



Published in final edited form as:

IEEE Trans Biomed Circuits Syst. 2021 October ; 15(5): 877–897. doi:10.1109/TBCAS.2021.3112756.

Closed-Loop Neural Prostheses with On-Chip Intelligence: A Review and A Low-Latency Machine Learning Model for Brain State Detection

Bingzhao Zhu [Student Member, IEEE],

School of Applied and Engineering Physics, Cornell University, Ithaca, NY, 14853 USA; Institute of Electrical Engineering and Center for Neuroprosthetics, EPFL, 1202 Geneva, Switzerland

Uisub Shin [Member, IEEE],

School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, 14853 USA; Institute of Electrical Engineering and Center for Neuroprosthetics, EPFL, 1202 Geneva, Switzerland

Mahsa Shoaran [Member, IEEE]

Institute of Electrical Engineering and Center for Neuroprosthetics, EPFL, 1202 Geneva, Switzerland

Abstract

The application of closed-loop approaches in systems neuroscience and therapeutic stimulation holds great promise for revolutionizing our understanding of the brain and for developing novel neuromodulation therapies to restore lost functions. Neural prostheses capable of multi-channel neural recording, on-site signal processing, rapid symptom detection, and closed-loop stimulation are critical to enabling such novel treatments. However, the existing closed-loop neuromodulation devices are too simplistic and lack sufficient on-chip processing and intelligence. In this paper, we first discuss both commercial and investigational closed-loop neuromodulation devices for brain disorders. Next, we review state-of-the-art neural prostheses with on-chip machine learning, focusing on application-specific integrated circuits (ASIC). System requirements, performance and hardware comparisons, design trade-offs, and hardware optimization techniques are discussed. To facilitate a fair comparison and guide design choices among various on-chip classifiers, we propose a new energy-area (E-A) efficiency figure of merit that evaluates hardware efficiency and multi-channel scalability. Finally, we present several techniques to improve the key design metrics of tree-based on-chip classifiers, both in the context of ensemble methods and oblique structures. A novel Depth-Variant Tree Ensemble (DVTE) is proposed to reduce processing latency (e.g., by 2.5× on seizure detection task). We further develop a cost-aware learning approach to jointly optimize the power and latency metrics. We show that algorithm-hardware co-design enables the energy- and memory-optimized design of tree-based models, while preserving a high accuracy and low latency. Furthermore, we show that our proposed tree-based models feature a highly interpretable decision process that is essential for safety-critical applications such as closed-loop stimulation.

Keywords

Neural prostheses; closed-loop neuromodulation; on-chip machine learning; symptom detection; decision trees

I. Introduction

Developing novel non-pharmacological treatments such as neurostimulation is becoming increasingly important to treat some of the most prevalent and intractable neurological disorders. Brain stimulation is currently the most common surgical treatment for movement disorders and has shown promise in epilepsy, neuropsychiatric disorders, memory, chronic pain, and traumatic brain injury, with new applications rapidly emerging. Despite promising proof-of-concept results, current clinical neurostimulators are limited in many aspects. For example, while deep-brain stimulation (DBS) can effectively control motor symptoms in most patients suffering from Parkinson's disease (PD), it causes persistent side effects (e.g., speech impairment and cognitive symptoms) [1], [2]. It is now widely known that this is due to the conventional 'open-loop' approach, which involves delivering constant high-frequency (~130Hz) stimulation regardless of the patient's clinical state. In addition, open-loop stimulation increases the power consumption and the need for surgical battery replacement. This simplistic open-loop approach is also a key limiting factor in designing clinically effective stimulation for more complex disorders such as depression [3], Alzheimer's disease [4], and stroke [5], [6], among others [5], [7], [8].

To further leverage the benefits of stimulation and address the aforementioned limitations, closed-loop neuromodulation techniques have been recently explored, such as the responsive neurostimulator for epilepsy [9] and PD [10], with promising results. In this approach, stimulation is dynamically controlled according to a patient's clinical state, either with a continuous (i.e., adaptive) or an on-off (i.e., on-demand) strategy. Through feedback from relevant biomarkers of a neurological symptom (e.g., a seizure event, tremor episode, or mood change), closed-loop stimulation can titrate charge delivery to the brain, thus reducing the side effects and the amount of stimulation delivered, enhancing the therapeutic efficacy and battery life compared to its open-loop counterparts [2]. However, several critical challenges remain to be addressed in order to fully exploit the potential of closed-loop therapies for neurological disorders. The existing closed-loop devices mainly rely on simple comparison of a pre-selected biomarker (typically from 1 out of 4 channels) against a fixed threshold. Such simplistic approaches are known to be suboptimal in terms of predictive accuracy, resulting in low sensitivity and high false alarm rates, while exacerbating other symptoms [8]. Multiple biomarkers and control loops may be necessary to reliably improve symptoms, leading to design complexity.

A promising solution to address this challenge is to implement a machine learning (ML) algorithm directly on the implant or wearable to predict the onset or severity of neurological symptoms, an approach that has gained significant interest in recent years [11]-[18]. Real-time symptom control can be achieved through on-chip biomarker extraction and ML-based disease state detection, followed by a closed-loop intervention (e.g., electrical, magnetic or

optical stimulation, drug delivery) to suppress the abnormal activity, as illustrated in Fig. 1. This approach offers significant advantages over the conventional wireless transmission and external processing methods [19], [20] that suffer from feedback loop latency, high power consumption due to continuous telemetry, security and privacy concerns [21], [22]. A number of clinical trials have recently shown the advantage of machine learning-based control for closed-loop stimulation in movement disorders, epilepsy, and memory [4], [23]. In addition, machine learning systems have been developed to forecast the onset of neurological symptoms during preictal phase, allowing sufficient time prior to seizure manifestation (e.g., in the order of several minutes) to provide early warnings to the patients and caregivers [24]-[26]. In closed-loop neural prostheses, however, both the ML decoder and neurostimulator are integrated on the implant, eliminating the need for excessively long symptom prediction horizons [27]. Therefore, most closed-loop devices train the classifier to differentiate ictal epochs from interictal period, several seconds prior to symptom onset [28]. Such systems detect the onset and termination (i.e., offset) of neurological symptoms to precisely control the delivery of stimulation [13].

Despite the benefits of using machine learning for closed-loop intervention, strict power and area requirements on an implantable or wearable device pose critical challenges for hardware realization of ML algorithms, particularly in the form of a miniaturized ASIC. The choice of learning algorithm and neural biomarkers affects the prediction accuracy and latency. Moreover, the prediction accuracy depends on the spatial resolution of the recording system and the number of input channels. Thus, there is a crucial need to develop high-performance, energy- and area-efficient biomarker extraction and ML solutions that are scalable to high channel counts and satisfy the implantable/wearable power budget and form factor.

In this paper, we review the state-of-the-art neural prostheses with embedded biomarker extraction and machine learning. We first discuss the closed-loop system components, requirements for the next-generation smart neural prostheses, their clinical applications, hardware techniques and trade-offs. Commercial and investigational closed-loop neuromodulation devices and a comparison of previously reported system-on-chips (SoCs) for neural signal classification are presented. In the second part of this paper, we discuss an emerging class of machine learning algorithms based on decision trees [12], [22], [29]-[31], including tree ensembles and oblique trees, that are particularly suitable for energy- and area-constrained platforms such as brain implants and wearables. We introduce novel techniques to improve the accuracy-latency trade-off in tree ensembles. A new class of tree-based models that effectively combine decision trees (DTs) with neural networks is further discussed. After presenting various techniques for energy, latency, and memory-efficient realization of oblique trees, we present the results of testing these models on two neural signal classification tasks relevant to closed-loop stimulation (epilepsy and PD).

