

Real-Time Seizure Prediction Using Low-Power Spiking Neural Networks on Wearable Edge Devices

Abdullah Ghanim Jaber*¹

¹ University of Information Technology and Communications, Baghdad, Iraq

Article Info	ABSTRACT
Article history: Received Apr., 27, 2026 Revised: May,30,2026 Accepted: Aug.,15,2026 Keywords: Spiking Neural Networks (SNNs) Seizure Prediction Wearable Edge Devices EEG Signal Processing Energy-Efficient Computing	Epilepsy is a condition that is present in 50 million individuals all over the world, but predicting seizures in real-time on wearable devices is a challenging task because of stringent power and latency requirements. The current paper introduces a wearable edge devices-based low-power seizure prediction system based on Spiking Neural Networks (SNNs). In contrast to traditional deep learning models, SNNs compute EEG signals as a result of events, and spikes, which consumes much less energy by a factor of ten. We suggest a three-layer SNN, leaky integrate-and-fire neurons, adaptive thresholding and sparse connectivity, trained using a surrogate gradient descent. The system based on CHB-MIT scalp EEG database and patient-wise splitting, has a sensitivity of 88.4, specificity of 93.2 and a false alarm rate of 0.08/hour, predicting seizures 5-30 minutes before they occur. The system is implemented on an ARM Cortex-M4 microcontroller, using 3.4 - 13 times less energy than a CNN-LSTM baseline to make an inference, and has 45 ms latency. Ablation experiments validate the work of adaptive thresholding, population coding and 16 ms timestep resolution. This is the first SNN-based seizure prediction system to the best of knowledge that has been tested on a publicly available EEG dataset and with wearable-edge hardware simulation. The framework shows that neuromorphic computing has the capabilities to address clinical and engineering needs of real-time neurological monitoring.
Corresponding Author: Abdullah Ghanim Jaber University of Information Technology and Communications (UoITC), Baghdad, Iraq Email: abdullah.ghanim@uoitc.edu.iq	

1. INTRODUCTION

Epilepsy is a disorder that is prevalent in the world with about 50 million individuals being affected by the disorder with close to one third of the affected individuals having uncontrolled seizures despite treatment [1]. The randomness of the seizures has a severe effect on the quality of life and there is a dire need to develop effective prediction mechanisms that will be able to give timely alerts. Conventional seizure detection methods based on artificial neural networks (ANNs) have also been promising but are limited in their application in the real world because they require a lot of computations and consume a lot of energy [2]. These limitations are especially troublesome in case of a continuous monitoring of wearable devices, where power efficiency and real-time processing are the most important factors.

The recent developments in neuromorphic computing have brought about spiking neural networks (SNNs) as a biologically realistic alternative to traditional ANNs [3]. SNNs are event-based, like biological neurons, in that information is computed using discrete spikes, instead of continuous activations. This basic distinction makes it possible to save a lot of energy since only when input stimuli reach some predetermined thresholds, calculations are made. Furthermore, the limited communication between neurons in SNNs lowers the bandwidth memory demands and hence they are especially appropriate with edge devices with limited resources [4].

The combination of SNNs and wearable technology offers a special chance to cope with the issues of real-time prediction of seizures. Wearable edge devices are capable of processing EEG data at the edge, and do not need to be connected to the cloud at all times, preserving patient privacy [5]. Recent developments in edge-based AI for

physiological monitoring have highlighted the need for both computational efficiency and data security [6]. Nevertheless, the current solutions tend to be weak in terms of accuracy in prediction and efficiency. There are those systems that are very sensitive at the expense of too many false alarms and others that are very sensitive but are low-power consuming [7]. This trade-off is particularly acute in the case of EEG signals variability in various patients and types of seizures.

The suggested SNN-based system is a new framework that is specifically developed to predict seizures in real-time on wearable edge devices. There are three main aspects in approach which is different as compared to the previous work. The first step is to create a lightweight preprocessing pipeline with signal fidelity and lower computational costs. Second, it presents a training methodology that is patient-specific and takes into consideration individual electrophysiological properties without the need to have too much calibration information. Third, the optimization of SNN architecture to be efficient in spike-based computation, allowing to make accurate predictions with the minimal energy usage.

The main contribution of this work is that SNNs can be used to attain clinically relevant prediction performance, and satisfy the power requirements of wearable devices, which are very stringent. The system is able to meet these competing requirements, whereas in the past researchers focused on one of these requirements, or sacrificed time resolution for energy efficiency. The experimental results show that the proposed method can predict seizures 5-30 minutes in advance with a high degree of accuracy, which is suitable for clinical application, and consumes orders of magnitude less power than the conventional deep learning algorithms.

A number of practical issues in implementing SNNs to predict seizures are also discussed in this paper. The discussion covers trade-offs between prediction horizon and false alarm rate, how patient-specific adaptation affects the system performance and the hardware implications of applying the framework to real edge devices. These findings can be applied to the overall area of neuromorphic healthcare application and can serve as a guide to future studies in the area of low-power neural signal processing.

The rest of this paper will be structured in the following way: Section 2 will be a review of related work in the fields of seizure prediction and neuromorphic computing. Section 3 gives background information needed on SNNs and characteristics of EEG signals. Section 4 outlines the suggested structure, and experimental setup in Section 5. Findings and discussion can be found in Section 6 and discussion and future work in Section 7. Section 8 is a conclusion of the paper.

2. RELATED WORK

Real-time seizure prediction systems have been developed over a number of technological generations, starting with the primitive signal processing methods, to the more recent machine learning methods. The initial techniques used mainly focused on handcrafted characteristics derived out of EEG signals including spectral power, entropy measures, and nonlinear dynamics [8]. Although these methods showed decent performance in the controlled environment, they were frequently not robust enough to be applied in the real-world since they were sensitive to noise and inter-patient variability.

