Aurevia:

Enabling Technologies, Predictive AI Architecture, and Clinical Translation Framework for a Personalized Wearable Seizure Prediction Device

Daniel KALU
· CORTEX LABs
June 2025

Abstract

Epilepsy affects 50 million people worldwide and remains inadequately managed for approximately one-third of patients due to drug resistance. A critical gap exists between the clinical necessity for advanced seizure warning and the technological capability to provide it: while FDA-cleared wearable devices can detect generalized tonic-clonic seizures after onset, no commercially available system reliably predicts seizures before they occur in a continuous, wearable, patient-individualized form. Aurevia is designed to address this gap — a predictive seizure-monitoring device that combines continuous scalp electroencephalography (EEG), multimodal physiological sensing, and edge-deployed personalized artificial intelligence (AI) to identify pre-ictal neural signatures and deliver timely alerts to users and caregivers. This paper presents the complete enabling technology and analytical framework for Aurevia, synthesizing the neurophysiological basis of EEG-based seizure prediction, the state of the art in machine- and deep-learning architectures for pre-ictal classification, the engineering requirements for real-time wearable deployment, and the data infrastructure required for personalized, continuously adaptive model performance. We introduce a formal mathematical framework that covers pre-ictal state modeling, a composite personalized prediction objective, a patient-adaptive confidence threshold, and an edge-deployment efficiency formulation. We analyze the comparative performance of six AI model classes — from classical ensemble methods to Spiking Neural Networks (SNNs) and Micro Tree-based Neural Additive Models (MT-NAMs) — against the specific constraints of a wearable, real-time, energy-constrained device. Federated Learning is identified as a strategic architecture for privacy-preserving personalization across a diverse user base. We survey the landscape of public EEG datasets and their limitations for pre-ictal model development, and argue for a proprietary longitudinal data strategy as a core competitive and scientific requirement. The paper includes an overview of the prototype hardware architecture and reserved figure spaces for device prototype images to be populated as development progresses. The Aurevia framework represents a rigorous, clinically grounded foundation for developing the first broadly deployable, truly predictive, personalized wearable epilepsy device.

Keywords: *epilepsy · seizure prediction · wearable EEG · deep learning · spiking neural networks · federated learning · edge computing · personalised medicine · pre-ictal detection · TinyML · Aurevia*

1 Introduction

1.1 The Unmet Need: Prediction vs Detection

Epilepsy is characterized by recurrent, unprovoked seizures arising from pathological neuronal hypersynchrony, affecting approximately 50 million individuals globally across all age groups [2,4]. Its clinical burden extends far beyond seizure events themselves: patients face elevated rates of cognitive decline, depression, social isolation, traumatic injury from unwarned falls, and significantly elevated mortality risk — including sudden unexpected death in epilepsy (SUDEP) [4]. Despite an expanding pharmacological armamentarium, approximately 30% of patients develop drug-resistant epilepsy (DRE), achieving neither seizure freedom nor an acceptable quality of life through medication [3].

Existing wearable technologies for epilepsy monitoring — including FDA-cleared devices such as EpiWatch (Apple Watch platform) and EpiMonitor (Empatica) — primarily address seizure detection: identifying that a seizure has occurred or is occurring based on motor manifestations and physiological changes. This is clinically valuable for post-event logging, caregiver alerting, and safety monitoring, but it does not deliver on the transformative promise of prediction: identifying that a seizure is likely to occur before it does, enabling the patient and caregivers to take proactive action [35,36].

Aurevia is designed to close this gap. Its core value proposition is not detection but prediction — identifying the pre-ictal neural signatures that precede seizure onset by minutes, delivering actionable, personalized alerts that transform the patient's relationship with their condition from reactive to proactive. This requires a fundamentally different technological stack from that of existing detection devices: more sophisticated AI architectures, deeper personalization, and a rigorous approach to the noisy, variable, high-dimensional challenge of continuous EEG prediction in a wearable form factor.

1.2 Why This is Hard: The Prediction Challenge

Seizure prediction from wearable EEG is among the most technically demanding applications of biosignal AI. The pre-ictal state — the period of altered neural dynamics preceding a clinically evident seizure — is subtle, variable, patient-specific, and evolves over a timescale of minutes to hours. It does not produce a discrete, easily identifiable signal; rather, it manifests as gradual changes in the spectral composition, non-linear dynamics, and functional connectivity of EEG across multiple frequency bands — changes that differ not only between patients but within the same patient over time as the disease evolves and medications are adjusted [3,8].

Compounding this, scalp EEG — the only viable modality for a non-invasive consumer device — has inherently limited spatial resolution and is highly susceptible to motion, electromyographic, and environmental artifacts in real-world use [7]. The class imbalance problem is severe: seizure events constitute less than 1% of continuously recorded EEG data, making it methodologically challenging to train robust predictive models on existing public datasets [21]. The personalization requirement is non-negotiable: a model trained on population-level data will not achieve clinically meaningful prediction accuracy for individuals whose pre-ictal signatures deviate from the population mean, which is most patients [23].

1.3 Scope and Contributions of This Paper

This paper makes the following contributions:

- A comprehensive synthesis of the neurophysiological basis of EEG frequency band dynamics and their relevance to pre-ictal state characterization.
- A structured comparative analysis of AI/ML architectures for seizure prediction, evaluated against the specific constraints of real-time wearable deployment.
- A formal mathematical framework for pre-ictal state modeling, personalized prediction scoring, adaptive confidence thresholding, and edge efficiency optimization.
- A hardware and system architecture overview for Aurevia, including reserved prototype image sections to be populated as device development progresses.
- A federated learning strategy for privacy-preserving personalization at scale.
- A critical analysis of public EEG dataset limitations and a proprietary data strategy for robust pre-ictal model development.
- A strategic development roadmap aligned with regulatory and clinical translation requirements.

---

2 Neurophysiological Foundations

2.1 EEG Principles and Clinical Utility

Electroencephalography measures the summed electrical activity of cortical neuronal ensembles via scalp electrodes, generating a continuous multi-channel time-series record of brain electrical dynamics [1]. The resultant waveforms are characterized by their frequency content, morphology, spatial distribution, amplitude, and synchrony across channels, collectively providing a functional map of brain state that is sensitive to both physiological and pathological changes [16].