It should be noted that closed-loop neural prostheses with on-chip intelligence are also being explored in the context of fully implantable brain-machine interfaces (BMI) [30], [32]-[35]. Such BMI systems can provide a sensory feedback to the brain and/or control prosthetic devices to restore lost motor or sensory function in paralyzed patients. However, the focus of this paper is on neural prostheses that directly record and modulate the brain

activity to treat neurological disorders, while motor neuroprosthetics (i.e., BMIs or brain-computer interfaces, BCI), peripheral [36] and spinal cord prostheses [37] (e.g., EMG-based interfaces) are beyond the scope of this paper. Furthermore, we limit our review to those systems that focus on ASIC implementation of neurological symptom detection algorithms (either validated on, or with a potential for closed-loop stimulation) due to similarity in design requirements. Thus, FPGA-based systems are not included in this review. While the focus of this review is on CMOS-based edge machine learning specifically for neural prostheses, a comprehensive review on embedded hardware (FPGA, neuromorphic, CMOS) for neural networks used in biomedical applications can be found in [38].

This paper is an extension of our conference paper [22] that presented a brief survey on closed-loop neural interface systems with on-chip machine learning and provides the following contributions:

- A comprehensive review on the latest developments in technology design for closed-loop stimulation, including novel electrodes for sensing and stimulation, emerging clinical applications, commercial, research-based and investigational devices for closed-loop stimulation.
- A detailed review of the reported neural interface SoCs with on-chip machine learning for neurological disease detection, either as a stand-alone chip or as part of a closed-loop system (implantable and wearable).
- Future directions for the next-generation closed-loop neural prostheses, including the integration of advanced design techniques, accommodating high channel counts and the need for online learning.
- Novel algorithm-hardware co-design techniques for next-generation energy-efficient neural prostheses. Specifically, we present a range of methods for cost-aware implementation of tree-based classifiers in brain implants and validate them on human neurophysiological datasets.

II. Closed-loop Neural Prostheses: Recent Trends, System Requirements, and Trade-offs

In a closed-loop neural prosthesis (Fig. 1), neurostimulation is triggered to suppress the impending signs of a neurological disease. Research on closed-loop neurostimulation has gained momentum in recent years, particularly with the success of proof-of-concept studies on epilepsy [43] and PD [1], [44], [45]. Closed-loop approaches are now being explored to treat a variety of medication refractory brain disorders where open-loop stimulation has been less effective. Yet, major technological challenges have limited the efficacy and clinical translation. These challenges include the low channel count of the current devices, the effect of stimulation artifacts on the sensing circuits, the need for miniaturization and improved energy efficiency, and the need for more advanced control algorithms [2], [8], [22]. Next-generation closed-loop neuromodulation systems will require significant improvements in the existing devices. For instance, higher numbers of recording and stimulation channels will be necessary for disorders that require multi-site neural recording and manipulation. More

sophisticated processing algorithms and complex stimulation patterns will be beneficial to improve therapeutic outcomes. However, this will increase the design complexity and required on-chip resources for symptom detection and stimulation, as well as the required processing time. Better localization of target regions for effective stimulation and improved stimulation artifact cancellation are also critical for bidirectional neural prostheses. In this paper, we discuss the major challenges and review the most recent advances in the field, with a particular focus on machine learning-embedded implantable and wearable systems.

A. Sensing and Stimulation

High-density neural recording and multi-site neurostimulation with low-power miniaturized circuits are crucial for the next-generation closed-loop neural prostheses. Particularly, complex disorders such as depression and Alzheimer's disease (AD) need multi-site rather than single-site recording that calls for more intelligent, data-driven closed-loop systems with high-density sensing and stimulation capabilities.

1) Conventional and Emerging Electrodes for Sensing and Stimulation: In a neural prosthesis, the electrophysiological activity of the brain can be recorded through various noninvasive, minimally-invasive, or invasive electrodes such as scalp EEG, subscalp EEG [39], electrocorticography (ECoG), also known as intracranial EEG (iEEG), stereo-EEG (sEEG) [41], [46], and deep-brain leads, providing various degrees of spatial and temporal resolution (Fig. 2). In some cases and predominantly in implantable prostheses, the same electrode can be used for delivering electrical stimulation to the brain to suppress disease symptoms.

The EEG electrodes have a cm-range distance and are noninvasive. Both scalp and subscalp EEG are suitable for wearable settings, with electrodes placed either above (scalp EEG) or under the scalp (subscalp EEG). Subscalp electrodes are particularly suitable for chronic (i.e., longer than one month) EEG recording in a home environment and require a minimally invasive surgery under general anesthesia to implant the subcutaneous electrodes [39]. The subscalp approach eliminates the need for constant electrode care (i.e., no need for an EEG cap or adhesives electrodes), providing a stable and less obtrusive recording modality compared to conventional EEG, Fig. 2(b). Furthermore, subscalp EEG has been shown to attenuate several types of artifacts and improve (or at least maintain) the signal quality compared to EEG. However, similar to scalp EEG, it is limited in temporal and spatial resolution compared to ECoG (i.e., <100Hz vs. several hundred Hz) and cannot monitor deep-brain structures. A number of subscalp EEG systems are currently certified or in development for long-term epilepsy monitoring (Section III).

The spacing of ECoG electrodes (epidural or subdural) is typically within mm-range, while state-of-the-art ECoG interfaces enable denser recording arrays for high-spatial-resolution recording of cortical activity [47]. For instance, it has been shown that high-density μ ECoG with a 400 μ m pitch outperforms lower density grids in classifying cognitive tasks in humans [48], highlighting its potential for future high-performance neuroprosthetic applications. High-frequency electrophysiological activity relevant to seizure prediction or epileptic foci localization can be captured on high-resolution ECoG from submillimeter scale cortical

regions [47], [49]-[51]. These novel electrodes are not yet adopted in diagnostic or closed-loop devices.

While ECoG provides a precise mapping technique at the level of cortical surface, stereo-EEG (sEEG) [41] is an alternative minimally-invasive method for identifying seizure onset zone in medically refractory focal epilepsy. Placement of stereo-EEG electrodes (typically 5–15 cylindrical shafts) requires small, localized burr holes to insert depth electrodes into the brain. Stereo-EEG enables a sparse sampling of localized brain regions, as opposed to the relatively large craniotomy required for strip/grid ECoG implantation [41].

The electrodes on a deep-brain lead (e.g., Medtronic 3387/3389 deep-brain stimulation lead with four cylindrical contacts) are placed several millimeters or even 100s of micrometers apart to capture the local field potential (LFP) activity (up to several 100 Hz) [42]. The leads employed in sEEG are similar to those used for deep brain stimulation (DBS). DBS is widely used as a treatment for essential tremor, PD and dystonia, with emerging applications in epilepsy, major depression, obsessive-compulsive disorder (OCD), and Tourette's syndrome. While DBS is primarily used for electrical stimulation, the chronic efficacy and stability of DBS leads suggest the use of long-term sEEG for sensing applications and closed-loop prostheses [41]. In rare cases, single-unit activity captured by μ DBS leads (100 μ m spacing [42]) or penetrating microelectrodes such as Utah array can be used to detect spike-based biomarkers (e.g., neuronal firing rates correlating with cognitive functions) for disease state prediction and guiding neurostimulation therapy [52], [53].

For stimulation, recent DBS electrodes employ directional leads with higher number of small contacts (e.g., 16, 40, 1760) as opposed to traditional leads with only four cylindrical contacts [8], [42], Fig. 2(e). Such directional leads with segmented electrodes can effectively steer the stimulation back toward a missed target structure, without exciting non-target regions and inducing adverse effects. Moreover, recent studies report the impact of using temporal patterns delivered via multiple contacts in enhancing plasticity and symptom relief [8], [42], highlighting the benefits of high-channel-count stimulation.

2) Concurrent Sensing and Stimulation: Accuracy and latency can be enhanced by measuring evolving disease state even as therapeutic stimulation is applied. This motivates the need for a new class of circuit and system techniques to enable detection of weak electrophysiological signals of interest in the presence of orders-of-magnitude stronger stimulus artifacts. This general problem of measuring weak signals in the presence of extreme self-interference represents a general challenge for modern mixed-signal circuit in various sensing and communication applications. The next generation 'full-duplex' neuromodulation devices must feature simultaneous sensing and stimulation for truly closed-loop operation.

The most common electrical approach is to use 'blanking' [54], [55] where recording amplifiers are disconnected from the electrode during and immediately after stimulation, and then reconnected after the stimulation artifact will no longer saturate the amplifier. Recent improvements allow the amplifier to be connected immediately after stimulation [56], using mixed-signal circuit realization of the analog front-end (AFE). However, this

method still suffers from its inability to record while stimulating, which is especially limiting in complex, multi-electrode stimulation patterns where extended stimulation blocks recording over much longer time stretches.