CNNs and RNNs emerged as the leading types of neural networks with the advent of deep learning and found their application in seizure prediction studies [9]. These models were able to automatically learn the relevant features of raw EEG data, without the need to manually engineer features. They were however computationally complex and high energy consuming, and not practical to use in wearable devices to do continuous monitoring. The research direction then moved to more efficient architecture exploration, such as hybrid architectures that were a combination of the traditional signal processing and shallow neural networks [10].

Recent studies have explored the possibility of spiking neural networks in predicting seizures, taking advantage of the event-driven nature and energy efficiency of the networks. A number of studies have shown that SNNs can be as accurate as traditional ANNs, and use much less power [11]. As an example, [12] used a seizure detector based on SNNs on a custom hardware with a power consumption of less than milliwatts. Their strategy was more of detection, and less of prediction and that more sophisticated architectures should be developed that can predict the seizures before they happen.

The use of SNNs in wearable devices has its own challenges in compressing the models and the processing in real-time. Other authors have studied the idea of quantization and pruning to decrease the memory footprint of SNN models [13]. Specialized hardware architectures have also been designed to compute using spikes, including neuromorphic processors which model biological neural networks [14]. The developments have made it possible to implement SNNs on resource-constrained systems, at the expense of worse prediction accuracy or higher latency.

2.1. Predicting seizures using EEG Signals

The basis of any seizure prediction system is to have effective EEG processing. PhysioNet database has been utilized as a useful tool to create and test algorithms and offer standardized EEG recordings of epilepsy patients [15]. Conventional methods used to use bandpass filtering to select the desired frequency content, and then extract features in time, frequency or time-frequency space. The more modern techniques use end-to-end learning, where neural networks directly take raw or minimally processed EEG signals [16].

The preprocessing steps are highly dependent on the choice of the preprocessing steps, which affect performance and computational requirements of the system. Downsampling, e.g., can be used to decrease the amount of data, but can eliminate high-frequency content that can be used to predict seizures [17]. Similarly, the temporal resolution and computational efficiency are affected by windowing, shorter windows have a finer temporal grain but require more frequent processing [18]. The work is based on these insights and optimizations specific to the processing of SNNs.

2.2. Spiking Neural Network Architectures

The SNNs are becoming a promising solution in edge computing because of their biological plausibility and energy efficiency. The model is based on the leaky integrate-and-fire (LIF) neuron model, a compromise between biological and computational complexity [19]. More recent developments in training algorithms, especially surrogate gradient algorithms, have overcome the non-differentiability of spiking neurons, and have made it possible to backpropagate through time [20].

Several architectural improvements have been made to improve the performance of SNN in temporal signal processing. The recurrent spiking networks are used to learn long-range dependencies in EEG, while the attention mechanisms are used to focus on clinically relevant patterns [21]. However, these improvements are generally more expensive in terms of computation and need to be optimized for wearable applications. The architecture that proposed is particularly sensitive to this trade-off with the use of efficient connectivity patterns and spike-based activation functions.

The system proposed is different in several aspects from the existing systems. First, the entire processing chain, from signal acquisition to classification, etc., is optimized for edge deployment, not changing the existing cloud-based models. Second, the patient-specific training approach takes into consideration individual variability without a large amount of calibration data. Third, it introduces a new set of architectural decisions and training methods that do not lower the accuracy of prediction and do not reduce the energy consumption. All these innovations allow real-time functioning of wearable devices with resource constraints, which is a serious gap in the existing seizure prediction technology.

2.3. Epilepsy and Seizure Basics

Epileptic seizures are caused by abnormal excessive electrical discharges in clusters of brain cells, which are temporary impairments in neurological functioning [22]. These events are usually between seconds and minutes and may have a wide range of clinical manifestation. The preictal period, which takes place minutes to hours prior to the onset of the seizure, has slight electrophysiological alterations that can be used as predictive biomarkers [23]. To detect these early signs, it is crucial to understand the complex spatiotemporal dynamics of neural activity, which can be captured by neural network models.

2.4. Electroencephalography (EEG) Signals

EEG is a voltage variation, brought about by the ionic current of the neurons and is typically observed by placing electrodes on the head. There are several common frequency bands of the signals, including delta (0.5-4 Hz), theta (4-8 Hz), alpha (8-13 Hz), beta (13-30 Hz) as well as gamma (>30 Hz) that refer to different brain states and functions [24]. These rhythms are usually disrupted during the epileptic seizures and the synchrony and amplitude of these rhythms is greater in specific areas. Preictal changes are however more elusive and they require sophisticated techniques of analysis in order to be identified.

An EEG signal $x(t)$ can be mathematically represented as:

$$x(t) = \sum_{i=1}^N a_i \sin(2\pi f_i t + \phi_i) + n(t) \quad (1)$$

where a_i , f_i , and ϕ_i represent the amplitude, frequency, and phase of each oscillatory component, while $n(t)$ accounts for noise and artifacts. This definition emphasizes the composite character of EEG signals, in which the appropriate information should be obtained out of a blend of neural activity and interference.

2.5. Neural Networks Fundamentals

The information is calculated using artificial neural networks with the assistance of network layers of computational units, known as neurons. The fundamental performance of a single neuron can be written as:

$$y = f(\sum_{i=1}^n w_i x_i + b) \quad (2)$$

In which x_i are the inputs, w_i are the weights, b is a bias term and f is a nonlinear activation function. Conventional neural networks operate with continuous-valued activations, which necessitate large amounts of computational resources to perform matrix multiplications and ensure high precision across the network.

Spiking neural networks are fundamentally different in that they use discrete, event-based communication between neurons. SNN neurons do not produce continuous activations, but rather binary spikes when their membrane potential $V(t)$ reaches a threshold V_{th} :

$$V(t) = \sum_i w_i \sum_{t_i} K(t - t_i) \quad (3)$$

where K a kernel function that represents the postsynaptic potential, and t_i the input spike times. This spatiotemporal processing is similar to biological neural dynamics, and allows more efficient computation with sparse, event-driven operations [25].