In epilepsy, EEG provides diagnostic information about seizure type (focal vs. generalized), origin (localization to a specific cortical region), and interictal characteristics (the pattern of abnormal discharges observed between seizures) [4,5]. The International 10-20 electrode placement system provides a standardized spatial framework for electrode positioning across 16–25 clinically used channels, with high-density configurations extending to 256 channels for research applications [17]. For wearable devices, the practical constraints of user comfort and miniaturization drive the electrode count down to 2–8 channels, typically positioned at temporal, frontal, or behind-the-ear (BTE) locations — a reduction that necessitates sophisticated AI to extract equivalent predictive information from sparser data [7,20].

2.2 EEG Frequency Bands and Pre-ictal Signatures

Seizure-related neural dynamics manifest across the full EEG frequency spectrum, from infra-slow oscillations below 0.5 Hz to high-frequency oscillations (HFOs) above 200 Hz [15]. Each canonical frequency band carries distinct physiological and pathological information:

Delta (0.5–4 Hz): Dominant in deep sleep and infancy; pathological in waking adults, where Temporal Intermittent Rhythmic Delta Activity (TIRDA) is specifically associated with temporal lobe epilepsy [15].

Theta (4–7 Hz): Physiological in drowsiness and childhood; sustained waking theta in adults suggests focal cortical dysfunction. Frontal theta enhancement is associated with heightened emotional arousal — relevant to seizure provocation states [15].

Alpha (8–12 Hz): Dominant rhythm of the relaxed, eyes-closed adult brain; attenuated with alerting or eye opening. Alpha band disruption is an early marker of cortical state change [16].

Beta (13–30 Hz): Prominent in alert, anxious, and task-engaged states; reduced over areas of cortical damage. Changes in beta activity may precede focal seizure onset [15].

High-Frequency Oscillations (HFOs, >80 Hz): Ripples (80–250 Hz) and fast ripples (250–500 Hz) at the seizure focus are established biomarkers of epileptogenicity. Their detection requires high-frequency EEG sampling and advanced signal processing [15].

A critical finding for Aurevia's design is that the frequency band in which prominent pre-ictal changes are most apparent varies across patients [8]. This is not a design failure — it is a biological reality that mandates a personalized, broad-spectrum feature extraction approach rather than a fixed-frequency-band prediction strategy.

2.3 The Pre-ictal State: What We Are Trying to Detect

The pre-ictal state is not a single, discrete phenomenon but a probabilistic gradation from interictal baseline toward ictal discharge, mediated by shifts in neuronal excitability, network synchrony, and metabolic state. Pre-ictal changes observed in EEG include: increases in signal energy across multiple bands; changes in non-linear dynamics (reduction in signal entropy, increase in Lyapunov exponent magnitude); shifts in inter-channel coherence and phase-locking; and the emergence of high-frequency oscillatory bursts at the seizure focus [3,15]. These changes occur on a timescale of minutes to hours before clinical seizure onset, providing the biological window that a prediction system must exploit. The challenge is that these changes are embedded in a high-noise, high-variability signal stream, and their discriminative power relative to non-pre-ictal baseline states varies substantially across individuals.

---

3 Methodology

3.1 Literature Review Protocol

A systematic narrative review was conducted using PubMed, Scopus, IEEE Xplore, arXiv, and medRxiv preprint servers, with searches executed through May 2025. Search terms combined variants of: (epilepsy OR seizure) AND (prediction OR forecasting OR pre-ictal) AND (EEG OR electroencephalography) AND ("machine learning" OR "deep learning" OR "neural network" OR "spiking neural network" OR "federated learning" OR "edge computing" OR "wearable"). Additional targeted searches covered specific model architectures (MT-NAM, ARNN, TinyML, SNNs), commercial device landscape (EpiWatch, EpiMonitor),

and public datasets (CHB-MIT, Kaggle Epilepsy, Neurovista). Primary empirical studies, systematic reviews, and authoritative technical reports were included; case reports were excluded.

3.2 Framework Scope

The mathematical formulations in Section 4 are developed de novo as analytical scaffolding for Aurevia's system design and future clinical evaluation rather than as empirically validated models. Performance metrics cited throughout are extracted verbatim from primary sources with their original sample characteristics; cross-study comparisons are made qualitatively, given the heterogeneity in the dataset, evaluation protocols, and patient populations. Clinical recommendations are not made; the paper is a technical framework document.

---

4 Mathematical Framework

We introduce four formal constructs that together define the computational architecture for Aurevia's personalized prediction system.

4.1 Pre-ictal State Probability Model

Let x(t) ∈ R^(C × T) denote the multi-channel EEG observation window of C channels and T time samples centered at time t. The pre-ictal state probability is modeled as:

$$P_i(t) = f_{\theta i}(\varphi(x(t))) \quad \text{(Eq. 1)}$$

where φ: R^(C×T) → R^d is a feature extraction function mapping the raw EEG window to a d-dimensional feature vector (encompassing time-domain, spectral, wavelet, non-linear, and cross-channel features as catalogued in Table 1), f_θi is a patient-specific predictive model parameterised by θ_i — learned from individual i's longitudinal data — and P_i(t) ∈ [0,1] is the instantaneous probability of being in the pre-ictal state. For hybrid deep learning architectures (Section 5), φ(·) and f_θi(·) are jointly learned end-to-end, with φ(·) implemented as convolutional feature extraction layers and f_θi(·) as recurrent or attention-based classification layers.

4.2 Personalized Prediction Objective

The patient-specific model parameters θ_i are optimized to balance sensitivity (capturing true pre-ictal epochs) against specificity (rejecting interictal and artifact epochs), subject to an edge-efficiency constraint:

$$\theta^*_i = \operatorname{argmin}_{\theta} [L_{CE}(\theta; D_i) + \lambda \cdot L_{FPR}(\theta; D_i) + \mu \cdot \Omega(\theta)] \quad \text{(Eq. 2)}$$

where L_CE is the cross-entropy loss on labeled pre-ictal/interictal windows from patient i's dataset D_i, L_FPR is an auxiliary false-positive-rate penalty (weighted by λ, reflecting the clinical cost of alarm fatigue — a primary driver of device abandonment), and Ω(θ) is a regularisation term (weighted by μ) that penalizes

model parameter count and computational complexity to enforce edge deployability. The false-positive penalty term is critical: Aurevia's 'peace of mind' value proposition is eroded not only by missed predictions but by excessive false alarms that erode user trust.

