Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AutoML-XAI convolutional neural network for hypnotic susceptibility classification from segmented multichannel EEG

Ali Walid Daher

ali.daher@aum.edu.kw

American University of the Middle East

Research Article

Keywords: hypnotic susceptibility, EEG classification, AutoML, convolutional neural network, explainable artificial intelligence, frequency bands

Posted Date: June 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10027260/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

AutoML-XAI convolutional neural network for hypnotic susceptibility classification from segmented multichannel EEG

Ali Walid Daher[1*]

[1*]College of Engineering and Technology, American University of the Middle East, Block 6, Building 1, Egaila, Kuwait.

Corresponding author(s). E-mail(s): ali.daher@aum.edu.kw;

Abstract

Hypnotic susceptibility classification from electroencephalography (EEG) remains underexplored compared with other EEG classification tasks. This paper presents an AutoML-XAI convolutional neural network (CNN) framework for hypnotic susceptibility classification using the public `nplresearch/hypnotic-state` EEG dataset. Each 16-channel EEG recording is filtered into five frequency-bands (rows), resampled, and split into ten non-overlapping observed EEG snippets. Each snippet is formatted as a $\mathbf{5 \times 1600}$ matrix, where rows represent frequency bands and columns preserve the ordered time-series structure. Unlike synthetic augmentation, the proposed sample expansion uses only observed EEG data and keeps all snippets from the same subject inside the same grouped fold. Optuna is used to optimize CNN architecture and hyperparameters, while SHAP and LIME are used to identify important frequency-band rows. In a single run, the XAI-row-weighted CNN achieved a best-fold subject-level accuracy, balanced accuracy, F1 score, and AUC of 1.00. Across folds, the top-two-band XAI model achieved the highest mean subject-level accuracy and balanced accuracy of 0.667, with alpha and gamma identified as the two most important bands. These findings suggest that AutoML and XAI can be combined to produce an interpretable and data-efficient EEG model for hypnotic susceptibility, while also emphasizing the need for larger independent validation datasets.

Keywords: hypnotic susceptibility, EEG classification, AutoML, convolutional neural network, explainable artificial intelligence, frequency bands

1 Introduction

Hypnotic susceptibility, also referred to as hypnotizability or hypnotic suggestibility, is commonly treated as an individual trait that influences the degree to which a subject responds to hypnotic suggestions. Although hypnosis has been studied using behavioral scales, EEG, fMRI, and other neurophysiological tools, EEG-based classification of hypnotic susceptibility remains a relatively small and heterogeneous research area. A recent bibliometric review identified EEG-based hypnosis research as a growing but still limited field and emphasized methodological inconsistency across EEG protocols and hypnotic induction procedures [1]. Reviews of hypnosis neuroimaging further report that EEG and neuroimaging findings have not fully converged, particularly for frequency-band and connectivity markers [2].

Existing hypnosis-related EEG machine learning studies provide important starting points but do not yet close this methodological gap. Alimardani et al. used differential entropy features and support vector machines for hypnotic susceptibility classification in a Hypnotrack BCI setting [3]. Keshmiri et al. reported that higher hypnotic suggestibility was associated with lower EEG signal variability in theta, alpha, and beta bands [4]. Landry et al. used machine learning to distinguish high and low hypnotic susceptibility from resting-state EEG and highlighted aperiodic neural activity as a discriminating feature [5]. Farahzadi et al. developed interpretable machine learning models for hypnotic depth, finding gamma power and beta power-envelope correlation to be strong predictors and using SHAP to interpret neural contributions [6]. These studies support the feasibility of EEG-based hypnosis modeling, but they do not combine AutoML-driven CNN architecture selection with XAI-guided band interpretation and model refinement for hypnotic susceptibility classification.

AutoML and XAI have each been used in EEG research outside hypnosis. AutoML has been applied to optimize EEG deep architectures and to compare tools such as AutoGluon, Auto-Sklearn, and Amazon SageMaker in EEG seizure detection [7, 8]. XAI methods such as SHAP and LIME have been used to investigate EEG epilepsy models and EEG-based BCI systems [9, 10]. To the best of the literature reviewed here, however, no previous EEG-hypnosis study has combined AutoML-based CNN optimization with SHAP/LIME-based XAI and XAI-guided frequency-band weighting for hypnotic susceptibility classification.