A popular model of SNN neurons is the leaky integrate-and-fire (LIF) model, which models the dynamics of membrane potentials as:

$$\tau_m \frac{dV}{dt} = -(V - V_{rest}) + RI(t) \quad (4)$$

where τ_m the time constant of the membrane, V_{rest} the resting potential, R the resistance of the membrane, and $I(t)$ the input current. When $V(t)$ reaches V_{th} , the neuron fires a spike and resets to V_{reset} . It is a simple but powerful model that captures important features of biological neurons and is computationally tractable to simulate on a large scale [26].

The time resolution of EEG and the event-based processing of SNNs are a natural synergy in seizure prediction. EEG offers milliseconds-level monitoring of brain activity, whereas SNNs are capable of effectively identifying and processing the temporal dynamics that are predictive of impending seizures. This correspondence between the properties of signals and the computational paradigm is the theoretical basis of the proposed framework.

3. METHOD

The framework proposed combines spiking neural networks with optimized signal processing to realize real-time seizure prediction on wearable edge devices. As shown in Figure 1, the system architecture consists of three main components: an EEG preprocessing module, an SNN-based feature extraction and classification module, and an alert generation mechanism. The design focuses on computational efficiency and has enough temporal resolution to detect early seizures, which is important in the context of continuous monitoring applications where there is a critical trade-off between prediction accuracy and power consumption.




Figure 1. System Architecture of the Wearable Seizure Prediction System

3.1. SNN Architecture for Seizure Prediction

The proposed SNN architecture consists of three layers designed to process EEG signals while maintaining low computational overhead. The preprocessed EEG data is converted into spike trains and fed into the input layer via a rate coding scheme. The input neurons I_j are associated with the EEG channels, and the firing rates are proportional to the normalized signal amplitude:

$$r_j = \frac{x_j - \mu_j}{\sigma_j} \cdot r_{\max} \quad (5)$$

where x_j being the raw EEG signal, μ_j and σ_j being channel-specific normalization factors, and $r_{\max}$ being the maximum firing rate. This transformation maintains the time dynamics of EEG signals but transforms them into a format that can be processed by SNN.

The concealed layers use leaky integrate-and-fire (LIF) neurons with adaptive thresholding to enhance noise resistance. $V_i^l(t)$ the membrane potential of neuron i in layer l is changing as:

$$\tau_m \frac{dV_i^l}{dt} = -(V_i^l - V_{\text{rest}}) + R \sum_j w_{ij}^l S_j^{l-1}(t) \quad (6)$$

where τ_m is the membrane time constant (20 ms), V_{rest} the resting potential (-70 mV), R the membrane resistance, w_{ij}^l the synaptic weight from neuron j in layer $l-1$, and $S_j^{l-1}(t)$ the incoming spike train. When $V_i^l(t)$ crosses the threshold V_{th} (-50 mV), the neuron emits a spike and resets to V_{reset} (-60 mV). The adaptive threshold mechanism prevents excessive firing by increasing V_{th} after each spike:

$$V_{\text{th}}(t) = V_{\text{th0}} + \alpha \sum_{t_k < t} e^{-(t-t_k)/\tau_{\text{th}}} \quad (7)$$

where V_{th0} is the baseline threshold, α controls the adaptation strength, and τ_{th} determines the decay time constant (100 ms). This adaptation assists in preserving sparse firing patterns that are necessary to achieve energy-efficient computation, and the output layer consists of two LIF neurons that model the preictal and interictal classes. The spikes more than 5 seconds are classified and the higher number is defined to predict the state. The connectivity of the network is sparsely random with a density of 30 percent to minimize the computational cost, but still has the adequate representational capacity. The Synaptic weights are initialized using a Gaussian distribution with mean 0 and variance $1/\sqrt{n_{\text{inputs}}}$ to prevent vanishing or exploding gradients during training.

3.2. Edge-Compatible Preprocessing Pipeline

The preprocessing step transforms raw EEG signals into a format that can be processed by SNN, while also optimizing the computational efficiency for deployment on the edge. The input signal $x(t)$ first undergoes bandpass filtering (0.5-40 Hz) to remove artifacts and isolate clinically relevant frequencies. A Butterworth fourth-order filter has adequate roll-off characteristics at a reasonable computational cost:

$$H(s) = \prod_{k=1}^2 \frac{\omega_0^2}{s^2 + \frac{\omega_0}{Q_k}s + \omega_0^2} \quad (8)$$

where ω_0 is center frequency and Q_k is quality factor of each second-order section. The filtered signal $x_f(t)$ is downsampled from 256 Hz to 128 Hz which reduces computational load but it's not lose any important temporal information, while downsampling operation D uses an anti-aliasing filter and then selects samples:

$$x_d[n] = x_f(n \cdot T_s'), \quad T_s' = 2T_s \quad (9)$$

where T_s is the original sampling period (1/256 s) and T_s' new sampling period (1/128 s). To tradeoff between temporal resolution and computational efficiency, the signal is divided into 5-second windows with 75% overlap (1.25 s step size). The samples in each window W_k are $N = 640$ (5 s $\times$ 128 Hz) and normalized to zero mean and unit variance:

$$\hat{W}_k = \frac{W_k - \mu_k}{\sigma_k} \quad (10)$$

where μ_k and σ_k are the window-specific mean and standard deviation. This normalization corrects for inter-window variability, but maintains relative amplitude relationships within each window.

The preprocessed windows are encoded into spike trains with a population coding scheme. The input neurons $I_{j,m}$ are associated with a range of amplitudes A_j and a time t_m in the window:

$$S_{j,m}(t) = \begin{cases} 1 & \text{if } \widehat{W}_k(t_m) \in A_j \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

The amplitude ranges A_j are equally spaced across the expected signal range (-4σ to $+4\sigma$) and 16 levels are employed to provide sufficient resolution without requiring too many levels in the input population. This encoding scheme maintains both time and amplitude data in a spike-based format that can be used in the SNN architecture.

The preprocessing pipeline is implemented in such a way that it can be run with fixed-point arithmetic to minimize the energy usage of edge devices. Filter coefficients and normalization parameters are all represented as 16-bit numbers, providing adequate signal fidelity with less memory and computation than floating-point counterparts. The whole pipeline is capable of operating at a latency of less than 10 ms per window, real-time on low-power hardware.

