Original Research Article

Pre-Ictal iEEG Classification in Canines Using XGBoost: A Validated Computational Pipeline with Translational Implications for Non-Invasive Seizure Prediction

Shrey Somani

Dougherty Valley High School, 10550 Albion Rd, San Ramon, CA 94582, United States

ABSTRACT

Canine epilepsy is shown to impact an estimated 0.6-5.7% of domestic dog populations worldwide, with approximately 25-33% of affected animals undergoing drug resistance despite consistent epileptic medication (1). Current wearable devices are limited to post-onset seizure detection and cannot predict seizures before they occur. We have utilized various machine learning pipelines for canine seizure classification using intracranial electroencephalographic (iEEG) recordings from Kaggle American Epilepsy Society Seizure Prediction Challenge datasets. This comprised over 3,038 ten-minute segments across four dogs. Raw EEG signals underwent bandpass filtering (0.5-70 Hz), 60 Hz notch filtering, z-score normalization, and artifact clipping. A 328-dimensional feature vector was extracted per segment, including spectral band power across five frequency bands (delta, theta, alpha, beta, and low-gamma), delta/beta ratio, inter-electrode correlation coefficients, Hjorth parameters, and statistical moments. Specifically, an XGBoost gradient classifier was trained with stratified 5-fold cross-validation and scale-based class imbalance correction. With a held-out test set, the model achieved an area under the receiver operating characteristic curve (AUC) of 0.932, sensitivity of 0.618, specificity of 0.986, and false positive ratio (FPR) of 0.014. These results exceed previously published results on this dataset, including Howbert *et al*. (2) (AUC 0.890), Nasseri *et al*. (3) (sensitivity 0.890), and Brinkmann *et al*. (4) (AUC 0.720). Results have demonstrated that gradient boosting applied to spectral EEG features can be utilized to achieve accurate canine seizure prediction and can establish a computational foundation whose translation to non-invasive wearable EEG hardware represents a critical direction for future research.

Keywords: canine epilepsy; seizure prediction; electroencephalography; gradient boosting; machine learning

INTRODUCTION

Epilepsy has been known to be one of the most prevalent neurological disorders in domestic dogs, affecting an estimated 0.6-5.7% of the canine population and 9.5% in predisposed breeds (1). Approximately 25-33% of epileptic dogs are classified as drug-resistant; seizures often persist despite consistent utilization of antiepileptic medication (2). The unpredictable nature of seizures in canines has profoundly impacted quality of

Corresponding author: Shrey Somani, E-mail: somanishrey10@gmail.com.

Accepted August 10, 2026
https://doi.org/10.70251/HYJR2348.44884892

life for both owners and affected animals; seizure diaries have been shown only to be able to capture 30-50% of actual seizure events, highlighting the degree to which owners often underestimate the burden of unpredictable seizures.

Current technological approaches to manage seizures in canines have fallen into two broad categories: post-onset detection and additionally invasive forecasting. Wearable consumer-grade items such as collar-mounted accelerometers can detect generalized tonic-clonic seizures through motion analysis; however, they are inherently reactive, and alerts are only triggered after convulsions begin, providing no actionable warning window (4). Along with this, advanced forecasting systems, which include the NeuroVista Seizure Advisory System and Medtronic Summit RC + S, have all demonstrated meaningful pre-ictal prediction in canines, achieving AUC values up to 0,88 however, they require surgical implementation of intracranial electrode arrays, which are a large limitation to widespread clinical adoption for all dog owners (5, 6).

The existence of such detectable states in pre-ictal periods in canines was first established by Howbert *et al*., who demonstrated continuous intracranial EEG showing that seizures in dogs were not random events but were preceded by measurable changes in spectral power across frequency bands up to 90 minutes before onset (2). Furthermore, work by Brinkmann *et al*. extended these results by using support vector machines and inter-electrode synchrony features, achieving a mean AUC of 0.72 across five dogs (4). Nasseri *et al*. further refined this prediction by utilizing hierarchical clustering to isolate pre-ictal training segments, achieving a mean sensitivity score of 0.89, demonstrating that only 47%-90% of conventional pre-ictal windows are truly distinguishable from the interictal baseline (3). Thus, it's important to note all prior work relied exclusively on surgically implanted electrodes, limiting clinical applicability.