The present work addresses this gap by proposing a grouped, non-overlapping, non-synthetic EEG snippet strategy coupled with an AutoML-optimized CNN and XAI-based frequency-band analysis. The main contributions are: (i) a 5×1600 EEG representation preserving band, channel, and time ordering; (ii) a leakage-controlled segmentation strategy in which each subject contributes ten observed snippets without cross-subject or synthetic recombination; (iii) an Optuna-based CNN search with grouped K-fold validation; and (iv) SHAP/LIME-guided identification of important EEG frequency bands and XAI-derived features.

2 Related work and research gap

2.1 Limited EEG-hypnosis data and findings

A central limitation in EEG-hypnosis modeling is the lack of available data. EEG studies are often limited by small subject counts, expensive acquisition, and inter-subject variability. This problem is particularly important in hypnosis because the available EEG-hypnosis literature remains small compared with EEG epilepsy, sleep staging, or BCI motor imagery [1]. Reviews also indicate that frequency-band findings remain inconsistent. Theta, alpha, beta, and gamma have all been implicated in hypnosis susceptibility, however findings vary by task, induction method, feature definition, and whether the analysis targets hypnotic state, depth, or trait susceptibility [2, 4, 6].

2.2 AutoML and XAI in EEG

AutoML is relevant for EEG because deep learning architectures have many hyperparameters and EEG signals are high-dimensional, noisy, and contain artifacts[7]. AutoML EEG studies have shown that automated model selection can produce strong results in domains such as epilepsy detection [8]. However, the epilepsy setting differs from hypnotic susceptibility because labels, signal morphology, and clinical targets are different. XAI is also increasingly important in EEG because clinicians and neuroscientists need to understand which channels, or frequency bands support a model decision [9, 10]. Farahzadi et al. demonstrated the value of SHAP for interpreting hypnosis-depth models [6], but that study did not perform AutoML-based CNN architecture selection or XAI-guided CNN reweighting for susceptibility classification.

Therefore, the main novelty of the present study is the integration of AutoML and XAI in a single EEG-hypnosis susceptibility pipeline. AutoML is used to search for a CNN design, while XAI is used not only after training but also to create interpretable reduced or reweighted inputs. This is distinct from prior hypnosis work based on classical features such as entropy, connectivity, or power alone [3–6].

2.3 Non-synthetic segmentation versus synthetic augmentation

Deep EEG models are often limited by small datasets. EEG augmentation studies have addressed this problem using synthetic or transformed examples, including time-domain, frequency-domain, and spatial transformations [11]. CutCat, for example, increases EEG data by cutting and concatenating portions of trials from within or across subjects [12]. Such approaches can improve performance in low-data regimes, but they also introduce assumptions about which transformations preserve labels and neurophysiology.

The present work uses a more conservative strategy: no synthetic EEG is generated, no artificial noise is injected, and no inter-subject recombination is used. Instead, each original EEG recording is divided into ten non-overlapping observed snippets. This increases the number of training samples while preserving the actual input signal. The method still requires careful validation because the ten snippets from a subject are not independent people. For that reason, all splitting is performed by subject using

grouped folds, ensuring that all snippets from the same subject remain together in either training or testing within a fold.

3 Materials and methods

3.1 Dataset and preprocessing

The EEG dataset was obtained from the public GitHub resource `nplresearch/hypnotic-state` [13]. The implementation reads EEGLAB `.set`/`.fdt` files from two class folders corresponding to high and low hypnotic susceptibility. Each original recording contains 16 EEG channels. The five frequency rows used in the CNN input correspond to delta ($1-4$ Hz), theta ($4-8$ Hz), alpha ($8-13$ Hz), beta ($13-30$ Hz), and gamma ($30-45$ Hz). Each subject recording is resampled to 1000 time points per channel and then divided into ten non-overlapping parts of 100 time points per channel. Therefore, each segment contains $16 \times 100 = 1600$ ordered values per frequency band.

Let the original recording-level dataset be

$$\mathcal{D} = \{(\mathbf{R}_i, y_i)\}_{i=1}^{N}, \tag{1}$$

where $\mathbf{R}_i \in \mathbb{R}^{16 \times L_i}$ is the ith multichannel EEG recording and $y_i \in \{0, 1\}$ denotes the susceptibility label. After filtering and segmentation, the CNN input for the jth snippet of subject i is

$$\mathbf{Z}_{i,j} \in \mathbb{R}^{5 \times 1600}, \qquad j = 1, \ldots, 10. \tag{2}$$

The 1600 columns preserve the ordered channel-time structure:

$$1600 = 16 \text{ channels} \times 100 \text{ time values per channel}. \tag{3}$$

