

Discovering Neurophysiological Characteristics of Pathological High-Frequency Oscillations in Epilepsy with an Explainable Deep Generative Model

Authors: Yipeng Zhang¹, Atsuro Daida², Lawrence Liu¹, Naoto Kuroda³, Yuanyi Ding¹, Shingo Oana², Tonmoy Monsoor¹, Shaun A. Hussain^{2,4}, Joe X Qiao⁵, Noriko Salamon⁵, Aria Fallah⁶, Myung Shin Sim⁷, Raman Sankar^{2,4}, Richard J. Staba⁸, Jerome Engel Jr.⁸⁻¹¹, Eishi Asano³, Vwani Roychowdhury¹*, Hiroki Nariai^{2,4}*

¹ Department of Electrical and Computer Engineering, University of California, Los Angeles, CA, USA

² Division of Pediatric Neurology, Department of Pediatrics, UCLA Mattel Children's Hospital, David Geffen School of Medicine, Los Angeles, CA, USA

³ Department of Pediatrics and Neurology, Children's Hospital of Michigan, Wayne State University School of Medicine, Detroit, MI, USA

⁴ The UCLA Children's Discovery and Innovation Institute, Los Angeles, CA, USA

⁵ Division of Neuroradiology, Department of Radiology, UCLA Medical Center, David Geffen School of Medicine, Los Angeles, CA, USA

⁶ Department of Neurosurgery, UCLA Medical Center, David Geffen School of Medicine, Los Angeles, CA, USA

⁷ Department of Medicine, Statistics Core, University of California, Los Angeles, CA, USA

⁸ Department of Neurology, UCLA Medical Center, David Geffen School of Medicine, Los Angeles, CA, USA

⁹ Department of Neurobiology, University of California, Los Angeles, CA, USA

¹⁰ Department of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles, CA, USA

¹¹ The Brain Research Institute, University of California, Los Angeles, CA, USA

Character count title:	145 characters
Word count abstract:	195 words
Word count text total:	4413 words (Introduction, Results, and Discussion)
Word count Methods:	3378 words
Number of Tables:	2 for print, 3 for supplementary data
Number of Figures:	6 for print, 10 for supplementary data
Running Title:	Self-supervised HFO classification with deep learning

*Corresponding author:

Hiroki Nariai, MD, PhD, MS (hnariai@mednet.ucla.edu) and Vwani Roychowdhury (vwani@g.ucla.edu)

ABSTRACT

Interictal high-frequency oscillation (HFO) is a promising biomarker of the epileptogenic zone (EZ). However, objective definitions to distinguish between pathological and physiological HFOs have remained elusive, impeding HFOs' clinical applications. We employed self-supervised deep generative variational autoencoders to learn such discriminative HFO features directly from their morphologies in a data-driven manner. We studied a large retrospective cohort of 185 patients who underwent intracranial monitoring and analyzed 686,410 candidate HFO events collected from 18,265 brain contacts across diverse brain regions. The model automatically clustered HFOs into distinct morphological groups in the latent space. One cluster consisted of putative morphologically defined pathological HFOs (mpHFOs): HFOs in that cluster were observed to be associated with spikes and exhibited high signal intensity both in the HFO band (>80 Hz) at detection and in the sub-HFO band (10-80 Hz) surrounding the detection and were primarily localized in the seizure onset zone (SOZ). Moreover, resection of brain regions based on a higher prevalence of interictal mpHFOs better predicted postoperative seizure outcomes than current clinical standards based on SOZ removal. Our self-supervised, explainable, deep generative model distills pathological HFOs and thus potentially helps delineate the EZ purely from interictal intracranial EEG data.

Key words: HFO, pathological HFOs, artificial intelligence, unsupervised learning, machine learning

Abbreviations: HFOs: high-frequency oscillations. EEG: electroencephalogram. mpHFOs: morphologically defined pathological HFOs. SOZ: seizure onset zone. EZ: epileptogenic zone. iEEG: intracranial EEG. SEEG: stereotactic EEG. DL: deep learning. VAE: variational autoencoder. mArtifacts: Artifacts defined based on morphological analysis using the VAE model. spkHFOs: HFOs with spikes. non-spkHFOs: HFOs without spikes. HS: hippocampal sclerosis. FCD: focal cortical dysplasia. t-SNE: t-distributed stochastic neighbor embedding. ROIs: regions of interest. AUC: area under the curve. CT: computed tomography. MRI: magnetic resonance imaging. RMS: root mean square. SD: standard deviation. STE: short-term energy.

INTRODUCTION

Over a third of people with epilepsy do not respond to medication, thereby becoming potential candidates for epilepsy surgery.¹ Presently, surgical plans are mainly informed by neuroimaging and electroencephalogram (EEG), including interictal spikes and the seizure onset zone (SOZ). Nonetheless, the success rate of surgical intervention in achieving seizure freedom is less than ideal, varying between 50% and 75%.^{2,3} Discovering a biomarker that can precisely define the boundaries of the epileptogenic zone (EZ: the brain regions responsible for generating seizures) would be a significant breakthrough. Epilepsy studies in both humans and animals have indicated that interictal high-frequency oscillations (HFOs) in intracranial EEG (iEEG) is a promising spatial biomarker for delineating the EZ.⁴⁻⁷ Numerous retrospective studies have shown a correlation between the excision of brain areas generating HFOs and postoperative seizure freedom.⁸⁻¹¹

A fundamental challenge arises because HFOs with similar frequency ranges appear despite significantly different mechanisms: pathological HFOs and physiological HFOs.¹² Pathological HFOs are conceptually associated with epileptogenesis, occurs exclusively in epileptogenic brain regions. Physiological HFOs are related to cognitive and sensorimotor function and occur in healthy brain regions.¹³ Indeed, it has been shown that sparing brain regions with physiological HFOs from resection can result in seizure freedom.¹⁴⁻¹⁷ A recent clinical trial investigating the effect of HFO-guided intraoperative resective surgery excluded patients with occipital lobe epilepsy from the study due to the concern of abundant physiological HFOs in the visual cortex.¹⁸ Hence, using HFOs as a spatial biomarker for directing epilepsy surgery requires developing a method to distinguish between pathological and physiological HFOs. The resection of brain regions generating pathological HFOs would be associated with favorable postoperative seizure outcomes. Studies have shown pathological HFOs generated by abnormal synchronous burst firing are morphologically distinct from physiological HFOs generated by inhibitory synchronous postsynaptic potentials when recorded with research microelectrodes.^{12,13} However, typical iEEG signals are recorded using clinical macro electrodes, and algorithms to characterize such morphological differences based on conventional signal processing have proven insufficient.^{19,20}

Supervised machine learning approaches, especially deep learning (DL), have proven effective when human annotators accurately label large-scale datasets. DL models can learn complex patterns representative of labeled categories and then perform automated classification on unseen signals and images, leading to significant advances in medical imaging applications.^{21,22} If such frameworks are used to classify pathological HFOs, certified experts would be required to label large repositories of HFOs from diverse groups of patients accurately and consistently. However, accurately annotating pathological HFOs is difficult because there is no consensus on what constitutes pathological HFOs. Our recent studies have shown that instead of expert labels, clinical evidence from a channel's resection status, postoperative seizure outcomes, or functional brain mapping results can be used

to design a weakly supervised framework that yields relatively accurate labels for a small cohort size, as verified by retrospective postoperative seizure outcome prediction.^{23,24} Since these labels are still noisy—not all identified HFOs in resected areas or the SOZ are necessarily pathological—extending such a weakly supervised approach to ensure surgical outcomes for unseen patients would require a much larger patient cohort with postoperative outcomes. However, this approach faces challenges: first, an increasingly significant percentage of patients are being recorded through stereotactic EEG (SEEG) for diagnostic purposes or neuromodulation, and only a small percentage of such patients go through the brain resection; thus, data from patients undergoing diagnostic SEEG cannot be used in the aforementioned weakly supervised method, limiting the availability of the training dataset; second, while such SEEG patients still need accurate predictions of pathological HFOs to delineate the EZ, they would no longer be represented in the training set, thereby potentially limiting the applicability of the weakly supervised model to these patients.

One effective way to overcome such difficulties in using a supervised learning approach is to investigate the original observation directly: since pathological and physiological HFOs indeed differ in their biological mechanisms, leading to distinct morphologies, a generative AI model for the HFOs, trained in a self-supervised manner, should be able to simulate these underlying processes if a sufficiently large set of HFOs is provided. In particular, one should be able to discover the underlying discriminative morphological patterns without any labels. Autoencoders, such as variational autoencoders (VAEs),²⁵ is a DL framework based on a self-supervised training paradigm, which can provide such a capability and has been proven to be successful in many natural language processing,²⁶ computer vision tasks,²⁷ medical data, including EEG analysis.^{28,29} VAEs are generative models that learn efficient low-dimensional representations (latent encodings) of unlabeled data by requiring that the original high-dimensional data be reconstructed from their latent codes. It is particularly effective in discovering clusters in the latent space, where each cluster corresponds to a different mechanism that led to the generation of data in it. Therefore, clusters automatically produced in the latent space of a VAE can potentially capture different underlying generative mechanisms. The HFO morphologies observed within these different clusters would have unique characteristics and can be interpreted through direct visualization and interpolations in the latent space.

In this study, we utilized a large multi-institutional cohort of 185 patients with epilepsy who underwent iEEG monitoring with grid or SEEG electrodes, providing comprehensive coverage of both deep and superficial brain regions. We developed a VAE framework that analyzes 686,410 HFOs to characterize three different classes of HFOs: 1) morphologically defined putative pathological HFOs (mpHFOs), 2) morphologically defined putative non-pathological HFOs (non-mpHFOs), and 3) morphologically defined artifacts (mArtifacts) of extra-cerebral origin. We leveraged the interpretability of VAEs to explore various neurophysiological characteristics of HFOs, such as (i) characterization of cross-frequency power distributions of HFOs, (ii) exploring the dependence of HFO morphologies on variables, including the origin of iEEG datasets and patient-specific variables such as age, sex,

and epilepsy etiology based on histopathology (referred to pathology in the rest of the manuscript), and (iii) investigating HFO morphologies based on anatomical locations of the brain.

Our results showed that mpHFOs, VAE-discovered putative pathological HFOs, generally originated within the SOZ and were associated with spikes. We also demonstrated that a latent space discovered by a VAE model exhibited refined characteristics of pathological HFOs within time-frequency plots. Finally, we showed that the proportion of the removal of mpHFOs outperformed the resection status of the SOZ in predicting postoperative seizure freedom with cross-validation. Our results demonstrate that deep generative models can characterize neurophysiological signatures of pathological HFOs. From a short duration of interictal iEEG data, one can potentially infer the EZ by differentiating pathological from physiological HFOs, thereby enhancing the use of these biomarkers in guiding epilepsy surgery.

RESULTS:

Building an explainable self-supervised DL algorithm for HFO analysis:

We studied 185 patients (91 females) from two centers who met the eligibility criteria (**Table 1**). The median age at surgery was 13 years (range: 2–44 years). A total of 18,265 artifact-free electrode sites (median: 106 per patient; range: 29–152) within 34 regions of interest (ROIs) were available for analysis (**Figure 1a**). There were 1,670 electrode sites sampled within the SOZ (mean: 9.42 per patient; frontal: 437; temporal: 464; parietal: 393; occipital: 135; limbic: 241), and 7,732 sites sampled within non-epileptogenic brain regions, defined as spared brain regions in patients with postoperative seizure freedom (frontal: 2943; temporal: 1806; parietal: 1853; occipital: 635; limbic: 495) (**Table 2**). The median duration of analyzed EEG data for the UCLA grid/strip dataset was 91.5 min [IQR: 90.6–94.8 min], and for the UCLA SEEG, it was 90.3 minutes [IQR: 87.1–96.8 min]. The median analyzed EEG recording duration for the Detroit grid/strip dataset was 5.3 minutes [IQR: 5.1–5.7 min]. In total, 686,410 putative HFOs were detected from all the datasets by the PyHFO, an automated HFO detection platform, using both STE and MNI detectors.³⁰ The median rate of HFOs (number of detections/min) within SOZ contacts was 3.11 (range: 0.73–7.17), and the median rate of HFOs within the non-epileptogenic contacts was 1.08 (range: 0.63–2.15) across the ROIs. All the detected HFOs were used for the unsupervised VAE model training (**Figure 1b**). Using the trained VAE model, such HFOs were automatically classified into artifacts (mArtifacts) and non-artifactual HFOs. A cluster for mArtifacts was determined based on detected HFOs' high reconstruction loss due to the high variability of their morphologies (**Figure 1c**). All the non-artifactual HFOs were further classified into two clusters—the cluster with a higher resection percentage in seizure-free patients after resection was deemed pathological. HFOs originating from such a pathological cluster were defined as

mpHFOs, while those from another cluster were defined as non-mpHFOs. All the detected HFOs were also classified into artifacts, HFOs with spikes (spkHFOs), and HFOs without spikes (non-spkHFOs) using our previous classification algorithm²³ (**Supplementary Table 1**) for the subsequent correlational analysis, including interpretability analysis (**Figure 1d**). The detailed workflow of the study is described in **Supplementary Figures 1-3**. There were 163 patients who underwent resective surgery, and 110 patients (67.5 %) achieved seizure freedom. Of the patients who had resection, pathology results were as follows: focal cortical dysplasia (FCD) (41.1%), hippocampal sclerosis (HS) (6.7%), tumor (19.0%), and others (33.1%).

Characterization of pathological HFOs based on the self-supervised VAE algorithm:

With five-fold subjective-wise cross-validation, the VAE classification was aligned with conventional knowledge that spkHFOs are more likely pathological: 93.56% of mArtifacts were predicted as artifactual HFOs, 80.1% of the mpHFOs were predicted as spkHFOs and 84.87 % of the non-mpHFOs were predicted as non-spkHFOs (p-value < 0.001). Furthermore, we visualized these findings by projecting the latent codes of each event onto a two-dimensional plane using t-distributed Stochastic Neighbor Embedding (t-SNE). (**Figure 2a, b**). Such clustering results were consistent throughout the five folds (**Supplementary Figure 4**). Moreover, the mpHFO rate (count/channel/min) was significantly higher than the non-mpHFO rate (count/channel/min) in the SOZ channels across three datasets (p-value < 0.01 for UCLA grid/strip, Detroit grid/strip, and UCLA SEEG dataset) and did not exhibit significant difference within the non-soz channels (p-value = 0.1 for Detroit grid/strip dataset, p-value = 0.49 for UCLA SEEG dataset, and p-value = 0.38 for UCLA grid/strip dataset) (**Figure 2c**). We then investigated the neurophysiological characteristics of mpHFO and non-mpHFO directly from the data. The morphological analysis of the time-frequency plot demonstrated that mpHFOs had higher amplitude values throughout the HFO band (≥ 80 Hz), around the center point (0 ms, where HFOs were detected) than non-mpHFOs (**Figure 2d, e**). Furthermore, there were statistically higher values of mpHFOs at the sub-HFO band (10-80 Hz) throughout the time window compared to non-mpHFOs. These bands together lead to a "hanging bell" template in the time-frequency plot (**Figure 2e**). Such a template showed consistency regardless of the variables, including sex, the origin of the dataset, pathology, and age categories (**Supplementary Figure 5**). mpHFOs also showed distinct characteristics in the frequency domain with FFT analysis. The mpHFOs demonstrated relatively high amplitude within the sub-HFO band (10-80 Hz), while non-mpHFOs, likely representing physiological HFOs, exhibited discrete peaks with the HFO range, maximal at 87.7 Hz (**Figure 2f, g**). To understand the sub-HFO band component of the "hanging bell" template in mpHFOs, we filtered mpHFOs with a 10-80 Hz bandpass to evaluate the corresponding time-series data. We found a spike-appearing EEG signal after this procedure (**Figure 2h-j**) but none in non-mpHFOs (**Figure 2k-m**).

HFO morphology analysis based on dataset origin, sex, age, and pathology:

We investigated whether HFO morphology was affected by variables including recording sites/type (UCLA grid/strip, UCLA SEEG, Detroit grid/strip), sex (female vs. male), age (0-5, 6-10, 11-15, 16-20, and 21+), and pathology (HS, FCD, Tumor, others). We first visualized the distribution of the latent space by projecting latent codes of HFOs into 2D space using t-SNE and color-coded with different subcategories within that specific demographic variable (**Figure 3a, c, e, g**). Three hundred projected codes were randomly sampled from each subject to achieve easier visualization. Then, the accuracy of a logistic regression model trained from actual data was compared with a model trained with label-permuted data. HFO morphology based on the dataset sources could not be differentiated by the classifier (p-value = 0.13, **Figure 3b**). Similarly, HFO morphology, based on demographic information such as sex, age group, and pathology, could not be differentiated (Sex: p-value = 0.44; Age Group: p-value = 0.10; Pathology: p-value = 0.83) (**Figure 3d, f, h**).

HFO morphology analysis based on ROIs (anatomical location):

We then investigated whether HFOs generated from different anatomic locations exhibit distinguishable morphologies (frontal, temporal, parietal, occipital, and limbic). Initially, we generated a plot of the latent space for HFOs derived from the preserved regions in patients who remained seizure-free after resection, presumably physiological HFOs. We found a cluster of HFOs corresponding to the occipital region, which were distinguishable from other areas (**Figure 4a, b**). The confusion matrix showed that such occipital HFOs could be successfully distinguished by the classifier, which achieved an average of 62% accuracy (**Figure 4c**). To further characterize such distinguishable features, we plotted the summation of the time-frequency plot of each HFO event in different anatomical locations (**Figure 4d**). The HFOs derived from the occipital area showed distinct peak frequency and power ratio compared to HFOs generated from other anatomical areas (**Supplementary Figure 6**). Nevertheless, HFOs generated from the SOZ did not demonstrate distinctive morphological differences across the anatomical locations using the same method (**Figure 4e-h**).