4.3 Adaptive Confidence Threshold

A real-time alert is triggered when the instantaneous pre-ictal probability exceeds a patient-specific and time-adaptive threshold $\tau_i(t)$:

$$\text{Alert}(t) = 1 \quad \text{iff} \quad P_i(t) \geq \tau_i(t) \tag{Eq. 3}$$

The threshold $\tau_i(t)$ is updated online using a lightweight exponential moving average of recent model performance:

$$\tau_i(t) = (1-\eta) \cdot \tau_i(t-1) + \eta \cdot \tau_{target,i} \tag{Eq. 4}$$

where η is a smoothing coefficient, and $\tau_{target,i}$ is a patient-specific target threshold updated by the Test-Time Template Adjuster (T3A) mechanism based on the model's recent sensitivity-specificity balance [24]. This formulation allows τ to respond to short-term changes in baseline brain state (e.g., medication adjustments, sleep deprivation, circadian variation) without requiring full model retraining.

4.4 Edge Deployment Efficiency

For a wearable device with constrained battery capacity B (mAh) and processing budget P (FLOPS/s), the effective monitoring duration M (hours) under continuous inference is:

$$M = B \cdot V_{ops} / \left(I_{base} + I_{model} \cdot FLOPs_{per\_window} / (T_{window} \cdot P) \right) \tag{Eq. 5}$$

where V_{ops} is the operating voltage, I_{base} is the baseline device current draw (EEG amplifier, radio, display), I_{model} is the current per FLOP of the inference engine, $FLOPs_{per\_window}$ is the computational cost of evaluating the prediction model on one EEG window, and T_{window} is the window duration in seconds. This formulation makes explicit the trade-off between model complexity ($FLOPs_{per\_window}$) and monitoring duration, directly motivating the selection of SNN or TinyML architectures that achieve 98–100x reductions in FLOPs relative to dense deep learning equivalents [11,24].

---

5 AI/ML Architectures for Seizure Prediction

5.1 Feature Engineering Pipeline

Effective seizure prediction from scalp EEG requires a multi-category feature-extraction pipeline that captures the full dimensionality of preictal neural dynamics. Six feature categories are identified in the literature as collectively discriminative [13,18]:

Feature Category	Example Features	Relevance to Pre-ictal Detection
Time-Domain	Mean amplitude, std deviation, skewness, kurtosis, zero-crossing rate, RMS, peak-to-peak amplitude, line length	Capture statistical properties and amplitude dynamics of the EEG signal over time
Frequency-Domain	Delta/Theta/Alpha/Beta/Gamma band power, continuous wavelet transform (CWT)	Highlight spectral content; essential for identifying seizure-related oscillatory signatures
Wavelet-based	Wavelet entropy, wavelet energy, DWT coefficients	Capture transient, non-stationary events across temporal and frequency scales simultaneously
Non-linear	Sample entropy, approximate entropy, Lyapunov exponent, Hurst exponent, Higuchi fractal dimension, Lempel-Ziv complexity	Quantify signal complexity, predictability, and chaotic dynamics characteristic of pre-ictal states
Seizure-Specific	Seizure duration, pre-seizure pattern, post-seizure recovery, interictal spike rate, seizure intensity index	Tailored features capturing the onset, progression, and recovery phases of seizure events
Cross-Channel	Cross-correlation between channels, Hjorth mobility, Hjorth complexity	Measure inter-electrode dependencies; essential for detecting network-level propagation patterns

Table 1. EEG feature categories for pre-ictal detection, with representative features and clinical rationale.

Dimensionality reduction via Principal Component Analysis (PCA) — which has been shown to compress 178 features to 53 while retaining 99.05% of the variance [23] — is applied downstream to reduce redundancy and computational load. Personalized channel selection, identifying the 2–4 most informative electrodes for each individual via variance-based ranking, further reduces input dimensionality while improving signal-to-noise ratio; a 2-channel personalized system has been shown to match the state-of-the-art performance of a 4-fixed-channel system with a 30% reduction in false-positive rate [20].

5.2 Model Architecture Comparison

Model Class	Architectures	Core Strengths	Performance	Relevance to Aurevia
Traditional ML	RF, SVM, Extra Trees, Gradient Boosting, DT	Interpretable; lower inference cost; good on engineered features	ET: 99.29%; RF: 96.52%; SVM: 95.28%	Suitable for benchmarking and lightweight on-device inference if features are pre-extracted; limited adaptability to raw dynamic signals
CNNs (1D/2D)	1D-CNN, 2D CNN, spectral CNN	Automated feature extraction; robust spatial/spectral pattern recognition	97.65% (validation); 94.29% (ECG-based	Core component for brainwave spectral pattern recognition; efficient with spectrogram representations

Model Class	Architectures	Core Strengths	Performance	Relevance to Aurevia
			seizure prediction)	
RNNs / LSTMs	LSTM, GRU	Strong temporal modeling; learns long-range dependencies from sequential EEG	LSTM: 90.17%; GRU: 97.59%	Essential for modeling pre-ictal temporal evolution; raw EEG can be fed directly without heavy preprocessing
Hybrid DL	CNN-LSTM, Wavelet+CNN+Attention, ARNN, HyEpiSeiD (CNN+GRU)	Combines spatial/spectral (CNN) and temporal (RNN) strengths; attention improves focus on critical windows	99.83% (Wavelet+CNN+Attention); 99.01% (HyEpiSeiD)	Optimal predictive core for Aurevia; captures both subtle spectral shifts and temporal seizure dynamics
MT-NAM / NAM	Neural Additive Model, Micro Tree-based NAM with T3A	High efficiency (~100x inference speedup); interpretable; adaptive via T3A test-time updates	NAM: 85% sens, 96% spec; MT-NAM: comparable after T3A	Critical for edge deployment; T3A enables backpropagation-free on-device adaptation to evolving seizure dynamics
SNNs	Spiking Conv-LSTM, CSNN, Liquid-Dendrite SNN	Ultra-low power (100x-1000x savings); neuromorphic hardware compatible; on-device training	98.58% complexity reduction; 1.28 µJ/classification; 3.1 ms inference on RPi5	Enables continuous, battery-efficient wearable monitoring; ideal for real-time edge inference on neuromorphic chips

Table 2. Comparison of AI/ML model classes for seizure prediction against Aurevia's specific requirements.