3.2 Non-overlapping segmentation

For each subject i, the segment length is

$$T_i = \left\lfloor \frac{L_i}{10} \right\rfloor. \tag{4}$$

The jth non-overlapping segment is

$$\mathbf{X}_{i,j} = \widetilde{\mathbf{R}}_i[:, (j-1)T_i : jT_i], \qquad j = 1, \ldots, 10, \tag{5}$$

where $\widetilde{\mathbf{R}}_i$ is the filtered recording. The label is inherited from the original subject-level recording:

$$y_{i,j} = y_i. \tag{6}$$

The segmented dataset is therefore

$$\mathcal{D}^{seg} = \{(\mathbf{Z}_{i,j}, y_i) : i = 1, \ldots, N,\ j = 1, \ldots, 10\}. \tag{7}$$

To avoid leakage, cross-validation is performed using grouped subject-level splits. If $g(i,j) = i$ denotes the subject group of snippet (i,j), then a valid fold satisfies

$$\{g(i,j) : (i,j) \in \mathcal{D}_{train}^{(k)}\} \cap \{g(i,j) : (i,j) \in \mathcal{D}_{test}^{(k)}\} = \emptyset. \tag{8}$$

Thus, snippets from the same subject never appear in both training and testing within the same fold.

3.3 AutoML CNN optimization

A CNN f_θ maps the 5×1600 input to a class-probability vector:

$$f_\theta : \mathbb{R}^{5\times 1600} \to [0,1]^2. \tag{9}$$

The predicted label for a snippet is

$$\widehat{y}_{i,j} = \arg\max_c f_\theta(\mathbf{Z}_{i,j})_c. \tag{10}$$

Optuna was used to search the CNN architecture and training parameters [14]. The search included the number of convolutional layers, filters, kernel dimensions, batch normalization, pooling, dropout, dense units, learning rate, and batch size. The objective maximized mean validation balanced accuracy across grouped folds:

$$J(a) = \frac{1}{K}\sum_{k=1}^{K} S\left(f_{\theta_{a,k}}, \mathcal{D}_{val}^{(k)}\right), \tag{11}$$

where a is an AutoML candidate and $S(\cdot)$ is the validation score. The best candidate is

$$a^\star = \arg\max_{a\in\mathcal{A}} J(a). \tag{12}$$

3.4 Recording-level prediction

Because each subject is represented by ten snippets, subject-level prediction aggregates the ten snippet probabilities:

$$\overline{p}_i = \frac{1}{10}\sum_{j=1}^{10} f_{\theta^\star}(\mathbf{Z}_{i,j})_1. \tag{13}$$

The binary subject-level prediction is

$$\widehat{y}_i = \begin{cases} 1, & \overline{p}_i \geq \tau, \\ 0, & \overline{p}_i < \tau. \end{cases} \tag{14}$$

3.5 XAI-guided band analysis

After training the best full-band CNN, SHAP and LIME were applied to estimate feature relevance [15, 16]. Let $\mathbf{E}_{i,j} \in \mathbb{R}^{5\times 1600}$ denote the XAI explanation map. Row-wise band relevance is computed as

$$\alpha_r = \frac{1}{|\mathcal{D}^{seg}| \times 1600} \sum_{(i,j)} \sum_{m=1}^{1600} |\mathbf{E}_{i,j}(r,m)|, \tag{15}$$

where $r \in \{1,\ldots,5\}$ indexes frequency bands. Channel-block relevance can similarly be computed over the 100 columns assigned to each channel:

$$\beta_q = \frac{1}{|\mathcal{D}^{seg}| \times 5 \times 100} \sum_{(i,j)} \sum_{r=1}^{5} \sum_{m=100(q-1)+1}^{100q} |\mathbf{E}_{i,j}(r,m)|. \tag{16}$$

The XAI stage produced three model variants: (i) the full-band CNN, (ii) a CNN using only the two most important XAI bands, and (iii) a CNN using all bands with XAI-derived row weights. A fourth stress-test variant multiplied the top two XAI bands by ten.

4 Results

The implementation used 5 grouped folds. In each fold, 24 subjects were used for training and 6 subjects for testing. Each subject contributed ten non-overlapping observed EEG snippets, producing 300 samples in the run. The best Optuna trial achieved a validation balanced accuracy of 0.7722. The selected architecture contained four convolutional layers, batch normalization, dropout of 0.1933, a dense layer of 256 units, learning rate of 5.53×10^{-5}, and batch size of 2.

Table 1 Grouped 5-fold subject-level results from one exploratory run. Mean values summarize all folds; best-fold values indicate the highest subject-level fold performance observed for that model.