Disentanglement of the latent space to establish neurophysiological characteristics of pathological HFOs:

We first plotted each dimension of the latent space and colored them with mpHFO and non-mpHFO for a specific fold (**Supplementary Figure 2: Latent dimension interpretation**); the mpHFO and non-mpHFO exhibited different distributions in some of the dimensions but did not in the others (**Supplementary Figure 7**). We then used the decoder of the VAE to understand the meaning of each latent dimension by perturbing each dimension in the latent space and visualizing the reconstructed image. Specifically, in each fold, we first took the

mean of the cluster of the mpHFO and non-mpHFO (**Supplementary Figure 2: Generation of representative mpHFO and non-mpHFO morphologies based on respective cluster centroids**) and then perturbed a single dimension from the smallest value (1 percentile) of that dimension to the largest value of that dimension (99 percentile) (**Supplementary Figure 2: Latent dimension interpretation**). We found a dimension that demonstrates the separation of mpHFO and non-mpHFO. The morphological transition of the generated images showed more power within the "hanging bell" template region for both mpHFO and non-mpHFO when we traversed the value of that dimension (**Figure 5a, b**). At a population level, we randomly sampled actual latent code from both mpHFO and non-mpHFO and traversed that specific dimension of each latent code from the smallest (1% tile) to the largest (99% tile); the decoded images showed an increase of the power within the template region (**Figure 5c**). We fed the decoded time-frequency plots through our inference pipeline and plotted the distribution of the model confidence (**Figure 5d**), demonstrating the increasing model confidence toward mpHFO as we increased the value of that dimension in the latent space. Furthermore, we also captured dimensions that, while not separating mpHFOs and non-mpHFOs, still manifested interpretable neurophysiological properties. Within the same fold, a dimension represented the HFO's peak frequency (**Figure 5e, f**), and subsequent population-level analysis indicated that an ascent in this dimension inversely correlated with the peak frequency in the time-frequency plots (**Figure 5g**) while the model confidence value did not change towards either the mpHFO and non-mpHFO (**Figure 5h**). Similarly, the slow-wave dimension (10-20Hz) was revealed (**Figure 5i, j**), where an increase in dimension value directly paralleled an increase in power within the slow-wave band in the decoded images (**Figure 5k**), but an increase in power within that dimension did not affect the model confidence towards either the mpHFO and non-mpHFO (**Figure 5l**).

Prediction of postoperative seizure outcomes using the classified HFOs:

Based on the hypothesis that mpHFOs are pathological and their resection status can forecast postoperative seizure freedom, we investigated two prediction models. We first applied a logistic regression model with five-fold cross-validation. Notably, utilizing the resection ratio of the mpHFOs (AUC = 0.63) demonstrated superior separation capability compared to using the resection ratio of unclassified HFOs (AUC = 0.53) and traditional expert-driven classification, spkHFO (AUC = 0.61) (**Figure 6a**). Further, a multivariable logistic regression model was fitted, incorporating subjects' demographic data as the baseline (age as a continuous variable and sex as a categorical variable) and the status of SOZ resection during surgery. The mpHFO resection ratio again showed enhanced classification performance over the spkHFO (AUC = 0.72 vs. 0.70) (**Figure 6b**). Additionally, we explored the feasibility of constructing a non-linear predictive model using a random forest classifier (**Figure 6c**). The random forest aimed to forecast post-surgical seizure freedom in the test set based on subject-specific features identified in the training set through a subject-wise five-fold cross-validation, ensuring

that our classifier was evaluated under consistent and rigorous conditions (**Supplementary Figure 3**). Our findings revealed that the random forest model, trained exclusively with the resection ratio of mpHFOs, exhibited better (p-value < 0.01, for both unclassified HFOs and spkHFO) predictive performance (F1 = 0.74) compared to using unclassified HFOs (F1 = 0.69) and spkHFO (F1 = 0.66). The model's predictive power improved when incorporating more features, such as subjects' demographic information and the SOZ resection status. Specifically, the combination of demographic data and mpHFO resection ratio (F1 = 0.77) outperformed (p-value < 0.01) traditional clinical predictions based on demographic information and SOZ resection status (F1 = 0.73). Furthermore, a comprehensive model including all features, including demographic data, SOZ resection, and mpHFO resection ratio (F1 = 0.81) demonstrated superior (p-value < 0.01) predictive power over the traditional expert-driven HFO classification including demographic data, SOZ resection, and spkHFO resection ratio (F1 = 0.80) (Detailed numbers are listed in **Supplementary Table 2**). We also conducted an ablation study on predicting surgical outcomes by only using subjects with a higher number of HFOs (**Supplementary Figure 10**): the mean F1 score remains consistent, demonstrating the robustness of our self-supervised framework.

DISCUSSION

Leveraging the innovative self-supervised deep learning approach, we aimed to characterize pathological HFOs using a cohort of 185 epilepsy patients who underwent intracranial monitoring with grid or stereotactic EEG electrodes across 18,265 brain contacts from 34 ROIs. We analyzed 686,410 HFOs obtained via fully automated detectors. The novel VAE method and extensive data allowed us to control for variables, including recording sites/types, sex, age, pathology, and anatomical location. By analyzing the latent space of VAEs, we identified morphological features of pathological HFOs, notably high signal intensity within the HFO band at detection, extending across the sub-HFO band (10-80 Hz). These features remained consistent across all studied variables. Incorporating the resection status of the identified pathological HFOs into prediction models markedly enhanced the accuracy of postoperative seizure freedom forecasts, outperforming traditional approaches that utilize the resection status of the SOZ.

Until now, the field of HFO research has encountered significant challenges. HFOs can be generated in healthy brain regions, adding complexity to their interpretation. Despite promising results from numerous retrospective studies,⁸⁻¹⁰ a clinical trial failed to demonstrate the utility of HFOs in improving postoperative seizure outcomes.¹⁸ The major challenge in HFO research is distinguishing between pathological HFOs that need to be targeted for resection and physiological HFOs that can be spared during resection. Attempts to differentiate these include identifying HFOs associated with spike-wave discharges,³¹⁻³³ which are generally considered indicative of pathology. However, there is no uniform method to define HFOs with spikes, and experts often

exhibit inconsistencies and poor inter-rater reliability when using visual annotations.³⁴ Moreover, simple analyses of HFO features such as frequency, amplitude, and duration have been ineffective in differentiating pathological from physiological HFOs.^{19,20} While fast ripples (250–500 Hz) may more accurately delineate epileptogenic zones than ripples (80–250 Hz), they are detected less frequently.^{35,36} Strategies like adjusting the detection rate of HFOs to match region-specific normative values can help estimate the degree of pathology in the corresponding brain regions,^{37,38} it cannot reliably classify each HFO event as pathological or physiological. This uncertainty complicates clinical decision-making, especially when treatment decisions hinge on a limited number of detected HFOs.

In this study, we adopted a fundamentally different approach. Instead of pre-selecting potential pathological HFO features, we implemented a data-driven, self-supervised strategy using VAE to define pathological HFOs. In the field of HFOs, prior studies have explored various types of neural networks, including convolutional neuronal networks,^{39,40} long short-term memory,^{41,42} and transformer⁴³, to classify HFOs into pathological and physiological based on human-annotated labels. This reliance limits scalability on large datasets and is constrained by the subjectivity inherent in expert labeling. Our method is contingent on the premise that pathological HFOs possess distinct morphological characteristics that are separable from their physiological counterparts. This hypothesis is biologically plausible, given that HFOs with similar frequency ranges emerge via markedly different generative mechanisms: Pathological HFOs arise from abnormal synchronous burst firing, whereas physiological HFOs result from inhibitory synchronous postsynaptic potentials.^{12,13} These morphological differences, potentially too subtle for detection with traditional signal processing methods using clinical macroelectrodes, may be effectively discerned through our innovative approach. The VAE is a DL framework capable of efficiently learning low-dimensional representations of unlabeled data.^{25,44} Without providing any labels associated with the input, an unsupervised classification can be achieved by reconstructing the original high-dimensional data from their latent codes. It excels at uncovering clusters in the latent space, with each cluster representing a distinct mechanism underlying the data generation. Consequently, clusters automatically formed in the latent space of a VAE have the potential to capture various morphological features and classify artifacts and HFOs with or without spikes, as demonstrated in our study. The diverse HFO morphologies generated by these clusters possessed distinct traits, allowing for interpretation through direct visualization and interpolations in the latent space, which enabled us to interpret what the algorithm has learned. The algorithm provided us with intuitive and interpretable latent spaces, such as a pathological template domain (the "hanging bell" template on a time-frequency plot), peak frequency domain, and slow-wave domain.

Our results are compelling, as the morphologically defined pathological HFOs identified by our deep generative model exhibited features commonly associated with expert-acknowledged pathological HFOs, such as spikes. While "spike" annotations can be performed by human experts or computationally, there is no standardized method for determining what constitutes pathological spikes. Our deep generative model objectively

identified distinctive features of the hanging bell template within the time-frequency plot. Specifically, pathological HFOs were associated with higher signal intensity within the sub-HFO band, peaking at 23 Hz. Such sub-HFO band EEG signals seemed to emerge as spikes in time series data. Thus, rather than arbitrarily define spikes to analyze HFOs with spikes, this data-driven approach to characterize the "spike" component associated with pathological HFOs seemed more unbiased. Another evidence of the pathological nature of mpHFOs was that they were primarily localized around the SOZ, although the model was trained without a SOZ label.

Our findings provide clinically significant insights; the resection ratio of mpHFOs outperformed traditional unclassified HFOs and spkHFOs in predicting postoperative seizure freedom. This supports the idea that mpHFOs are more indicative of pathological HFOs than unclassified HFOs and spkHFOs. The potential clinical utility is evident as the resection ratio of mpHFOs exceeded the current clinical standards regarding the removal status of SOZ in predicting postoperative outcomes. The mpHFOs can be analyzed from short interictal EEG data, offering the potential to reduce the duration of EEG monitoring, hospital stays, and associated costs for the patient. Furthermore, combining the resection status of the SOZ with the proportion of mpHFO resection further enhanced the prediction, suggesting an additive effect when mpHFOs are combined with the current clinical standard.

Another notable discovery was that non-pathological HFOs originating from the occipital lobe displayed distinct morphological features compared to HFOs from other brain regions. Physiological HFOs were reportedly abundant in the occipital lobe^{15,16,45} and also showed distinct coupling with slow waves.⁴⁶ This study added essential findings in the literature to establish the unique morphology of HFOs originating in the occipital lobe. These findings could potentially overcome the limitation faced by the HFO trial, which necessitated the exclusion of subjects with occipital lobe epilepsy due to the likelihood of abundant physiological HFOs in such patients.¹⁸

There are several factors to consider when interpreting these results. The study was conducted using only macroelectrode recordings, and while we hypothesized that the different biological mechanisms of pathological and physiological HFOs would manifest differently in these recordings, we could not verify the actual neurophysiological mechanisms at the single-neuron level. Although our sample included 185 patients, only 18 were studied using SEEG, limiting our ability to sample from deeper brain areas. With more balanced coverage of both superficial and deep areas, network analysis to account for HFO propagation may provide better accuracy in the prediction of postoperative surgical outcomes.⁴⁷ Since we aimed to investigate the HFOs' morphology as a function of age, we primarily included the pediatric population in the study. Expanding the inclusion of the adult population will be needed to generalize our findings. Additionally, the majority of our EEG data comprised short recordings, typically around five minutes from the initial night during sleep. Also, the sampling frequency of the Detroit dataset was limited up to 1,000 Hz, limiting the analysis of the fast ripple band (250-500 Hz). Although our results suggest that the peak frequency of HFOs did not affect pathological classification, further investigation

of the fast ripple band will be needed. Longer, multi-day recordings might reveal HFOs with potentially varying morphologies over time.⁴⁸ The vigilance state should also be considered,⁴⁹ as morphological differences between various sleep stages and wakefulness remain under-investigated. Finally, our sample size may not be sufficient to fully characterize the subtle differences in HFO morphology associated with different types of epilepsy pathologies. Conditions like focal cortical dysplasia (FCD) and tumors have distinct subtypes whose morphological differences might become clearer with a larger number of cases.

In the foreseeable future, we plan to collaborate with additional institutions to test the generalizability of our approach. Analyzing publicly available datasets will also be considered. Analyzing multi-day recordings from different sleep stages presents challenges, yet such studies could be feasible with our fully automated algorithm. Once we establish definitions for pathological HFOs across different anatomical regions, we can potentially integrate this data with other types of information, such as neuroimaging studies, to guide surgical resections. We aim to test the utility of this approach prospectively and may eventually consider a clinical trial to verify whether HFO-guided resections truly improve postoperative seizure outcomes.

METHODS:

Patient cohort: This was a multi-institutional retrospective cohort study. The inclusion criteria consisted of [a] simultaneous video-iEEG recording for epilepsy surgery between August 2016 and December 2023 at UCLA Mattel Children's Hospital or between January 2007 and May 2018 at Children's Hospital of Michigan, Detroit, [b] iEEG sampling rate of at least 1,000Hz, [c] iEEG contained at least an artifact-free 20 min slow-wave sleep epoch at least two hours apart from clinical seizure events, and [d] known postoperative seizure outcomes over one year. The exclusion criteria included [a] undergoing hemispherectomy or hemispherectomy, and [b] the presence of massive brain malformations (such as megalencephaly and perisylvian polymicrogyria) or previous surgeries that make it difficult to identify brain anatomy during the iEEG study. The institutional review board at UCLA and Wayne State University have approved the protocol. We obtained written informed consent from patients or the guardians of pediatric patients.

Patient evaluation: All patients with medically refractory epilepsy referred during the study period underwent a standardized presurgical evaluation, which—at a minimum—consisted of inpatient video-EEG monitoring, high resolution (3.0 T) brain magnetic resonance imaging (MRI), and 18 fluoro-deoxyglucose positron emission tomography (FDG-PET), with MRI-PET co-registration. The margins and extent of resections were determined mainly based on the SOZ, clinically defined as regions initially exhibiting sustained rhythmic waveforms at the

onset of habitual seizures. In some cases, the seizure onset zones were incompletely resected to prevent an unacceptable neurological deficit.

iEEG recording: Macroelectrodes, including platinum grid electrodes (10 mm intercontact distance) and depth electrodes (platinum, 5 mm intercontact distance), were surgically implanted. The placement of intracranial electrodes was guided by the results of scalp video-EEG recording and neuroimaging studies. All electrode plates were stitched to adjacent plates, the edge of the dura mater, or both to minimize movement of subdural electrodes after placement. Regarding the SEEG placement, BrainLab Elements software was used for planning the electrodes to the intended targets using T1-weighted sequences, and the trajectories were guided by a gadolinium-enhanced T1-weighted MRI. Both institutions obtained iEEG recordings using Nihon Kohden Systems (Irvine, California, USA). The sampling frequency was set at 1,000 Hz in Detroit and at 2,000 Hz in UCLA upon acquisition.

Acquisition of three-dimensional (3D) brain surface images: We obtained preoperative high-resolution 3D magnetization-prepared rapid acquisition with gradient echo (MPRAGE) T1-weighted image of the entire head. Using the FreeSurfer scripts, we created the averaged surface image to which all electrode locations were spatially normalized.^{23,38} In cases where the software failed to detect the pial surface accurately due to insufficient cerebral myelination, we manually delineated the pial surface using the Control Point function. For patients younger than two, we used the Infant FreeSurfer software package to reconstruct the surface image.⁴⁶ The averaged surface image functioned as the template for the analysis of anatomical location.

Anatomical labeling and determination of ROIs: For the dataset from UCLA, each implanted contact was labeled visually according to the Desikan-Killiany-Tourville atlas.⁵⁰ The location of electrodes was directly defined within a FreeSurfer-based 3D surface image using post-implant computed tomography (CT) images using Brainstorm software.²³ For the dataset from Detroit, all implanted subdural contacts were coregistered with 3D surface images within the FreeSurfer with an FSaverage vertex label.³⁸ We defined 34 ROIs for further analysis (**Table 2**). For the data harmonization between the two institutions, the FSaverage vertex of Detroit datasets was converted to MNI coordinates.⁵¹ Finally, these data were combined with UCLA patients, which were projected to the MNI normalized space under Brainstorm for the co-registration image (**Figure 1a**).

iEEG data pre-processing: We used a customized common average reference for the grid²³ and a bipolar montage for SEEG data. The EEG was first resampled to be the same sampling frequency of 1,000 Hz, and a band-reject filter was to reject the 60 Hz and its harmonics with a stopband of 2 Hz. iEEG channels not recording from grey matter (e.g., located outside of the brain) or otherwise deemed 'bad' (e.g., excessively noisy or

artificial) by the clinicians were discarded from the analysis. All EEG data pre-processing and analysis were performed using Python 3.9.1 (Python Software Foundation, Wilmington, DE).

Automated HFO detection: HFOs were detected by the PyHFO platform using both STE and MNI detectors to enhance the sensitivity of detection.³⁰ For STE and MNI detectors, the detection parameters were set to default (**Supplementary Table 3**), except that the frequency band for filtering the signal was set to be 80-300 Hz to accommodate a minimum 1,000 Hz sampling frequency.