An alternative approach based on high dynamic range (DR) AFE incorporating amplifiers and analog-to-digital converters (ADC) can reliably record the neural signal along with the persistent artifacts without saturation [57], [58]. Alternatively, the design in [59] proposes a linear-interpolation-based artifact cancellation implemented on an FPGA. Another approach employs a front-end cancellation technique that avoids using a high DR AFE [60]. However, this method requires a significant convergence time (impractical for closed-loop systems). Artifact cancellation generally poses additional hardware overhead on the AFE and on the back-end for digital cancellation, which limits the area and energy efficiency of the closed-loop system.

B. Disease Biomarkers and Machine Learning

While artificial intelligence and machine learning can contribute to various aspects of neurotechnology (e.g., optimizing the programming of stimulation to activate target regions, offline analysis of chronic neural recordings, understanding the underlying disease mechanism), our focus in this paper is on real-time on-device disease state prediction using machine learning. This is inspired by the unique potential of ML techniques in classifying high-dimensional electrophysiological signals, typically outperforming conventional methods in various applications [2], [11], [12], [61]-[64]. Accurate and timely detection of symptoms in brain disorders is critical to enable closed-loop neuromodulation, and it typically requires the use of correlating biomarkers (i.e., features) of an underlying disease state along with a machine learning algorithm. The widely used features in electrophysiological studies include the spectral power (or bandpower) in various frequency bands relevant to the neurological symptom of interest, time-domain and statistical features (e.g., line-length [65], the Hjorth parameters of activity, mobility, and complexity [2], [63], [66], number of peaks, peak-to-peak amplitude and peak latency [63]), biomarkers that measure connectivity between different brain regions such as phase-amplitude coupling and phase locking value [2], [64], [67]-[69], and the correlation structure of multi-channel neural data [70].

Some initial steps have been taken recently toward embedding biomarkers and machine learning algorithms on brain implants or wearables for disease monitoring and closed-loop therapy, and in investigational neuromodulation systems such as Medtronic's Summit RC+S [71] and Percept PC systems, as summarized in the next sections.

1) Classifier requirements – High accuracy, low latency: Symptom detection requires high accuracy and low latency. The classification algorithms should be robust in handling the typically small amounts of training data in such applications, due to the lack of chronic recordings. In some cases, the recording length could be limited to the duration of surgery for device implantation (e.g., up to 30 minutes for DBS surgery in PD, several days for epilepsy patients undergoing pre-surgery evaluation at the hospital). With the increasing interest in devices with chronic recording capability (e.g., the NeuroPace RNS

and Medtronic Percept), it is expected that more long-term human data will be available in near future, enabling data-driven algorithm and hardware developments.

Depending on the distribution of different classes in a neurophysiological dataset, the appropriate measure of accuracy may be used to evaluate the classifier's performance. Sensitivity (i.e., True Positive rate), specificity (i.e., selectivity or True Negative rate), accuracy, F1 score, the area under the ROC curve (AUC), and the false alarm rate (FAR) are among the commonly used metrics in ML studies on neural datasets. The F1 score (i.e., the harmonic mean of sensitivity and precision: $2 \times (\text{precision} \times \text{sensitivity}) / (\text{precision} + \text{sensitivity})$, where precision represents the positive predictive value) is particularly useful in dealing with imbalanced datasets (i.e., datasets with non-uniform distribution of classes), such as EEG or iEEG recordings in epilepsy [12]. Balanced accuracy (i.e., the average of sensitivity and specificity) is another metric used for imbalanced datasets [64].

Most closed-loop systems rely on external computing for feature extraction and classification, which suffers from long loop latency, thus jeopardizing the real-time feedback. The on-chip integration of ML can significantly speed up the closed-loop therapy and enable feedback loops of msec-range latency. If the feedback is too slow, the detector may miss the window of opportunity to trigger or adjust stimulation, resulting in poor therapeutic outcomes. More sophisticated processing algorithms may improve the decoding accuracy at the cost of increased processing latency.

While 'latency' has been used to represent various types of 'processing delay' in literature (e.g., feature extraction and classification delay resulting from window-based processing), the detection latency of a closed-loop system is typically defined as the delay between the electrographic, expert-marked, or externally labeled symptom onset and the onset declared by the on-chip processor, for instance in detecting seizures in epilepsy [12], [61], [72]-[76] or tremor onset in PD [2], [68], [77]. In disorders such as epilepsy, the onset of clinical symptoms could be several seconds (in some cases, up to 30 seconds [28]) after the time of earliest detectable changes in neural activity. Therefore, therapeutic feedbacks within that time frame can be still beneficial for the patients. In other cases, e.g., in movement disorders with more rapid changes in electrophysiological state, a low latency (i.e., negative latency or lead [2]) is preferred to enable closed-loop stimulation.

2) Classifier requirements – Low power and small area: To enable efficient local processing in a brain implant, silicon-realizable ML algorithms that can precisely predict a neurological symptom are essential. Neural prostheses with on-device ML do not require continuous wireless telemetry. Yet, low-power realization of machine learning algorithms is crucial to avoid excessive power dissipation. Optimized use of memory and computational resources and compact silicon area are further required to process multiple channels. The computational complexity of the classifier (and features) could set a limit on the number of input channels, thus hindering its application in more complex disorders.

The conventional implementation of most classification algorithms is resource intensive such that devices in existence today [43] sacrifice the classification accuracy and latency to

meet the power and size constraints [12]. Some limited processing is embedded in recently developed neuromodulation devices, but this applies to 1–4 channels only, requiring external classifiers for more accurate symptom detection [71]. There is a crucial need for energy- and area-efficient machine learning algorithms via co-design of algorithm and hardware, as discussed in the next sections.

3) Neurophysiological Datasets: In contrast to computer vision tasks that benefit from standard datasets for direct benchmarking of machine learning models, the electrophysiological datasets used in disease prediction tasks are diverse and not directly comparable. Furthermore, these datasets include different numbers of patients with various levels of symptom detection complexity, making it challenging to compare classifiers evaluated on the same dataset but on different patients. Another critical challenge is the lack of data sharing and open-source datasets in emerging applications beyond epilepsy (e.g., movement disorders, depression, Alzheimer’s disease), which greatly limits the development of biomarkers and ML solutions and subsequent device implementation for such disorders.

III. Commercial and Investigational Closed-loop Devices

One of the few platforms currently available for closed-loop stimulation is the NeuroPace’s Responsive Neurostimulator (RNS) for medication-refractory epilepsy (Fig. 3(a)). RNS continuously analyzes cortical activity to detect and halt seizure events from 4 channels, by comparing a simple pre-selected feature (signal intensity, line-length, or half-wave) against a threshold [43], [78], and it is currently in clinical use in patients. Both cortical and deep-brain stimulation are enabled in RNS (8 channels). The recently published results of a nine-year, multi-center chronic study of RNS device on 230 patients in 34 epilepsy centers [79] showed significant reductions in seizure rates: 75% median reduction, at least 50% reduction in 73% of patients. The sudden unexpected death in epilepsy (SUDEP) was also significantly reduced. The responsive neurostimulation was a well-tolerated treatment, with a similar safety profile to other epilepsy procedures.

Similarly, Medtronic’s investigational Activa PC+S, Summit RC+S [71] and Percept PC system (Fig. 3(c)) are capable of sensing and closed-loop stimulation for movement disorders such as essential tremor and PD. Compared to RNS, the Medtronic devices implement slightly more complex spectral analysis and a linear classifier, relying on only 4 sensing channels with 2–8 features in total, and 8–16 stimulation channels. For both RNS and Medtronic devices, external algorithms with advanced machine learning capabilities may be necessary for more accurate symptom tracking [71], [80], at the cost of long loop latency and high power demands to support continuous wireless streaming [12], [31].