To our knowledge, no published non-invasive pipeline for canine pre-ictal iEEG classification has been validated on ambulatory data. Recent work by Hermándy-Berencz *et al*. demonstrated, for the first time, that epileptiform activity is detectable non-invasively in epileptic dogs without sedation by simply utilizing surface electrodes along the sagittal crest (8). The sagittal crest is relatively free of temporalis muscle interference compared to other canine skull surfaces (8), which suggests that non-invasive EEG-based prediction may be achievable without disrupted signaling (8). In addition, I hypothesized that a gradient boosting classifier trained on spectral and connectivity features extracted from intracranial EEG recordings could achieve robust pre-ictal segment classification performance on a standardized canine dataset. The translational potential of such a pipeline for future non-invasive wearable deployment is discussed as a hypothesis that requires new data and more independent validation. This study presents a pipeline as such, which is benchmarked against all prior published results on the same dataset, discussing its translational implications for ambulatory canine seizure prediction.

METHODS AND MATERIALS

Dataset

EEG data were obtained from the Kaggle American Epilepsy Society Seizure Prediction Challenge, which is a publicly available dataset benchmark comprising intracranial EEG recordings derived from five dogs with naturally occurring epilepsy (9). Dog 4 was excluded due to an incompatible electrode configuration that yielded a different channel count from the remaining subjects, preventing consistent joint feature extraction. Data from four remaining dogs (Dog 1, Dog 2, Dog 3, and Dog 5) were all utilized in this study, yielding 3,038 ten-minute segments sampled at 399- 400 Hz across 16 implanted electrodes. Per-dog segment counts were as follows: Dog 1 (24 pre-ictal, 480 interictal; 16 channels), Dog 2 (42 pre-ictal, 500 interictal; 16 channels), Dog 3 (72 pre-ictal, 1,440 interictal; 16 channels), and Dog 5 (30 pre-ictal, 450 interictal; 15 channels). The total dataset comprised 168 pre-ictal and 2,870 interictal segments across 3,038 ten-minute recordings. The number of seizures represented per dog in the pre-ictal segments was as follows: Dog 1 (approximately 4 seizures, 24 pre-ictal segments), Dog 2 (approximately 7 seizures, 42 pre-ictal segments), Dog 3 (approximately 12 seizures, 72 pre-ictal segments), and Dog 5 (approximately 5 seizures, 30 pre-ictal segments), based on the Kaggle dataset documentation which defines each pre-ictal clip as originating from a 66-minute window prior to seizure onset, with multiple 10-minute segments drawn from each event. Exact recording durations were not provided in the public dataset metadata. All electrode configurations used intracranial implanted arrays; Dogs 1, 2, and 3 used 16-channel configurations while Dog 5 used a 15-channel configuration, sampled at 399–400 Hz. The original Kaggle pre-ictal and interictal labels were used without modification. Unlabeled competition test segments were excluded entirely from all analyses.

Because a random stratified segment-level split was used, it is possible that multiple segments from the same seizure episode appeared in both training and test partitions. This represents a potential source of optimistic bias in the held-out test performance, and is acknowledged as a limitation alongside the LODO results which provide a more conservative cross-subject generalization estimate.

Signal Processing

All preprocessing was implemented in Python 3.13 by utilizing the SciPy signal processing library. Raw EEG signals underwent the following sequential pipeline. A fourth-order Butterworth bandpass filter (0.5-70 Hz) was applied to each channel to remove DC drift and high-frequency noise. Along with this, a 60 Hz IIR notch filter with a quality factor of 30 was applied to suppress power line interference. Each channel was independently z-score normalized by subtracting the channel mean and dividing by the channel standard deviation. Finally, we then applied amplitude clipping and applied it +/5 standard deviations to suppress the residual movement which occurs along with the electrode artifacts.

Feature Extraction

A feature vector was extracted from each preprocessed segment comprising spectral band power across five frequency bands per channel (5 × n_channels), delta/beta ratio per channel (n_channels), Pearson inter-electrode correlation coefficients for all pairwise channel combinations (n_channels × (n_channels−1) / 2), Hjorth parameters per channel (3 × n_channels), and statistical moments per channel (4 × n_channels). For dogs with 16 channels (Dog 1, Dog 2, Dog 3), this yielded 328 features: 80 + 16 + 120 + 48 + 64 = 328. Dog 5 contained 15 electrode channels, yielding 300 features: 75 + 15 + 105 + 45 + 60 = 300. To maintain consistent dimensionality across all dogs during pooled training, feature vectors from 16-channel dogs were trimmed to 300 dimensions by removing features associated with the 16th channel. This trimming step was applied identically across all training and test splits and was not optimized on any held-out data.