Overall training and inference method: The overall study flow is outlined in **Figure 1** and **Supplementary Figure 1-3**.

Morphology-based pathological HFO classification by deep generative model

Subject-wise k-fold Cross-validation: To thoroughly test our method, we use a subject-wise five-fold cross-validation. In this setup, for each fold, we set aside 20% of the subjects as a test set (controlling sampling from UCLA grid/strip, UCLA SEEG, and Detroit grid/strip, ensuring that over five folds, every subject is tested once. All remaining data became the training set. Within the training set, we then randomly sampled ten subjects uniformly from three datasets for the validation set.

Feature Representation: Each event detected by the automatic detector was represented by a time-frequency plot (Morlet wavelet transform), as the time-frequency plot is known for representing spatial-temporal information for each HFO event. The time-frequency plot spans ± 285 ms (centered in the middle of the HFO event) in time and 10-290 Hz in frequency. This plot was resized to a 64 x 64 resolution, and the value was normalized within the range of 0 to 1 before being sent to the VAE encoder.

VAE architecture: VAE emerges as a prominent deep learning mode. Distinct from traditional autoencoders, VAE introduces a probabilistic approach; the encoder encodes inputs as distributions. During the decoding process, the model stochastically samples from these distributions, enabling it not only to accurately reproduce known data but also requires the decoder to generate new similar data. The loss of training the VAE, reconstruction loss, and variational regularization, the reconstruction loss measures the distance from the input to the output, and the variational regularization ensures the distribution generated from the encoder follows the normal distribution. We've chosen ResNet⁵² as the backbone for both the encoder and decoder, drawing inspiration from its widespread application in the community of computer vision, as evidenced by sources, as well as the successful application of using ResNet in capturing morphological information within the time-frequency plot.^{23,24} The designated size for the latent space was fixed at 8. As an ablation study, we set the latent dimension to 16 and demonstrated that it would produce more redundant latent space by visualization (**Supplementary Figure 8**).

Self-supervised VAE training: Our VAE underwent training over 80 epochs with the main goal of reconstructing the input image (**Figure 1a**). During the training, we used Adam optimizer with a learning rate of $3e-4$ and batch size of 512. Furthermore, the time-frequency plot was augmented by randomly flipping at the time axis to improve the model's generalization ability. To ensure an even representation of each subject during every epoch, we employed a stratified sampling method, capping the sample number at 2,500 per subject in each epoch of training. Moreover, we selected perceptual loss as our reconstruction loss criterion to capture the morphological discrepancies more effectively between input and output images. We adopted the beta-VAE⁴⁴ in training to better disentangle the latent space; therefore, the final loss function $loss = perceptual\ loss + \beta * KL\ divergence$, where $\beta = 0.1$. The model iteration (epoch) exhibiting the minimum validation loss was the one selected for subsequent utilization in the unsupervised classifier assembly and inference.

Unsupervised discovery of HFO clusters: We designed a hierarchical (two-stage) Gaussian Mixture Model (GMM) based clustering pipeline to discover morphologically different classes of HFOs (**Figure 1c**). Once the VAE had been trained, we processed the entire training dataset through the VAE encoder, thereby extracting the associated latent codes (**Figure 1b**). After this extraction, at the first stage, we employed the GMM clustering algorithms, setting the cluster number, k , to 2. We adopted stratified sampling to balance the contribution of information from each subject. We capped the number of events from each patient at 10,000. The cluster with high reconstruction loss was defined as the cluster of artifacts of extra cerebral origin, as the artifact is diverse in morphology and the VAE cannot properly decode. Then, at the second stage, within the cluster that was classified as non-artificial events, we trained another GMM to cluster them further into two clusters. We capped the number of events from each patient at 2000. To allocate either a pathological (mpHFO) or physiological (non-mpHFO) label to each cluster, we employed a minimalistic use of clinical data. The cluster with a higher resection percentage in seizure-free patients after resection was deemed pathological. As part of an ablation study, we also explored different stratified sampling methods, which demonstrated similar performance in predicting surgical outcomes (**Supplementary Figure 9**).

HFO morphology inference pipeline: To predict the HFO events from a new subject, time-frequency plots were first computed. Then, these time-frequency plots were sent into the VAE encoder to extract the latent codes. The trained hierarchical GMMs predicted each latent code and assigned class labels: mpHFO, non-mpHFO, and mArtifact for each detected event.

mpHFO characterization

VAE-based HFO morphology inference pipeline:

Time-frequency plot characteristics of pathological and physiological HFOs: We investigated whether the time-frequency scalogram of mpHFOs identified by the VAE exhibited a distinction from that of non-mpHFOs (**Figure 2d, e**). We used similar methodologies outlined in preceding studies.^{23,24} Specifically, we executed a t-test across every pixel in the time-frequency scalogram comparing the predicted mpHFOs and non-mpHFOs. The foundational null hypothesis is that scalogram values associated with mpHFOs are higher than those linked with non-mpHFOs. We binarized the p-value less than 0.05 (if the p-value is less than 0.05, the value is 1, and vice versa). The characteristics map displays mpHFOs' morphological distinction within frequency and time domains.

Latent space interpretability analysis:

Latent space 2D visualization: The encoder extracts latent vectors from the time-frequency plot of each event, and the latent code could contain essential morphological information that the decoder can reconstruct. To visually demonstrate the topology of the latent space trained from VAE better, we used t-SNE to project the latent space acquired from the VAE into 2D space. We used the TSNE accelerated by GPU (cuml package),⁵³ the parameters of fitting the TSNE using default parameters as the released package except for the n_components as 2 (**Supplementary Figure 2: Clinical interpretation of latent space clustering**). We projected all of the latent code into the 2D space. For better visualization, we randomly sampled a maximum of 200 data points from each subject for visualization and color-coded each 2D point by different types of HFOs: mpHFO, non-mpHFO, mArtifact (**Figure 2a**); spkHFO, non-spkHFO, the artifact (**Figure 2b**); different subcategories in subjects demographic information included sex, recording sites/type, age, pathology (**Figure 3**) and different anatomical locations (**Figure 4**).

Statistical tests to evaluate the dependence of HFO morphology on the significance of different demographic subject-level information and anatomic locations of the HFO events variability:

We evaluated if HFO morphology is dependent on demographic and anatomical factors. Suppose HFO morphology depended on a particular variable (for example, the subject's sex, **Figure 3c**). In that case, the latent codes of the HFOs labeled with the corresponding subcategories (female or male) may be distinguishable. Consequently, a classifier designed to identify true labels from the latent codes of HFOs should achieve higher accuracy compared to a classifier trained with randomly assigned labels. The methodology can be illustrated using sex as an example. For each fold of the cross-validation, we randomly selected three subjects from each subcategory (male, n=3; female, n=3). We randomly sampled 100 HFO events from each subcategory (100 HFO events for male subjects and 100 HFO events for female subjects). This resulted in 200 (n_subcategories * 100) latent codes for training a logistic regression classifier. For computing accuracy, we repeated the sampling process for the left-out subjects for that fold. This resulted in one accuracy datapoint. We repeated this sampling, training,

and testing process five times (trials) for each fold. Thus, we have 25 trails in total (5 folds * 5 trails) for a classifier to determine sex based on actual sex labels. Next, we describe the design of a surrogate classifier based on randomizing the labels. For each fold and each trial, the same latent codes sampled for the true-label case are used. However, the sex label is randomized. That is, we randomly shuffle the subjects and assign the label of male to the first three and female to the last three. Then, a surrogate logistic classifier is trained with the randomized sex labels. For testing, we again used the same data samples from the true-label case and randomized the sex assignment. This leads to another 25 accuracy samples for the surrogate classifier. A one-tailed t-test was employed to assess whether the true-label classifier's accuracy was significantly better than that of the surrogate.

We performed two such tests, one for the dependence of the morphology of HFOs from preserved regions in seizure-free patients (presumed physiological HFOs) on anatomical locations. The second one was for the dependence of the morphology of the HFOs from SOZ regions (presumed pathological HFOs) on anatomical locations (**Figure 4**). Note that each channel in each subject was assigned a unique anatomical location. Given an HFO type, for each fold and each trial, the HFO events were sampled equally ($n=100$, in total 500) from each anatomical location across all patients. A multi-class logistic regression model was trained, and the confusion matrix was computed based on 500 samples acquired from the test subject under the same fold. Five such trials were done for each fold. Thus, for each HFO type, we obtained 25 confusion matrices, and the average of the confusion matrices and the accuracy were computed. Then, for the surrogate classifier, the randomization of labels of the 500 samples was done by shuffling samples on the channel level; for example, all samples from a particular channel that were initially labeled as "Frontal" would be reassigned to the same randomly selected label such as "Occipital". This led to 25 surrogate accuracy samples. A one-tailed t-test was employed to assess whether the true-label classifier's accuracy was significantly better than that of the surrogate.

Latent space disentanglement on latent dimension perturbation:

By manipulating specific dimensions of the latent space and observing the changes in the reconstructed images, we can infer the significance of each dimension in relation to the data characteristics it influences (**Figure 5, Supplementary Figure 2: Clinical interpretation of latent space interpolation**). This approach not only can aid in the interpretability of complex models but also provide insights into the underlying data structure that the VAE has learned to encode. The latent code was a one-dimensional vector sampled from a distribution generated by the encoder using a time-frequency plot. The distribution of latent codes was plotted on each dimension and classification type (mpHFO and non-mpHFO) to identify the dimension that contributes the most to the decision boundary. To further understand the morphological characteristics contributed from each dimension, we

manipulated each dimension of the latent code and then observed the resultant images generated from the decoder. The trend of the resultant images corresponding to the manipulation could reveal the morphological characteristic for that dimension. Specifically, for one latent code, we moved the value of a specific dimension to the 1st percentile and 99th percentile of all of the values within this dimension (we avoided extreme values to exclude outliers) while keeping the values in other dimensions unchanged. For visualization, we focused on the morphological impact on the two most representative latent codes, which were determined by averaging the latent vectors of all identified mpHFO and non-mpHFO events, respectively. To evaluate the effect of each dimension at a population level, we first randomly selected 500 high-confident mpHFO (model confidence > 0.99) and 500 high-confident non-mpHFO (model confidence < 0.01). We then manipulated the latent codes of 1000 randomly selected samples, irrespective of the event types (mpHFO/non-mpHFO). To quantitatively evaluate the morphological change of resultant images, we designed various metrics based on simple image processing techniques on the time-frequency plot as a proxy to represent the morphological characteristic of the corresponding HFO events. To identify if an event was more pathological (determining pathological dimension), we calculated the ratio of power in the "hanging bell" template (**Figure 2e**) to the overall power within the time-frequency plot. To calculate the peak frequency of an event (determining the peak-frequency dimension), we averaged the information across -90 ms to 90 ms (as the center of the event to be 0 ms) of the time-frequency plot to construct a frequency-power vector and found the frequency that had the max power within 80 Hz to 290 Hz. To measure the slow-wave power of an event (determining the slow-wave dimension), we computed the power in the 10-20 Hz range on the time-frequency plot to the total power. These analyses were visualized using box plots. Furthermore, we used the inference pipeline (encoder and GMM) to classify each generated event and visualized the model confidence through the violin plot.

Clinical correlation: Predicting surgical outcomes: Once the VAE and GMM classifiers were trained on the training set, this pipeline was able to make a prediction on the test set. In the following analysis, we used all subjects who underwent resection. We evaluated the effectiveness of a discovered biomarker by predicting the surgical outcomes of the resected subjects. Specifically, features associated with each subject were extracted, and then a predictive model was trained using these subject-specific features to predict whether or not one subject would be seizure-free. In particular, we used the resection ratio (No. resected events / No. overall detected events), which was the same as in previous studies,²³ and subjects' demographic features such as sex and age. Furthermore, we also considered the resection status of the SOZ (whether the SOZ assigned by clinicians was removed or not), which is a current clinical standard to guide epilepsy surgery. We applied two approaches to validate and predict the surgical outcomes. First, we evaluated the separation of post-surgical seizure-free subjects and post-surgical non-seizure-free subjects in the feature space and also evaluated the balance of the trained unsupervised models in the five-fold cross-validation, i.e., if a model trained in one fold is too conservative or aggressive on predicting

the biomarkers. Therefore, we trained a single/multivariable logistic regression model to predict post-surgical outcomes on features generated from all subjects in all test sets among the five-fold cross-validation and compared the AUC of the logistic regression. Second, we evaluated the forward predictive power of our discovered biomarker. In each fold of five-fold cross-validation, we trained a random forest model from the training set and validation set, using subject-associated features as input of random forest and surgical status as labels. We used the trained random forest to predict all subjects underwent resection in the test set. Then, one random forest model was trained for each fold, and a set of performance metrics (accuracy, precision, recall, and F1 score) were reported; we reported the mean and SEM of the F1 score based on the unbalanced dataset.

Statistical analysis: The above-mentioned statistical calculations were carried out using Python (version 3.9.1; Python Software Foundation, USA). The deep neural network was developed using PyTorch (version 2.1.0; Facebook's AI Research lab). Quantitative measures are described by medians with interquartile or mean with standard deviations. Comparisons between groups were performed using chi-square for comparing two distributions and Student's t-test for quantitative measures (in means with standard deviations). All comparisons were two-sided, and significant results were considered at $p < 0.05$ unless stated otherwise. Specific statistical tests performed for each experiment were described in each section. Machine learning model performance was evaluated using accuracy ($(TP + TN)/(TP+TN+FP+FN)$), recall ($TP/(TP+FN)$), precision ($TP/(TP+FP)$), and F-1 score ($2/[1/recall + 1/precision]$).

Data sharing and availability of the methods: All data produced in the present study are available upon reasonable request to the corresponding authors. Anonymized EEG data and metadata, including labels (channel's resection status, SOZ, patients' demographics, seizure outcomes, and pathology) used in this study will be available on the OpenNEURO website (<https://openneuro.org>) in the near future. The Python-based code used in this study is available at (<https://github.com/roychowdhuryresearch/HFO-VAE>). One can train and test the deep learning algorithm from their data and confirm our methods' validity and utility.

FUNDING

The authors have no conflict of interest to disclose. HN is supported by the National Institute of Neurological Disorders and Stroke (NINDS) (K23NS128318), the Sudha Neelakantan & Venky Harinarayan Charitable Fund, the Elsie and Isaac Fogelman Endowment, and the UCLA Children's Discovery and Innovation Institute (CDI) Junior Faculty Career Development Grant (#CDI-SEED-010124). AD is supported by the Uehara Memorial Foundation and the SENSHIN Medical Research Foundation. EA is supported by the National Institute of Health (R01NS064033). RS serves on scientific advisory boards and speakers bureaus and has received

honoraria and funding for travel from Eisai, Greenwich Biosciences, UCB Pharma, Sunovion, Supernus, Lundbeck Pharma, Liva Nova, and West Therapeutics (advisory only); receives royalties from the publication of Pellock's Pediatric Neurology (Demos Publishing, 2016) and Epilepsy: Mechanisms, Models, and Translational Perspectives (CRC Press, 2011). RS is also supported by the Sudha Neelakantan & Venky Harinarayan Charitable Fund. RJS is supported by the NINDS R01NS106957, R01NS033310, and R01NS127524. JEJ is supported by NINDS U54NS100064 and R01NS033310. RJS and JEJ are supported by the Christina Louise George Trust. The research described was also supported by NIH/National Center for Advancing Translational Science (NCATS) UCLA CTSI Grant Number UL1TR001881.

COMPETING INTERESTS

The authors have no conflict of interest to disclose.

FIGURE LEGENDS

Figure 1. Study flow. (a) Spatial distribution of intracranial electrodes: Electrode contact locations within the standardized MNI brain space from various views (Lateral left, Lateral right, Anterior, Posterior). (b) Variational Autoencoder (VAE) training: Time-frequency EEG data representing HFOs serve as input into the VAE, which outputs a reconstructed image. The VAE's encoder generates a latent distribution of mean and variance while the decoder reconstructs the time-frequency plot from a sampled latent vector. The loss function is a combination of perceptual loss (to capture morphological differences) and KL divergence (measuring the latent distribution's deviation from a normal distribution). (c) HFO classification pipeline: A two-stage, morphology-based classification process uses Gaussian Mixture Models (GMMs) for unsupervised learning. The first stage identifies artifacts (mArtifact) by latent codes and reconstruction loss; the second stage distinguishes putative pathological HFOs (mpHFOs) cluster; the cluster with a higher resection percentage in seizure-free patients after resection was deemed pathological. This process trains two unsupervised classifiers to be used on the test set. (d) Interpretability Analysis. Left: Latent space clustering, visualized via t-SNE, groups clinically relevant HFO classes, facilitating interpretation. Right: Dimensional perturbations in latent space discover specific neurophysiological features in reconstructed images, providing a bridge between clinical insights and the VAE's feature representations.

