256 to match *Phot_Data*. The class proportion in each dataset is as follows:

- Total Number of EEG windows in *Phot_Data*: 540,200.
  - Number of non-PPR instances: 539,886 (**99.94%**).
  - Number of PPR instances: 314 (**0.06%**).
- Total Number of EEG windows in *NONPhot_Data*: 290,721.
  - Number of normal instances: 261,635 (**90.00%**).
  - Number of EEG anomaly instances: 29,086 (**10.00%**).

2.2 Experimentation Setup

The Short-Time Fourier Transform algorithm was applied to each window to extract frequency bands corresponding to brain rhythms: *Delta* (0.5–4 Hz), *Theta* (4–8 Hz), *Alpha* (8–16 Hz), *Beta* (16–40 Hz), *Gamma* (40–80 Hz), and *High-Gamma* (80–120 Hz). Then, the band power ratios were calculated for each band across both cross-head channel pairs, resulting in 12 features per window. All values were standardized (Mean = 0, Standard Deviation = 1) for each signal.

The architecture of the VAE model is presented in 2. The encoder receives the input feature vectors and compresses them into a latent space represented by a pair of vectors: mean and standard deviation. The decoder then samples a value from this latent space and reconstructs the original input. Recurrent layers using GRU cells are included before the encoder and the decoder to keep track of the temporal dependencies between EEG windows. Moreover, a stack

of normalizing flow layers is created between the encoder and the decoder to perform contractions and expansions on the latent space, increasing its flexibility.

Training involves minimizing the difference between the original data and its reconstruction. Both the encoder and decoder are preceded by an RNN-based layer to capture temporal dependencies between consecutive instances. For AD tasks, the model is trained solely on "normal" data, creating a "normal" latent space. As a result, anomalous instances are expected to produce latent vectors that deviate from the existing latent space, making the decoder unable to reconstruct the original anomaly accurately and thus yielding higher reconstruction errors.

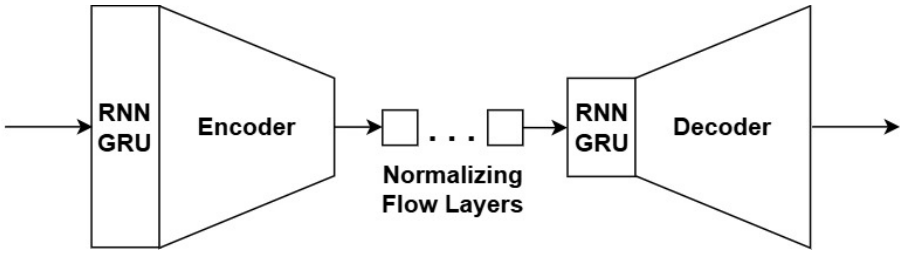

Fig. 2: VAE model architecture: both encoder and decoder are preceded by a GRU-based recurrent layer to consider the information from consecutive instances over time, and normalizing flow layers are added in between to make the latent space more flexible.

Each recording gathered in *Phot_ Data* and *NONPhot_ Data* was split into two segments: the 1^{st} hour –or the first 20 minutes for the shorter sessions–; and the remaining data. The IPS sessions are located at the end of the second segments, so the first ones only present non-PPR activity. In a previous study [7], the VAE model was trained following a Leave-One-Subject-Out (LOSO) cross-validation using the first segment from 8 patients and then tested with the second segment of the remaining patient. In this research, the previously obtained 9 models are now calibrated to each new non-photosensitive patient from *NONPhot_ Data* using the corresponding first segment and evaluated with the second segment. For this purpose, the the EEG anomalies from *NONPhot_ Data* are considered as "non-PPR" instances following the same labeling system from the training set from *Phot_ Data*, thus a Positive instance is regarded as PPR activity. The data segmentation and the experimentation workflow are shown in 3 and 4, respectively.