5.3 Recommended Architecture: Hybrid Edge-Optimized DL

Based on the analysis in Table 2, Aurevia's optimal predictive architecture is a tiered hybrid system:

Edge tier (on-device): A TinyML-optimized MT-NAM or Liquid-Dendrite SNN provides ultra-low power, sub-5ms inference for continuous real-time prediction. The T3A mechanism enables backpropagation-free on-device adaptation to evolving seizure dynamics without requiring cloud connectivity, preserving both responsiveness and privacy [24,25].

Cloud tier (periodic): A full CNN-LSTM or ARNN model hosted on cloud infrastructure receives periodic data uploads (with user consent via Federated Learning) to perform comprehensive model retraining, circadian pattern analysis, long-term trend identification, and the generation of updated edge model weights, which are pushed back to the device [6,30,31].

This hybrid architecture ensures that the real-time, battery-constrained primary alert function operates independently of connectivity, while the cloud layer continuously improves predictive accuracy over the device lifetime.

6 Aurevia System Architecture and Hardware Prototype

6.1 System Architecture Overview

Aurevia's hardware system comprises four integrated functional layers: (i) a bio-signal acquisition layer (EEG amplifier, multi-modal sensors); (ii) a signal processing and feature extraction layer (dedicated DSP or microcontroller); (iii) an on-device AI inference layer (neuromorphic chip or TinyML microprocessor); and (iv) a communication and alert layer (BLE/Wi-Fi module, vibration motor, LED alert, companion app). The EEG acquisition subsystem targets a minimum 4-channel configuration with personalized channel selection from a predefined 8-channel layout, sampled at 256 Hz with 24-bit ADC resolution. Multi-modal sensors — including 3-axis accelerometer, photoplethysmography (heart rate/SpO2), and single-lead ECG — provide corroborating physiological signals for artifact rejection and multi-modal seizure characterization [33,34].

6.2 Prototype Hardware — Image Section

The following figure spaces are reserved for prototype device images to be inserted as hardware development progresses. These sections will document the physical evolution of the Aurevia wearable from initial breadboard prototype through enclosure design to clinical-grade pre-production unit.

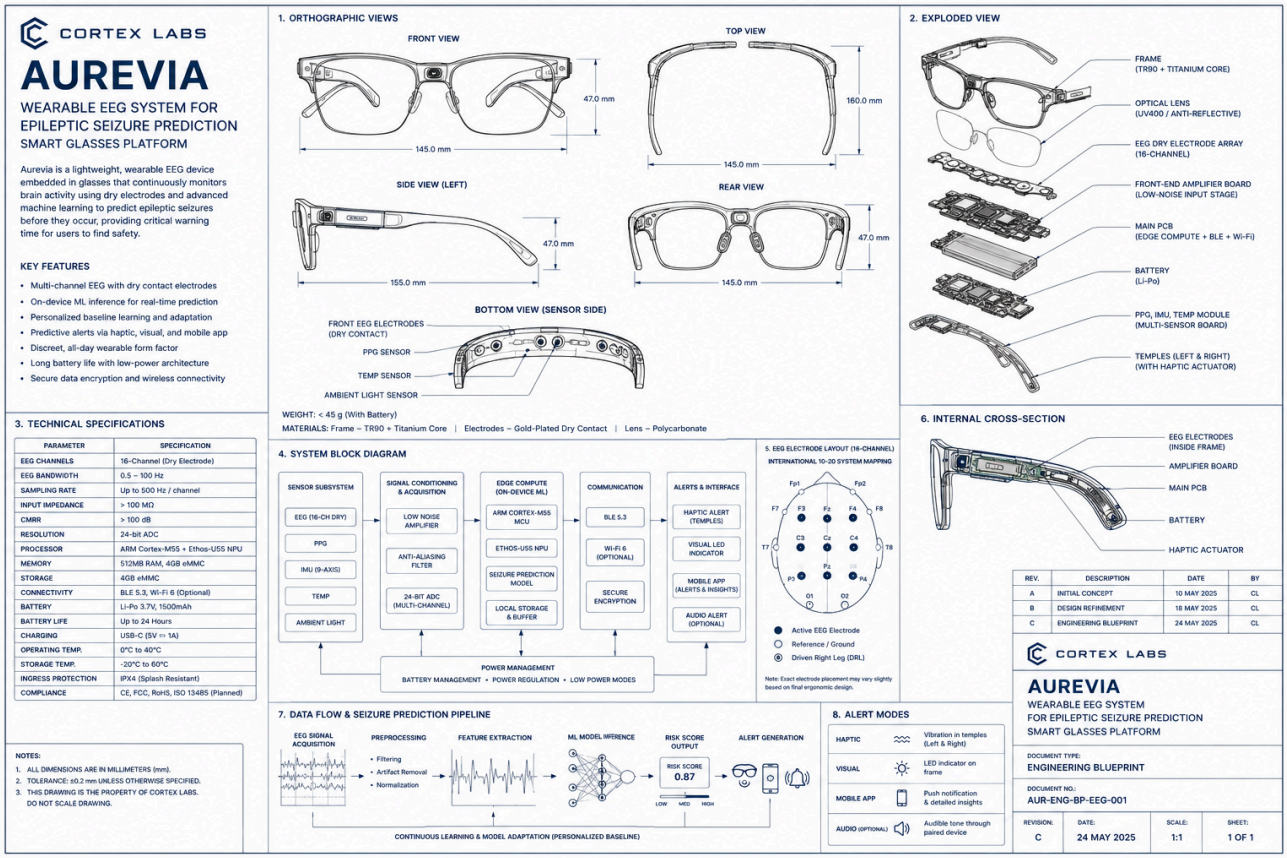

Figure 6.1. Aurevia prototype v1 — full device assembly, Electrode array, and wearable headband configuration.