3.3. Surrogate Gradient Learning for SNN Training

The non-differentiable characteristic of spiking neurons is a fundamental problem in training SNNs by backpropagation. To overcome this, I use a surrogate gradient method that estimates the derivative of the spiking function in the backward pass and the actual LIF dynamics in the forward pass. The spike generation function $S(t) = \Theta(V(t) - V_{th})$, where Θ is substituted in the backpropagation process with a fast-sigmoid surrogate:

$$\frac{\partial S(t)}{\partial V(t)} \approx \frac{1}{(1 + \beta|V(t) - V_{th}|)^2} \quad (12)$$

In this case, β regulates the steepness of the surrogate gradient and use 5 in implementation, which gives enough gradient information to learn effectively, but is close to the actual spiking behavior. This approximation allows gradient-based optimization using the temporal dynamics of the SNN without compromising the energy efficiency of event-driven computation.

Backpropagation through time (BPTT) with truncated horizon of 20 timesteps (320 ms at 16 ms resolution) is used in the training process.

$$\mathcal{L} = -\sum_c y_c \log p_c + \lambda \sum_{l,i} \left(\frac{1}{T} \sum_{t=1}^T S_i^l(t) - r_{\text{target}} \right)^2 \quad (13)$$

where y_c represents the true class probabilities, p_c the predicted probabilities based on output spike counts, λ the regularization strength (0.01), and r_{target} the target firing rate (10 Hz). The second term encourages sparse firing patterns, reducing energy consumption during inference.

The AdamW optimizer updates network parameters with weight decay separation:

$$\theta_{t+1} = \theta_t - \eta \frac{\widehat{m}_t}{\sqrt{\widehat{v}_t + \epsilon}} - \eta \lambda \theta_t \quad (14)$$

where η is the learning rate (initialized at 1×10^{-3}), $\widehat{m}_t$ and $\widehat{v}_t$ the bias-corrected first and second moment estimates, and λ the weight decay factor (5×10^{-4}). The learning rate follows a cosine decay schedule with warmup:

$$\eta_t = \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min})(1 + \cos(\pi t/t_{\max})) \quad (15)$$

for $t \leq t_{\text{warmup}}$, where t_{warmup} is 10% of total training steps, this scheduling strategy provides stable convergence while allowing fine-grained parameter updates in later training stages.

3.4. Patient - Wise Data Splitting

In order to guarantee the generalizability of the proposed seizure prediction framework, I use a patient-wise data splitting approach that strictly divides training, validation, and test sets by patients. This method eliminates data leakage that may artificially inflate performance measures when samples of the same patient are used in more than

one split [27]. Where $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$ denotes the set of N patients in the dataset, P_i is the set of M_i EEG recordings of patient with labeled preictal and interictal segments. The data is divided in such a way that:

$$\mathcal{P} = \mathcal{P}_{\text{train}} \cup \mathcal{P}_{\text{val}} \cup \mathcal{P}_{\text{test}}, \quad \mathcal{P}_{\text{train}} \cap \mathcal{P}_{\text{val}} \cap \mathcal{P}_{\text{test}} = \emptyset \quad (16)$$

The training set $\mathcal{P}_{\text{train}}$ comprises 70% of patients ($\lfloor 0.7N \rfloor$), while the validation $\mathcal{P}_{\text{val}}$ and test sets $\mathcal{P}_{\text{test}}$ each contain 15%. This division makes sure that the model is performing based on its capability to generalize to unknown patients and not memorizing patient-specific artifacts.

To process the recordings in batches, the recordings subdivide of each patient $P_i \in \mathcal{P}_{\text{train}}$ into non-overlapping segments. Where $S_{i,j}$ is the j -th 5-second window of patient P_i , and $y_{i,j} \in \{0,1\}$ is the label (0 interictal, 1 preictal). The training batches are built by sampling randomly segments of various patients:

$$\mathcal{B}_k = \{S_{i,j} | i \sim U(\mathcal{P}_{\text{train}}), j \sim U(1, M_i)\} \quad (17)$$

where U denotes uniform sampling and $|\mathcal{B}_k|$ represents the batch size of 32. This patient-wise shuffling ensures that the model does not overfit to temporal dependencies in individual recordings, but still has diversity in the training data.

The validation and test sets are run in a sequence of patient to replicate real world deployment conditions. For $P_i \in \mathcal{P}_{\text{test}}$, each segment $\{S_{i,1}, \dots, S_{i,M_i}\}$ is tested sequentially to determine the performance of the model on continuous EEG streams. This assessment plan is indicative of the time dynamics of seizure prediction where the system is expected to have low latency and handle real-time incoming data.

3.5. Timestep Resolution and Its Role

One of the trade-offs in the prediction accuracy and computational efficiency of the SNN is the time resolution of the SNN. The define of timestep Δt as the time interval between updates of the network state, with smaller values providing a finer time resolution at the cost of increased computation. The empirically chose 16 ms as the optimal timestep to predict seizures, as it is the best timestep in ablation experiments, and can still be run in real-time on edge hardware. The forward Euler method is used to discretize the membrane potential update equation 6 with the following time step:

$$V_i^l[t+1] = \beta V_i^l[t] + (1 - \beta)(RI_i^l[t] + V_{\text{rest}}) \quad (18)$$

where $\beta = e^{-\Delta t/\tau_m}$ represents the decay factor (approximately 0.45 for $\tau_m = 20$ ms and $\Delta t = 16$ ms). The input current $I_i^l[t]$ is the sum of incoming spikes during the timestep:

$$I_i^l[t] = \sum_j w_{ij}^l \sum_{t_k \in [t, t+\Delta t)} \delta(t - t_k) \quad (19)$$

where δ is the Dirac delta function. This formulation ensures that each spike contributes to the postsynaptic current only at its own timestep, maintaining time accuracy and enabling computation to be performed efficiently. The 16 ms timestep has a number of benefits. First, it complies with the Nyquist criterion of recording EEG gamma band activity (up to 40 Hz after filtering), which does not lose any clinically relevant information. Second, it enables the network to run every 5-second window in 312 timesteps (5000 ms / 16 ms), which is a reasonable amount to run in real-time. Third, this resolution allows accurate timing of preictal pattern detection and maintains the computational cost within the power budget of wearable devices.