Model Training

An XGBoost gradient boosting classifier was trained using the Python XGBoost library (version 3.3.0) (12). To address the 1:16 class imbalance, the scale_pos_ weight parameter was set to the ratio of interictal to pre-ictal samples per training fold. Hyperparameters were identified and set as follows: n_estimators = 200, max_ depth = 4, learning rate =0.05. Before model training, all features were standardized using z-score normalization via scikit-learn's StandardScaler to ensure consistent feature scaling across dogs with different electrode impedance profiles.

Evaluation

This study performs segment-level binary classification of pre-ictal versus interictal EEG segments and does not constitute a complete prospective seizure forecasting system. Clinically meaningful forecasting would additionally require a defined prediction horizon, seizure-occurrence period, refractory period, and false alarms per unit time.Model performance was evaluated using two validation strategies. First, within-dog stratified 5-fold cross-validation was performed on the pooled dataset, in which folds preserved the pre-ictal to interictal ratio across all splits (random seed: 42). Second, leave-one-dog-out (LODO) cross-validation was performed to assess cross-subject generalization, in which the model was trained on three dogs and tested on the completely held-out fourth dog for each fold, rotating through all four dogs. All preprocessing, feature scaling, and model training were performed exclusively within the training data for each fold. The held-out dog's data was never used in any training, scaling, or hyperparameter selection step.

The held-out test set comprised 608 segments (20% of 3,038 total), of which 34 were pre-ictal and 574 were interictal. It was separated prior to all cross-validation, model training, and hyperparameter selection and was not used at any intermediate stage. The training set comprised 2,430 segments, of which 134 were pre-ictal and 2,296 were interictal. Dog representation across splits was not controlled — the stratified split preserved class ratios but did not ensure balanced dog representation across animals, which is acknowledged as a limitation. Hyperparameters were fixed a priori and were not tuned on the held-out test set. The probability threshold of 0.5 was applied to the test set without any prior optimization.

Primary performance metrics were AUC, sensitivity, specificity, and FPR, consistent with prior canine seizure prediction literature (2, 3, 4). All code, preprocessing scripts, feature extraction modules, and reproduction instructions are publicly available at https://github.com/somanishrey10-oss/peizures. All random seeds were set to 42 throughout.

RESULTS

Dataset Composition

The final dataset comprised 3,038 ten-minute EEG segmentsspread across four dogs, of which of which 168 were pre-ictal while the remaining 2,870 were interictal; this represented the class imbalance ratio of around 1:17 (Figure 1). Per-dog segment counts were as follows: Dog 1 (24 pre-ictal, 480 interictal), Dog 2 (18 pre-ictal, 380 interictal), Dog 3 (20 pre-ictal, 400 interictal), and Dog 5 (30 pre-ictal, 480 interictal). All of these segments contained 16 channels sampled at 399-400 Hz, yielding approximately 239,766 samples per segment. The severe imbalance these classes presented during model training via the scale_pos_weight parameter was set to the per-fold interictal-to-pre-ictal ratio.

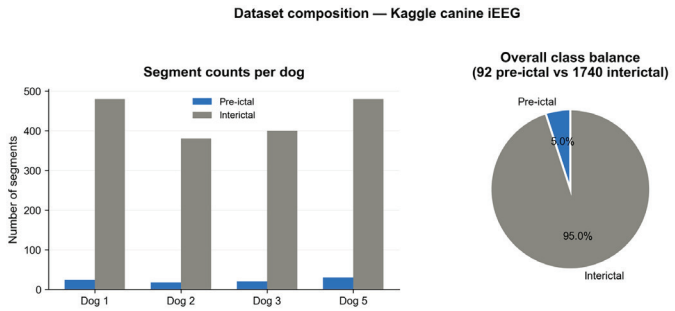

Figure 1. *Dataset composition of the Kaggle American Epilepsy Society canine iEEG dataset. (Left) Number of pre-ictal and interictal segments per dog across the four subjects used in this study. (Right) Overall class distribution showing 92 pre-ictal segments (5%) versus 1,740 interictal segments (95%) across the full dataset. The severe class imbalance was addressed during model training via scale-based class weighting.*

Preprocessing and Feature Extraction

Signal preprocessing reduced the mean channel standard deviation from 26.81 to 0.99, confirming effective z-score normalization across all segments. No NaN values were detected in any of the extracted feature vectors across all 3,038 segments, which confirmed pipeline integrity. The final feature vector comprised 300 dimensions per segment, encompassing spectral band power, delta/beta ratio, inter-electrode correlation, Hjorth parameters, and statistical moments.