Model	Mean Acc.	Mean F1	Mean AUC	Best-fold Acc./AUC
All bands CNN	0.633	0.518	0.578	0.833/1.000
Top-two XAI bands CNN	**0.667**	**0.527**	**0.689**	0.833/1.000
All bands with XAI row weights	0.600	0.505	0.556	**1.000/1.000**
Top-two bands multiplied by 10	0.533	0.448	0.511	0.667/0.778

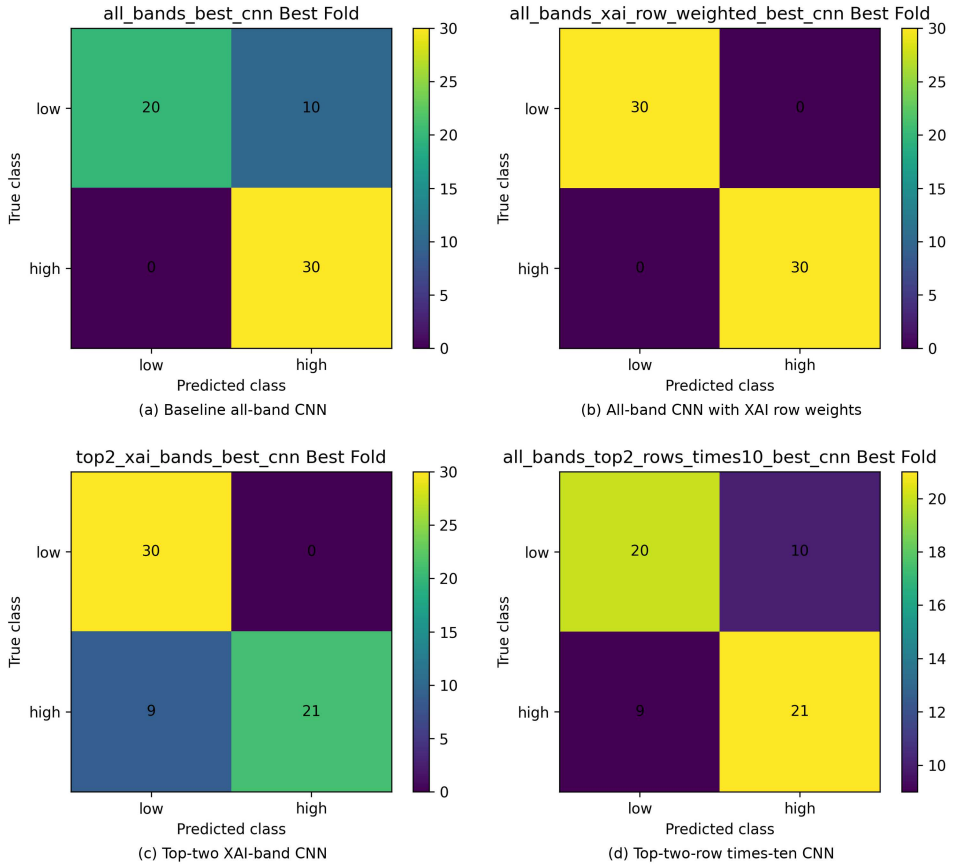


Fig. 1 Best-fold confusion matrices for the four CNN variants evaluated in the run: (a) baseline all-band CNN, (b) all-band CNN with XAI-derived row weights, (c) CNN trained using only the top two XAI-selected frequency bands, and (d) stress-test CNN with the top two XAI-selected rows multiplied by ten. Rows indicate the true hypnotic susceptibility class and columns indicate the predicted class. The XAI-row-weighted all-band CNN produced the clearest best-fold result, with 30/30 low-susceptibility snippets and 30/30 high-susceptibility snippets classified correctly.

Figure 1 presents the best-fold confusion matrices for all CNN variants on a dedicated figure page for readability. The all-band CNN and the top-two-row-times-ten stress-test model both correctly identified all high-susceptibility snippets in their best folds but retained errors in the low-susceptibility class. The XAI-row-weighted all-band CNN achieved the clearest best-fold separation, with no false low or false high predictions, while the top-two-XAI-band CNN preserved perfect low-class detection but retained errors for high-susceptibility snippets.

The strongest result was obtained by the all-band XAI-row-weighted CNN, which reached subject-level accuracy, balanced accuracy, F1 score, and AUC equal to 1.000 in the best fold. This result is important because it shows that the XAI-derived row weighting can produce a perfect fold-level separation in this run. However, because the mean across folds was lower, the perfect fold should be interpreted as a promising exploratory result rather than as conclusive evidence of generalization (To avoid overfitting).