The effect of timestep on energy consumption is nonlinear. Let E_{op} be the energy per operation (e.g. membrane potential update or spike propagation). The overall energy E_{total} to process one window is:

$$E_{\text{total}} \propto N_{\text{neurons}} \cdot N_{\text{timesteps}} \cdot E_{\text{op}} \quad (20)$$

where N_{neurons} is the total number of neurons (input + hidden + output layers) and $N_{\text{timesteps}}$ the number of timesteps per window. Reducing Δt from (32 ms to 16 ms) doubles $N_{\text{timesteps}}$ but allows earlier seizure detection, which forms a trade-off between energy consumption and prediction horizon. The experimentally determined that 16 ms is the best compromise, with 5-30 minute advance warnings and sub-milliwatt power consumption. The time step also affects the time coding of the network. The SNN is capable of detecting the

difference in spike timing resolution of 1-2 timesteps with 16 ms resolution, which is enough to detect the slow changes in synchronization that are typical of preictal states. This accuracy would be compromised with larger timesteps (e.g., 50 ms), and smaller resolutions (e.g., 1 ms) would yield diminishing returns in predicting seizures at much greater computational cost. The selected timestep is thus a tradeoff between biological plausibility and engineering.

4. EXPERIMENTAL SETUP

4.1. Dataset Description and Preprocessing

The analysis is conducted on the CHB-MIT Scalp EEG Database [28], which is a common source of reference in the research of seizure prediction. The data set contains EEG recordings of 23 children with intractable epilepsy, sampled at 18-23 EEG channels at 256 Hz. In total, it provides approximately 980 hours of continuous monitoring and 198 clinically annotated seizure events. According to the standard procedures [29], five frontal and temporal channels are selected: F7-T7, F8-T8, T7-P7, T8-P8, and Cz-Pz. These channels are usually linked to robust ictal patterns as they assist in minimizing the effects of muscle and movement artifacts.

A standardized preprocessing pipeline is used to process the raw EEG signals. The first step is to use a fourth-order Butterworth design to filter out low-frequency drift and high-frequency noise with a 0.540 Hz bandpass filter as illustrated in Equation (8). The filtered signals are then downsampled to 128 Hz, as specified in Equation (9), and divided into non-overlapping 5-second windows to be analyzed initially. The z-score of each window is normalized by Equation (10) and the normalization parameters are calculated based on the first 30 minutes of every recording to determine patient-specific baselines. In the case of SNN input, the windows are further subdivided into 16 ms sub-windows, which generate 312 timesteps, and coded into spike trains by population coding, as in Equation (11).

4.2. Baseline Methods

The proposed SNN is compared with four representative baselines from traditional and deep learning.

1. **1D CNN:** A temporal convolutional network with three 1D convolutional layers (64, 128, 256 filters) followed by two dense layers, implementing the architecture from [30]. The model processes raw EEG windows after downsampling to 128 Hz.
2. **CNN-LSTM:** A hybrid model that uses 2D convolutional layers to extract spatial features and LSTM layers to model time, as per the state-of-the-art design in [31]. The input is spectrograms calculated with 250 ms Hanning windows and 125 ms overlap.
3. **Random Forest:** A traditional machine learning model based on 100 decision trees trained on handcrafted features such as bandpower (delta, theta, alpha, beta, gamma), Hjorth parameters, and line length features derived on 5-second windows [32].
4. **Amplitude Threshold:** A basic detector that raises alarms when any channel is more than 3 standard deviations above the baseline mean over a period of more than 2 seconds, using the technique in [33].

To compare the SNN with all baseline models, all baseline models are trained and evaluated on the same patient-wise split (Equation 16). The deep learning models (CNN, CNN-LSTM) use the same training protocols (AdamW optimizer, cosine learning rate schedule) as the SNN to ensure consistency.

4.3. Evaluation Metrics

Five clinically relevant measures are used to evaluate the system performance:

1. **Sensitivity (Recall):** The percentage of true seizures that are correctly predicted, as $\frac{TP}{TP+FN}$, where TP and FN are true and false positives respectively. The target threshold is >85% according to clinical usability requirements [8].
2. **Specificity:** The percentage of interictal periods correctly identified, calculated as $\frac{TN}{TN+FP}$, aiming at a target of >90% to reduce unnecessary interventions.
3. **False Alarm Rate (FAR):** The mean false predictions per hour, and less than 0.1/hour is regarded as clinically acceptable in wearable devices [34].
4. **Prediction Horizon:** The interval between the alarm and the seizure onset, in minutes. We consider 5-30 minutes horizons to determine early warning ability.
5. **AUC-ROC:** The region beneath the receiver operating characteristic curve, which gives a complete account of the classification performance at all decision thresholds.

In the case of temporal measures (sensitivity, FAR), i use the time-based evaluation protocol of [35], where a prediction is said to be correct when it is made within 30 minutes before the onset of a seizure, and false alarms are only counted during interictal intervals (at least 4 hours between seizures).

4.3. Training Protocol

The SNN is trained on patient-specific data with leave-one-patient-out cross-validation. The model is trained on all other patients with weights, and then fine-tuned on the first 20% of the interictal data of P_i (at least 30 minutes) and any available preictal data. This method is a balance between generalization and personalization and is clinically practical to implement.

The training is continued until 100 epochs or early stopping in case the validation loss is not decreasing after 20 epochs. The 32 batch size offers consistent gradient estimates and fits within the memory limits of edge devices. i use gradient clipping (max norm 1.0) to avoid exploding gradients in the temporal processing. The learning rate is based on a cosine decay schedule with 10-epoch warmup in equation (15), starting with 1×10^{-3} and decreasing to 1×10^{-5} .