3 Results and Discussion

The results obtained by each K_i model trained with their corresponding training sets from *Phot_ Data* following LOSO, then calibrated and evaluated for each P_i

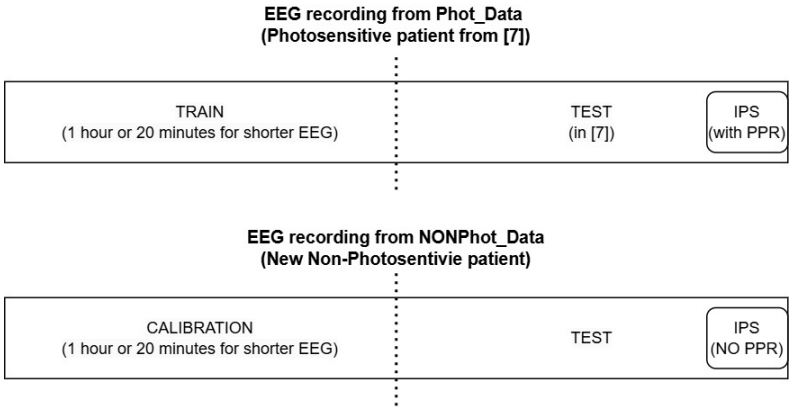

Fig. 3: Segmentation of EEG recordings in both datasets: the first segment presents only non-PPR data, while the second segment contains the IPS session. The VAE model was previously trained with the first segments from *Phot_ Data* in [7], and now they are calibrated for each patient with their first segment and evaluated with their second one.

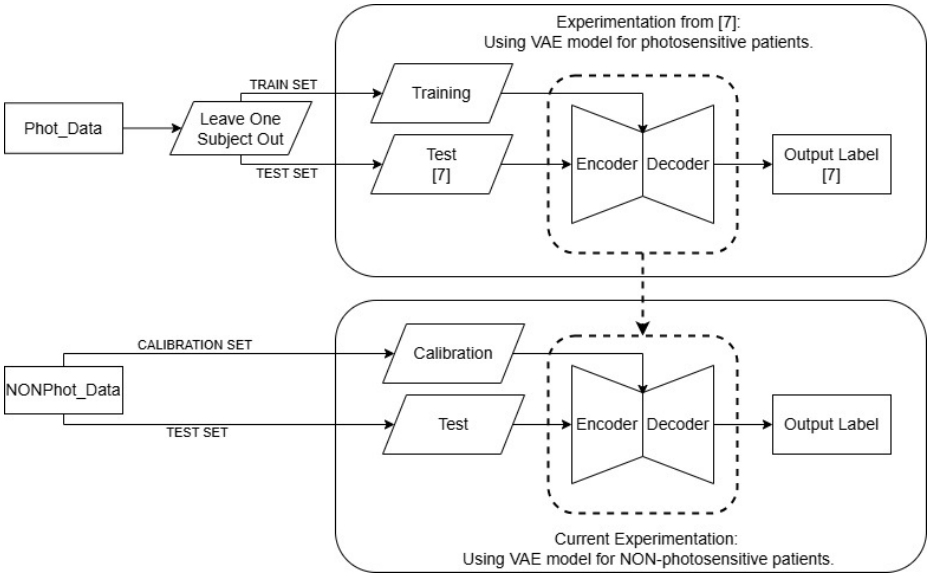

Fig. 4: Experimentation workflow: the VAE model was trained using the training segments from *Phot_ Data* following the LOSO scheme, then calibrated for and evaluated with each non-photosensitive patient from *NONPhot_ Data*.

non-photosensitive patient from *NONPhot_Data* for the PPR detection task are gathered in Table 1. The metrics used for evaluating the detection performance are Accuracy (ACC), Sensitivity (SENS) and Specificity (SPEC).