The AspireSR 106 (LivaNova) is an implantable Vagus Nerve Stimulator (VNS) with an optional AutoStim mode in which the VN stimulation can be adjusted in response to ictal heart rate changes which are potentially associated with an impending seizure (Fig. 3(b)) [81]. In a study on the efficacy of open-loop VNS on 5554 patients [82], a growing increase in seizure freedom was observed post therapy, with 49% responding to treatment 0–4 months after implantation (i.e., >50% seizure frequency reduction). The efficacy of closed-loop AspireSR versus the preceding open-loop device was recently studied, where 4

(from 11) patients who were less responsive to the open-loop VNS achieved >50% seizure reduction [83]. Of note, there have been reports on the Federal Drug Administration (FDA) device recall for different models of VNS due to concerns on device malfunctions.

In addition to the devices described above, there is an increasing effort in developing novel closed-loop stimulation devices for a variety of brain disorders. One example is the AlphaDBSTM system [84], which recently received the CE mark approval to treat Parkinson's disease (Fig. 3(d)). This closed-loop system developed by Newronika (S.p.A, Milan, Italy) can record deep-brain local field potentials and adjust the stimulation amplitude and frequency. DyNeuMo (Bioinduction, Bristol, UK) is a closed-loop neuromodulation research device that can titrate stimulation according to the current motor state (e.g., posture and activity) [85] (Fig. 3(e)). The device uses off-the-shelf consumer technology and embeds three-axis accelerometer sensors and 8-channel programmable neurostimulators, and is currently in preparation for first-in-human research trials.

Minimally-invasive signal modalities such as subscalp EEG are also being considered for long-term epilepsy monitoring. For instance, the Epios device (Wyss Center for Bio and Neuroengineering, Geneva, Switzerland) [39] enables both focal recording and full-montage coverage using subscalp EEG for chronic seizure analysis and forecasting (Fig. 2(b)). The EEG data is wirelessly transmitted to a wearable unit and temporarily stored, supporting multimodal ECG, audio, and accelerometry recording. Signals are then transmitted to the cloud for long-term data analysis and visualization. The Epios device is currently in preparation for clinical trial phase. The Minder device (Epi-Minder, Melbourne, Australia) [39] implants an electrode lead across the skull to cover both hemispheres (Fig. 3 (g)). This subscalp system provides continuous long-term measurement of EEG for chronic epilepsy diagnosis and monitoring (clinical trial in progress). Alternatively, in the EASEE system by Precisis (Heidelberg, Germany) five subscalp electrodes are implanted above the seizure focus for sensing and closed-loop neurostimulation with a personalized setting (clinical trial in progress) [39].

IV. Neural Prostheses with On-chip ML

In recent years, the application of machine learning techniques in closed-loop neuromodulation and its CMOS implementation have received considerable interest. Machine learning has been used to more accurately predict optimal stimulation times [2], [13], [16], [64], [86] and several clinical studies have shown the advantage of ML-based closed-loop therapy in movement disorders [23], epilepsy [87], and memory [4]. The most prominent benefits of integrating machine learning algorithms on a brain implant include:

- Eliminating the need for excessive wireless transmission for external processing, thus allowing design miniaturization, lower power dissipation, and higher mobility.
- Increasing patient independence and alleviating security concerns by avoiding the transmission of private data.
- Improving symptom prediction accuracy and latency.

The latter advantage largely depends on the number of sensing channels, the quality of neural signal (e.g., its sampling rate and signal-to-noise ratio), the choice of machine learning algorithm and neural biomarkers, and the chronic robustness of the algorithm. As mentioned in the previous section, current clinical devices do not offer sufficient embedded biomarker extraction and ML, relying on telemetry and cloud-based processing for accurate symptom prediction.

Various hardware implementations of machine learning algorithms have been reported for neurological symptom detection, as discussed below. Here, we limit our review to state-of-the-art neural prostheses with an ASIC implementation, validated on animal or human datasets (acute and/or chronic, either diagnostic only or closed-loop).

A. Implants and Wearables for Epilepsy

The most common application of on-chip classification in a neural prosthesis is in the context of seizure detection for medically refractory epilepsy, where a supervised ML algorithm is typically used to detect the onset of seizure events from multi-channel neural recordings. Neurostimulation offers an attractive treatment for intractable epilepsy (approximately one third of epileptic patients). Due to severity of refractory epilepsy, open-source epileptic EEG datasets (both scalp and intracranial) are largely available, as well as established animal models for device validation and preclinical studies. Therefore, several groups have integrated various biomarkers and machine learning algorithms on ASIC for automated seizure detection [11], [12], [14], [73], [90]-[95] and for controlling an on-chip stimulator [13], [16], [17], [29], [88], [89].

Most ML-embedded SoCs for epilepsy have adopted classifiers based on support vector machines (SVMs), as shown in Fig. 4 and Fig. 5(a). Several variants of SVM kernels including linear, second-order polynomial, and radial basis function (RBF) have been reported for on-chip implementation. An SVM classifier generates weighted feature matrices using multiply-and-accumulate (MAC) blocks and separates them into different classes via linear or non-linear separation boundaries. For example, [11] reported an 8-channel linear SVM classifier with digital bandpower features implemented using a distributed quad-LUT architecture, Fig. 5(a). The system was verified on the MIT PhysioNet EEG database from the Children's Hospital Boston (CHB-MIT). This dataset includes 906 hours of recordings from 24 patients with epilepsy with ~190 registered seizures, and is commonly used in EEG-based seizure detection SoCs (Table. I). Alternatively, the design in [14] implemented a Gaussian basis function (GBF) SVM classifier to account for linearly non-separable seizure patterns, Fig. 4(b). A natural log operator was employed to linearize the GBF equation and replace multiplications with additions. Time-division multiplexing was used to implement the bandpower features in an area- and energy-efficient manner. The non-linear SVM typically requires sufficient seizure patterns for training, which might be impractical for patients with limited training sets. Later, a combination of two linear SVMs was introduced [13] to address this limitation, Fig. 4(a). The two SVMs were trained separately to achieve high sensitivity and specificity, and the classification results were combined to generate final decisions. This noninvasive closed-loop SoC integrates a transcranial

electrical stimulator (tES) to suppress impending seizures. The classification performance and ASIC specifications are summarized in Table. I.

A 32-channel closed-loop neuromodulation system integrating frequency and phase-domain features, a 32-to-4 autoencoder for dimensionality reduction, and an exponentially decaying memory SVM (EDM-SVM) was proposed for seizure control [16], Fig. 4(f). This system was validated on 500 hours of iEEG data (4 patients, 44 seizures) provided by the EU dataset. The design proposed in [89] is an 8-channel closed-loop neuromodulation system for DBS, that was verified using stereo-EEG (sEEG) electrodes. The classifier is composed of a two-level coarse/fine detector, in which the DSP chip (separate from the core sensing chip) is only activated in case of suspected seizures raised by the coarse detector. In this mode, maximum-modulus discrete wavelet transform (MODWT) and kernel density estimation (KDE) are computed and classified by a least-squares SVM (LS-SVM) for fine classification, Fig. 4(h). Furthermore, [90] reported a configurable SVM processor with various kernels (RBF, polynomial, linear), validated on the MIT EEG dataset.

It should be noted that in addition to the machine learning processor, the feature extraction circuits can be highly power- and area-demanding, particularly in systems with many input channels. Minimizing the number of extracted features and their hardware complexity without jeopardizing the classification accuracy is essential to reduce the overall energy consumption and area. The required computational resources in SVM linearly scale with the number of neural channels, making such optimizations more critical in practice.

An 8-channel closed-loop iEEG-based seizure control SoC was presented in [17], computing frequency spectrum and time-domain entropy along with a linear least-square classifier, Fig. 4(c). This system was acutely verified in Long-Evans rats. Similarly, the closed-loop 16-channel design in [88] integrated a biosignal processor to extract approximate entropy (ApEn) and FFT-based bandpower features, passed to a ridge regression classifier (RRC). The system was verified on ECoG data from five patients (duration not reported), and acutely for closed-loop seizure suppression in mini-pigs, Fig. 4(d).