Random channel dropout ($p=0.1$) and additive Gaussian noise ($\sigma=0.1$) are used to augment the data to enhance robustness. The issue of class imbalance is solved by synthetic minority oversampling (SMOTE) [36] on the preictal class, with a 1:1 ratio of interictal samples in training.

4.4. Hardware Implementation

The SNN is deployed on an ARM Cortex-M4 microcontroller (STM32L4 series) that is commonly used in wearable devices, featuring 320 KB of SRAM and 1 MB of flash. The simulation and deployment of SNNs is implemented using Intel's Lava framework [37] that compiles the trained model into optimized C code with fixed-point arithmetic (16-bit weights, 8-bit activations). Continuous EEG streams are processed and power consumption is measured using a Monsoon Power Monitor [38]. The energy per prediction E_{pred} is calculated as:

$$E_{pred} = P_{avg} \cdot t_{window} \quad (21)$$

where P_{avg} is the average power of window processing and t_{window} is the 5 seconds window time. The ARM Memory Protection Unit is used to profile memory usage during inference. To compare, CNN and CNN-LSTM baselines are run on the same hardware with 8-bit quantization using TensorFlow Lite Micro [39]. All models are tested in the same conditions with peripherals turned off to isolate processing costs.

4.5. Statistical Analysis

The performance measures are presented in the form of mean and standard deviation of all test patients. Paired t-tests with Bonferroni correction of multiple comparisons ($\alpha=0.05$) are used to test significance. The sensitivity and FAR differences are measured with the help of Cohen d:

$$d = \frac{\mu_1 - \mu_2}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}}} \quad (22)$$

where μ and σ are the metric means and standard deviations of two compared methods. Bootstrap resampling (1000 times) is used to compute the confidence intervals (95) to take into account the variability of patients.

5. RESULTS AND ANALYSIS

The proposed SNN-based seizure prediction system is efficient in numerous ways as shown in the experimental study. This section provides quantitative comparisons to baseline methods, time prediction properties, and computational efficiency measures that are important in wearable implementation.

5.1. Prediction Performance Comparison

Table 1 provides an overview of the comparative performance of the proposed SNN with four baseline methods in terms of important evaluation metrics. The SNN has a high sensitivity 88.4% and specificity 93.2% which is better than all the traditional methods. It is important to note that the false alarm rate (0.08/hour) is within clinical acceptability limits and offers 5–30-minute advance warnings.

Table 1. Comparative performance of seizure prediction methods

Model	Sensitivity	Specificity	AUC	False Alarms (/h)	Prediction Lead Time (min)
Random Forest	72.3%	85.1%	0.82	0.15	8.2 ± 3.1
CNN (1D)	81.5%	88.3%	0.89	0.12	10.5 ± 4.3
CNN-LSTM	86.2%	91.7%	0.93	0.10	12.1 ± 5.0
SNN (Proposed)	88.4%	93.2%	0.95	0.08	12.7 ± 4.8

Figure 2 shows that the receiver operating characteristic (ROC) curves indicate that the SNN has a steady advantage at all decision thresholds. The 0.95 area under the curve (AUC) is much better than the CNN-LSTM baseline (AUC=0.93, $p < 0.05$) and has a more desirable sensitivity-false positive rate trade-off. This is enhanced by the fact that the SNN is able to detect fine temporal dynamics in preictal states due to its event-based processing.




Figure 2. Receiver operating characteristic curves comparing classification performance across methods

Individual patient analysis shows significant differences in prediction performance. Figure 3 presents the sensitivity and false alarm rates of each test patient. Although the majority of patients have a sensitivity of $>85\%$ with a false alarm rate of less than 0.1 false alarms/hour, there are outliers such as chb03 (71.4% sensitivity, 0.12 FA/h) that demonstrate the difficulty of inter-patient variability in EEG patterns. Such cases can be improved with longer calibration times or adaptive threshold changes at deployment.




Figure 3. Patient-wise distribution of sensitivity and false alarm rates

5.2. Temporal Prediction Characteristics

The proposed system has clinically relevant prediction times with 65% of correct predictions 5-10 minutes before the start of seizures, 25% 10-20 minutes and 10% 20-30 minutes. The cumulative distribution of prediction times is shown in Figure 4, where 90% of alarms were triggered at least 5 minutes before the onset of a seizure. This prediction system is beneficial for preventive measures and low false alarms.




Figure 4. Cumulative distribution of seizure prediction lead times

The preictal dynamics analysis shows that the confidence of the SNN prediction, defined as the difference in the number of spikes in the output, usually starts to increase 15-20 minutes before the seizure onset, as shown in Figure 5. This rise in confidence indicates that the network is capable of detecting subtle electrophysiological changes and not abrupt signal changes. Moreover, the timing of spikes in the hidden layers, as shown in Figure 6, allows the detection of subtle preictal dynamics that are usually ignored by traditional approaches.




Figure 5. Evolution of SNN Prediction Confidence During Preictal Period




Figure 6. Hidden Layer Spike Raster Plot During Preictal Detection

5.3. Computational Efficiency

The event-based architecture of the SNN is incredibly energy-efficient in contrast to traditional deep learning methods. The proposed approach uses 3.4 mJ per prediction -13 times less than the CNN-LSTM baseline (78.6 mJ) and has lower latency (45 ms vs 310 ms) as indicated in Table 2. This efficiency allows constant operation of wearable devices using standard coin-cell batteries.

Table 2. Computational resource requirements

Model	Energy (mJ/pred)	Latency (ms)	Memory (KB)	MCU Utilization (%)
CNN (1D)	45.2	120	286	78
CNN-LSTM	78.6	310	412	92
SNN (Proposed)	3.4	45	128	35

Figure 7 shows that 68% of the power of the SNN is used in spike propagation between layers, with only 22% of the power used in membrane potential updates. This distribution implies that additional optimization can be

achieved by sparse connectivity patterns or event-based memory access while 10% is assigned to peripheral overhead (ADC sampling, preprocessing).