Figure 2. Characterization of pathological HFOs (mpHFOs) based on the self-supervised VAE algorithm. (a). Projected latent space, color-coded by predicted results (mpHFO, non-mpHFO, and mArtifact) from VAE-based HFO morphology inference pipeline on one example fold, shows the 2D projected eight-dimensional latent codes by the t-SNE. (b). The same latent space is color-coded by conventional HFO classification (spkHFO, non-spkHFO, and artifact). (c) HFO rates (number of HFO detections[count]/min/channel) of mpHFO and non-mpHFO are plotted in box plots based on the location (SOZ versus non-SOZ) across three datasets (outliers were removed for better visualization quality). The rate of mpHFO (count/min/channel) was higher in the SOZ than in the non-SOZ. The rates of non-mpHFOs (count/min/channel) did not differ between the SOZ and non-SOZ (*: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$). (d) Morphological analysis of the time-frequency plot for an example subject. The pathological counterparts (mpHFOs) have higher values throughout the HFO band (> 80 Hz), around the center point (0 ms, where HFOs were detected) than non-mpHFOs; furthermore, higher values of mpHFOs at the sub-HFO band (10-80 Hz) throughout the time window compared to non-mpHFOs are exhibited. (e) The overall template (mean) of all subjects resembles a "hanging bell" shape (pixel comparisons that were significantly higher in mpHFOs than non-mpHFOs in more than 50% of patients were colored orange). (f) FFT amplitude is plotted as a function of frequency from an example patient (shaded areas indicate one SD). Note the FFT peak (around 80-90 Hz) is seen in non-mpHFOs but not in mpHFOs. The mpHFO has a higher sub-

HFO frequency signal intensity, so the HFO peak may be masked. (g) FFT amplitude ratio of (mpHFO - non-mpHFO / non-mpHFO) is plotted as a function of frequency, demonstrating the presence of relatively high FFT amplitude in mpHFO within the sub-HFOband, peaked around 23 Hz. (h) Time-frequency plot of an example predicted mpHFO. (i) EEG tracing of the same mpHFO with the detected HFO part colored in red. (j) EEG tracing bandpassed between 10 and 80 Hz of the same mpHFO with HFO detection colored in red. (k/l/m) An example predicted non-mpHFO presented in the same fashion as the mpHFO example. Note the presence of a spike-wave activity in EEG in the mpHFO but not in the non-mpHFO.

Figure 3. Investigating potential heterogeneity of HFO morphology based on variables. (a) Visualization of the latent space for HFOs, color-coded by different recording sites/types (UCLA SEEG, UCLA grid/strip, and Detroit grid/strip) on a specific fold. (b) Across five-fold, the classifiers trained using the actual recording sites/type labels (Real) did not show significantly better accuracy than those trained by using the permuted labels (Random). (c) Visualization of the latent space for HFOs color-coded by different sexes (male and female) on a specific fold. (d) Across five-fold, the classifiers trained using the actual sex labels (Real) did not show significantly better accuracy than those trained by permuted labels (Random). (e) Visualization of the latent space for HFOs color-coded by different age groups (0-5, 6-10, 11-15, 16-20, and 21+) on a specific fold. (f) Across five-fold, the classifiers trained using the actual age group labels (Real) did not show significantly better accuracy than those trained by label permuted data (Random). (g) Visualization of the latent space for HFOs color-coded by different pathologies (HS, FCD, Tumor, and Others) on a specific fold. (h) Across five-fold, the classifiers trained using actual pathology labels (Real) did not show significantly better accuracy than those trained by label-permuted data (Random).

Figure 4. Morphological investigations of HFOs from various anatomical regions. (a) Visualization of the latent space for HFOs from preserved brain regions of post-surgical, seizure-free subjects, color-coded by anatomical locations (frontal, temporal, parietal, occipital, and limbic regions) on a specific fold. (b) Across five-fold, the classifiers trained using actual anatomical location label (Real) showed significantly better ($p = 6.67e-10$) accuracy than those trained by permuted label (Random) (Real mean=0.286, std= 0.046, Random mean = 0.230, std= 0.033). (c) Averaged confusion matrix on the test set across five trials and five-fold ($n=25$ trials) using actual anatomical locations for HFOs from preserved brain regions of subjects who achieved postoperative seizure freedom. Note that HFOs from the occipital region were distinguishable. (d) Averaged time-frequency plots for each anatomical location for HFOs from each brain region. Note that HFOs from the occipital region exhibited distinct features on the time-frequency plot. (e) Visualization of the latent space for HFOs from SOZ channels, color-coded by anatomical locations. (f) Across five-fold, for HFOs from SOZ channels, the classifiers trained using the anatomical location label (Real) did not show significantly better ($p = 0.090$) accuracy than those trained

by permuted label (Random) (Real mean=0.241, std=0.043, Random mean = 0.23, std=0.027). (g) Averaged confusion matrix on the test set across five trials and five-fold (n=25) using actual anatomical locations for HFOs from SOZ channels. Note that HFOs from the SOZ were indistinguishable from any anatomical origin. (h) Averaged time-frequency plots for each anatomical location for HFOs from the SOZ channels. Note that HFOs from the SOZ exhibited similar features within time-frequency plots across the anatomical regions.

Figure 5. Disentanglement of the latent space to establish neurophysiological characteristics of pathological HFOs (mpHFOs). (a) Pathological dimension visualization. This panel illustrates VAE-identified latent space, which enables separation between mpHFOs and non-mpHFOs. (b) A traversal from the lower to the upper percentile of this dimension revealed morphological evolution from non-mpHFO to mpHFO characteristics for both mpHFO (orange) and non-mpHFO (blue) representatives in the decoded output, depicted in the image sequence. (c) Power trend within the pathological dimension at the population level. The box plot aggregates the percentage of power within the designated pathological template (the "hanging bell shape") region of decoded images, showing an ascending trend with higher values, as indicated by the fitted median line. (d) Distribution of model probability scores for each sample. The red circles indicate the mean probability scores, showing increased confidence in the model as the value of the pathological dimension increases. (e) Peak frequency dimensional visualization. A similar visualization of another dimension is shown. This latent space represents the peak frequency of HFOs. (f) The output of the decoder traversing the dimension, displayed in the image sequence, showed a descending trend in peak frequency from upper to lower percentiles of the value of that dimension. (g) At the population level, the box plots indicated a negative correlation between the peak frequency dimension value and the peak frequency in decoded images, with a trend line fitted from the median of each box. (h) Distribution of model probability scores for each sample. The red circles indicate the mean probability scores, showing the average confidence of perturbed events was around 0.5 (unchanged). (i) Slow-wave dimension visualization. This latent space represented the slow-wave component of HFOs at 10-20 Hz, separated by mpHFO and non-mpHFO prediction. (j) The output of the decoder traversing the dimension, displayed in the image sequence, showed an increased trend in slow wave power from lower to upper percentiles of the value of that dimension. (k) At the population level, the box plot demonstrated a positive correlation between slow-wave dimension values and slow-wave power in decoded images, with a line fit illustrating the median trend. (l) Model probability scores distribution corresponding to each sample, where the mean of the probability marked as red circles, showed the average confidence of perturbed events was around 0.5.

Figure 6. Clinical validation: resection status of pathological HFOs (mpHFOs) helps predict postoperative seizure outcomes. (a) AUCs (area under the curve) of logistic regression (a single-variable classifier) using resection ratios as a variable to predict postoperative seizure freedom are shown based on different types of HFO

resection ratios: HFO (unclassified HFO detection), spkHFO (HFO with spike), and mpHFO (pathological HFO defined by VAE). (b) AUCs from multivariable logistic regression models with different types of HFO resection ratios as the main predictive variable while incorporating subject-wise demographic information (baseline demographics: age and sex) are shown. As an additional predictive variable, the resection status of SOZ was also included. Note that combining the baseline demographics, SOZ resection status, and resection ratio of mpHFO provided the most favorable classification performance (base+soz+mpHFO, AUC=0.72). (c) The mean F1 scores (F1) and standard error of the mean for the random forest models, trained on the training subjects using five-fold cross-validation and evaluated on the test subjects across the five folds and using different features, are shown. Note that the resection ratio of mpHFO achieved a comparable predictive performance compared to the current clinical standard (SOZ resection status). By combining all the features (baseline demographics, mpHFO resection ratio, and SOZ resection status), the model achieved high performance with an F1 score of 0.812 ± 0.078 .

REFERENCES

1. Kwan P, Brodie MJ. Early identification of refractory epilepsy. *The New England journal of medicine*. Feb 03 2000;342(5):314-9. doi:10.1056/nejm200002033420503
2. Wiebe S, Blume WT, Girvin JP, Eliasziw M. A randomized, controlled trial of surgery for temporal-lobe epilepsy. *The New England journal of medicine*. Aug 2 2001;345(5):311-8. doi:10.1056/nejm200108023450501
3. Dwivedi R, Ramanujam B, Chandra PS, et al. Surgery for Drug-Resistant Epilepsy in Children. *The New England journal of medicine*. Oct 26 2017;377(17):1639-1647. doi:10.1056/NEJMoa1615335
4. Bragin A, Engel J, Jr., Wilson CL, Fried I, Mathern GW. Hippocampal and entorhinal cortex high-frequency oscillations (100--500 Hz) in human epileptic brain and in kainic acid--treated rats with chronic seizures. *Epilepsia*. Feb 1999;40(2):127-37.
5. Staba RJ, Wilson CL, Bragin A, Jhung D, Fried I, Engel J, Jr. High-frequency oscillations recorded in human medial temporal lobe during sleep. *Annals of neurology*. Jul 2004;56(1):108-15. doi:10.1002/ana.20164
6. Worrell GA, Parish L, Cranstoun SD, Jonas R, Baltuch G, Litt B. High-frequency oscillations and seizure generation in neocortical epilepsy. *Brain*. Jul 2004;127(Pt 7):1496-506. doi:10.1093/brain/awh149
7. Jirsch JD, Urrestarazu E, LeVan P, Olivier A, Dubeau F, Gotman J. High-frequency oscillations during human focal seizures. *Brain*. Jun 2006;129(Pt 6):1593-608. doi:10.1093/brain/awl085
8. Jacobs J, Zijlmans M, Zelmann R, et al. High-frequency electroencephalographic oscillations correlate with outcome of epilepsy surgery. *Annals of neurology*. Feb 2010;67(2):209-20. doi:10.1002/ana.21847
9. Wu JY, Sankar R, Lerner JT, Matsumoto JH, Vinters HV, Mathern GW. Removing interictal fast ripples on electrocorticography linked with seizure freedom in children. *Neurology*. Nov 09 2010;75(19):1686-94. doi:10.1212/WNL.0b013e3181fc27d0
10. Akiyama T, McCoy B, Go CY, et al. Focal resection of fast ripples on extraoperative intracranial EEG improves seizure outcome in pediatric epilepsy. *Epilepsia*. Oct 2011;52(10):1802-11. doi:10.1111/j.1528-1167.2011.03199.x
11. van 't Klooster MA, van Klink NEC, Zweiphenning W, et al. Tailoring epilepsy surgery with fast ripples in the intraoperative electrocorticogram. *Annals of neurology*. May 2017;81(5):664-676. doi:10.1002/ana.24928
12. Zijlmans M, Jiruska P, Zelmann R, Leijten FS, Jefferys JG, Gotman J. High-frequency oscillations as a new biomarker in epilepsy. *Annals of neurology*. Feb 2012;71(2):169-78. doi:10.1002/ana.22548
13. Jefferys JG, Menendez de la Prida L, Wendling F, et al. Mechanisms of physiological and epileptic HFO generation. *Prog Neurobiol*. Sep 2012;98(3):250-64. doi:10.1016/j.pneurobio.2012.02.005
14. Engel J, Jr., Bragin A, Staba R, Mody I. High-frequency oscillations: what is normal and what is not? *Epilepsia*. Apr 2009;50(4):598-604. doi:10.1111/j.1528-1167.2008.01917.x

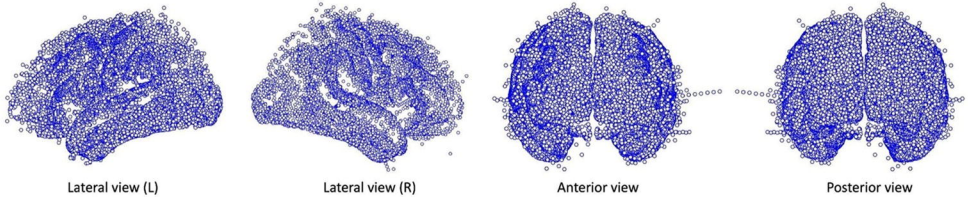
15. Frauscher B, von Ellenrieder N, Zelmann R, et al. High-Frequency Oscillations in the Normal Human Brain. *Annals of neurology*. Sep 2018;84(3):374-385. doi:10.1002/ana.25304
16. Guragain H, Cimbalnik J, Stead M, et al. Spatial variation in high-frequency oscillation rates and amplitudes in intracranial EEG. *Neurology*. Feb 20 2018;90(8):e639-e646. doi:10.1212/wnl.0000000000004998
17. Jacobs J, Wu JY, Perucca P, et al. Removing high-frequency oscillations: A prospective multicenter study on seizure outcome. *Neurology*. Sep 11 2018;91(11):e1040-e1052. doi:10.1212/wnl.0000000000006158
18. Zweiphenning W, Klooster MAV, van Klink NEC, et al. Intraoperative electrocorticography using high-frequency oscillations or spikes to tailor epilepsy surgery in the Netherlands (the HFO trial): a randomised, single-blind, adaptive non-inferiority trial. *Lancet neurology*. Nov 2022;21(11):982-993. doi:10.1016/s1474-4422(22)00311-8
19. Matsumoto A, Brinkmann BH, Matthew Stead S, et al. Pathological and physiological high-frequency oscillations in focal human epilepsy. *J Neurophysiol*. Oct 2013;110(8):1958-64. doi:10.1152/jn.00341.2013
20. Cimbalnik J, Brinkmann B, Kremen V, et al. Physiological and pathological high frequency oscillations in focal epilepsy. *Ann Clin Transl Neurol*. Sep 2018;5(9):1062-1076. doi:10.1002/acn3.618
21. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. May 28 2015;521(7553):436-44. doi:10.1038/nature14539
22. Litjens G, Kooi T, Bejnordi BE, et al. A survey on deep learning in medical image analysis. *Med Image Anal*. Dec 2017;42:60-88. doi:10.1016/j.media.2017.07.005
23. Zhang Y, Lu Q, Monsoor T, et al. Refining epileptogenic high-frequency oscillations using deep learning: a reverse engineering approach. *Brain Commun*. 2022;4(1):fcab267. doi:10.1093/braincomms/fcab267
24. Zhang Y, Chung H, Ngo JP, et al. Characterizing physiological high-frequency oscillations using deep learning. *J Neural Eng*. Dec 7 2022;19(6)doi:10.1088/1741-2552/aca4fa
25. Kingma DP, Welling M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*. 2013;
26. Kenton JDM-WC, Toutanova LK. Bert: Pre-training of deep bidirectional transformers for language understanding. 2019:2.
27. Razavi A, Oord Avd, Vinyals O. Generating Diverse High-Fidelity Images with VQ-VAE-2. *arXiv preprint arXiv:1906.00446*. 2019;
28. Dai M, Zheng D, Na R, Wang S, Zhang S. EEG Classification of Motor Imagery Using a Novel Deep Learning Framework. *Sensors*. Jan 29 2019;19(3)doi:10.3390/s19030551
29. Tian C, Ma Y, Cammon J, Fang F, Zhang Y, Meng M. Dual-Encoder VAE-GAN With Spatiotemporal Features for Emotional EEG Data Augmentation. *IEEE Trans Neural Syst Rehabil Eng*. 2023;31:2018-2027. doi:10.1109/tnsre.2023.3266810
30. Zhang Y, Liu L, Ding Y, et al. PyHFO: lightweight deep learning-powered end-to-end high-frequency oscillations analysis application. *J Neural Eng*. May 9 2024;doi:10.1088/1741-2552/ad4916

31. Weiss SA, Orosz I, Salamon N, et al. Ripples on spikes show increased phase-amplitude coupling in mesial temporal lobe epilepsy seizure-onset zones. *Epilepsia*. Nov 2016;57(11):1916-1930. doi:10.1111/epi.13572
32. Song I, Orosz I, Chervoneva I, et al. Bimodal coupling of ripples and slower oscillations during sleep in patients with focal epilepsy. *Epilepsia*. Nov 2017;58(11):1972-1984. doi:10.1111/epi.13912
33. Cai Z, Sohrabpour A, Jiang H, et al. Noninvasive high-frequency oscillations riding spikes delineates epileptogenic sources. *Proceedings of the National Academy of Sciences of the United States of America*. Apr 27 2021;118(17)doi:10.1073/pnas.2011130118
34. Nariai H, Wu JY, Bernardo D, Fallah A, Sankar R, Hussain SA. Interrater reliability in visual identification of interictal high-frequency oscillations on electrocorticography and scalp EEG. *Epilepsia open*. Dec 2018;3(Suppl Suppl 2):127-132. doi:10.1002/epi4.12266
35. Roehri N, Pizzo F, Lagarde S, et al. High-frequency oscillations are not better biomarkers of epileptogenic tissues than spikes. *Annals of neurology*. Jan 2018;83(1):84-97. doi:10.1002/ana.25124
36. Nevalainen P, von Ellenrieder N, Klimeš P, Dubeau F, Frauscher B, Gotman J. Association of fast ripples on intracranial EEG and outcomes after epilepsy surgery. *Neurology*. Oct 20 2020;95(16):e2235-e2245. doi:10.1212/wnl.0000000000010468
37. Zweiphenning W, von Ellenrieder N, Dubeau F, et al. Correcting for physiological ripples improves epileptic focus identification and outcome prediction. *Epilepsia*. Feb 2022;63(2):483-496. doi:10.1111/epi.17145
38. Kuroda N, Sonoda M, Miyakoshi M, et al. Objective interictal electrophysiology biomarkers optimize prediction of epilepsy surgery outcome. *Brain Commun*. 2021;3(2):fcab042. doi:10.1093/braincomms/fcab042
39. Zhao B, Hu W, Zhang C, et al. Integrated Automatic Detection, Classification and Imaging of High Frequency Oscillations With Stereoelectroencephalography. *Front Neurosci*. 2020;14:546. doi:10.3389/fnins.2020.00546
40. Besheli BF, Sha Z, Gavvala JR, et al. A sparse representation strategy to eliminate pseudo-HFO events from intracranial EEG for seizure onset zone localization. *J Neural Eng*. Aug 24 2022;19(4)doi:10.1088/1741-2552/ac8766
41. Medvedev AV, Agoureeva GI, Murro AM. A Long Short-Term Memory neural network for the detection of epileptiform spikes and high frequency oscillations. *Sci Rep*. Dec 18 2019;9(1):19374. doi:10.1038/s41598-019-55861-w
42. Nejedly P, Kremen V, Sladky V, et al. Exploiting Graphoelements and Convolutional Neural Networks with Long Short Term Memory for Classification of the Human Electroencephalogram. *Sci Rep*. Aug 6 2019;9(1):11383. doi:10.1038/s41598-019-47854-6