In addition to the above models, machine learning algorithms that exploit decision trees, either as base estimators in ensemble methods such as bagging and boosting [12], [29], [96] or as stand-alone classifiers [31] have been used in neural signal classification tasks. While Random Forests [97] apply a bagging technique to DTs in order to reduce variance, boosting is a bias reduction technique in which individual trees are incrementally added to the ensemble to correct the previously misclassified samples. Popular implementations of boosting methods include gradient boosting [98] and AdaBoost [99]. Both bagging and AdaBoost use classifiers as base estimators, while gradient boosting requires regressors. Particularly, ensembles of gradient-boosted DTs have recently emerged as an accurate [24], yet hardware-efficient [12], [92], [96] machine learning solution for neural SoC platforms and for applications with limited training sets. DT ensembles avoid hardware-intensive MAC operations and enable low-complexity hardware architectures for neural prosthesis applications.

In [12], a gradient-boosted DT ensemble achieved a record energy efficiency (41.2nJ/class, 32-channel) and a compact area (1 mm²) for seizure detection, Fig. 4(e). The system was validated on iEEG from 26 epilepsy patients (3074 hours, 393 seizures), available on the iEEG portal [102], a collaborative platform for sharing large iEEG datasets. An on-demand feature extraction approach was adopted by sequentially using a single feature extraction unit in each tree, thus substantially reducing the number of extracted features and the overall hardware cost for inference. As opposed to other classifiers that compute all features for each input channel, this unique property of DTs allows the classifier to selectively extract a limited number of features to minimize the loss function, thus accommodating a higher number of input channels (Table. I). Another CMOS implementation of tree-based models used AdaBoost with 1024 trees of depth one for seizure detection and closed-loop stimulation [29], Fig. 4(g). Thanks to a bit-serial processing scheme, this 8-channel SoC reported state-of-the-art energy efficiency (36nJ/class) for 8-channel iEEG classification. Recent work replaced axis-aligned splits with logistic regression to construct powerful oblique trees as an efficient combination of neural networks and DTs [31] (Section VI) for epileptic seizure and PD tremor detection.

B. Implants for Movement Disorders

Multiple feasibility studies using closed-loop DBS devices like Medtronic's Percept and Summit have demonstrated additional benefits using closed-loop versus open-loop DBS in movement disorders [103], [104]. Closed-loop DBS in PD [1], [10], [45] has led to improvements in tremor control, reduced stimulation time and power consumption, and reduced speech side effects compared to open-loop DBS. However, wider adoption of this approach is awaiting advances in implantable hardware, control algorithms, and chronic validation. Current systems predominantly use single-biomarker thresholding, which precludes the optimized control of tremor.

Recently, ML approaches have been used for detecting motor symptoms (e.g., tremor) in patients with PD and essential tremor [2], [23], [68], [77], [105], [106] to control DBS in closed loop. An approach based on feature engineering and tree boosting [2], [68] used various correlating features of tremor such as bandpower in multiple frequency bands, the ratio of high-frequency oscillations, phase-amplitude coupling, and tremor power to detect the onset of rest-state tremor episodes in PD. Using only five selected features, the system was able to predict tremor with a 89.2% sensitivity and detection lead of 0.52s in 12 patients, significantly better than conventional beta-thresholding approach. Fixed-point quantization and power-aware inference were later used to enable low-power gradient boosting, achieving 55.4% energy reduction compared to conventional tree ensemble [105]. A method based on resource-efficient oblique trees (ResOT) was recently applied to PD tremor detection, enabling significant energy and memory reduction by various hardware-algorithm co-design techniques [31]. A similar study was recently done on patients with essential tremor (ET) [23] who suffer from tremor during voluntary movements. Using a binary classifier, postural tremor and voluntary movements were detected from LFP features recorded via the DBS lead, achieving an average sensitivity of 80% in 7 patients with ET. Such machine learning techniques hold the promise to enable accurate symptom detection in

closed-loop neural prostheses for various movement disorders. More developments in SoC design for such applications are expected in near future.

C. Implants for Neuropsychiatric Disorders and Memory

Neuromodulation, particularly invasive technologies like DBS, has been recently explored for treating psychiatric disorders such as major depressive disorder (MDD) and obsessive compulsive disorder (OCD) [8], [107]. However, despite promising early results, the high-profile clinical trials have shown inconsistent effects. One major limiting factor is the open-loop approach used in conventional DBS, which has been shown to be inefficient in engaging target brain regions in complex disorders such as depression and OCD [3]. While the application of neurostimulation techniques has made a significant impact on the lives of patients with movement disorders, major advances are needed to treat more prevalent conditions such as depression. Closed-loop patient-specific stimulation appears to be the most viable solution.

Development of algorithms for automated detection of emotional states and shifts in arousal, vigilance, and wakefulness has received considerable attention in EEG-based human studies, with some recent reports on SoC design. For example, a deep neural network (DNN) classifier was implemented for emotion detection in children with Autism [100]. The valence/arousal binary classification by the 4-layer DNN was used to detect four-state emotions. A reduction in energy consumption was achieved through a pipelined DNN architecture with a central arithmetic logic unit, Fig. 5(c). This DNN processor can analyze two EEG channels with an accuracy of 85.2% and energy efficiency of 10.1 $\mu\text{J}/\text{class}$. In another design, a convolutional neural network (CNN) was proposed for emotion detection [101], offering an online training feature, Fig. 5(d). To minimize area and memory overhead due to batch processing, hardware re-use and mini-batch data were employed for training and acceleration, at the expense of longer training time. Using an external feature extraction engine, this system obtained a 83.36% accuracy in binary classification of emotions (Table. II). Machine learning has also been explored in sleep stage classification [108], task engagement [86] and mental fatigue prediction [69] to potentially trigger a neurostimulation therapy.

Disorders such as Alzheimer's disease exhibit network abnormalities, necessitating the need for multi-site electrophysiological recordings. The closed-loop stimulation approach in [4] used a patient-specific logistic regression classifier to decode the brain-wide electrocorticography (ECoG) signals, and subsequently triggered stimulation in response to the predicted periods of poor memory encoding to enhance memory. The results suggest a predictive role of increased high-frequency as well as decreased low-frequency activity for memory recall, and that responsive neuromodulation in the lateral temporal cortex could improve recall performance. More developments in neural prosthesis design for mental and memory disorders are expected in the coming years.

D. Wearables for Migraine

While most current devices have been developed for epilepsy and movement disorders, there is an increasing demand for novel therapeutic devices for other medication-resistant

neurological disorders. Migraine, for instance, is the most common neurological disorder that affects millions around the world. Migraine patients suffer from episodic headaches lasting hours to days and often move from a stage of low-frequency attacks into chronic migraine. The diagnosis mainly relies on patient diaries and clinical interviews [63], [109]. As an emerging alternative, neurophysiological monitoring techniques have shown to be beneficial in assessing migraine progression [109]. The automated detection of migraine state using continuous brain recordings could help in early and more effective treatment, either with medications or neurostimulation.

A machine learning approach was recently proposed for noninvasive migraine state detection from somatosensory evoked potential (SEP) biomarkers in 42 migraine patients, as described in [63]. The results suggest the potential use of SEP as a feedback signal for migraine attack prediction. Based on this idea, [18] reported a low-power feature extraction and ML processor for migraine state prediction, using single-channel SEP as input. Multiple features such as bandpower, time-domain and statistical features of high-frequency oscillations [63] were integrated with a multi-class artificial neural network (ANN), achieving a predictive accuracy of 76%, Fig. 5(b) (chip layout post place-and-route).

E. Implants for Stroke and Traumatic Brain Injury

Neurostimulation can be used to facilitate post-stroke plasticity and functional recovery. Compared to noninvasive methods such as transcranial magnetic or direct-current stimulation (TMS, tDCS), invasive tools such as direct cortical stimulation offer a higher temporal and spatial resolution. However, current cortical stimulation approaches for stroke are limited by the poor localization of stimulation targets and open-loop operation [110], urging the need for advanced data analysis and machine learning techniques.

In addition, patients with severe-to-moderate traumatic brain injury (smTBI) suffer from persistent cognitive dysfunction and chronic mental fatigue that significantly impacts all aspects of their functioning. Despite extensive efforts to develop rehabilitation and medication-based therapies, there are no effective therapeutic options for these patients. In a break-through study, it was shown that therapeutic DBS in the central thalamus (CT-DBS) could restore executive function, fluent communication and motor control in a patient who remained in a minimally conscious state for six years following a TBI [111]. Similar improvements have been observed in individuals with chronic mental fatigue. In a recent study, the ECoG activity from two healthy non-human primates (NHPs) during a sustained attention task was used to predict the onset of mental fatigue [64], [69]. Using spectrotemporal and connectivity biomarkers and a tree ensemble classifier, the decline in animal's performance was predicted, seconds prior to NHP's behavioral response. This approach could potentially be used for closed-loop neurostimulation in patients with TBI.