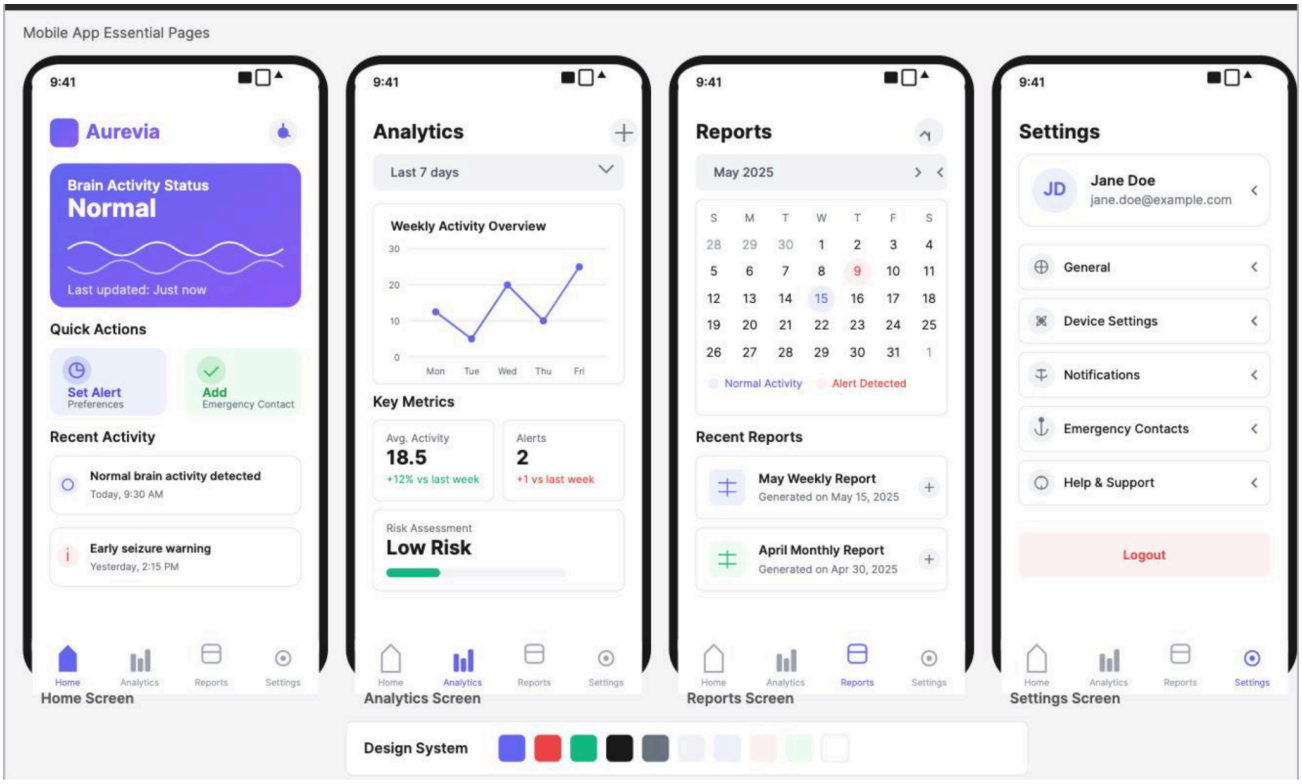

Figure 6.2. Aurevia companion app — dashboard and alert interface.

Note: Blueprint-level design drawings and detailed schematic documentation are retained as proprietary intellectual property and are not included in this publication. The prototype images above are intended to document physical hardware evidence of system development for research and funding purposes only.

6.3 Signal Processing and Inference Pipeline

The on-device signal processing pipeline operates as follows: raw multi-channel EEG is acquired at 256 Hz and segmented into 1-second non-overlapping windows with 50% overlap for transition continuity. Each window undergoes artifact detection (threshold-based checks of amplitude and gradient) and wavelet-based noise suppression [12]. The 75-dimensional feature vector (per Table 1) is extracted per window, PCA-compressed to 53 dimensions, and passed to the on-device MT-NAM or SNN inference engine. The resulting $P_i(t)$ probability estimate is compared against the adaptive threshold $\tau_i(t)$ (Eq. 3–4). If $P_i(t) \geq \tau_i(t)$ for three consecutive windows (a 3-second confirmation window to reduce single-window false positives), the alert pipeline is triggered: a vibrotactile notification on-device, a push notification via BLE to the companion app, and an optional SMS/call to registered caregivers.

---

7 Personalization and Federated Adaptive Learning

7.1 The Inter- and Intra-Patient Variability Problem

Epilepsy is among the most inter-individual variable neurological conditions in its EEG manifestations: the frequency band, spatial distribution, temporal dynamics, and pre-ictal lead time of seizures differ substantially between patients and within the same patient across different seizure clusters, medication changes, and disease progression stages [8,23]. A static, population-trained model will fail to generalize to a large proportion of individual users — either missing their specific pre-ictal signature (false negatives, undermining the safety value) or generating high false alarm rates from non-specific pattern matches (false positives, undermining the peace-of-mind value). Patient-specific, individualized models — trained and evaluated exclusively on data from the same patient — have demonstrated significantly superior performance, achieving a 73.6% F1 score for seizure cluster prediction compared with population-level baselines [8].

7.2 On-Device Adaptation: T3A Mechanism

The Test-Time Template Adjuster (T3A) is a backpropagation-free on-device adaptation mechanism developed for MT-NAM that dynamically updates the model's classification templates during inference, improving sensitivity to individual patients' current ictal patterns without requiring a full retraining cycle [24]. This makes it suitable for lightweight, battery-constrained edge devices. Combined with the adaptive threshold formulation in Eq. 4, T3A provides a two-layer personalization mechanism: threshold adaptation for rapid response to short-term state changes, and template adaptation for medium-term model refinement as new pre-ictal examples are observed.

7.3 Federated Learning for Privacy-Preserving Scale

Federated Learning (FL) enables model improvement across Aurevia's user base without centralizing sensitive brainwave data. In the FL paradigm, each device trains a local model update using its own patient-specific data, transmits only the gradient update (not raw EEG) to a central aggregation server, and receives the aggregated global model update in return [30]. This provides two complementary benefits for Aurevia: (i) global model generalization improves as the federated dataset diversifies across the user base, improving cold-start performance for new users; and (ii) personalized local fine-tuning can occur on each device using patient-specific data, without that data ever leaving the device. The Kaggle Epilepsy Dataset [18] is specifically structured to support federated deep learning research, providing a publicly available validation benchmark for Aurevia's FL architecture.