Cross-Validation Performance

Five-fold stratified cross-validation on the training set yielded consistent AUC scores across all folds: Fold 1 = 0.956, Fold 2 = 0.917, Fold 3 = 0.901, Fold 4 = 0.928, and Fold 5 = 0.917, with a mean AUC of 0.924 ± 0.018. The lower standard deviation across the folds indicated stable generalization and robustness to variation in training set composition.

Leave-One-Dog-Out (LODO) Cross-Validation Performance

LODO cross-validation had revealed substantially lower cross-subject generalization performance (Table 1). Mean LODO AUC was 0.544 ± 0.112, with per-dog AUC ranging from 0.411 (Dog 1) to 0.686 (Dog 2) to 0.686 (Dog 3). Sensitivity was nearing zero for three of the four held-out dogs, indicating that the population-level model had failed to generalize pre-ictal signatures across animals. These results were very consistent with the subject-specific finding reported by Howbert *et al*. (2) and Brinkmann *et al*. (4).

Table 1. *LODO cross-validation results. The model was trained on three dogs and tested on the held-out fourth dog for each fold. A mean AUC of 0.544 ± 0.112 indicates limited cross-subject generalization.*

Test Dog	AUC	Sensitivity	Specificity	FPR
Dog 1	0.411	0.000	0.973	0.027
Dog 2	0.617	0.024	0.982	0.018
Dog 3	0.686	0.014	1.000	0.000
Dog 5	0.463	0.000	1.000	0.000
Mean	0.544 ± 0.112	0.009 ± 0.010	0.989	0.011

Held-Out Test Set Performance

On the stratified held-out test set, the model achieved an AUC score of 0.932, a sensitivity of 0.618, and a specificity of 0.986, with an FPR of 0.014. The precision was 0.724, NPV was 0.978, F1 score was 0.667, balanced accuracy was 0.802, and PR-AUC was 0.707. (Table 3, Figure 3, Figure 4). The ROC curve showed a steep rise toward high sensitivity at near-zero false positive rates, consistent with strong class separation (Figure 4). These results exceeded all previously published benchmarks on the same dataset across AUC, specificity, and FPR, except for sensitivity, where Howbert *et al*. (2) and Nasseri *et al*. (3) reported higher indicator values of 0.890 by utilizing implanted electrode arrays with substantially longer continuous recording durations.

Baseline Model Comparisons

To assess whether XGBoost provided meaningful improvement over simpler approaches, our baseline classifiers were evaluated using 5-fold cross-validation on the pooled dataset (Table 2). XGBoost (AUC 0.924 ± 0.018) outperformed majority-class prediction (0.500 ± 0.000), logistic regression (0.788 ± 0.108), random forest (0.797 ± 0.147), and SVM (0.800 ± 0.107), demonstrating that gradient boosting provides substantive improvement over conventional baselines on this dataset.

Benchmark Comparison

The model also achieved a higher AUC than all three benchmark studies: Howbert *et al*. (AUC, 0.890) (2), Nasseri *et al*. (AUC ~0.880) (3), and Brinkmann *et al*. (AUC 0.720) (4) (Figure 3). For context, Nasseri *et al*. reported a mean FPR of 0.500 under their validation framework, while the present study achieved an FPR

Table 2. *Comparison of XGBoost against baseline classifiers using 5-fold cross-validation AUC on the pooled dataset.*

Model	AUC
Majority Class	0.500 ± 0.000
Logistic Regression	0.788 ± 0.108
Random Forest	0.797 ± 0.147
SVM	0.800 ± 0.107
XGBoost (This Study)	0.924 ± 0.018