In a proof-of-concept study [53], a closed-loop neural SoC was used to facilitate recovery after brain injury in a rat model of brain injury. The action potentials detected in premotor cortex were used to trigger neurostimulation in somatosensory cortex for several weeks. This spike-triggered stimulation led to significantly improved reaching and grasping functions, enhancing the functional connectivity between the two brain regions. These findings

motivate the design of novel closed-loop neural prostheses to treat brain injury and similar neurological indications.

F. Comparison of ML-embedded SoCs

A comparison on hardware specifications and classification performance of state-of-the-art neural prostheses with on-chip machine learning is presented in Table. I (for epilepsy) and Table II (for other applications). When comparing the performance and hardware cost of different ML SoCs, one should consider various factors that affect the overall predictive performance and design complexity, such as the input signal modality and dataset, the number of analyzed patients, the duration of recording and seizure count, and the metrics used to evaluate the algorithm/hardware performance (e.g., accuracy, F1 score, sensitivity, power vs. energy efficiency, detection vs. system latency). In addition, the number of processed channels should be taken into account to fairly compare various architectures and assess their scalability.

Energy efficiency has been a common metric to compare different ML-embedded biomedical SoCs in literature. However, we note that the energy efficiency is not being reported in a unified manner (e.g., total power consumption/sampling rate [11]-[14], [29] or total power consumption/classification rate [16], [17], [88] has been used), which may hinder appropriate design choices. Furthermore, the number of channels is not taken into account, which is particularly important in modern neural prostheses. Here, we define a new energy-area efficiency figure of merit (E-A FoM) as follows:

$$\text{E-A FoM} = \frac{P_{Ch} \cdot A_{Ch}}{f_s} \quad (1)$$

where P_{Ch} and A_{Ch} indicate per-channel power and area of the ML SoC, respectively, and f_s is the per-channel sampling rate of the signal processing circuits. Similar FoMs have been used in AFE and ADC design for multi-channel neural recording [112]. The E-A FoM fairly represents the energy-area efficiency of the system while also factoring in the multi-channel scalability. Other performance metrics such as accuracy and latency are excluded as those metrics can vary among different datasets and applications. Table. I and II report the E-A FoM of the state-of-the-art neural prostheses along with their per-channel area and energy consumption. Only the power and area of the ML processor (i.e., feature extractor, classifier, and memory for parameter storage) have been considered. This FoM indicates that the tree-based models achieve orders of magnitude superior energy-area efficiency compared to SVM classifiers, while providing comparable classification accuracy and latency. With cost-aware hardware-algorithm co-design, we aim to improve the efficiency of tree-based classifiers even further, as discussed in Sections VI-VII.

The predictive power and hardware efficiency of different SoCs are greatly affected by their selection of ML algorithms. For example, DT-based ML models feature a lightweight inference scheme where we simply compare feature values to thresholds to route samples through the tree. On the contrary, the inference of kernelized SVM involves vector multiplications and the calculation of Gram matrix, which partially explains the E-A superiority of DTs over SVMs in Table. I. Moreover, inspired by the recent success of deep

learning algorithms, there is an increasing interest in deploying CNNs and DNNs on neural SoCs [18], [100], [101]. However, compared to conventional approaches, deep learning models generally require more training data and consume higher power consumption [113]. The benefits of using deep learning in neural SoCs need to be further investigated in the future.

G. Limitations of the Current SoCs and Future Directions

High-density electrode arrays have shown promise in both neurophysiological monitoring [48] and therapeutic neurostimulation [114]. However, the channel count of state-of-the-art ML SoCs is limited to 32, which could hinder their clinical application. The most critical challenges to realizing high-channel-count ML-embedded neural prostheses lie in the AFE, the back-end signal processing, and the memory for parameter storage. Over the past years, the field has witnessed a growth of channel count in neural prostheses, such as Neuralink's BMI platform with 3072 channels [115]. Recently, a 1024-channel closed-loop BMI SoC was presented with a successful demonstration of motor intention decoding (performed offline) in a macaque monkey [116]. Novel area- and power-efficient AFE design techniques (such as mixed-signal [117] and time-division multiplexing [118], [119]) should continue to be explored. This will enable advanced neural prostheses with high resolution, reduced invasiveness, longer lifetime, and minimized heat-induced tissue damage. In addition to area-power constraints on the AFE, the burden of the back-end signal processing (i.e., feature extraction and classification) is a major bottleneck to next-generation high-channel-count prostheses. The amount of computation in the current ML SoCs grows linearly with channel count, posing a major challenge on the energy consumption. The on-demand feature computation scheme in [12] could be a viable solution to realizing a scalable ML SoC. Only relevant features from a subset of channels are computed in each processing window, achieving a substantial reduction in hardware cost. Similar techniques will pave the way for the integration of novel high-density electrodes (Section II. A) in future diagnostic and closed-loop devices. Another on-demand processing approach was adopted in an SVM-based two-level (coarse/fine) classifier to reduce the system power consumption [89]. Exploiting the sparseness of seizures, the otherwise power-demanding SVM classifier (fine) in a separate chip is only activated upon seizure declaration by the coarse detector. The two-level SVM classifier performs 266 classifications/hour with 1.16 μW average power, improving $>135\times$ over the conventional SVM. Single-chip integration and multi-channel scalability have yet to be addressed with this approach.

In addition, most current classifiers integrated on neural prostheses use an offline training scheme with fixed parameters, thus neglecting the non-stationary dynamics of neural signals. The next generation ML-embedded neural SoCs are expected to perform active, incremental learning to account for the previously unseen changes in neurological patterns. In online machine learning, the model parameters are updated with the sequential arrival of data, thus dynamically adapting to new signal patterns. Online learning algorithms have shown promise in stable chronic neural decoding [120]. Yet, the deployment of such models on ASIC with minimal area and power consumption remains an open direction. Current on-chip systems based on SVM [91], [121] are highly energy and memory demanding, while off-

chip recalibrations pose security risks and reduce patient independence. More developments in this area are expected in near future.

The ML-embedded neural prostheses, like other edge AI devices in IoT and healthcare, may greatly benefit from developments in algorithm and circuit design that could lead to higher performance, lower energy and more compact area. For instance, future ML SoCs are expected to benefit from emerging techniques in CMOS design such as analog, mixed-signal [122], and approximate computing, as well as in-memory computing techniques. Particularly, in-memory computing has shown the potential to achieve remarkable improvements over conventional digital implementations [32], [38], [123]. Compared to current SoCs, neuromorphic hardware integrates spiking neural networks (SNN) and in-memory computing to avoid the communication overhead between processors and memory, and allows unsupervised online learning via Spike Timing Dependent Plasticity (STDP). As discussed in [38], currently the memristor-based designs are rarely used in the biomedical domain. Moreover, it should be noted that the decoding performance of SNN is relatively low due to the lack of maturity of the training algorithms [38]. Deploying high performance SNN and memristor-based designs in neural prostheses remains as a future direction.

V. Hardware-Algorithm Co-Design of Decision Tree Ensembles

Designing machine learning models that consume little energy and area, while providing a high classification accuracy and low detection latency is essential to the next-generation smart neural prostheses. As discussed in Section IV, decision trees are widely used in edge applications and neural decoding tasks thanks to their low inference complexity, easy and fast training, as well as high predictive power in ensemble methods or oblique structures [12], [22], [24], [29], [31]. These advantages are essential for extremely resource-constrained platforms such as a brain implant or wearable with high channel counts. In this section, we present novel approaches to optimize the key design metrics of an on-chip DT ensemble, including the power consumption and processing latency, in the context of neural signal classification tasks. Some of these techniques are broadly applicable to other machine learning algorithms for various implantable and edge applications.