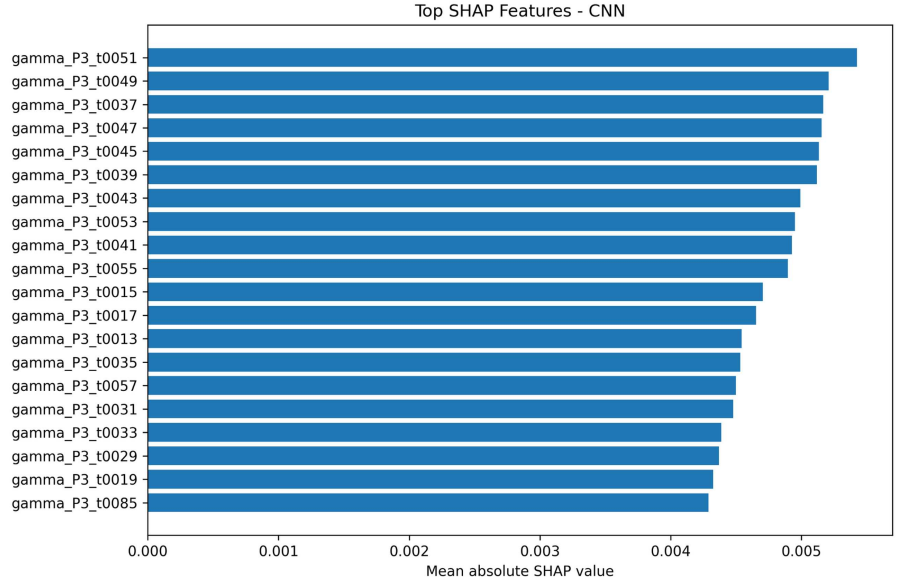
The most robust mean performance in this run was obtained by the top-two XAI-band CNN, which achieved the highest mean subject-level accuracy and balanced accuracy of 0.667 and the highest mean AUC of 0.689. This suggests that the XAI-selected bands retained discriminative information while reducing the input from five frequency rows to two.

The combined SHAP and LIME analysis ranked the frequency bands as follows:

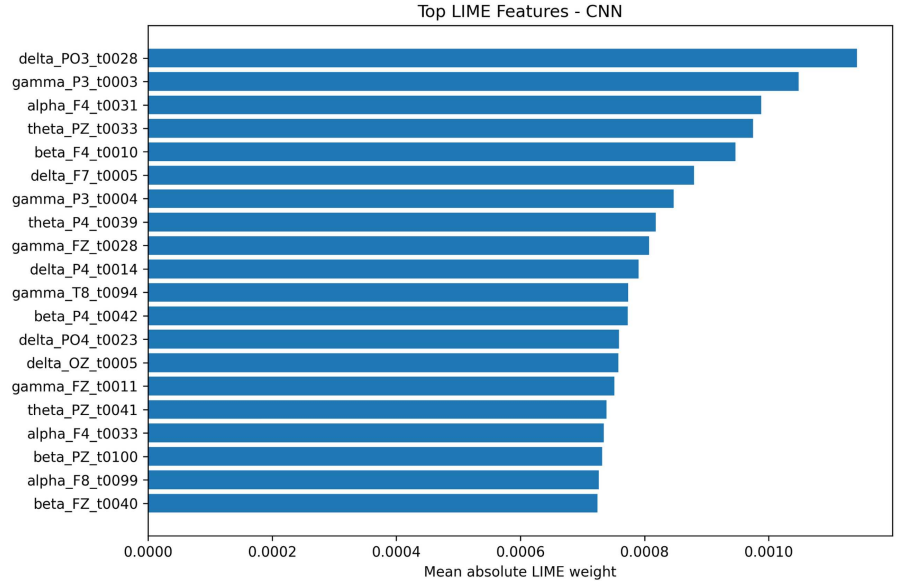
$$\alpha : 221.64, \quad \gamma : 207.16, \quad \beta : 180.09, \quad \theta : 145.99, \quad \delta : 82.32. \tag{17}$$

Thus, alpha and gamma were selected as the top two XAI frequency bands. The normalized XAI row weights used for retraining were

$$w_\delta = 0.492, \quad w_\theta = 0.872, \quad w_\alpha = 1.324, \quad w_\beta = 1.076, \quad w_\gamma = 1.237. \tag{18}$$



(a) SHAP feature ranking



(b) LIME feature ranking

Fig. 2 Feature-level XAI rankings for the CNN: (a) top SHAP feature ranking and (b) top LIME feature ranking. SHAP concentrates the strongest attributions in gamma-band P3 time features, whereas LIME assigns relevance to a broader set of band-channel-time features, including delta, gamma, alpha, theta, and beta components.