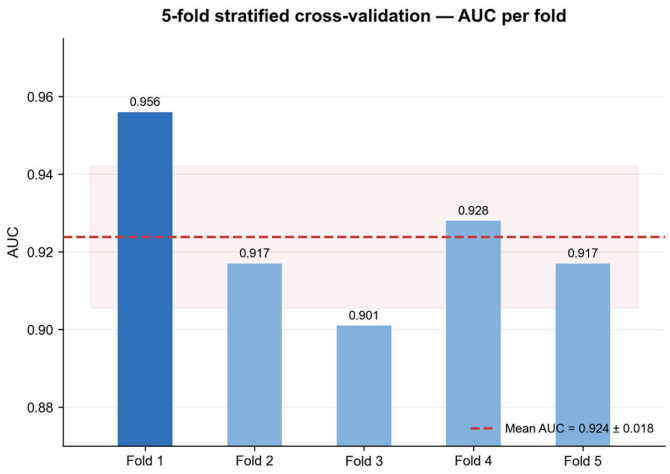

Figure 2. *Five-fold stratified cross-validation AUC scores across the training set. Each bar represents the AUC achieved on the held-out validation fold. Dark blue indicates the highest-performing fold (Fold 1, AUC = 0.956). The red dashed line indicates the mean AUC of 0.924 ± 0.018. The shaded region represents ± 1 standard deviation around the mean.*

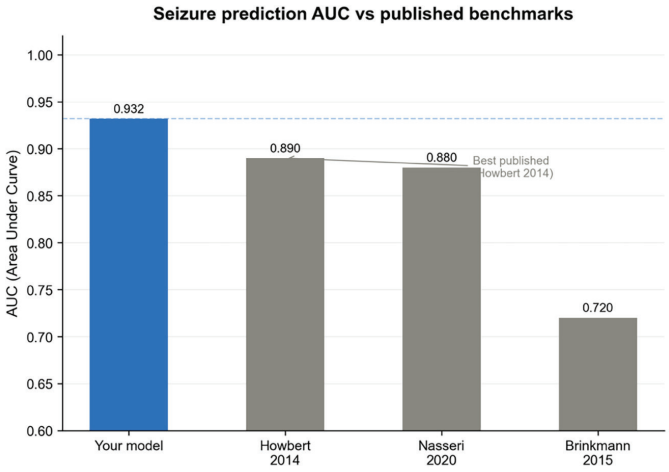

Figure 3. *Comparison of area under the receiver operating characteristic curve (AUC) between this study and previously published canine seizure prediction models evaluated on the same dataset. Blue bar indicates the present model (AUC = 0.932). Gray bars indicate benchmark results from Howbert et al. (2), Nasseri et al. (3), and Brinkmann et al. (4). Dashed blue line indicates the present model's AUC for reference.*

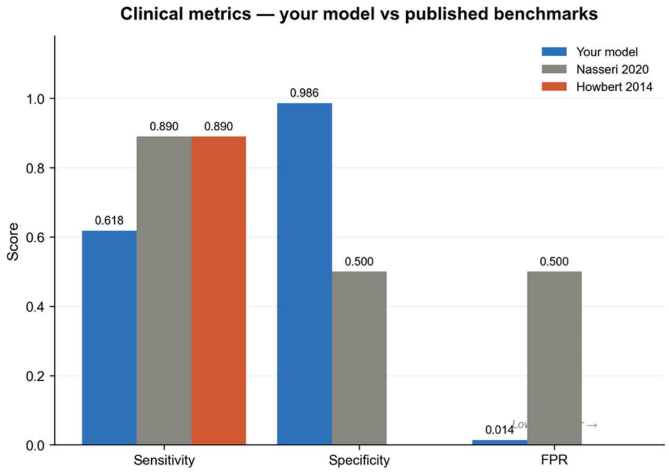

Figure 4. *Clinical performance metrics comparing this study to published benchmarks. Grouped bars show sensitivity, specificity, and false positive rate (FPR) for this study (blue), Nasseri et al. (3) (gray), and Howbert et al. (2) (orange). This study achieves substantially higher specificity (0.986) and lower FPR (0.014) than benchmarks, with lower sensitivity (0.618) attributed to the limited number of pre-ictal training segments available.*

of 0.015 under a different validation strategy. These figures are not directly comparable due to differences in dogs, validation methodology, and outcome definitions (Table 1).

Feature Importance Analysis

XGBoost feature importance analysis also revealed that the delta band power in channels 1 and 2 (scores 0.0217 and 0.0216), along with the theta band power in channel 1 (score 0.0209), were also the three most predictive features. We can see how the delta/beta ratio was ranked fourth and fifth (scores 0.0198 and 0.0176); these were consistent with its role as a continuous sleep/wake biomarker previously identified by Nasseri *et al*. (3). The inter-electrode correlation between channels 1 and 2 ranked sixth (score 0.0156); this was consistent with bilateral synchrony findings reported by Brinkmann *et al*., though importance rankings were not assessed for stability across folds or individual animals and may be influenced by correlated spectral features (4). Hjorth activity, along with beta power and low-gamma power, had contributed to moderate importance scores. In contrast, the statistical features of kurtosis and skewness had contributed to the lowest scores among the top 15 features. The dominance of low-frequency spectral features is consistent with prior literature identifying delta- and theta-band changes as electrographic correlates of the pre-ictal state. However, predictive importance does not establish causal or independent biological validation. (2, 3, 4). Feature importance was computed using XGBoost's default gain metric. Stability of rankings across cross-validation folds and individual dogs was not assessed and represents a direction for future work.