Figure 7. Power consumption breakdown during SNN inference

5.4. Ablation Studies

The perform systematic ablation experiments to assess the effects of important design elements. Table 3 shows that removing adaptive threshold mechanism in equation (7) decrease 6.2% in sensitivity and a 0.03 false alarm rate per hour. Likewise, when the population coding input is substituted with direct analog values, the AUC decreases by 0.04, which is mainly because of the loss of temporal encoding.

Table 3. Ablation study of SNN components

Configuration	Sensitivity	Specificity	AUC	FA Rate (/h)
Full Model	88.4%	93.2%	0.95	0.08
w/o Adaptive Threshold	82.2%	90.1%	0.91	0.11
Analog Input Encoding	84.7%	91.5%	0.91	0.09
Fixed Learning Rate	85.3%	92.0%	0.92	0.10
Dense Connectivity (100%)	87.1%	92.8%	0.94	0.09

The low-density connectivity pattern 30% density is essential in energy efficiency without compromising accuracy. The 100% density fully connected variant has a negligible performance improvement +0.3% sensitivity and 2.8 times higher energy per prediction. This result confirms that sparse architectures can be effective and reduce the cost of computation significantly when designed properly.

The timestep resolution of 16 ms is a trade-off between time accuracy and efficiency. Figure 8 shows that the response time can be increased by 8 ms, which results in a 1.1% increase in sensitivity, but a doubling of power consumption but a 32 ms increase in response time saves 35% of power and at the same time decreases sensitivity by 4.7%, which again 16 ms in optimum operating point.




Figure 8. Impact of timestep resolution on performance and energy consumption

6. CONCLUSION

The paper presented a real-time, ultra-low-power seizure prediction system using spiking neural networks (SNNs) specifically designed for wearable edge devices. The proposed system has successfully demonstrated that neuromorphic computing can provide a viable compromise between the accuracy of prediction and energy consumption, one of the most critical issues in ongoing healthcare monitoring. The experimental results of the CHB-MIT EEG data show that the model has a high predictive accuracy, with a sensitivity of 88.4, a specificity of 93.2 and a low false alarm rate of 0.08/hour and a prediction horizon of 5-30 minutes. Furthermore, the system requires much less energy (3.4 mJ per prediction) than traditional deep learning algorithms, making it much more practical for wearable devices. These are satisfactory results, but there are some restrictions. It is also found that the system is variable from patient to patient, especially when there are atypical preictal patterns, and more powerful and comprehensive feature extraction methods are required. Moreover, the current architecture treats EEG channels separately, which could be a limitation of the architecture for detecting complex spatial relationships between brain regions. The 16 ms temporal resolution employed is effective, but may not be capable of resolving very fast neural dynamics in certain forms of epilepsy.

Moreover, the model is founded on patient-specific adaptation that is not dynamic and continuous learning, which can reduce the performance in the long-term as the conditions of patients vary. The proposed system has a broad scope of real-life deployment applications in terms of applications. It may be used to monitor systems in the home, assist in the management of pediatric epilepsy and enhance clinical decision-making in hospitals. However, it should be used ethically including privacy of data, bias in the algorithms, and psychological impacts of being predicted across the board. It will be required to ensure safe on-device processing, comprehensible system behavior, and fair performance in a wide population to ensure real-world adoption.

The next area of research should be on how to enhance the strength and flexibility of the systems. Multi-modal physiological cues, such as heart rate variability and skin conductance, may be used to increase the accuracy of predictions. This would be facilitated by incorporation of online or incremental learning processes to make sure that patient specific changes are continuously adjusted to. Moreover, the multi-scale temporal processing and spatial attention mechanisms might be taken into account to further enhance the model in the representation of the complex dynamics of EEG. The proposed approach also needs to be tested at large scale and on a diverse group of patients to guarantee the generalizability and clinical utility of the proposed approach. To sum up, this paper has shown that the neuromorphic systems using SNNs have a massive potential in energy efficient and real time seizure prediction. The suggested framework is not just a step towards the next generation of the wearable healthcare technology, but also a base to the future intelligent, low-power biomedical monitoring systems.

ACKNOWLEDGEMENTS

Thanks to the University of Information Technology and Communications (UoITC), Baghdad, Iraq.