A. Depth-Variant Tree Ensemble for Latency Reduction

In a decision tree, test sample traverses a single root-to-leaf path during inference [31], [124]. Despite being lightweight and area-efficient, the single-path scheme requires conditional computation and evaluates nodes in a sequential order [125]. As a result, DT-based classifiers impose a latency that increases proportionally with the decision path length. However, early symptom detection is critical to effectively treat neurological disorders, and it is directly affected by the processing latency. Previous work reduced seizure detection latency by either using shorter windows [126] or replacing the widely used bandpower biomarkers with new features such as neuronal potential similarity [127]. However, such methods may suffer from a degraded classification performance (since low-frequency features that require a longer window could be critical in symptom detection [126]) or poor generalizability due to the use of specific biomarkers [127]. To the best of our knowledge, this study is the first to address latency reduction from an algorithmic perspective.

Tree ensembles have shown promise in various neural classification tasks [2], [12], [24], [31]. However, conventional ensembles impose a uniform maximum-depth constraint on all base-estimators in the ensemble, such that the system latency is similar across different trees. In this work, we propose the Depth-Variant Tree Ensemble (DVTE), a novel low-latency variation of conventional ensemble methods. As shown in Fig. 6, DVTE consists of decision trees with different maximum depths, resulting in non-uniform latencies across trees. In a DVTE, shallow trees perform fast inference to reduce system latency, while deep trees are trained to compensate for misclassification errors caused by shallow trees.

We trained the proposed DVTE model using the popular gradient boosting framework [98], [128]. In the first two boosting rounds, we initialized the ensemble with decision stumps (i.e., decision trees with a single internal node). In the third and fourth rounds, two DTs with a max depth of two were trained to compensate for the residual errors from previous rounds. Deeper trees were gradually added to DVTE in later boosting rounds to better fit on training data. During inference, all decision trees in a DVTE run freely in parallel, with no need for synchronization. Therefore, shallow trees can update the decision outcome more frequently than deeper trees. If the current inference in a deep tree is incomplete (i.e., test samples have not yet reached the leaf nodes), we used the most recent output of that tree. The final prediction of DVTE is calculated as the sum of the outputs of all trees in the ensemble, which can be updated at the same rate as the shortest tree (i.e., $d=1$). While shallow trees make predictions with low latency (trees of $d=1$ in Fig. 7), they often have a limited predictive power and may not fit well on training data. To tackle this problem and achieve the best trade-off between latency and classification accuracy, DVTE incorporates deeper trees in the gradient boosting framework to reduce bias.

Unlike DVTE which effectively combines shallow and deep trees in the gradient boosting ensemble to jointly optimize the latency and accuracy, previous work either used a few deep trees (e.g., 8 trees with a max depth of 4 [12]) with potential latency concerns as discussed above, or implemented a large number of shallow trees (1024 decision stumps in [29]), requiring many parallel feature processing units. The aim of DVTE is to benefit from both shallow and deep trees and enable low-latency inference with a small tree ensemble. This is particularly critical in time-sensitive classification tasks such as PD tremor detection with strict latency requirements.

As an example, we built a DVTE with 8 trees and various depths from 1 to 4 (Fig. 6). This model was benchmarked against conventional ensemble (8 trees, max depth: 4 [12]). We used a learning rate of 0.3 for both models and implemented them using the lightGBM library in Python [128]. We tested our classifier on epileptic seizure detection using iEEG recordings (11 patients, 106 annotated seizures over 1255 hours). The number of channels varied from 47 to 128. This dataset can be accessed via iEEG portal [102]. Handcrafted features were extracted over various window lengths as detailed in Table III. It should be noted that both EEG and iEEG have been widely used in on-chip seizure detectors [11]-[14], [16], [17], [29], [88], [89]. However, iEEG is more commonly used in closed-loop prostheses, as it can be easily combined with invasive neuromodulation techniques for improved symptom control [16], [17], [29], [88], and it has been used in our study.

Figure 8 compares the proposed DVTE and the conventional ensemble method in terms of classification performance (sensitivity, specificity) and latency. The performance was evaluated using bit-accurate classifier models in MATLAB and Python. We estimated the processing latency by calculating the average time to traverse a root-to-leaf decision path in the trees. Compared to the conventional ensemble, DVTE caused a marginal performance reduction (<3% in sensitivity and <1% in specificity). On the other hand, DVTE achieved an average latency of 0.86s, significantly lower than the latency of a conventional ensemble (2.12s, 2.5× reduction).

B. Cost-Aware Learning for Latency and Power Reduction

The inference phase of tree-based models is relatively simple. In axis-aligned decision trees, we compare a feature value to a threshold in order to select the child node at each internal node. The leaf node contains a constant weight indicating the prediction result. Given the lightweight inference in tree-based models, the hardware cost (e.g., power, latency) is largely affected by the feature extraction process [12].

Table III summarizes the biomarkers used in our seizure detection task and their power and latency cost. We implemented digital feature extraction hardware in a TSMC 65 nm LP process using Synopsys Design Compiler and Cadence Innovus. The power cost of each feature was simulated under a 1.2-V supply using Synopsys PrimeTime. Line-length, a widely used feature in epilepsy studies, is hardware-friendly and low-power. Bandpower features, on the other hand, consume higher power since they require an FIR filtering stage. The latency associated with a feature depends on the window size used to compute that feature. Long windows were used to extract low-frequency bandpower, while short windows were used for time-domain features and high-frequency bandpowers. Specifically, we used 1s windows to extract Delta (δ), 0.5s for Theta (θ) and Alpha (α), and 0.25s for other features.

We apply the cost-aware learning approach to tree-based classifiers (e.g., DVTE) to reduce the inference hardware cost. Specifically, we use the total power consumption and latency along the decision path as a regularization term in the training process. The training of cost-aware decision trees attempts to minimize the following expression:

$$\min \sum_i L(y_i, f(\mathbf{x}_i)) + C(\Psi_{pow}(f, \mathbf{x}_i) + \Psi_{lat}(f, \mathbf{x}_i)), \quad (2)$$

where $L(y_i, f(\mathbf{x}_i))$ is the loss function that measures the misclassification error as the difference between groundtruth y_i and prediction $f(\mathbf{x}_i)$, Ψ_{pow} and Ψ_{lat} indicate the estimated power consumption and latency along the decision path, respectively, and C is the regularization coefficient that determines the trade-off between hardware cost and performance. The effect of varying C on latency and power in DVTE is shown in Fig. 9. For a greater regularization coefficient, cost-aware decision trees achieve a lower hardware cost. Since power and latency span over different ranges, we standardized the cost by removing the mean value and normalizing both power and latency to their unit variance.

We applied the cost-aware inference approach to DVTE to reduce both power and latency on seizure detection task. Figure 10 shows the classification performance (sensitivity, specificity) as a function of the cost metrics (latency, power). We adjusted the regularization coefficient C to achieve different trade-offs between power/latency and performance. For both the low-power (Fig. 10(a)) and low-latency (Fig. 10(b)) DVTEs, the best trade-off is observed at a point where a maximum reduction in latency or power can be achieved with only a marginal performance loss.

Figure 11 shows the number of extracted features for the cost-aware DVTE. The number of feature extractions are normalized to each 0.25s window. Thus, the normalized feature count is upper bounded by the number of trees. For $C > 0$, we used the hardware cost to regularize the model and as a result, DVTE was trained to minimize the inference power and latency. As the regularization coefficient increases, the model further penalizes inefficient features. With $C = 0.01$, we achieved the best trade-off between performance and hardware cost (Fig. 10), reducing the power by $3\times$ and latency by $1.7\times$ compared to DVTE without cost-aware learning.

C. Hardware Implementation of DVTE Classifier

We implemented the DVTE classifier in hardware to demonstrate the efficacy of the proposed cost-aware learning approach. Figure 12(a) presents the system architecture of the DVTE classifier, which supports 32-channel 500-S/s 10-bit input data. Each of the 8 decision trees consists of a feature extraction unit (FEU), a comparator, and a tree control unit (TCU). A 32-tap programmable FIR bandpass filter was implemented to extract the bandpower feature in a selected frequency band, and a single FIR coefficient memory was shared between 8 trees. The FEU extracts only one feature during each window, which allows us to clock- and data-gate unused feature blocks for dynamic power saving. The extracted feature is then compared to a threshold to decide the decision path in the tree. The TCU reads the trained tree information (i.e., feature type, channel index, threshold, and leaf value) from memory and controls the FEU based on the current node information and comparison result. When a leaf node is reached, the tree sends out a leaf value and repeats the process starting from the root node. Leaf values from the 8 trees are summed to make a final decision. The proposed lightweight DVTE classifier utilizes a 0.4kB on-chip memory.