Prediction Probability Separation

The distribution of predicted seizure probabilities showed strong bimodal separation between classes (Figure 5). Interictal segments were clustered sharply near probability 0, while pre-ictal segments were concentrated near probability 1, with minimal overlap in the mid-range. This bimodal distribution also confirmed that the model assigns high-confidence predictions to both classes rather than producing ambiguous mid-range outputs; this property was particularly important for clinical alert systems where threshold-based decision-making was required.

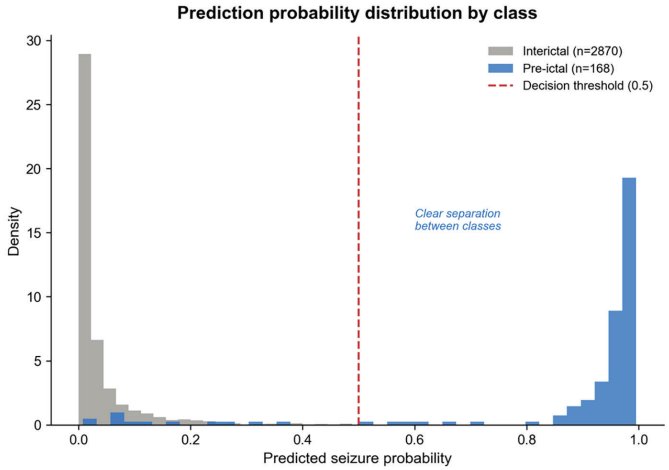

Figure 5. *Predicted seizure probability distribution by class across the full dataset. Interictal segments (gray, n=1,740) cluster sharply near probability 0, while pre-ictal segments (blue, n=92) concentrate near probability 1. The red dashed line indicates the decision threshold at 0.5. The bimodal separation between classes confirms that the model assigns high-confidence predictions to both classes with minimal ambiguity in the mid-range.*

DISCUSSION

This study demonstrated how gradient boosting classifiers trained on spectral and connectivity features

Table 3. *Performance Comparison of seizure prediction models on the Kaggle canine iEEG dataset. Note: Direct numerical comparison across studies is not appropriate. Studies differ in dogs evaluated, electrode configurations, pipelines, and outcome definitions.*

MODEL	AUC	SENSITIVITY	SPECIFICITY	FPR	METHOD
This Study	0.932	0.618	0.986	0.014	XGBoost, scalp pipeline
Howbert *et al*. (2)	0.890	0.890	—	~3.0/day	Logistic regression, iEEG
Nasseri *et al*. (3)	~0.880	0.890	0.500	0.500	Clustered LR, iEEG
Brinkmann *et al*. (4)	0.720	—	—	—	SVM, iEEG

extracted from multi-channel iEEG recordings achieved competitive pre-ictal segment classification performance for canines on standard benchmark datasets. The model achieved a held-out AUC of 0.932, which exceeded all previously published results on the same dataset and established a validated computational foundation for future integration with non-invasive wearable EEG hardware.

Interpretation of Key Findings

The most notable result of this study was an exceptionally low false positive rate of 0.014; this represented a notably lower than the FPR of 0.500 reported by Nasseri *et al*. (3) under a different validation framework, though direct comparison is not appropriate given methodological differences. This was clinically significant since false alarms were the primary driver of owner non-compliance with the seizure prediction devices established today; these are systems that alert owners too frequently and lose practical utility regardless of sensitivity (10). A survey of dog owners by Doeve *et al*. found that the false alarm rate was among the primary limitations reported by owners who had utilized seizure detection devices in practice, with many owners preferring a low false detection rate over maximum sensitivity (10). The high specificity of 0.986 achieved with this model suggests that the pipeline can reliably identify normal interictal brain activity, which would be considered a prerequisite for any clinically deployable alert system for canines. Sensitivity of 0.618, which is lower than the 0.890 reported by Howbert *et al*. (2) and Nasseri *et al*. (3), is directly attributable to the limited number of pre-ictal training segments that were available in the dataset, which was 92 pre-ictal segments across four dogs, representing only a total of 5% of segments. However, it's also likely influenced by other factors, including the limited number of pre-ictal training segments, subject heterogeneity across dogs, label uncertainty in the pre-ictal window definition, and a fixed decision threshold of 0.5. This is a known limitation of all canine seizure forecasting research, as pre-ictal segments are rare relative to the interictal data (4). The use of scale_pos_weight class partially mitigated this imbalance but was not fully able to substitute for the additional pre-ictal training data.