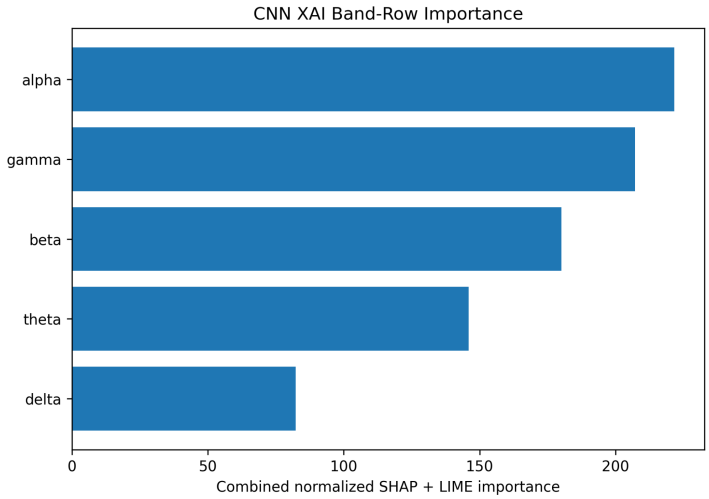


Fig. 3 Combined SHAP and LIME frequency-band row importance for the CNN. Alpha and gamma are the two highest-ranked rows, supporting their selection for the top-two-band CNN and their emphasis in the XAI-guided interpretation.

The detailed feature-level explanations are visualized in Figure 2, while the combined frequency-band row importance is summarized in Figure 3. The SHAP ranking is dominated by gamma-band features around the P3 channel, while LIME distributes high relevance across delta, gamma, alpha, theta, and beta features. When SHAP and LIME are combined at the row level, alpha and gamma receive the highest total importance, followed by beta, theta, and delta. The convergence of alpha and gamma in the combined ranking supports the use of XAI as an interpretability layer rather than relying only on validation accuracy.

5 Discussion

5.1 Main finding and novelty

The present results support the feasibility of an AutoML-XAI approach for hypnotic susceptibility classification. The best-fold perfect result obtained by the XAI-row-weighted CNN demonstrates that the proposed representation can separate high and low susceptibility subjects under at least one grouped train-test partition. More importantly, the model is not a black-box-only pipeline: it identifies alpha and gamma as the most important frequency rows and uses those XAI results to construct reduced and reweighted CNN variants.

The proposed workflow differs from previous EEG-hypnosis machine learning studies. Earlier susceptibility work used differential entropy and classical SVM classification [3, 4], while more recent hypnosis-depth work used interpretable models and SHAP but did not perform AutoML CNN search or XAI-guided CNN reweighting [6]. AutoML and XAI have been explored in other EEG domains such as BCI and epilepsy [7–10], but the reviewed literature does not show a prior combination of

AutoML-optimized CNNs and XAI-guided band selection for hypnotic susceptibility classification.

5.2 Reliability of the data expansion strategy

The segmentation strategy is intentionally conservative. EEG augmentation methods can be useful in low-data regimes, but many generate transformed or synthetic samples by applying time, frequency, spatial, noise, or cut-and-concatenate operations [11, 12]. In contrast, the present work uses only original recorded EEG snippets. No artificial EEG signals are generated, no cross-subject snippets are recombined, and no overlapping windows are used. Therefore, the enlarged dataset is based on observed EEG evidence rather than synthetic augmentation.

This choice improves interpretability and reduces classification assumptions, however it does not remove the dependence among snippets from the same subject. The ten snippets produced from one EEG recording are not ten independent participants. The use of grouped K-fold splitting directly addresses this issue by keeping all ten snippets from a subject in the same fold. This makes the validation design more reliable than segment-level random splitting, which would risk leakage and unreliable or even inflated performance.

5.3 Frequency-band interpretation

The XAI analysis selected alpha and gamma as the two most important bands, with beta ranked third. This partly aligns with previous hypnosis EEG findings but also adds a CNN/XAI-specific perspective. Farahzadi et al. found gamma power and beta PEC to be the top predictors of hypnotic depth and reported SHAP evidence that reduced gamma power in midline frontal and parietal regions and increased beta DAN connectivity contributed to deeper hypnosis [6]. Keshmiri et al. found that high hypnotic suggestibility was associated with lower EEG signal variability in theta, alpha, and beta bands [4]. De Pascalis reviewed evidence that theta, alpha, beta, and gamma oscillations and connectivity patterns can all change with hypnotic susceptibility, but also emphasized that findings are not fully convergent [2].