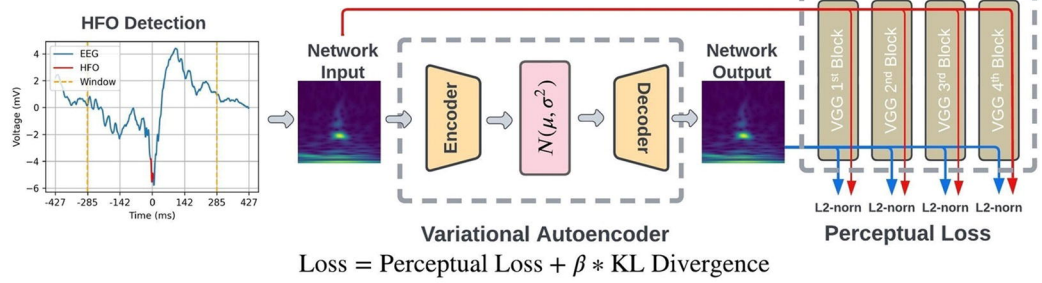
43. Guo J, Xiao N, Li H, et al. Transformer-Based High-Frequency Oscillation Signal Detection on Magnetoencephalography From Epileptic Patients. *Front Mol Biosci.* 2022;9:822810. doi:10.3389/fmolb.2022.822810
44. Irina Higgins, Loic Matthey, Arka Pal, et al. VAE: LEARNING BASIC VISUAL CONCEPTS WITH A CONSTRAINED VARIATIONAL FRAMEWORK. *ICLR 2017.* 2017:1-22.
45. Alkawadri R, Gaspard N, Goncharova, II, et al. The spatial and signal characteristics of physiologic high frequency oscillations. *Epilepsia.* Dec 2014;55(12):1986-95. doi:10.1111/epi.12851
46. Sakakura K, Kuroda N, Sonoda M, et al. Developmental atlas of phase-amplitude coupling between physiologic high-frequency oscillations and slow waves. *Nature communications.* Oct 13 2023;14(1):6435. doi:10.1038/s41467-023-42091-y
47. Tamilia E, Park EH, Percivati S, et al. Surgical resection of ripple onset predicts outcome in pediatric epilepsy. *Annals of neurology.* Sep 2018;84(3):331-346. doi:10.1002/ana.25295
48. Gliske SV, Irwin ZT, Chestek C, et al. Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings. *Nature communications.* Jun 1 2018;9(1):2155. doi:10.1038/s41467-018-04549-2
49. von Ellenrieder N, Dubeau F, Gotman J, Frauscher B. Physiological and pathological high-frequency oscillations have distinct sleep-homeostatic properties. *NeuroImage Clinical.* 2017;14:566-573. doi:10.1016/j.nicl.2017.02.018
50. Klein A, Tourville J. 101 labeled brain images and a consistent human cortical labeling protocol. *Front Neurosci.* 2012;6:171. doi:10.3389/fnins.2012.00171
51. Wu J, Ngo GH, Greve D, et al. Accurate nonlinear mapping between MNI volumetric and FreeSurfer surface coordinate systems. *Human brain mapping.* Sep 2018;39(9):3793-3808. doi:10.1002/hbm.24213
52. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition.* 2016:770-778.
53. Raschka S, Patterson J, Nolet C. Machine learning in python: Main developments and technology trends in data science, machine learning, and artificial intelligence. *Information.* 2020;11(4):193.

Figure 1

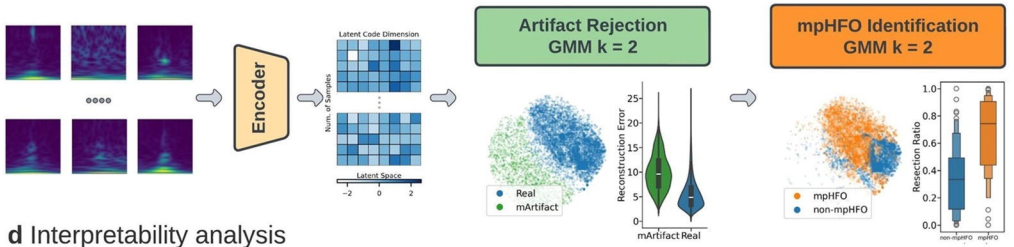
a Spatial distribution of intracranial electrode sampling



b Self-supervised variational autoencoder training

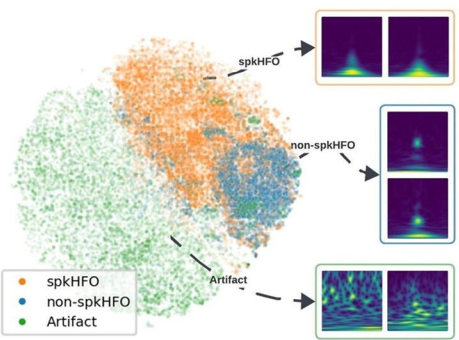


c Morphology-based HFO classification



d Interpretability analysis

Clinical interpretation of latent space clustering



Neurophysiological interpretation of latent space disentanglement

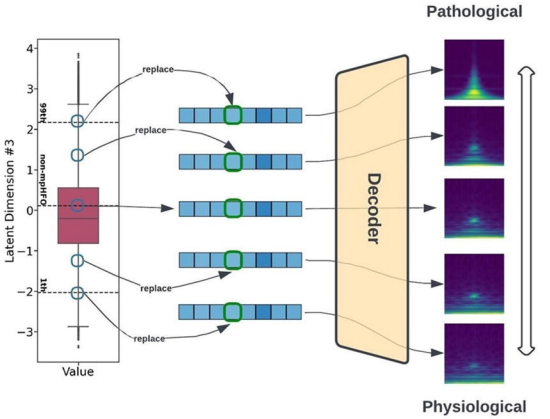


Figure 2

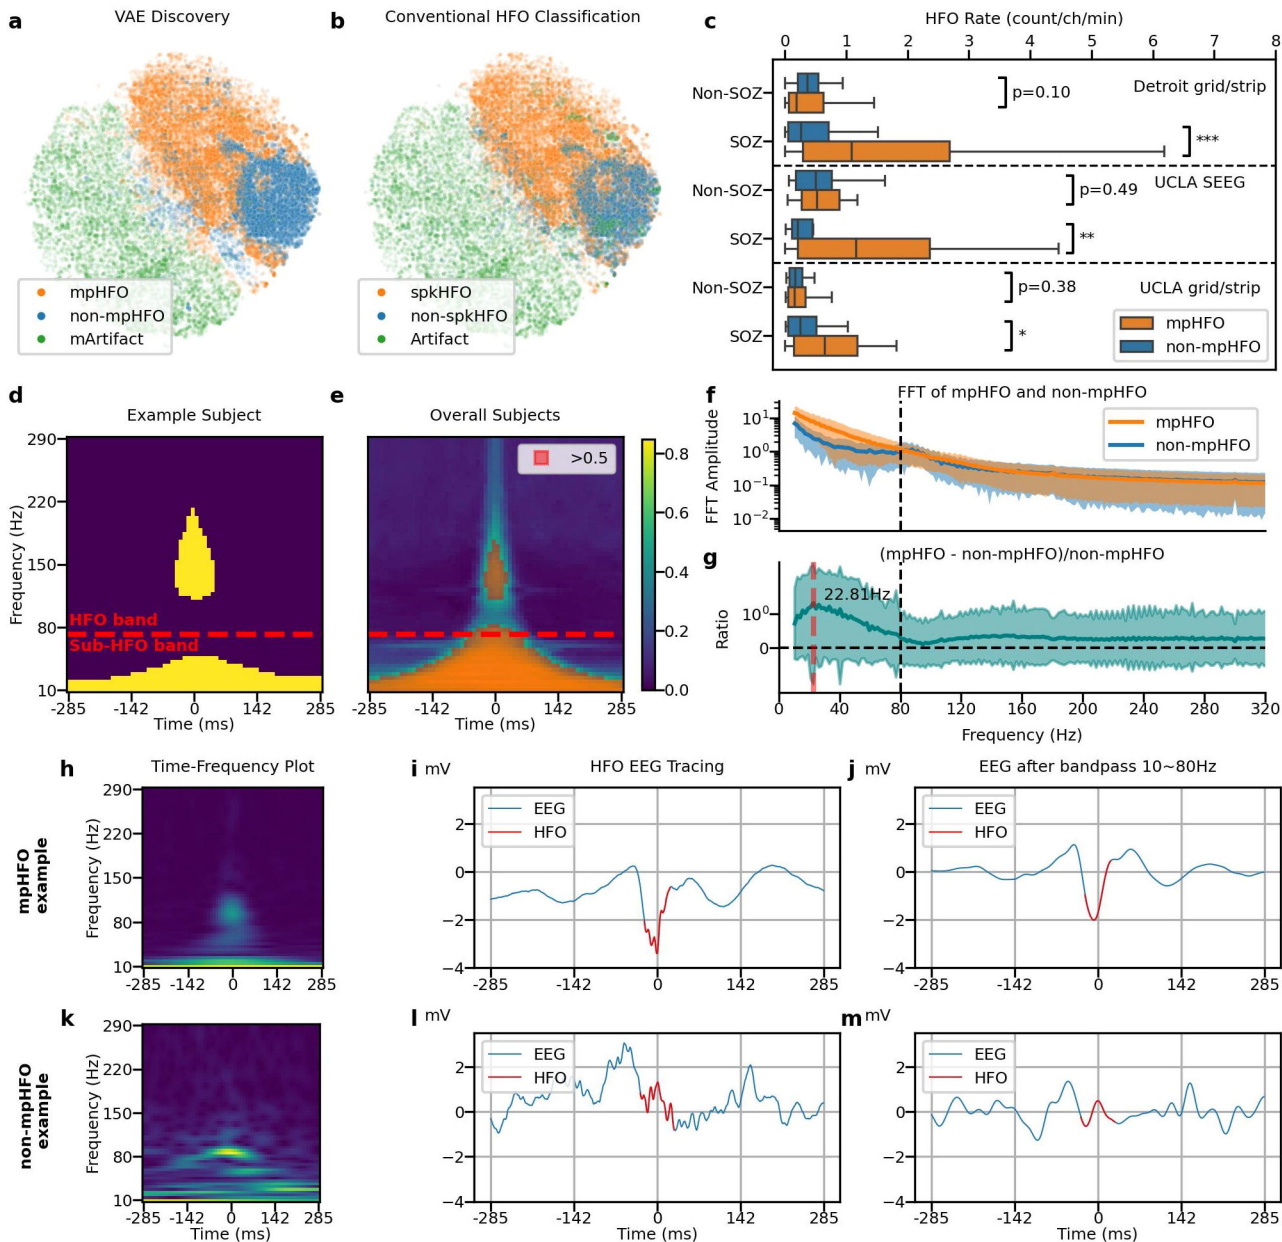
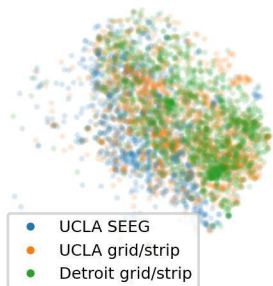


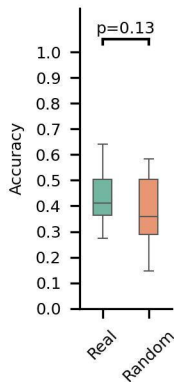
Figure 3

Recording sites/types

a

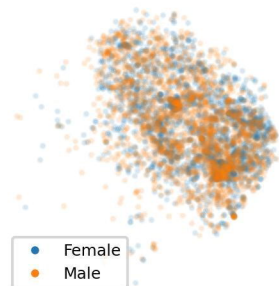


b

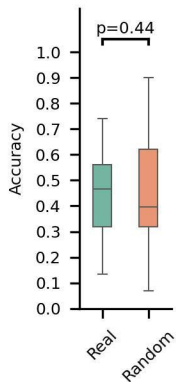


Sex

c

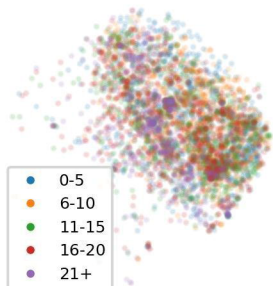


d

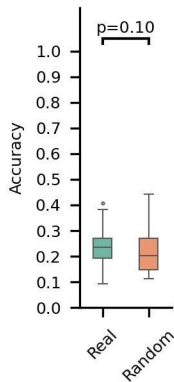


Age groups

e

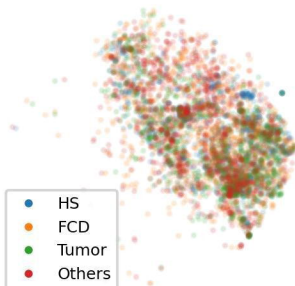


f



Pathologies

g



h

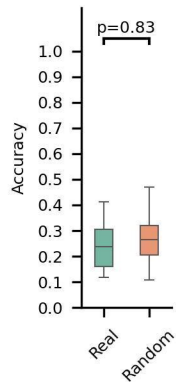


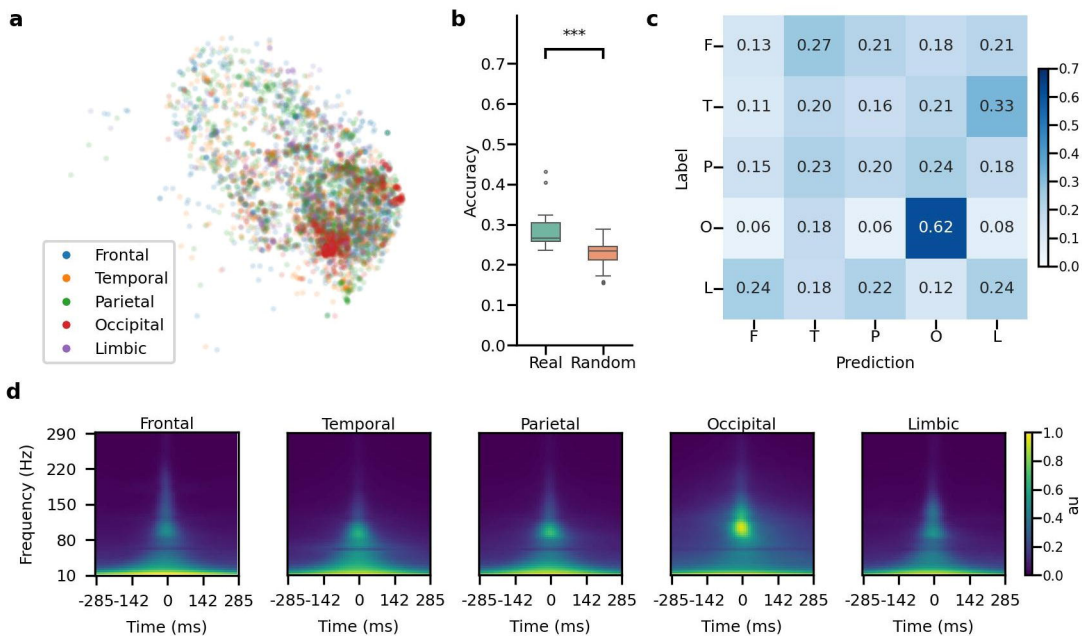
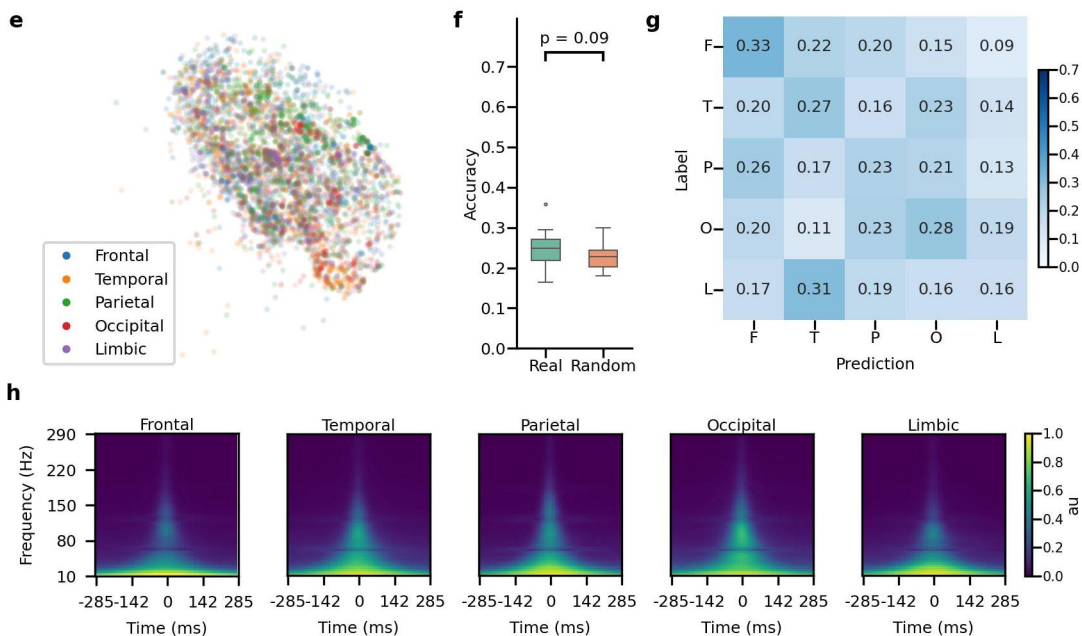
Figure 4**HFOs from preserved brain regions of post-surgical, seizure-free subjects****HFOs from SOZ channels**