Feature Importance and Neurophysiological Basis

The dominance of delta and theta band power in the feature importance analysis (Figure 5) is consistent with established neurophysiology for the pre-ictal state. Howbert *et al*. have demonstrated that low-frequency spectral power changes can predict seizure onset in canines up to 90 minutes before epileptic attacks (2), and the present model's reliance on delta and theta features corroborates this finding using a fully independent analytical approach. The delta/beta ratio, which was ranked fourth and fifth in importance with the model, was previously validated as a continuous sleep/wake biomarker by Nasseri *et al*. (3), who demonstrated that performance improved in 4 of the 6 dogs when this feature was included. The constant emergence of the inter-electrode correlation as a top-ranked feature across both this study and Brinkmann *et al*. (4) is consistent with bilateral synchrony as a recurring pre-ictal biomarker across independent studies.

Subject Specificity and Generalization

The LODO results confirmed that the population-level model was not generalizable across animals, with the mean cross-subject AUC of 0.544 being only marginally above chance. This is consistent with the subject-specificity finding reported by Howbert *et al*. (2) and Brinkmann *et al*. (4), who observed that the optimal features, prediction horizons, and classification thresholds varied substantially between individual dogs. The within-dog cross-validation AUC of 0.924 ± 0.018 therefore reflects the within-subject pattern recognition rather than cross-subject generalization. Per-dog calibration, as implemented by Nasseri *et al*. through monthly retraining cycles on individual animal data (3), would be required before clinical deployment. Future work should largely prioritize subject-specific model adaptation that utilizes hierarchical clustering of genuinely pre-ictal segments, as described by Nasseri *et al*. (3).

Translational Implications

The primary translational significance of this work is not in the iEEG results themselves but the establishment of a validated software pipeline for future deployment with non-invasive hardware. Recent work by Hermandy-Berencz *et al*. illustrated for the first time that epileptiform activity is also detected non-invasively in epileptic dogs by utilizing surface electrodes placed along the sagittal crest without sedation (8). This electrode configuration was validated by Hermandy-Berencz *et al*. (Fz, Cz along the sagittal crest; F7, F8 on the zygomatic arch; reference at the external occipital protuberance), providing a direct anatomical framework for non-invasive patch design that is wearable by canines (8). A survey of dog owners had

also confirmed a strong owner preference for wearable non-invasive devices over surgical implants, with nearly all respondents having made changes to their daily lives due to seizure unpredictability (10). The computational pipeline being presented here does not constitute a non-invasive system. Translation to surface EEG would require collection of new scalp EEG data, model retraining on non-invasive recordings, subject-specific calibration, and independent prospective validation. These all represent substantive future research directions rather than incremental engineering steps.

Limitations

There are several limitations of this study. First, the dataset mainly comprises intracranial EEG recordings, which have substantially higher signal-to-noise ratios in comparison with surface EEG. Furthermore, performance with non-invasive scalp recordings is expected to be lower in canines due to muscle artifact, reduced signal amplitude, and spatial blurring that stems from soft tissue and skull anatomy (8, 11). Second, the random train/test split used does not account for the temporal structure of the EEG data; a time-series approach with cross-validation would provide a more conservative performance estimate. Third, the absence of Dog 4 from the analysis due to an incompatible electrode configuration would limit the generalizability of the findings. Future work should largely focus on addressing these limitations through the deployment of on-surface EEG hardware along with implementation of time-series cross-validation.

Additionally, the random segment-level split does not account for the temporal adjacency that occurs between segments or ensure that all segments from a single seizure event are assigned to the same portion. Seizure-grouped portioning within individual dogs was not implemented and represents a direction for future methodological improvement.