The present alpha-gamma result is therefore plausible but should not be interpreted as replacing prior theta-, beta-, or connectivity-based findings. Instead, it suggests that when the raw ordered band-channel-time matrix is processed by an AutoML-optimized CNN, alpha and gamma rows carry the strongest combined SHAP/LIME relevance in this run. Gamma agrees strongly with the hypnosis-depth literature, while alpha may reflect the importance of attention, or connectivity-related differences often discussed in hypnosis EEG studies. The beta band remained important but ranked below alpha and gamma in this pipeline. These differences may arise from the use of non-overlapping snippets, CNN filters, row-wise weighting or even the lack of available training samples.

6 Conclusion

This work presents an AutoML-XAI CNN framework for hypnotic susceptibility classification from segmented 16-channel EEG. The method uses a 5×1600 frequency-band representation, ten non-overlapping observed snippets per subject, grouped K-fold validation, Optuna-based CNN search, and SHAP/LIME-based band interpretation. In one exploratory run, the XAI-row-weighted CNN reached a perfect best-fold subject-level result, while the top-two-band XAI model produced the strongest mean subject-level performance across folds. The findings support alpha and gamma as the most important bands in this pipeline and position AutoML-XAI as a promising direction for interpretable EEG-based hypnosis classification.

Acknowledgements. The authors acknowledge the availability of the public `nplresearch/hypnotic-state` dataset and the open-source Python ecosystem used for model development and interpretation.

Declarations

- **Funding** No funding information is provided in this draft.
- **Conflict of interest** The authors declare no competing interests.
- **Ethics approval** This study uses a public dataset. Any required ethics statement should be completed according to the dataset documentation and institutional requirements.
- **Data availability** The EEG data used in this work are cited as the `nplresearch/hypnotic-state` GitHub repository [13].
- **Code availability** The implementation used Python with Optuna for AutoML-style CNN optimization [14], SHAP and LIME for explainability [15, 16], and comparative AutoML context from Auto-Sklearn, AutoGluon, and Amazon SageMaker [17–19].

References

[1] Rinaldi, M.R., Luijtelaar, G., Ul Hasanat, N., Yuniarti, K.W.: Mapping eeg-based hypnosis research: A bibliometric study. American Journal of Clinical Hypnosis (2025) https://doi.org/10.1080/00029157.2025.2532452 . Summary: Bibliometric review of EEG-hypnosis studies; highlights limited literature, methodological inconsistencies, and emerging AI themes.

[2] De Pascalis, V.: EEG Oscillations and Neural Functional Connectivity Underpinning Hypnosis and Hypnotizability. Preprints.org. Summary: Review of EEG/fMRI/PET hypnosis findings; emphasizes lack of convergence and frequency-band/connectivity patterns. (2023). https://doi.org/10.20944/preprints202312.0382.v1

[3] Alimardani, M., Keshmiri, S., Sumioka, H., Hiraki, K.: Classification of eeg signals for a hypnotrack bci system. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 240–245 (2018). https://doi.org/10.

1109/IROS.2018.8594136 . Summary: SVM classification of high/low hypnotic susceptibility based on differential entropy features.

[4] Keshmiri, S., Alimardani, M., Shiomi, M., Sumioka, H., Ishiguro, H., Hiraki, K.: Higher hypnotic suggestibility is associated with the lower eeg signal variability in theta, alpha, and beta frequency bands. PLOS ONE **15**(4), 0230853 (2020) https://doi.org/10.1371/journal.pone.0230853 . Summary: Differential entropy in theta, alpha, and beta; high suggestibility linked to lower EEG signal variability.

[5] Landry, M., Silva Castanheira, J., Sackur, J., Raz, A., Ogez, D., Rainville, P., Jerbi, K.: Unravelling the Neural Dynamics of Hypnotic Susceptibility: Aperiodic Neural Activity as a Central Feature of Hypnosis. bioRxiv preprint. Summary: Machine learning for high/low hypnotic susceptibility using resting-state EEG before and after induction; highlights aperiodic exponent. (2023). https://doi.org/10.1101/2023.11.16.567097

[6] Farahzadi, Y., Alldredge, C., Kekecs, Z.: Gamma power and beta envelope correlation are potential neural predictors of deep hypnosis. Scientific Reports **14**, 6329 (2024) https://doi.org/10.1038/s41598-024-56633-x . Summary: Interpretable EEG machine learning for deep hypnosis; gamma power and beta PEC were top predictors; SHAP localized gamma and beta network contributions.