Figure 5

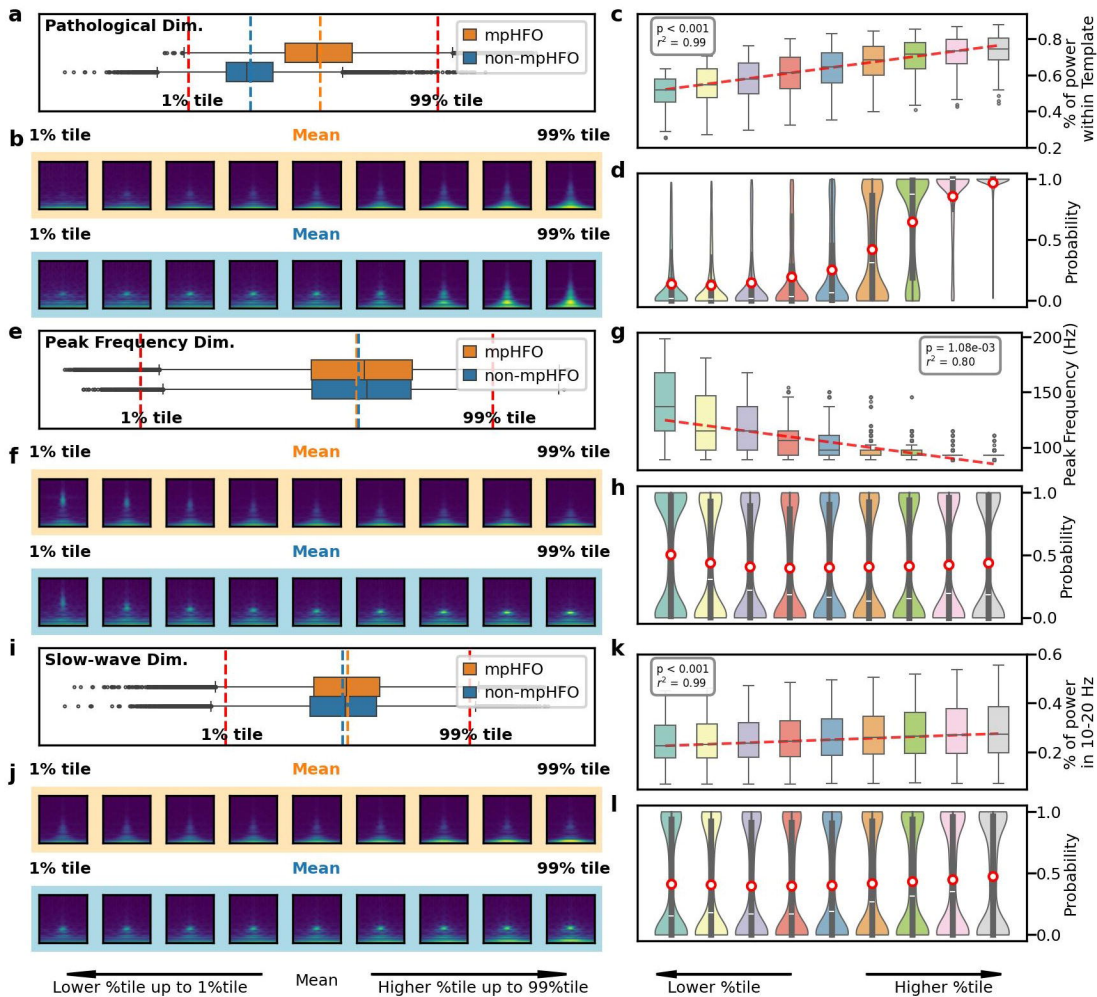


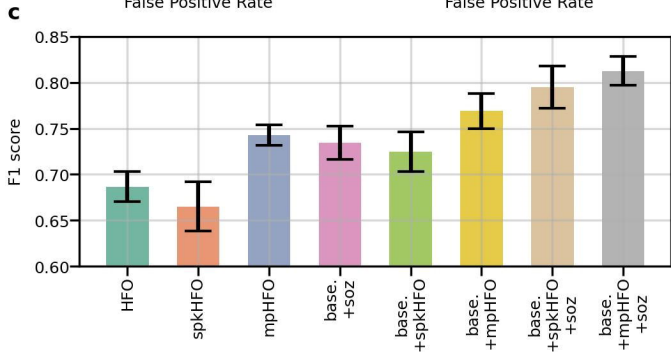
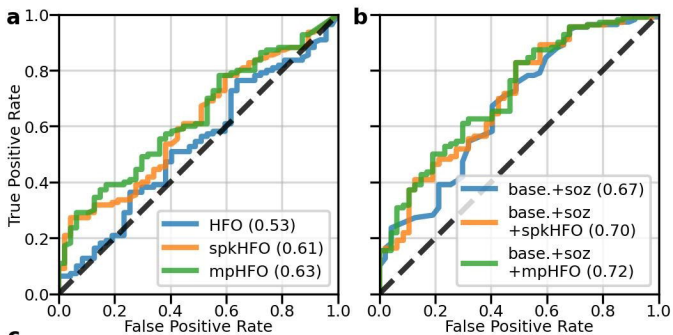
Figure 6

Table 1: Patient Demographics

	UCLA	Detroit	All
Number of patients	50	135	185
Median age in years (range)	14.5 (2-28)	12.0 (4-44)	13 (2-44)
Proportion of female (%)	24 (48.0%)	67 (49.6%)	91 (49.2%)
Proportion of grid case (vs. SEEG) (%)	32 (64.0%)	135 (100.0%)	167 (90.3%)
Sampled hemisphere (%)			
Left	21 (42.0%)	61 (45.19%)	82 (44.3%)
Right	11 (22.0%)	51 (37.78%)	62 (33.5%)
Both	18 (36.0%)	23 (17.04%)	41 (22.2%)
Seizure onset zone			
Frontal	24	42	66
Temporal	16	63	79
Parietal	22	49	71
Occipital	4	23	27
Limbic	18	59	77
Patients who underwent resection [#]	28 (56.0%)	135 (100%)	163 (88.1%)
Patients with postoperative seizure-freedom (%) [#]	15 (53.6%)	95 (70.4%)	110 (67.5%)
Pathology (%) [*]			
Focal cortical dysplasia	18 (64.29%)	49 (36.3%)	67 (41.1%)
Hippocampal sclerosis	0 (0.0%)	11 (8.15%)	11 (6.7%)
Tumor	3 (10.71%)	28 (20.74%)	31 (19.0%)
Others	7 (25.0%)	47 (34.81%)	54 (33.1%)

*Pathology was considered only in resected patients.

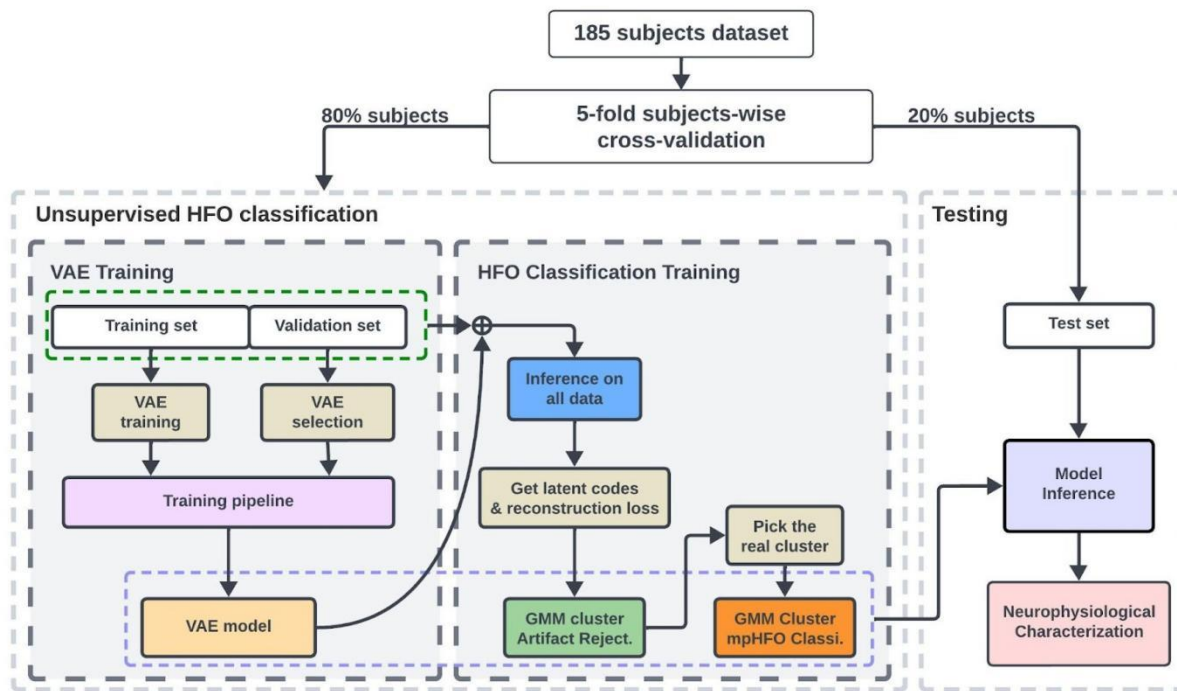
#Only three patients in SEEG were resected. Therefore, they were excluded for the outcome analysis.

Table 2: Spatial distribution of intracranial electrode sampling

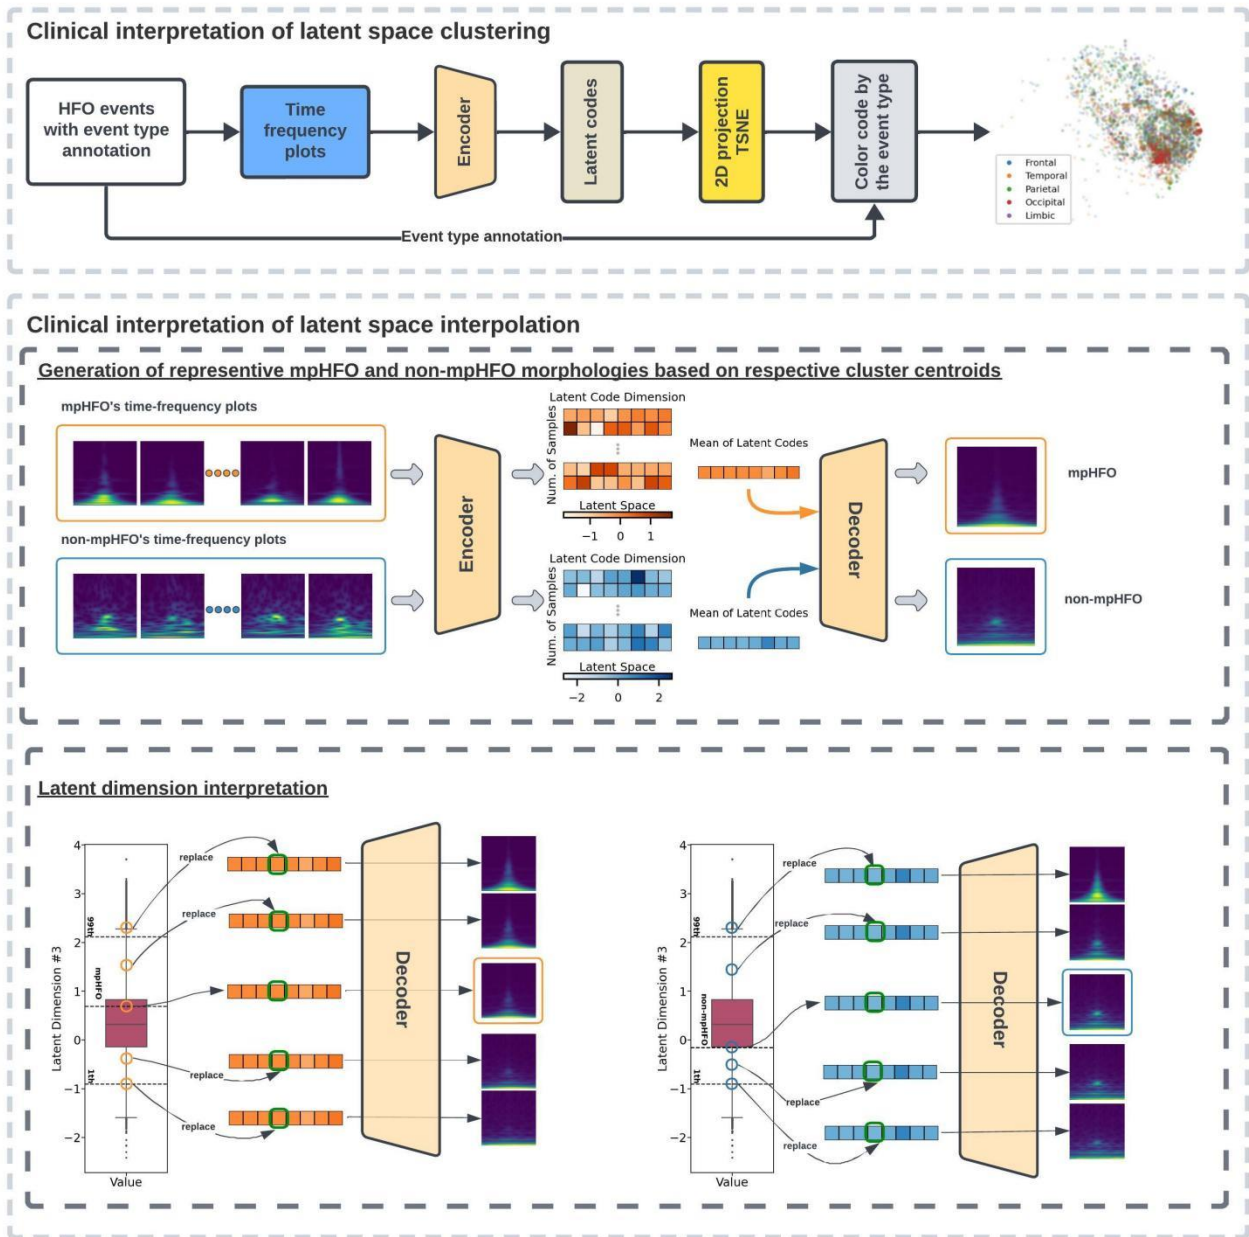
Region of Interest (ROI)	No. of contacts	No. of contacts	No. of contacts	Proportion of Grid contacts	Proportion of SEEG contacts (%)	No. of SOZ contacts	No. of non-epileptogenic contacts*	HFO rate at SOZ (/min) [#]	HFO rate at non-epileptogenic contacts (/min) [#]
	Left	Right	Total	(%)	(%)				
Frontal:	3197	3074	6271	95.39	4.61	437	2943	3.60	1.08
caudalmiddlefrontal	477	400	877	97.38	2.62	52	433	2.77	1.10
frontalpole	2	11	13	100	0	1	3	NA	NA
paracentral	71	99	170	98.82	1.18	12	80	5.68	1.31
parsopercularis	325	203	528	95.27	4.73	33	286	3.14	1.31
parsorbitalis	137	162	299	97.99	2.01	31	118	2.24	1.28
parstriangularis	265	295	560	92.32	7.68	51	250	3.09	1.05
precentral	920	797	1717	95.92	4.08	71	933	6.35	1.03
rostralmiddlefrontal	613	629	1242	95.09	4.91	116	508	2.84	1.04
superiorfrontal	387	478	865	93.18	6.82	70	332	3.26	1.00
Temporal:	2707	2128	4835	93.9	6.1	464	1806	2.99	1.05
fusiform	481	391	872	97.82	2.18	108	358	2.56	1.02
inferiortemporal	469	395	864	97.11	2.89	92	296	3.76	1.06
middletemporal	720	526	1246	89.41	10.59	99	440	2.51	0.98
superiortemporal	876	652	1528	92.87	7.13	115	618	3.35	1.04
temporalpole	158	154	312	100	0	46	94	2.74	1.52
transversetemporal	2	10	12	16.67	83.33	4	0	NA	NA
Parietal:	2231	1936	4167	96.02	3.98	393	1853	5.40	1.22
inferiorparietal	314	353	667	95.95	4.05	81	208	3.67	1.19
postcentral	832	661	1493	97.19	2.81	91	783	7.17	0.94
precuneus	137	142	279	96.42	3.58	31	129	4.91	1.14
superiorparietal	237	173	410	92.2	7.8	48	125	4.42	1.54
supramarginal	711	607	1318	95.83	4.17	142	608	5.83	1.44
Occipital	806	670	1476	98.98	1.02	135	635	3.19	0.81
cuneus	79	62	141	100	0	17	54	4.04	0.87
lateraloccipital	504	356	860	98.84	1.16	78	355	3.13	0.87
lingual	219	249	468	99.79	0.21	40	226	2.95	0.68
Limbic	808	708	1516	82.19	17.81	241	495	2.07	1.28
Hippocampus	32	28	60	0	100	9	7	1.65	1.23
amygdala	18	14	32	0	100	0	4	NA	NA
caudalanteriorcingulate	26	28	54	55.56	44.44	5	18	3.51	0.66
entorhinal	164	169	333	99.7	0.3	83	89	1.86	2.15
insula	40	52	92	0	100	17	21	1.22	0.63
isthmuscingulate	66	59	125	97.6	2.4	8	55	0.73	1.62
lateralorbitofrontal	226	179	405	94.81	5.19	50	155	2.42	1.09
medialorbitofrontal	49	36	85	94.12	5.88	8	37	1.66	1.25
parahippocampal	104	80	184	96.2	3.8	44	57	1.69	1.37
posteriorcingulate	68	48	116	89.66	10.34	10	49	4.77	0.93
rostralanteriorcingulate	15	15	30	56.67	43.33	7	3	2.67	NA