---

8 Real-Time Wearable Deployment

8.1 Computational and Power Constraints

The fundamental tension in wearable AI device design is between model sophistication — which drives prediction accuracy — and computational and power efficiency — which determine battery life and form-factor feasibility. Traditional dense deep learning models (e.g., full CNN-LSTM) require gigaflop-scale inference operations that are untenable on milliamp-hour battery budgets for continuous monitoring. The solution space comprises three complementary approaches [9,11,24,25]:

- **SNNs:** Spiking Convolutional LSTM variants achieve 100x-1000x power reduction relative to equivalent dense architectures. The Liquid-Dendrite SNN consumes 1.28 μJ per classification,

processes 1 second of EEG in 3.1 ms on a Raspberry Pi 5 (a proxy for comparable embedded SoC), and requires only 535 KB memory with 130K trainable parameters — within the budget of current neuromorphic microcontrollers [11,25].

- **TinyML:** TinyML-optimized models (e.g., quantized 1D-CNN) have demonstrated 98–99% seizure detection accuracy with 1.4 KB RAM and 17.3 KB FLASH usage and 4 ms execution time — fully compatible with ARM Cortex-M4 class microcontrollers [9].
- **MT-NAM + T3A:** The Micro Tree-based NAM achieves approximately 100x inference speedup relative to full NAM with comparable sensitivity (recoverable with T3A), providing an interpretable, computationally efficient model with intrinsic adaptability [24].

8.2 Multi-Modal Sensing for Robustness

While EEG is Aurevia's primary prediction modality, multi-modal sensor fusion provides significant complementary value. Convolutional Neural Networks applied to single-lead ECG data have demonstrated 94.29% seizure prediction accuracy [27], suggesting cardiac autonomic signatures are partially predictive. Accelerometry-based motor activity monitoring enables post hoc seizure-type classification (tonic-clonic vs. non-motor) and provides an independent channel for artifact rejection: movement artifacts in EEG that coincide with the absence of accelerometry-detectable motor activity can be confidently suppressed [33,34]. Heart rate and SpO2 provide autonomic indicators useful for detecting the post-ictal state and assessing caregiver severity.

8.3 Challenges and Engineering Solutions

Challenge	Impact on Aurevia	Technical Solution(s)
Computational efficiency	Prevents continuous complex monitoring on miniaturized hardware	SNNs; TinyML (98-99% accuracy, 4ms exec, 1.4KB RAM); MT-NAM (100x inference speedup); edge computing for local processing
Power consumption	Limits device uptime; undermines the continuous monitoring promise	SNNs (100x-1000x power savings); energy-efficient dLIF neuron models; avoid STFT/FFT transforms; Liquid-Dendrite SNN (1.28 µJ/classification)
Data volume and latency	Large EEG streams delay alert delivery; breaks real-time promise	Edge-first architecture; SNN inference in 3.1 ms (RPi5); optimised 1-second non-overlapping windowing; cloud only for ambiguous cases
Data privacy and security	Brainwave data is maximally sensitive; centralization is a liability	Federated Learning (model updates only, no raw data transfer); on-device TinyML/SNN inference minimizes external transmission
User comfort and wearability	Electrode bulk/discomfort leads to non-adherence; defeats continuous monitoring	Reduced electrode count (2-channel, BTE configurations); personalised channel selection; compact lightweight form factor
Inter- and intra-patient variability	Static models fail for individual users; seizure patterns evolve	Personalised model adaptation; T3A test-time updates (backpropagation-free); individualized ML; federated learning for diverse data
Real-world artifacts	Motion, EMG, and environmental noise generate false positives	Wavelet-based preprocessing; multi-modal sensor corroboration (ECG, accelerometry); models trained on noisy real-world data

Table 3. Wearable deployment challenges and corresponding technical solutions for Aurevia.

---

9 Data Strategy and EEG Dataset Landscape

9.1 Public Dataset Overview and Limitations

Dataset	Signal / Scale	Sampling / Channels	Labels / Annotations	Key Limitations for Aurevia
CHB-MIT Scalp EEG	Scalp EEG, 22 subjects	256 Hz, 23 channels	198 annotated seizures; continuous long-term; 10-20 system	Seizure:non-seizure ratio 1:344; gaps between files; many studies ignore non-seizure data — biases model generalization
Kaggle Epilepsy Dataset	EEG-derived features, 289,010 records	75 features (6 categories)	Multi-class: Normal, Pre-Seizure, Seizure, Post-Seizure; Generalized vs Focal labels	Designed for federated DL; rich feature set; data source heterogeneity may introduce inconsistencies
University of Bonn	Mixed scalp EEG + iEEG, 500 subjects	Varies	5 brain-state categories, including seizure, healthy (eyes open/closed)	Non-continuous random segments; generalized (not patient-specific); limited for true prediction tasks
Neurovista Ictal	iEEG, 15 patients	16 electrodes (bipolar)	Seizure events + 60s pre-ictal windows	Limited pre-ictal window duration; invasive data source; small cohort
Siena Scalp EEG	Scalp EEG, multiple subjects	Varies	Short-term continuous multi-session recordings	Segmented sessions may have discontinuities; limited for long-term circadian pattern analysis

Table 4. Key public EEG datasets for seizure prediction research — characteristics and limitations for Aurevia development.

The common limitations of publicly available EEG datasets — severe class imbalance (seizure:non-seizure ratio up to 1:344 in CHB-MIT [21]), absence of standardized pre-ictal annotations, lack of patient-specific longitudinal continuity, and heterogeneous acquisition protocols [38] — collectively preclude their use as the sole training data source for Aurevia's personalized predictive models. They are valuable for initial model prototyping, architecture benchmarking, and feature engineering validation, but cannot substitute for longitudinal, patient-specific, real-world data.

9.2 Proprietary Data Collection Strategy

Aurevia's long-term predictive performance requires a proprietary data infrastructure built around the device itself. The following strategy is proposed:

1. Collect continuous multi-channel EEG and multi-modal physiological data from consenting users through the Aurevia device, with passive background data acquisition augmented by user-triggered event logging (seizure confirmation, false alarm dismissal).