	P_1			P_2			P_3		
K_i	ACC	SENS	SPEC	ACC	SENS	SPEC	ACC	SENS	SPEC
K_1	0.9287	0.0000	0.9287	0.9368	0.0000	0.9368	0.7498	0.0000	0.7498
K_2	0.9287	0.0000	0.9287	0.9414	0.0000	0.9414	0.7025	0.0000	0.7025
K_3	0.9274	0.0000	0.9274	0.9377	0.0000	0.9377	0.7499	0.0000	0.7499
K_4	0.9288	0.0000	0.9288	0.9429	0.0000	0.9429	0.7010	0.0000	0.7010
K_5	0.9284	0.0000	0.9284	0.9423	0.0000	0.9423	0.7140	0.0000	0.7140
K_6	0.9181	0.0000	0.9181	0.9507	0.0000	0.9507	0.6828	0.0000	0.6828
K_7	0.9248	0.0000	0.9248	0.9353	0.0000	0.9353	0.6920	0.0000	0.6920
K_8	0.9271	0.0000	0.9271	0.9393	0.0000	0.9393	0.6934	0.0000	0.6934
K_9	0.9255	0.0000	0.9255	0.9317	0.0000	0.9317	0.7152	0.0000	0.7152
Mean	0.9264	0.0000	0.9264	0.9388	0.0000	0.9397	0.7034	0.0000	0.7092
Median	0.9274	0.0000	0.9274	0.9393	0.0000	0.9393	0.7025	0.0000	0.7025
StD	0.0032	0.0000	0.0032	0.0035	0.0000	0.0050	0.0114	0.0000	0.0194

	P_4			P_5		
K_i	ACC	SENS	SPEC	ACC	SENS	SPEC
K_1	0.9705	0.0000	0.9705	0.9173	0.0000	0.9173
K_2	0.9723	0.0000	0.9723	0.9184	0.0000	0.9184
K_3	0.9717	0.0000	0.9717	0.9167	0.0000	0.9167
K_4	0.9701	0.0000	0.9701	0.9155	0.0000	0.9155
K_5	0.9577	0.0000	0.9577	0.9173	0.0000	0.9173
K_6	0.9726	0.0000	0.9726	0.9166	0.0000	0.9166
K_7	0.9728	0.0000	0.9728	0.9179	0.0000	0.9179
K_8	0.9698	0.0000	0.9698	0.9182	0.0000	0.9182
K_9	0.9725	0.0000	0.9725	0.9161	0.0000	0.9161
Mean	0.9700	0.0000	0.9700	0.9171	0.0000	0.9171
Median	0.9717	0.0000	0.9717	0.9173	0.0000	0.9173
StD	0.0045	0.0000	0.0045	0.0009	0.0000	0.0009

Table 1: Results achieved by the 9 VAE models trained with LOSO in detecting PPR discharges. Accuracy (ACC), Sensitivity (SENS), and Specificity (SPEC) were obtained for each non-photosensitive patient, and Mean, Median, and Standard Deviation (StD) were computed for each metric.

Since *NONPhot_Data* contains no PPR activity, no instance should be labeled as positive. Therefore, the expected results for the PPR detection task are 0% SENS and ACC and SPEC values of 100%. However, the results reveal that the VAE model detected PPR instances in non-photosensitive patients in the form of False Positives. While the models achieved ACC and SPEC values between 90%–97% for most patients, they showed significantly lower performance for patient 3, reaching only 70%. These results further support the previous hypothesis that there are EEG anomalies that the VAE model can still distinguish.

When training the model with PPR instances as the sole type of anomaly, it is expected that all False Positives in non-photosensitive patients correspond

to EEG anomalies at least. To prove this hypothesis, a second experiment is carried out: the prediction labels generated by each model are now compared to the original *NONPhot_Data* labeling system ("anomaly" VS "non anomaly"), so an EEG anomaly detection task where a Positive will correspond to an EEG anomaly instead of PPR activity is executed. In this way, it is possible to determine if the models trained with a training set that presented no PPR activity but multiple unlabeled EEG anomalies (*Phot_Data*) were still able to identify them, affecting previous PPR detection performances. The results are presented in Table 2.