*Non-epileptogenic contacts were defined as EEG contacts covering spared brain regions in patients who achieved postoperative seizure freedom

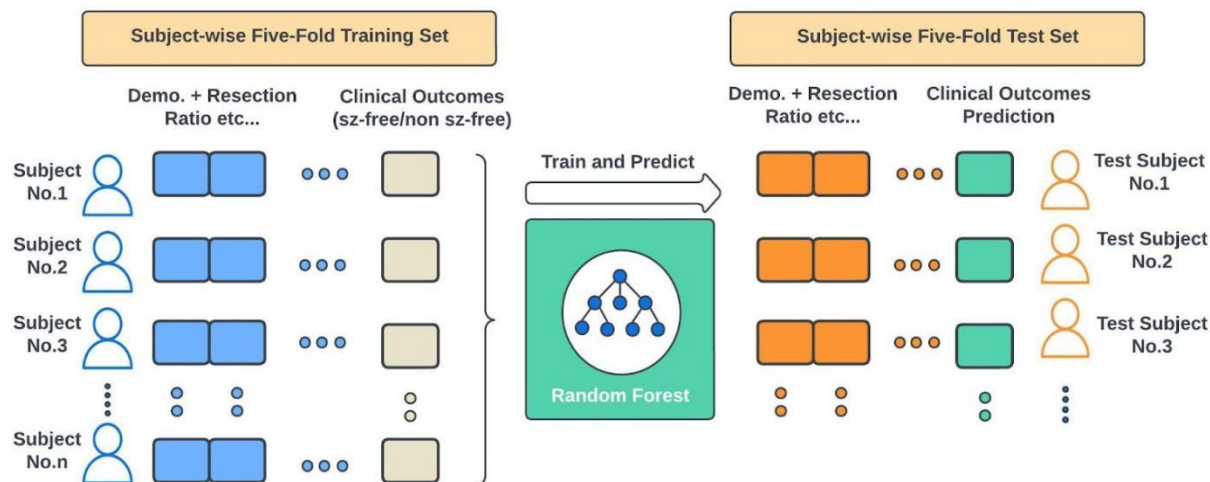
ROIs with contacts less than 5 were removed from the analysis



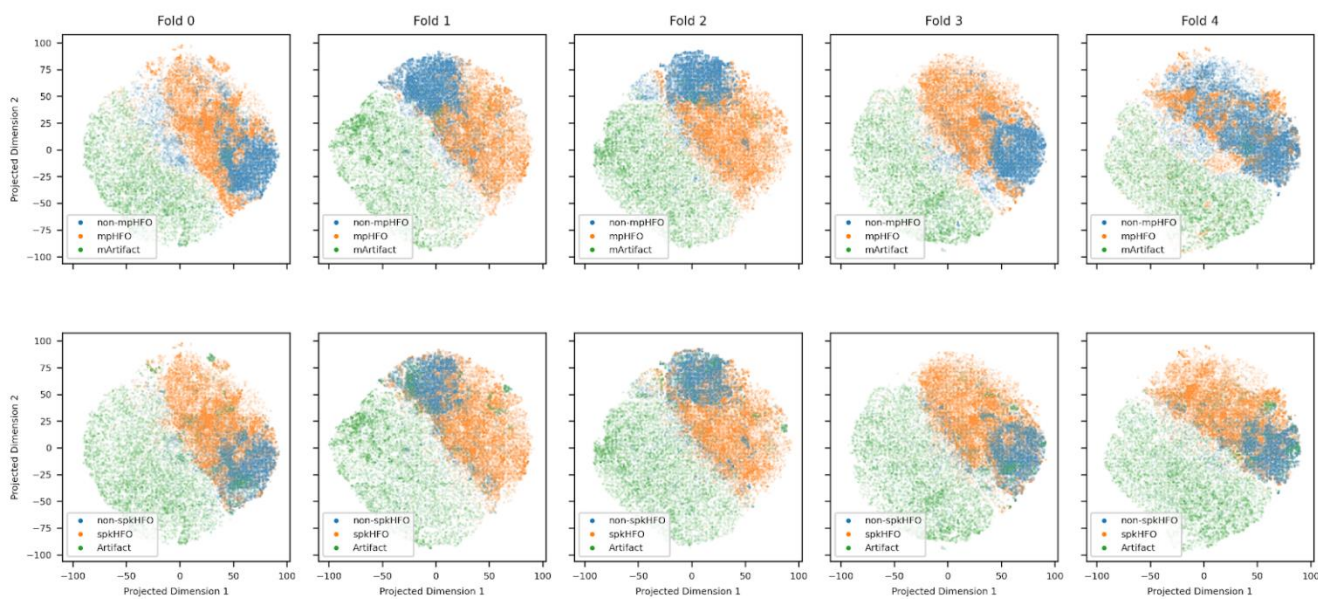
Supplementary Figure 1. VAE training and HFO classification workflow.



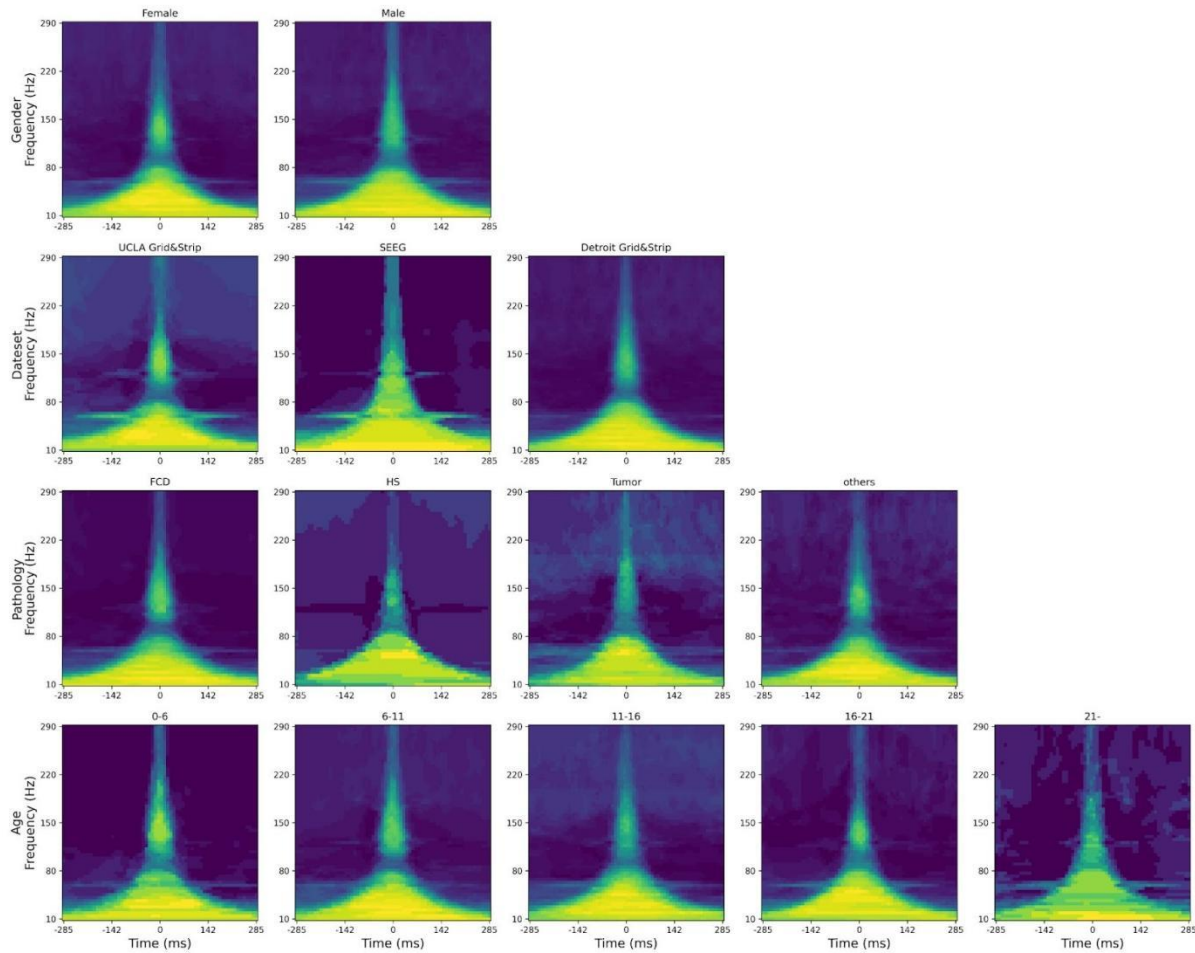
Supplementary Figure 2. Model interpretation analysis workflow.



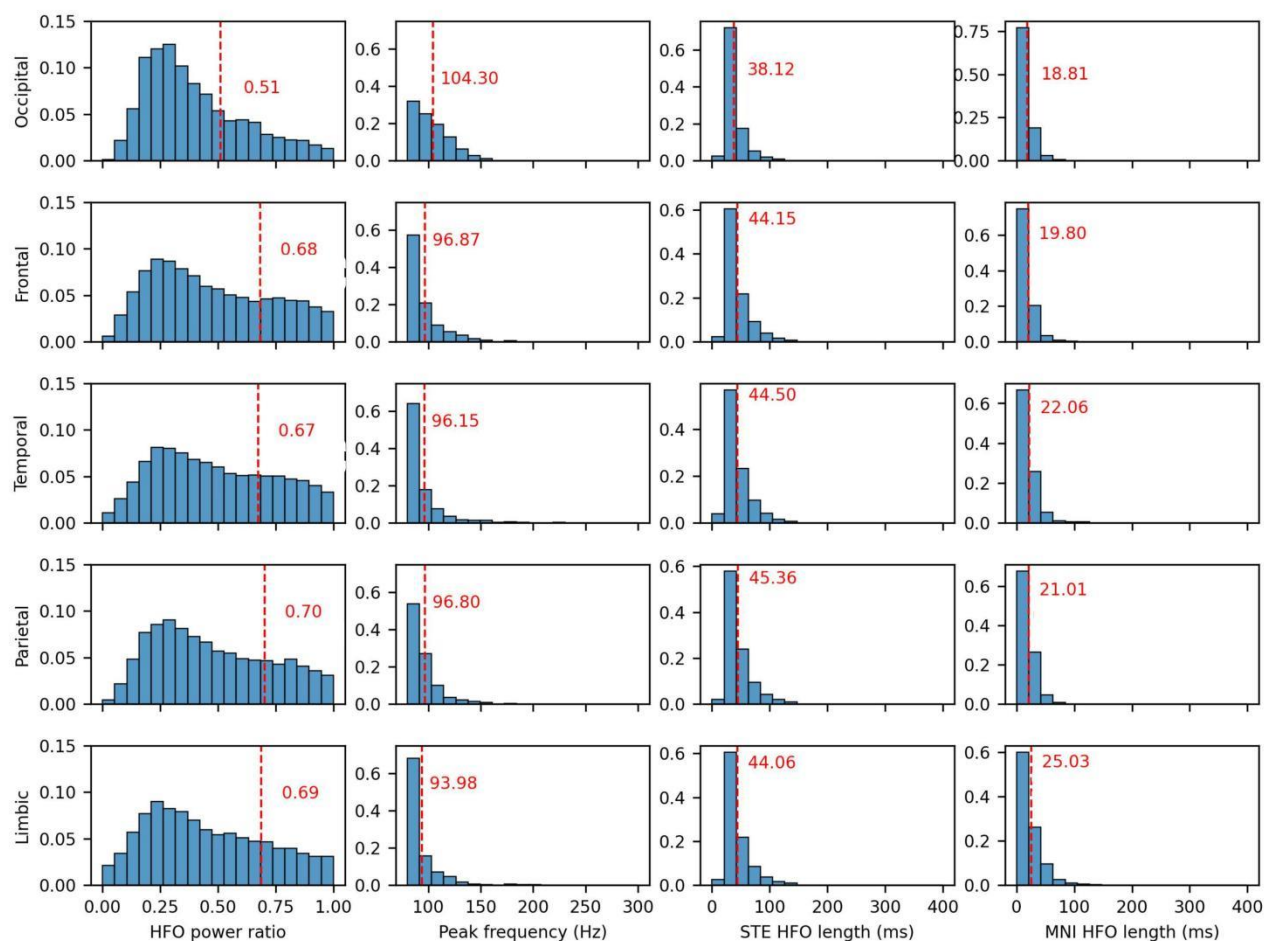
Supplementary Figure 3. Building postoperative seizure outcome prediction model using HFO information



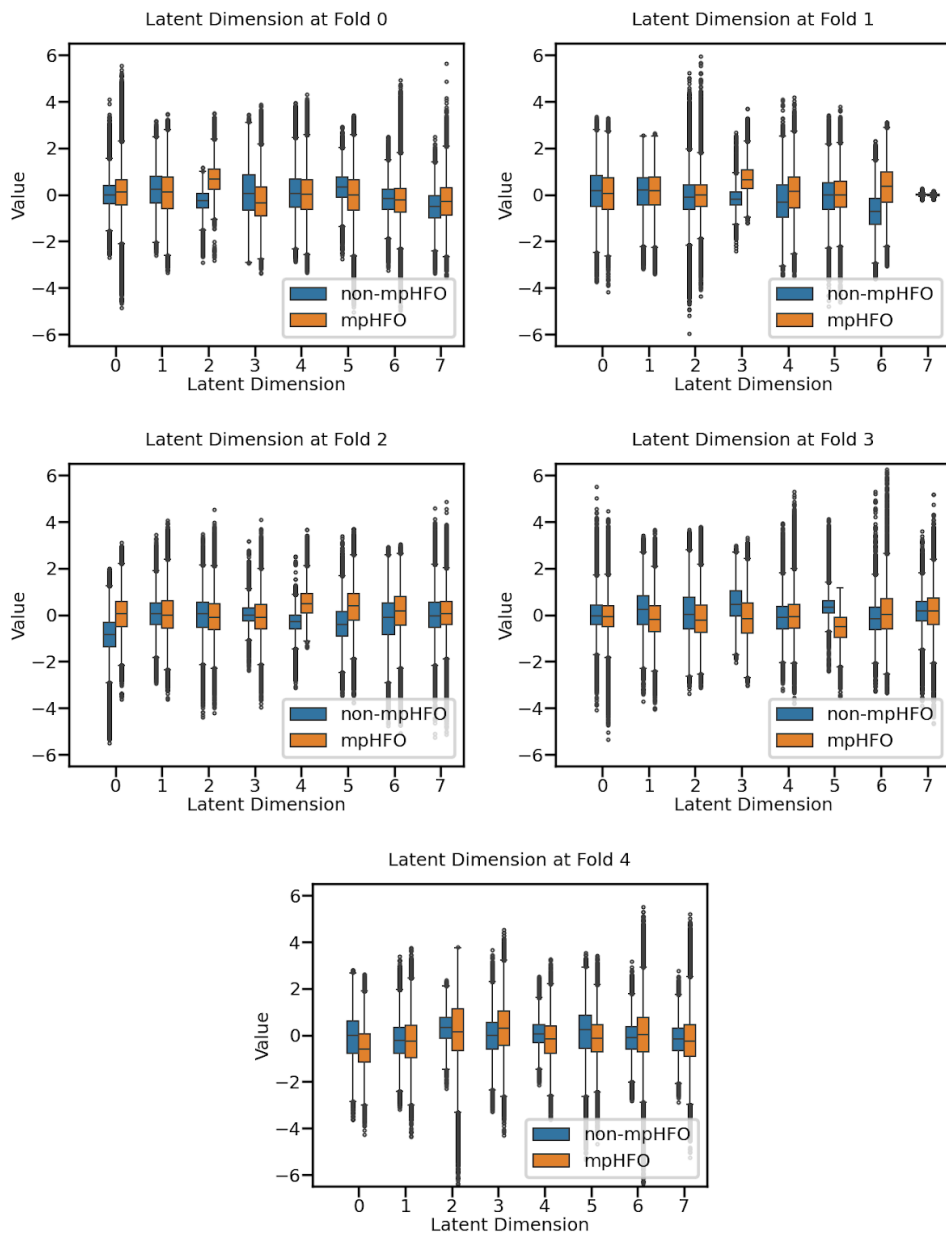
Supplementary Figure 4. Latent space comparison: This figure illustrates the 2D projection of the latent space, color-coded according to classifications from a VAE (mpHFO, non-mpHFO, and mArtifact) and from conventional supervised methods (spKHFO, non-spKHFO, and Artifact), evaluated across five folds. It highlights the alignment between VAE predictions and conventional supervised classifications, demonstrating that classifiers developed from unsupervised learning can achieve performance comparable to those developed from supervised methods.