2. Implement structured annotation of the pre-ictal, ictal, and post-ictal phases using clinician-verified seizure logs and the companion app's seizure diary functionality.
3. Address class imbalance through: (a) window overlap oversampling of pre-ictal epochs; (b) generative augmentation using GAN-synthesized pre-ictal EEG segments to supplement rare positive examples [23]; and (c) cost-sensitive loss weighting in model training (the $\lambda \cdot L\_FPR$ term in Eq. 2).
4. Apply Federated Learning to aggregate model updates across the Aurevia user base without centralizing raw EEG, building a continuously improving global model while maintaining data sovereignty on each device.
5. Adopt NWB-compatible data formatting internally to ensure compatibility with research collaboration partners and potential future dataset contributions to the scientific community.

---

10 Competitive Landscape and Differentiation

10.1 Existing Devices

EpiWatch (FDA 510(k)-cleared, Apple Watch platform, Johns Hopkins) detects tonic-clonic seizures using motion and heart rate data from the Apple Watch accelerometer and optical heart rate sensor [35]. Its prediction capability is limited to the minutes immediately before a tonic-clonic event based on heart rate changes — not the multi-minute or multi-hour pre-ictal window that EEG-based prediction can potentially access. EpiMonitor (Empatica, FDA-cleared) achieves 98% detection accuracy for generalized tonic-clonic seizures using wrist-worn multi-modal sensing, with 7-day battery life and automated caregiver alerting [36]. Again, the emphasis is post-onset detection rather than pre-ictal prediction.

US Patent 10,596,377 B2 describes a system for seizure detection, prediction, and prevention via neurostimulation, employing two sequential ML models — one for prediction, one for neuromodulation parameter optimization — with adaptive self-learning based on stimulation outcomes [37]. This patent demonstrates active interest in ML-driven seizure prediction within the patent landscape and establishes prior art for adaptive ML approaches, while fundamentally differing from Aurevia in its invasive neurostimulation component.

10.2 Aurevia's Differentiation

Aurevia's competitive differentiation rests on four pillars: (i) genuine pre-ictal prediction (not post-onset detection) from continuous EEG; (ii) personalised, continuously adaptive AI that learns each individual user's seizure signature; (iii) non-invasive scalp EEG in a comfortable wearable form factor designed for daily use; and (iv) a privacy-preserving federated learning architecture that improves model performance at population scale without compromising individual data sovereignty. No currently FDA-cleared consumer device combines all four of these capabilities.

---

11 Discussion

11.1 The Personalization Imperative

The consistent theme across the technical evidence reviewed is that personalization is not a desirable feature enhancement for Aurevia — it is a prerequisite for clinical meaningfulness. Population-level models trained on existing public datasets achieve impressive aggregate accuracy metrics that obscure profound variability in individual-level performance. The patient whose pre-ictal EEG change occurs in a frequency band not well-represented in the training population will receive either no alert or a stream of false alarms — both outcomes that are worse than no device at all. The 30% false-positive reduction achieved by personalized channel selection alone [20], and the superiority of patient-specific seizure cluster prediction models [8], make the scientific case for individualized architecture emphatically.

11.2 From Detection to Prediction: The Technical Hurdle

The transition from seizure detection (identifying that a seizure is occurring) to seizure prediction (identifying that a seizure will occur) represents a qualitative escalation in technical difficulty that current public benchmarks understate. Detection can be validated on ictal EEG segments; prediction requires accurate labeling of pre-ictal windows that are, by definition, clinically indistinguishable from normal interictal periods at the time of recording. This labeling can only be done retrospectively (after a seizure has confirmed that a pre-ictal state was present), creating a fundamental annotation dependency on longitudinal real-world seizure data from individual patients — precisely the data that public datasets lack in sufficient volume and annotation quality [38].

11.3 Regulatory Pathway Considerations

Aurevia's regulatory strategy will be shaped by whether it claims a diagnostic, monitoring, or predictive function. FDA-cleared detection devices (EpiWatch, EpiMonitor) achieved 510(k) clearance as monitoring aids rather than diagnostic tools. An AI-based predictive device will likely require demonstration of clinical sensitivity and specificity in a prospectively enrolled, clinically validated patient cohort — a rigorous bar that necessitates planning the clinical validation study infrastructure in parallel with technical development. The FDA's Digital Health Center of Excellence and Software as a Medical Device (SaMD) framework provides the regulatory scaffolding; early Pre-Submission (Q-Sub) engagement is advisable.

11.4 Limitations of This Framework

The mathematical formulations in Section 4 are theoretical constructs; the parameters (η, λ, μ in Eqs. 2–4; I_{model}, $FLOPs_{per\_window}$ in Eq. 5) require empirical determination from hardware characterization and clinical validation data not yet available. Performance metrics cited throughout are drawn from heterogeneous studies with varying datasets, patient populations, and evaluation protocols; direct numerical comparisons should be interpreted with caution. The prototype hardware architecture described is conceptual; hardware build specifications will be refined through iterative prototyping.

---

12 Conclusions

Aurevia addresses one of the most consequential unsolved problems in epilepsy care: providing individuals with drug-resistant epilepsy with reliable, personalized advanced warning of seizures before they occur. The

technical foundation for this is now within reach — but only through the deliberate integration of the right AI architecture, the right hardware constraints, the right personalization strategy, and the right data infrastructure, working together as a coherent system rather than as independently optimized components.

This paper has provided a coherent system framework. The key findings are: (i) hybrid deep learning architectures combining CNN-based spectral pattern recognition with LSTM or attention-based temporal modelling achieve the highest prediction accuracy on benchmark datasets, but must be paired with SNN or TinyML optimisation for edge deployment; (ii) patient-specific personalisation — implemented through on-device T3A adaptation and federated learning at scale — is non-negotiable for clinical usefulness; (iii) the class imbalance and annotation limitations of public EEG datasets necessitate a proprietary longitudinal data collection strategy built into the device itself; and (iv) multi-modal sensing (ECG, accelerometry, heart rate) materially enhances robustness and enables seizure-type characterisation beyond what EEG alone provides.

The prototype image sections in Section 6 are provided as reserved spaces to be populated as hardware development progresses, establishing a documentation framework that tracks the physical realization of the Aurevia system from first prototype to clinical device. The mathematical framework in Section 4 provides quantitative scaffolding for model evaluation, device optimization, and clinical trial endpoint design.