	P_1			P_2			P_3		
K_i	ACC	SENS	SPEC	ACC	SENS	SPEC	ACC	SENS	SPEC
K_1	0.9998	1.0000	0.9998	0.9936	1.0000	0.9933	0.9637	0.8769	0.9981
K_2	0.9999	1.0000	0.9999	0.9982	1.0000	0.9981	0.9863	1.0000	0.9809
K_3	0.9985	1.0000	0.9984	0.9945	1.0000	0.9942	0.9638	0.8769	0.9983
K_4	0.9999	1.0000	0.9999	0.9995	1.0000	0.9994	0.9848	1.0000	0.9788
K_5	0.9996	1.0000	0.9996	0.9991	1.0000	0.9991	0.9978	1.0000	0.9969
K_6	0.9892	1.0000	0.9884	0.9901	0.8472	0.9987	0.9667	1.0000	0.9534
K_7	0.9959	1.0000	0.9956	0.9921	1.0000	0.9916	0.9759	1.0000	0.9663
K_8	0.9982	1.0000	0.9981	0.9961	1.0000	0.9958	0.9772	1.0000	0.9681
K_9	0.9966	1.0000	0.9964	0.9885	1.0000	0.9878	0.9990	1.0000	0.9987
Mean	0.9975	1.0000	0.9973	0.9943	0.9830	0.9953	0.9795	0.9726	0.9822
Median	0.9985	1.0000	0.9984	0.9945	1.0000	0.9958	0.9772	1.0000	0.9829
StD	0.0032	0.0000	0.0035	0.0037	0.0480	0.0037	0.0128	0.0512	0.0159

	P_4			P_5		
K_i	ACC	SENS	SPEC	ACC	SENS	SPEC
K_1	0.9969	1.0000	0.9968	0.9988	1.0000	0.9988
K_2	0.9987	1.0000	0.9987	1.0000	1.0000	1.0000
K_3	0.9981	1.0000	0.9980	0.9983	1.0000	0.9981
K_4	0.9965	1.0000	0.9964	0.9971	1.0000	0.9968
K_5	0.9840	1.0000	0.9836	0.9989	1.0000	0.9988
K_6	0.9990	1.0000	0.9990	0.9982	1.0000	0.9980
K_7	0.9992	1.0000	0.9992	0.9995	1.0000	0.9994
K_8	0.9962	1.0000	0.9961	0.9998	1.0000	0.9997
K_9	0.9989	1.0000	0.9989	0.9977	1.0000	0.9975
Mean	0.9964	1.0000	0.9963	0.9987	1.0000	0.9986
Median	0.9981	1.0000	0.9980	0.9989	1.0000	0.9988
StD	0.0045	0.0000	0.0046	0.0009	0.0000	0.0010

Table 2: Results achieved by the 9 VAE models trained with LOSO in detecting EEG anomalies. Accuracy (ACC), Sensitivity (SENS), and Specificity (SPEC) were obtained for each non-photosensitive patient, and Mean, Median, and Standard Deviation (StD) were computed for each metric.

Surprisingly, this second table shows that although the VAE model was not trained for this purpose, it accurately identified all EEG anomalies present in the EEG recordings, effectively distinguishing them from normal brain activity. This was consistent across all trained instances, with a low number of false

positives and achieving 100% SENS and 99% SPEC and ACC values. Further analysis confirmed that Patient 3 exhibited a high number of EEG anomalies, which greatly increased the number of False Positives found in the PPR detection experiment, explaining the reduced ACC and SPEC values of 70%.

These results confirm that these unaccounted EEG anomalies negatively impacted the performance achieved in our previous research, where only PPR discharges were considered anomalies. Furthermore, despite being trained on a dataset where the EEG anomalies were not labeled, the model still detected them. This suggests that even without explicit labeling, these phenomena are distinct enough to be recognized as anomalies among themselves. However, further research is needed to ensure the correct and appropriate application of the model for each of the evaluated tasks.

4 Conclusions

In this research, we proposed the use of VAE-based AD models trained and used in our previous work [7] for the detection of EEG anomalies under the hypothesis that these anomalies, often present in EEG recordings, were negatively affecting the PPR detection performance of these models. Two datasets were used in this study: the first one was used to train the models on "non-PPR" data from photosensitive patients without considering EEG anomalies, while the second was used to verify if the models were able to detect those anomalies using recordings from non-photosensitive data without PPR discharges at all but with all EEG anomalies correctly labeled. Two cross-head channels were created and then divided into overlapping windows. The power band ratio was calculated from the frequency bands associated with the brain rhythms and extracted by applying the Short-Time Fourier Transform. Following a Leave-One-Subject-Out scheme, the VAE model was trained on "non-PPR" data from the first dataset, treating PPR as the sole anomaly. Then, each model instance was calibrated and evaluated with each non-photosensitive recording.