REFERENCES

- [1] KM Fiest, KM Sauro, S Wiebe, SB Patten, CS Kwon, et al. (2017) Prevalence and incidence of epilepsy: a systematic review and meta-analysis of international studies. *Neurology*.
- [2] R Sahu, SR Dash, LA Cacha, RR Poznanski, et al. (2020) Epileptic seizure detection: a comparative study between deep and traditional machine learning techniques. *Journal of Integrative Neuroscience*.
- [3] G Bugmann (1997) Biologically plausible neural computation. *Biosystems*.
- [4] W Wei, M Zhang, J Zhang, A Belatreche, J Wu, et al. (2024) Event-driven learning for spiking neural networks. *arXiv preprint arXiv:2403.00270*.
- [5] E Covi, E Donati, X Liang, D Kappel, et al. (2021) Adaptive extreme edge computing for wearable devices. *Frontiers in Neuroscience*.
- [6] Jaber, A.G. (2026). Privacy-Preserving and Secure AI-based Stress Monitoring Federated Learning on Wearable Edge Devices Using Homomorphic Encryption. *International Journal of Research and Innovation in Applied Science (IJRIAS)* 11(2).
- [7] A Yadollahpour & M Jalilifar (2015) Seizure prediction methods: a review of the current predicting techniques. *Biomedicine And Pharmacology Journal*.
- [8] K Gadhoumi, JM Lina, F Mormann & J Gotman (2016) Seizure prediction for therapeutic devices: A review. *Journal of Neuroscience Methods*.
- [9] H Daoud & MA Bayoumi (2019) Efficient epileptic seizure prediction based on deep learning. *IEEE Transactions On Biomedical Engineering*.
- [10] S Ryu & I Joe (2021) A hybrid DenseNet-LSTM model for epileptic seizure prediction. *Applied Sciences*.
- [11] S Ghosh-Dastidar & H Adeli (2007) Improved spiking neural networks for EEG classification and epilepsy and seizure detection. *Integrated Computer-Aided Engineering*.
- [12] R Li, G Zhao, DR Muir, Y Ling, K Burelo, M Khoe, et al. (2024) Real-time sub-milliwatt epilepsy detection implemented on a spiking neural network edge inference processor. *Computers in Biology and Medicine*.
- [13] L Deng, Y Wu, Y Hu, L Liang, G Li, X Hu, et al. (2021) Comprehensive snn compression using admm optimization and activity regularization. *IEEE Transactions On Neural Networks And Learning Systems*.
- [14] A Javanshir, TT Nguyen, MAP Mahmud, et al. (2022) Advancements in algorithms and neuromorphic hardware for spiking neural networks. *Neural Computation*.
- [15] AL Goldberger, LAN Amaral, L Glass, JM Hausdorff, et al. (2000) PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *circulation*.
- [16] W Sun, Y Su, X Wu & X Wu (2020) A novel end-to-end 1D-ResCNN model to remove artifact from EEG signals. *Neurocomputing*.
- [17] Y Pan, F Dong, J Wu & Y Xu (2023) Downsampling of EEG signals for deep learning-based epilepsy detection. *IEEE Sensors Letters*.
- [18] K Belwafi, S Gannouni & H Aboalsamh (2020) An effective zeros-time windowing strategy to detect sensorimotor rhythms related to motor imagery EEG signals. *IEEE Access*.
- [19] E Orhan (2012) The leaky integrate-and-fire neuron model. *no*.
- [20] EO Neftci, H Mostafa & F Zenke (2019) Surrogate gradient learning in spiking neural networks. *Ieee Signal Processing Magazine*.
- [21] W Ponghiran & K Roy (2022) Spiking neural networks with improved inherent recurrence dynamics for sequential learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [22] S Engelborghs, R D'hooge & PP De Deyn (2000) Pathophysiology of epilepsy. *Acta neurologica belgica*.
- [23] M Le Van Quyen, J Soss, V Navarro, R Robertson, et al. (2005) Preictal state identification by synchronization changes in long-term intracranial EEG recordings. *Clinical Neurophysiology*.
- [24] JJ Newson & TC Thiagarajan (2019) EEG frequency bands in psychiatric disorders: a review of resting state studies. *Frontiers in human neuroscience*.
- [25] J Wu, Y Wang, Z Li, L Lu & Q Li (2024) A Review of Computing with Spiking Neural Networks. *search.ebscohost.com*.

- [26] S Lian, J Shen, Z Wang & H Tang (2024) IM-LIF: Improved neuronal dynamics with attention mechanism for direct training deep spiking neural network. *IEEE Transactions On Circuits And Systems I: Regular Papers*.
- [27] L Billeci, D Marino, L Insana, G Vatti & M Varanini (2018) Patient-specific seizure prediction based on heart rate variability and recurrence quantification analysis. *PloS one*.
- [28] J Prasanna, MSP Subathra, MA Mohammed, et al. (2021) Automated epileptic seizure detection in pediatric subjects of CHB-MIT EEG database—a survey. *Journal of Personalized Medicine*.
- [29] E Gordon, N Cooper, C Rennie, et al. (2005) Integrative neuroscience: the role of a standardized database. *Clinical EEG And Neuroscience*.
- [30] H Aldawsari, S Al-Ahmadi & F Muhammad (2023) Optimizing 1D-CNN-based emotion recognition process through channel and feature selection from EEG signals. *Diagnostics*.
- [31] M Shahbazi & H Aghajan (2018) A generalizable model for seizure prediction based on deep learning using CNN-LSTM architecture. In *2018 IEEE Global Conference On Signal And Information Processing*.
- [32] M Iftikhar, SA Khan & A Hassan (2018) A survey of deep learning and traditional approaches for EEG signal processing and classification. In *2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference*.
- [33] D Selvathi & H Selvaraj (2017) FPGA implementation for epileptic seizure detection using amplitude and frequency analysis of EEG signals. In *2017 25th International Conference on Electronics, Circuits and Systems (ICECS)*.
- [34] G Segal, N Keidar, RM Lotan, Y Romano, et al. (2023) Utilizing risk-controlling prediction calibration to reduce false alarm rates in epileptic seizure prediction. *Frontiers in Neuroscience*.
- [35] M Winterhalder, T Maiwald, HU Voss, et al. (2003) The seizure prediction characteristic: a general framework to assess and compare seizure prediction methods. *Epilepsy & Behavior*.
- [36] GA Pradipta, R Wardoyo, A Musdholifah, et al. (2021) SMOTE for handling imbalanced data problem: A review. In *International Conference on Informatics and Analytics*.
- [37] DL Manna, A Vicente-Sola, P Kirkland, TJ Bihl, et al. (2023) Frameworks for SNNs: a review of data science-oriented software and an expansion of spyketch. *Applications of Neural Networks*.
- [38] A Schulman, T Schmid, P Dutta & N Spring (2011) Phone power monitoring with battor. In *Proc. ACM MobiCom*.
- [39] R David, J Duke, A Jain, et al. (2021) Tensorflow lite micro: Embedded machine learning for tinyml systems. In *Proceedings of Machine Learning and Systems*.

BIOGRAPHIES OF AUTHORS



ABDULLAH GHANIM JABER received the B.Sc. degree in computer science from Infrastructure University Kuala Lumpur (IUKL), Malaysia, and the M.Sc. degree from the Faculty of Information Science and Technology (FTSM), Universiti Kebangsaan Malaysia (UKM), in 2021. He is currently pursuing the Ph.D. degree in Cybersecurity with the Research Centre for Cyber Security, Faculty of Information Science and Technology, UKM, Bangi, Selangor, Malaysia. Since 2022, he has been a Lecturer at University of Information Technology and Communications (UOITC), Baghdad, Iraq. He has published multiple articles in reputable international journals and conferences. His research interests include cryptography, secure data processing, encryption, cybersecurity frameworks, and secure computation. He can be contacted at email: abdullah.ghanim@uoitc.edu.iq.