[7] Aquino-Britez, D., Ortiz, A., Ortega, J., Leon, J., Formoso, M., Gan, J.Q., Escobar, J.J.: Optimization of deep architectures for eeg signal classification: An automl approach using evolutionary algorithms. Sensors **21**(6), 2096 (2021) https://doi.org/10.3390/s21062096 . Summary: AutoML approach optimizing deep EEG architectures; relevant to EEG AutoML but not hypnosis susceptibility.

[8] Lenkala, S., Marry, R., Gopovaram, S.R., Akinci, T.C., Topsakal, O.: Comparison of automated machine learning (automl) tools for epileptic seizure detection using electroencephalograms (eeg). Computers **12**(10), 197 (2023) https://doi.org/10.3390/computers12100197 . Summary: Compared AutoGluon, Auto-Sklearn, and Amazon SageMaker for EEG epilepsy; supports AutoML in EEG but outside hypnosis.

[9] Sanchez-Hernandez, S.E., Torres-Ramos, S., Roman-Godinez, I., Salido-Ruiz, R.A.: Evaluation of the relation between ictal eeg features and xai explanations. Brain Sciences **14**(4), 306 (2024) https://doi.org/10.3390/brainsci14040306 . Summary: Used SHAP and LIME for EEG epilepsy channel relevance; demonstrates XAI in EEG but not hypnosis/AutoML.

[10] Apicella, A., Isgro, F., Pollastro, A., Prevete, R.: Toward the Application of XAI Methods in EEG-Based Systems. Summary: XAI for EEG/BCI dataset shift and feature relevance; useful for positioning XAI-based feature selection. (2024)

[11] Rommel, C., Paillard, J., Moreau, T., Gramfort, A.: Data augmentation for

learning predictive models on eeg: A systematic comparison. Journal of Neural Engineering **19**(6), 066020 (2022) https://doi.org/10.1088/1741-2552/aca220 . Summary: Formal comparison of 13 EEG augmentations; useful contrast to non-synthetic segmentation.

[12] Al-Saegh, A., Dawwd, S.A., Abdul-Jabbar, J.M.: Cutcat: An augmentation method for eeg classification. Neural Networks **141**, 433–443 (2021) https://doi.org/10.1016/j.neunet.2021.05.032 . Summary: Cut-and-concatenate augmentation for EEG; illustrates synthetic/trial recombination methods.

[13] nplresearch: hypnotic-state: EEG Hypnosis Dataset and Analysis Repository. GitHub repository. Accessed 2026-06-11. Summary: Dataset repository used for hypnotic-state EEG analyses. (2026). https://github.com/nplresearch/hypnotic-state

[14] Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M.: Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 2623–2631 (2019). https://doi.org/10.1145/3292500.3330701 . Summary: Hyperparameter optimization framework used for AutoML-style CNN search.

[15] Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems, vol. 30 (2017). Summary: SHAP method used for feature and frequency-band importance.

[16] Ribeiro, M.T., Singh, S., Guestrin, C.: "why should i trust you?" explaining the predictions of any classifier. In: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations, pp. 97–101 (2016). https://doi.org/10.18653/v1/N16-3020 . Summary: LIME method used for local EEG feature explanations.

[17] Feurer, M., Klein, A., Eggensperger, K., Springenberg, J.T., Blum, M., Hutter, F.: Efficient and robust automated machine learning. In: Advances in Neural Information Processing Systems, vol. 28 (2015). Summary: Introduced auto-sklearn for robust AutoML.

[18] Erickson, N., Mueller, J., Shirkov, A., Zhang, H., Larroy, P., Li, M., Smola, A.: AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data. Summary: AutoML platform using stacking and ensembling; cited as comparative AutoML context. (2020)

[19] Liberty, E., Karnin, Z.S., Xiang, B., Rouesnel, L., Coskun, B., Nallapati, R., Delgado, J., Sadoughi, A., Astashonok, Y., Das, P., Balioglu, C., Chakravarty, S., Jha, M., Gautier, P., Arpin, D., Januszewski, R.: Elastic machine learning algorithms in amazon sagemaker. Proceedings of the VLDB Endowment **13**(12), 3198–3211 (2020) https://doi.org/10.14778/3415478.3415545 . Summary: Amazon SageMaker algorithms/platform reference.