The results gathered show that the VAE model with RNN layers, which was trained with a dataset that did not distinguish EEG anomalies, can still perform a high-precision detection of EEG anomalies. The findings highlight the negative impact that ignored EEG anomalies had on earlier research.

Incorporating these additional labels into the validation of the models allows us to assess their AD capabilities more precisely. Moreover, fully labeled EEG recordings are needed for further research in multiclass models that distinguish between all types of EEG anomalies or multilabeling classifiers that assign multiple labels to a single window with a certain degree of certainty. These insights pave the way for the refinement of anomaly detection in EEG analysis, ultimately enhancing diagnostic and clinical applicability.

For future work, refining the available dataset by improving the marks and labels to include all possible EEG anomalies is essential for further research on AD approaches in PPR activity or other epileptiform discharges. Since unsupervised models were already tested in the previous study, future evaluations will

also explore supervised and semi-supervised methods. Additionally, a multiclass classification step could be incorporated after AD detection to recognize the different types and sources of EEG anomalies, as well as studying multilabeling models that detect the anomalies and provide multiple labels at once.

Ethical approval

The study was carried out according to the Declaration of Helsinki and approved by the Ethics Committee of the University Hospital of Burgos (Protocol Code CEIm2467, 23 February 2021).

Informed consent was obtained from all subjects involved in the study. Written informed consent was obtained from the patients to publish this paper.

Author Contributions

Fernando Moncada Martins, Víctor M. González, and José R. Villar designed the methodology and experiments and executed all technical and computer work. They all wrote the paper.

Clinical specialists Antonia Gutiérrez, Pablo Calvo, Sara Urdiales, Ricardo Díaz, and Alinne Dalla-Porta Acosta collected the data at the Cabueñes University Hospital's Neurophysiology Service, recorded and anonymised the data, and also studied the VAE results and interpreted them.

References

1. Dong, H., Supratak, A., Pan, W., Wu, C., Matthews, P.M., Guo, Y.: Mixed neural network approach for temporal sleep stage classification. IEEE Transactions on Neural Systems and Rehabilitation Engineering **26**(2), 324–333 (2018). https://doi.org/10.1109/TNSRE.2017.2733220
2. Fisher, R.S., Acharya, J.N., Baumer, F.M., French, J.A., Parisi, P., Solodar, J.H., Szaflarski, J.P., Thio, L.L., Tolchin, B., Wilkins, A.J., Kasteleijn-Nolst Trenité, D.: Visually sensitive seizures: An updated review by the epilepsy foundation. Epilepsia **n/a**(n/a) (Feb 2022). https://doi.org/10.1111/epi.17175, https://doi.org/10.1111/epi.17175
3. Han, C., Rundo, L., Murao, K., Noguchi, T., Shimahara, Y., Milacski, Z.A., Koshino, S., Sala, E., Nakayama, H., Satoh, S.: Madgan: unsupervised medical anomaly detection gan using multiple adjacent brain mri slice reconstruction. BMC Bioinformatics **22** (2021). https://doi.org/10.1186/s12859-020-03936-1
4. Kalitzin, S., Parra, J., Velis, D., Lopes da Silva, F.: Enhancement of phase clustering in the eeg/meg gamma frequency band anticipates transitions to paroxysmal epileptiform activity in epileptic patients with know visual sensitivity. IEEE transactions on bio-medical engineering **49**, 1279–86 (12 2002). https://doi.org/10.1109/TBME.2002.804593
5. Karpov, O.E., Khoymov, M.S., Maksimenko, V.A., Grubov, V.V., Utyashev, N., Andrikov, D.A., Kurkin, S.A., Hramov, A.E.: Evaluation of unsupervised anomaly detection techniques in labelling epileptic seizures on human eeg. Applied Sciences **13**(9) (2023). https://doi.org/10.3390/app13095655, https://www.mdpi.com/2076-3417/13/9/5655