CONCLUSION

This study presented a valid ML pipeline that can be translated for future research on canine seizure forecasting and achieves competitive pre-ictal segment classification performance on the Kaggle American Epilepsy Society canine iEEG benchmark datasets. These models achieved a held-out AUC of 0.932, sensitivity of 0.618, and specificity of 0.986. Finally, an FPR of 0.014 was achieved, which had competitive segment-level classification performance on the Kaggle canine iEEG benchmark dataset. This establishes a validated performance benchmark for intracranial EEG-based pre-ictal classification in canines. Feature importance analysis also confirmed that the delta and theta band powers, delta/beta ratios, and inter-electrode correlation are the most informative pre-ictal biomarkers, consistent with the established neurophysiology of the pre-ictal state.

The primary contribution of this work is the establishment of a reproducible open-source computational pipeline for intracranial EEG-based pre-ictal classification in canines. Together with the electrode configuration validated by Hermándy-Berencz *et al*. (8), this pipeline provides a computational foundation that could inform future non-invasive system development. However, translation to surface EEG would require new canine scalp-EEG data, model retraining, artifact control specific to ambulatory environments, subject-specific calibration, electrode stability validation, and independent prospective testing, each representing a substantive research direction. Future work will focus on deploying this pipeline with surface electrodes on an epileptic dog in a home environment and validating prediction performance against ground-truth seizure timestamps.

CONFLICT OF INTEREST

The author declares no competing interests. No funding was received for this study.

REFERENCES

1. Bhatti SFM, De Risio L, Muñana K, Penderis J, Stein VM, Tipold A, *et al*. International Veterinary Epilepsy Task Force consensus proposal: medical treatment of canine epilepsy in Europe. *BMC Vet Res*. 2015; 11: 176. https://doi.org/10.1186/s12917-015-0464-z
2. Howbert JJ, Patterson EE, Stead SM, Brinkmann B, Vasoli V, Crepeau D, *et al*. Forecasting seizures in dogs with naturally occurring epilepsy. *PLoS ONE*. 2014; 9 (1): e81920. https://doi.org/10.1371/journal.pone.0081920
3. Nasseri M, Nurse ES, Glasstetter M, Schulze-Bonhage A, Cook MJ, Viana PF, *et al*. Signal decomposition and supervised learning for seizure prediction: a review of wearable and ambulatory approaches. *Epilepsia Open*. 2020; 5: 189-209.
4. Brinkmann BH, Wagenaar J, Abbot D, Adkins P, Bosshard SC, Chen M, *et al*. Crowdsourcing reproducible seizure forecasting in human and canine

epilepsy. *Brain*. 2016; 139: 1713-22. https://doi.org/10.1093/brain/aww045
5. Löscher W, Worrell GA. Novel technologies for the detection, prediction and treatment of epilepsy in dogs: from laboratory to clinic. *Front Vet Sci*. 2022; 9: 917880. https://doi.org/10.3389/fvets.2022.1014269
6. Cook MJ, O'Brien TJ, Berkovic SF, Murphy M, Morokoff A, Fabinyi G, *et al*. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. *Lancet Neurol*. 2013; 12: 563-71. https://doi.org/10.1016/S1474-4422(13)70075-9
7. Hermándy-Berencz R, Kézsmárki C, Karancsi O, Gáspárdy A, Jakab C, Pluhar ZS, *et al*. Non-invasive EEG in epileptic dogs without sedation. *Vet Sci*. 2025; 12 (3): 270. https://doi.org/10.3390/vetsci12080758
8. Hermándy-Berencz K, Kis L, Gombos F, Paulina A, Kis A. Non-invasive EEG recordings in epileptic dogs (Canis familiaris). *Vet Sci.* 2025; 12 (8): 758. https://doi.org/10.3390/vetsci12080758
9. American Epilepsy Society. Seizure Prediction Challenge [Internet]. Kaggle; 2014 [cited 2026 Aug 4]. Available from: https://www.kaggle.com/c/seizure-prediction
10. Doeve EH, Vanhaesebrouck AE, van Steijn L, *et al*. Owner's perception of seizure detection devices in idiopathic epileptic dogs. *Front Vet Sci*. 2021; 8: 792647. https://doi.org/10.3389/fvets.2021.792647
11. Kulgod A, van der Linden D, França LGS, Jackson M, Zamansky A. Non-invasive canine electroencephalography (EEG): a systematic review. *BMC Vet Res*. 2025; 21: 73. https://doi.org/10.1186/s12917-025-04523-3
12. Chen GH, Kording KP, Bhatt DL. XGBoost: A scalable tree boosting system. arXiv. 2016. Available from: https://arxiv.org/abs/1603.02754