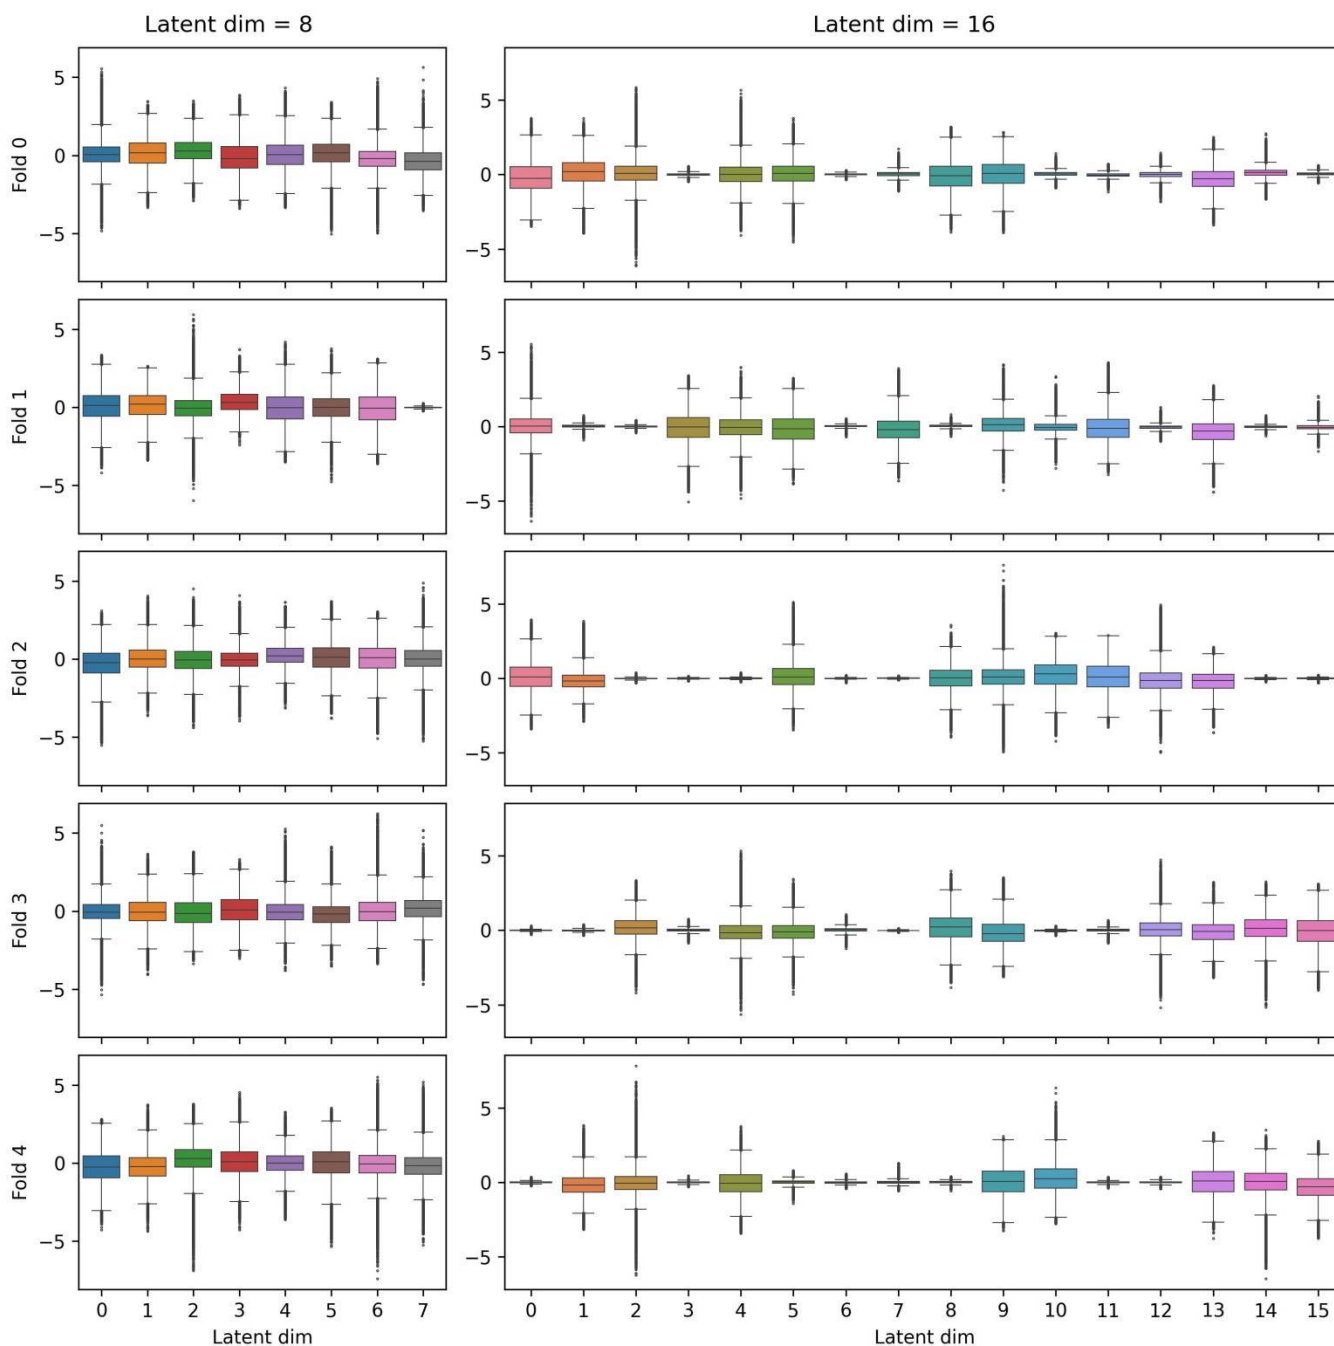
Supplementary Figure 5. mpHFO shares similar morphology across the variables. The morphological analysis of mpHFOs' time-frequency plot is shown based on subgroups, each corresponding to a specific variable. Consistent morphological features of mpHFOs are seen across the subgroups. mpHFOs have higher values throughout the HFO band (> 80 Hz), around the center point (0 ms, where HFOs were detected) than non-mpHFOs; furthermore, higher values of mpHFOs at the sub-HFO band (10-80 Hz) throughout the time window compared to non-mpHFOs are seen.



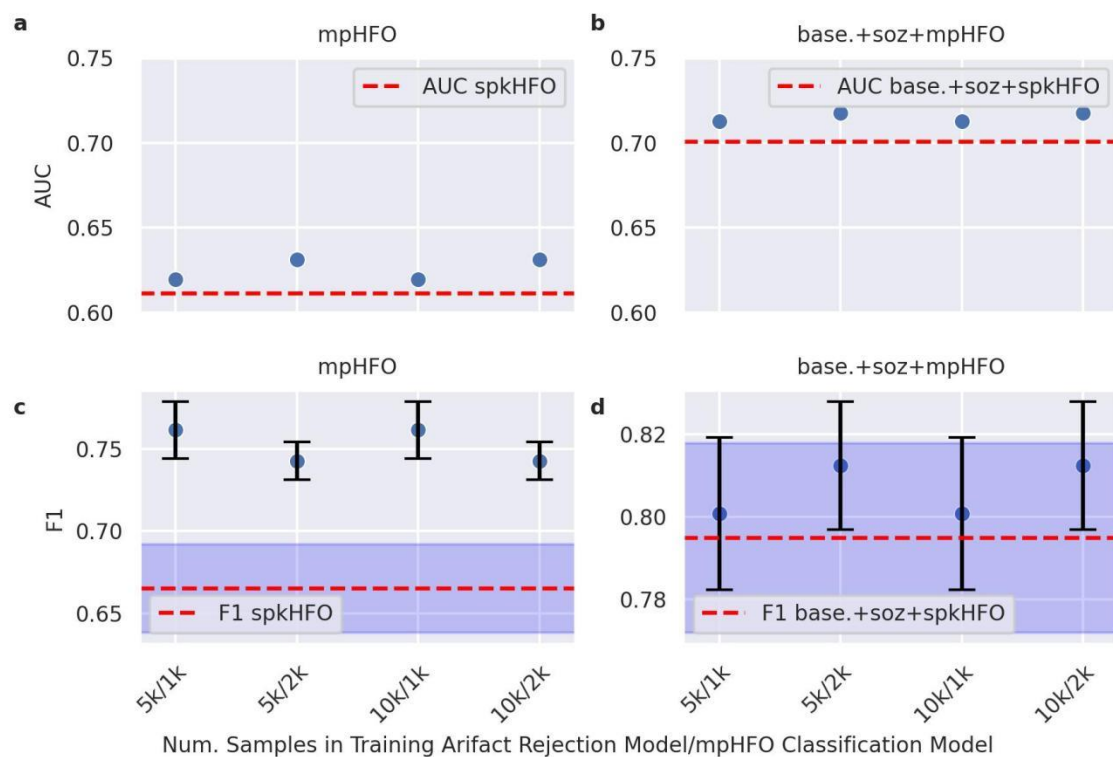
Supplementary Figure 6. Conventional physiological non-mpHFO characteristics across different anatomical locations. This figure displays multiple aspects of HFOs: the HFO power ratio in the first column, HFO peak frequency in the second column, STE HFO length in the third column, and MNI HFO length in the fourth column. The HFO power ratio is defined as the ratio of the signal power within ± 50 ms of the HFO center to the power within the entire time window of ± 285 ms. The peak frequency is identified as the frequency component that has the most power within the 80-300 Hz range. Both STE and MNI HFO lengths are measured in milliseconds as the duration of the HFO event. Each histogram's mean value is annotated within its respective panel. Additionally, statistical comparisons revealed that physiological HFOs in the occipital region exhibited lower power ($p < 0.001$, one-sided t-test), higher peak frequency ($p < 0.001$, one-sided t-test), and shorter durations in both STE and MNI HFOs ($p < 0.001$ for all, except in the frontal location where $p < 0.01$, one-sided t-test) compared to those in other anatomical locations.



Supplementary Figure 7. Full latent space visualization across five folds. This visualization presents the distribution of each latent dimension, trained within five folds and depicted by separating mpHFO (in orange) and non-mpHFO (in blue). It is evident that mpHFO and non-mpHFO exhibit significant separation in some dimensions while others show less distinction.

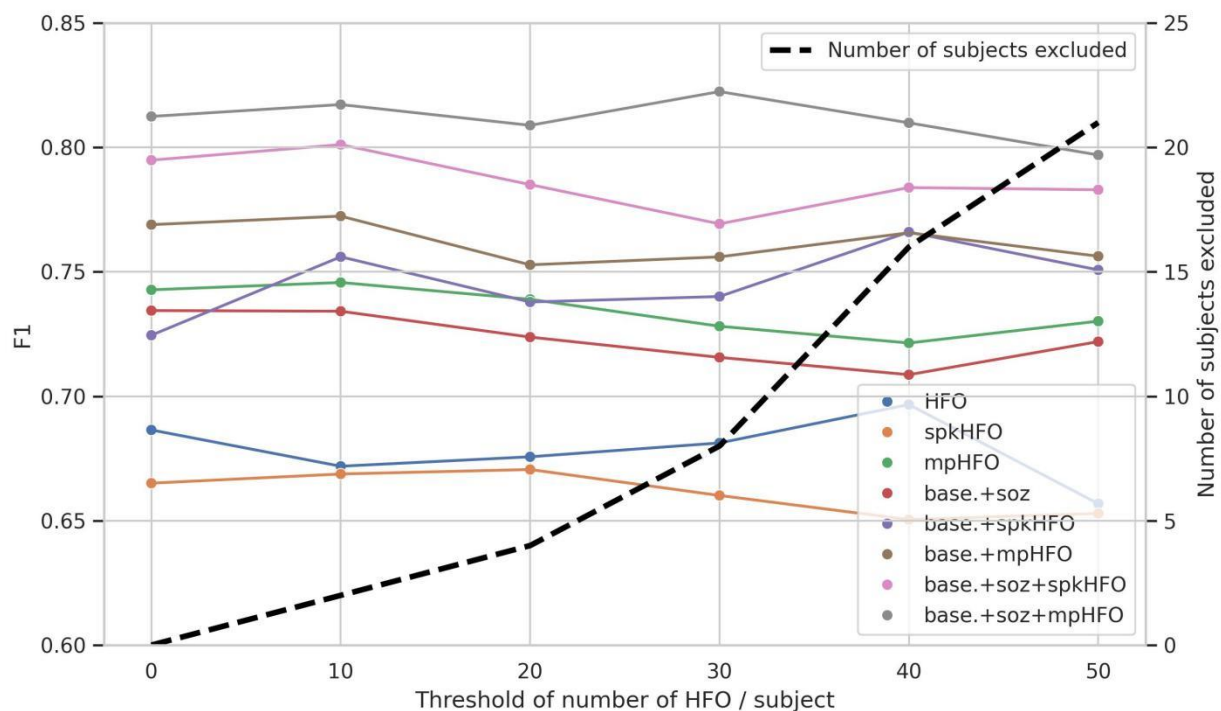


Supplementary Figure 8. Ablation study on different latent dimensions. The left panel shows the distribution of dimensions for a VAE model trained with a latent dimension of 8, across five folds. The right panel displays the distribution for the same model configuration but with a latent dimension of 16. Several dimensions in the 16-dimension model do not spread well, indicating that these dimensions are underutilized during training. This compactness suggests a potential reduction in dimensionality. Based on these observations and the more effective utilization of dimensions seen in the boxplot for the 8-dimension model, we chose to proceed with a latent dimension of 8, as it demonstrates optimal use of all available dimensions during training.



This is the new figure using threshold = 0

Supplementary Figure 9. Ablation study on predicting surgical outcomes using different sample sizes for training GMM models. We conducted experiments using different stratified sampling to train GMM models and investigated how capping the number of samples per subject used to train the artifact rejection and mpHFO classification impacts surgical outcomes prediction performance metrics such as AUC and F1 score. We compared models capped at 10,000 or 5,000 samples per subject for artifact rejection and 2,000 or 1,000 samples for mpHFO classification. (a) Comparison of AUC scores for mpHFO models trained from different stratified sampling sizes against those from spkHFO models. (b) A similar comparison for base.+soz+mpHFO against base.+soz+spkHFO. (c) Comparison of F1 scores for mpHFO models trained from different stratified sampling sizes against those from spkHFO models. (d) Similar procedure as in (c), Similar comparison for base.+soz+mpHFO against base.+soz+spkHFO.



Supplementary Figure 10. Ablation study on predicting surgical outcomes by excluding patients with fewer HFOs: We conducted an ablation study to assess the impact of excluding patients with a smaller number of high-frequency oscillations (HFOs) on the performance of our Random Forest model in predicting post-surgical outcomes. In addition to the results presented in the main manuscript, which included all patients, we compared the F1 scores for predicting post-surgical outcomes by selectively including patients with a minimum of 10, 20, 30, 40, and 50 HFOs. The F1 scores were plotted for different feature sets used in constructing the Random Forest model. The results indicate that the F1 score remains consistent even when including only patients with more than 50 HFOs, demonstrating the robustness of our self-supervised framework.

Supplementary Table 1

Dataset	No. Event	No. Artifact	No. spkHFO	No. non-spkHFO
STE Detector				
Detroit grid/strip	35610	4374	20481	10755
UCLA grid/strip	103360	14578	54455	34327
UCLA SEEG	35861	7886	21748	6227
All STE HFOs	174831	26838	96684	51309
MNI Detector				
Detroit grid/strip	36279	6494	26953	2832
UCLA grid/strip	124064	42848	69266	11950
UCLA SEEG	351236	275550	50434	25252
All MNI HFOs	511579	324892	146653	40034

Supplementary Table 2

Feature (s)	Accuracy	F1
HFO	0.559(0.089)	0.686(0.081)
spkHFO	0.553(0.141)	0.665(0.134)
mpHFO	0.657(0.051)	0.743(0.056)
base.+ soz	0.637(0.106)	0.734(0.090)
base.+ spkHFO	0.624(0.113)	0.725(0.108)
base.+ mpHFO	0.668(0.111)	0.769(0.096)
base.+ soz + spkHFO	0.719(0.126)	0.795(0.115)
base.+ soz + mpHFO	0.732(0.084)	0.812(0.078)

Supplementary Figure 3

	STE		MNI
Filter freq (Hz)	[80, 300]	Filter freq (Hz)	[80, 300]
RMS window (s)	0.003	Epo CHF (s)	60
Min window (s)	0.006	Per CHF (%)	0.95
Min gap (s)	0.01	Min win (s)	0.01
Min osc (count)	6	Min gap (s)	0.01
RMS thres (in SD)	5	Thrd perc (%)	0.999999
Peak thres (in SD)	3	Base seg (s)	0.125
Epoch len (ms)	600	Base shift (0-1)	0.5
		Base thrd (0-1)	0.67
		Base min (%)	5
		Epoch time (ms)	10
		Seed	0