6. Martins, F.M., Suárez, V.M.G., Flecha, J.R.V., López, B.G.: Data augmentation effects on highly imbalanced eeg datasets for automatic detection of photoparoxysmal responses. Sensors **23**(4) (2023). https://doi.org/10.3390/s23042312, https://www.mdpi.com/1424-8220/23/4/2312
7. Moncada, F., González, V.M., Villar, J.R., Gutiérrez, M.A., Calvo, P., Urdiales, S., Díaz, R., Dalla-Porta, A.: Anomaly detection comparison for photo-paroxysmal response detection. Logic Journal of the IGPL. –Aceptado pero no publicado–. (2024)
8. Moncada, F., Martín, S., González, V.M., Álvarez, V.M., García-López, B., Gómez-Menéndez, A.I., Villar, J.R.: Virtual reality and machine learning in the automatic photoparoxysmal response detection. Neural Computing and Applications (2022). https://doi.org/10.1007/s00521-022-06940-z
9. Moncada Martins, F., González, V.M., Villar, J.R., García López, B., Isabel Gómez-Menéndez, A.: Inception networks, data augmentation and transfer learning in eeg-based photosensitivity diagnosis. Machine Learning: Science and Technology **6**(1), 015034 (feb 2025). https://doi.org/10.1088/2632-2153/adb008, https://dx.doi.org/10.1088/2632-2153/adb008
10. Pilcevic, D., Djuric Jovicic, M., Antonijevic, M., Bacanin, N., Jovanovic, L., Zivkovic, M., Dragovic, M., Bisevac, P.: Performance evaluation of metaheuristics-tuned recurrent neural networks for electroencephalography anomaly detection. Frontiers in Physiology **14** (2023). https://doi.org/10.3389/fphys.2023.1267011
11. S, Y., B, H.C., YM, S., DW, S., IY, K.: Semi-supervised automatic seizure detection using personalized anomaly detecting variational autoencoder with behind-the-ear eeg. Comput Methods Programs Biomed (2022). https://doi.org/10.1016/j.cmpb.2021.106542
12. S, Y., BH, C., S, Y., JY, K., YM, S., DW, S., IY, K.: Unsupervised automatic seizure detection for focal-onset seizures recorded with behind-the-ear eeg using an anomaly-detecting generative adversarial network. Comput Methods Programs Biomed (2020). https://doi.org/10.1016/j.cmpb.2020.105472
13. Shvetsova, N., Bakker, B., Fedulova, I., Schulz, H., Dylov, D.V.: Anomaly detection in medical imaging with deep perceptual autoencoders. IEEE Access **9**, 118571–118583 (2021). https://doi.org/10.1109/ACCESS.2021.3107163
14. Strigaro, G., Gori, B., Varrasi, C., Fleetwood, T., Cantello, G., Cantello, R.: Flash-evoked high-frequency eeg oscillations in photosensitive epilepsies. Epilepsy Research **172**, 106597 (5 2021). https://doi.org/10.1016/j.eplepsyres.2021.106597, https://linkinghub.elsevier.com/retrieve/pii/S0920121121000504
15. International Federation of Clinical Neurophysiology: Report of the committee on methods of clinical examination in electroencephalography. Electroencephalography and Clinical Neurophysiology **10**(2) (1958). https://doi.org/10.1016/0013-4694(58)90053-1
16. Trenité, D.G.N., Binnie, C.D., Harding, G.F., Wilkins, A.: Photic stimulation: Standardization of screening methods. Epilepsia **40**(9) (1999). https://doi.org/10.1111/j.1528-1157.1999.tb00911.x
17. Trenité, D.K.N.: Photosensitivity in epilepsy. electrophisiological and clinical correlates. Acta Neurologica Scandinavica **125**(Suppl.), 3–149 (1989)
18. Waltz, S., Christen, H.J., Doose, H.: The different patterns of the photoparoxysmal response - a genetic study. Electroencephalography and Clinical Neurophysiology **83**(2) (1992). https://doi.org/10.1016/0013-4694(92)90027-F