For the millions of people living with drug-resistant epilepsy and the constant unpredictability it imposes on every aspect of daily life, Aurevia's vision — genuine, personalized, wearable seizure prediction — represents not just a product opportunity but a meaningful human need. This paper establishes the scientific and technical foundation from which that vision can be rigorously built.

Conflict of Interest

The authors declare no conflicts of interest. Aurevia is a proprietary development project of CORTEX LABs. This paper does not constitute a commercial offer or warranty of device performance.

Data Availability

This paper presents a theoretical and design framework. No primary datasets were generated. Public datasets referenced are accessible via CHB-MIT at PhysioNet (https://archive.physionet.org/physiobank/database/chbmit/) and the Kaggle Epilepsy Dataset (https://www.kaggle.com/datasets/datasetengineer/epilepsy-dataset). Proprietary Aurevia development data will be subject to institutional data governance protocols and participant consent frameworks.

References

[1] Davenport G. Electroencephalography: Principles and Clinical Applications. Amazon / Academic Press, 2023.

[2] Mayo Clinic. EEG (electroencephalogram). https://www.mayoclinic.org/tests-procedures/eeg/about/pac-20393875. Accessed May 2025.

[3] Siddiqui MK et al. Machine and Deep Learning-Based Seizure Prediction: A Scoping Review. Applied Sciences 2025;15(11):6279. https://doi.org/10.3390/app15116279

[4] CURE Epilepsy. EEG for Epilepsy. https://www.cureepilepsy.org/understanding-epilepsy/diagnosing-epilepsy/eeg. Accessed May 2025.

[5] Epilepsy Foundation. Electroencephalography (EEG) for Epilepsy. https://www.epilepsy.com/diagnosis/eeg. Accessed May 2025.

[6] Zhu H et al. ARNN: Attentive Recurrent Neural Network for Multi-channel EEG Seizure Detection. arXiv 2024. https://arxiv.org/html/2403.03276v2

[7] Nasri M et al. Patient-Independent Epileptic Seizure Detection with Reduced EEG Channels. Information 2025;16(1):20. https://doi.org/10.3390/info16010020

[8] Blum S et al. Individualized seizure cluster prediction using iEEG. OpenReview 2024. https://openreview.net/pdf?id=MmSWvH2bmd

[9] Billeci L et al. High Accuracy of Epileptic Seizure Detection Using Tiny Machine Learning. Neurol Int 2025;5(1):14. https://doi.org/10.3390/neurolint5010014

[10] Alotaibi FM et al. Seizure Onset Zone Detection Based on CNN. Brain Sciences 2024;14(11):1090.

[11] Liquid-Dendrite Spiking Neural Network for Edge Devices: A 130K Parameter Model. medRxiv 2025. https://doi.org/10.1101/2025.02.09.25321975

[12] Albawi S et al. Hybrid Deep Learning Model for Epileptic Seizure Classification using 1D-CNN with Multi-head Attention. arXiv 2025. https://arxiv.org/html/2501.10342v1

[13] Elias A et al. Automated Seizure Diagnosis System Based on Feature Extraction. PMC 2021. PMC7881082.

[14] Gomez C et al. Effective Early Detection of Epileptic Seizures through EEG Signals. Diagnostics 2023;13(11):1957.

[15] Tatum WO et al. EEG Normal Waveforms. StatPearls / NCBI Bookshelf, 2024. https://www.ncbi.nlm.nih.gov/books/NBK539805

[16] McGill University. EEG Signal Characteristics. Medicine.mcgill.ca, 2024.

[17] LearningEEG. EEG Montages and Technical Components. https://www.learningeeg.com/montages-and-technical-components. Accessed May 2025.

[18] DatasetEngineer. Epilepsy Dataset. Kaggle 2024. https://www.kaggle.com/datasets/datasetengineer/epilepsy-dataset

[20] Rashwan A et al. Personalizing Seizure Detection by Optimal Channel Selection. Sensors 2025;25(9):2715. https://doi.org/10.3390/s25092715

[21] Boonyakitanont P et al. Epileptic seizure detection using CHB-MIT dataset: Overlooked perspectives. PMC 2024. PMC11286169.

[22] Comparative Analysis of Epileptic Seizure Prediction Algorithms. arXiv 2023. https://arxiv.org/abs/2308.05176

[23] Shoeibi A et al. AI-augmented digital care pathway for epilepsy. PMC 2024. PMC11459578.

[24] MT-NAM: An Efficient and Adaptive Model for Epileptic Seizure Detection. arXiv 2025. https://arxiv.org/pdf/2503.08251

[25] Liquid-Dendrite SNN for Edge Devices: 130K Parameter Model — Full Text. medRxiv 2025. https://doi.org/10.1101/2025.02.09.25321975

[27] Poh MZ et al. Detecting and Predicting Seizures with High Accuracy from ECG. PubMed 2024. PMID 38359446.

[30] Advancements in Personalized Federated Learning for Epilepsy Detection. SciTePress 2024. https://www.scitepress.org/Papers/2024/133115

[31] An EEG Signal Recognition Algorithm During Epileptic Seizure Using Edge-Cloud Framework. IJIMAI 2022.

[33] Mabrouk R et al. Artificial Intelligence Guided Wearable Seizure Detection Device (EpiPatch). ResearchGate 2023.

[34] Loonen ICM et al. Seizure Detection Devices. J Clin Med 2025;14(3):863. https://doi.org/10.3390/jcm14030863

[35] AHA Center for Health Innovation. FDA Clears Apple Watch-Powered Platform for Seizure Monitoring. April 2025.

[36] Empatica. EpiMonitor — The Future of Epilepsy Monitoring. https://www.empatica.com/epimonitor. Accessed May 2025.

[37] Google Patents. US10596377B2 — Seizure Detection, Prediction and Prevention Using Neurostimulation. https://patents.google.com/patent/US10596377B2/en

[38] Asif U et al. EEG Datasets for Seizure Detection and Prediction — A Review. ResearchGate 2023.

[39] PhysioNet. CHB-MIT Scalp EEG Database. https://archive.physionet.org/physiobank/database/chbmit