The DVTE classifier was implemented in a TSMC 65 nm 1P9M LP process. Figures 12(b) and (c) show the chip layout occupying 0.31 mm^2 and its area breakdown, respectively. Using Synopsys PrimeTime, the power consumption of the system was simulated at $2.8 \mu\text{W}$ under a 1.2-V supply. The parallel implementation of 8 trees allowed a low system clock (500 Hz). In addition, the use of high-Vt transistors saved both dynamic and static power consumption. The energy efficiency and E-A FoM of the DVTE classifier are 5.6 nJ/class. and $1.7 \text{ pJ}\cdot\text{mm}^2/\text{S}$, respectively, achieving $>6.4\times$ and $>23.7\times$ improvements over the state-of-the-art designs in Table. I.

In this advanced technology node with a low operating frequency and efficient clock- and data-gating, the static power consumption acts as the dominant source of power, as indicated in the breakdown of Fig. 12(d). Here, 83.8% of system power is consumed by leakage currents in the ensemble. Therefore, power-gating of the unused feature extraction blocks

can further improve the energy efficiency of the proposed cost-aware DVTE classifier. This is possible thanks to the on-demand feature extraction scheme of DVTE. To estimate the potential power savings, we performed post-layout simulations for each feature extraction block with power-gating header switches [129]. The results showed that the static power consumption of each feature substantially reduced to 30 pW with the supply power gated. For the best trade-off case in Fig. 11 ($C = 0.01$) with power-gating applied, the overall system power is estimated to be 0.68 μ W. It is our ongoing work to implement the power-gating technique reliably at the system level with a minimal area overhead, to potentially achieve sub- μ W total power consumption.

VI. Hardware-Algorithm Co-Design of Oblique Trees

In the previous section, we proposed a novel tree ensemble, DVTE, and a cost-aware learning approach to improve latency and power. However, tree ensembles may require a large number of axis-aligned DTs for non-trivial classification tasks [29], [63], [130], resulting in a large model size and on-chip memory. Different from conventional trees that use axis-aligned decision boundaries, oblique trees calculate a weighted sum of multiple features and compare the result to a threshold [31]. Thanks to their powerful split functions, oblique trees are capable of generating accurate predictions using a single tree with a reduced model size. Moreover, in our previous work, we built a new class of oblique trees that are compatible with model compression techniques to further reduce the hardware complexity and memory needs [31]. In this section, we present the hardware-algorithm co-design of oblique trees to simultaneously achieve low power consumption, low latency and small model size.

We built oblique DTs using a probabilistic routing scheme [31], where the i -th internal node sends samples to a child according to the probabilistic distribution, as follows

$$P_i(\mathbf{x}_n) = \text{softmax}(\mathbf{x}_n^\top \boldsymbol{\theta}_i), \quad (3)$$

where \mathbf{x}_n indicates the feature vector and $\boldsymbol{\theta}_i$ is the trainable weight vector of the same shape as \mathbf{x}_n . The *softmax* function normalizes the output space into a probability distribution within (0,1) interval. Here, \mathbf{x}_n visits the left child with a probability of $P_i(\mathbf{x}_n)$ and the right child with $1 - P_i(\mathbf{x}_n)$. In the probabilistic routing scheme, samples arrive at multiple leaf nodes with different probabilities and the final prediction is given by

$$\hat{y}_n = \sum_l P_l(\mathbf{x}_n) \omega_l, \quad (4)$$

where $P_l(\mathbf{x}_n)$ indicates the probability of sample \mathbf{x}_n visiting the leaf node l and ω_l is the constant leaf predictor. For classification tasks, we measured the cross-entropy loss $\sum_n L(\hat{y}_n, y_n)$ using the groundtruth (y_n) and the prediction of the oblique tree \hat{y}_n .

A. Model Compression and Cost-Aware Learning

Various compression techniques have been applied to DNNs, including fixed-point quantization [131], weight pruning and sharing [132]. Interestingly, within the probabilistic

training scheme, oblique trees are compatible with gradient descent-based optimization, similar to the training of a neural network. Therefore, we propose to combine oblique trees with DNN-based compression techniques to reduce model size and hardware cost. We trained the oblique tree by minimizing the loss $\sum_n L(\hat{y}_n, y_n)$ on training data. During the training process, we applied weight pruning to slim the oblique tree and weight sharing to further reduce model size. Specifically, we used a simple neural network with input and output layers to represent the oblique decision functions in the internal nodes. Weight pruning/sharing were applied to 2-layer NNs for creating sparse connections and reducing the model size. We pruned the oblique tree by iteratively setting small values to zero and retraining the remaining parameters. For weight sharing, we uniformly clustered the weights into k shared values, requiring only $\lceil \log_2 k \rceil$ bits to store the index. It should be noted that oblique trees are compatible with the aforementioned cost-aware learning framework, by simply replacing the loss function in Eq. 2 with oblique tree training loss. In cost-aware learning, oblique trees assign smaller weights to costly features so that they hardly survive the pruning process.

We compared the hardware efficiency of oblique trees against axis-aligned tree ensembles. Specifically, we built resource-efficient oblique trees (ResOT) [21] by combining cost-aware learning with model compression. We used the conventional lightGBM ensemble [128] and gradient boosting with power-efficient training (PEGB [105]) as baseline. In addition to seizure detection, we tested our model on LFPs recorded from 12 PD patients via DBS leads (3-channel, 2048 Hz sample rate, 16 recordings) to detect the tremor onset [2]. For both tasks, a single ResOT was built (max depth: 4) with 16 shared weights (4 bits). Hyperparameters of oblique trees including the number of parameters post pruning and the regularization coefficient were optimized for each patient. We used 5-fold chronological cross-validation to measure the F1 score, and leave-one-out for epilepsy patients with <5 seizures. Cross-validation has been widely used in previous studies [12], [16], [61]. It allows testing on multiple train-test splits to fairly assess the model performance on unseen data. Compared to the hold-out method [13], cross-validation is less dependent on a specific train-test split and could provide a reliable measure of performance for patients with few seizure events. We employed a block-wise data splitting method, where each block includes a complete seizure event and its neighbouring non-seizure period, to avoid information leakage during training [12]. Cross-validation was performed on pre-recorded data to estimate the model performance and optimize the hyperparameters. In a clinical setting, the final set of parameters (i.e., feature index, threshold, leaf weights, and feature weights for oblique trees) will be trained using the entire pre-recorded data of each patient and loaded to the chip to predict future seizure events. We estimated the memory requirements of various models using the size of the trainable weight matrix. Compressed sparse column and delta encoding were used to store the sparse matrix after weight pruning. For power comparison, we considered the power consumption for extracting features along the decision path during inference (Table III). In our simulations, ResOT achieved an average saving of 7.0 \times in model size and 10.7 \times in power cost compared to lightGBM, as shown in Fig. 13. It also outperformed the hardware-efficient ensemble (PEGB) (3.1 \times in model size and 3.0 \times in power cost).

The oblique node evaluation time is set by the longest feature computed in that node. Here, we pruned the oblique tree to use a maximum of 8 features per internal node to fairly compare it against DVTE. The hardware cost in an oblique tree (e.g., power, latency) can also benefit from the introduced cost-aware learning scheme, by including a cost regularization in the oblique tree objective. Figure 14 plots the distribution of features in ResOT on seizure detection task. The latency was reduced from 2.67s to ~1s via cost-aware training, while the power cost was reduced from 696nW to 241nW.

B. Parallel Node Evaluation for Latency Reduction

Previous work on oblique trees employed a single-path inference scheme by visiting the most probable path, which suffers from a latency proportional to the length of the decision path [31]. An alternative is to evaluate multiple nodes in parallel to reduce latency, as shown in Fig. 15. The single-path inference scheme is presented in Fig. 15(a), where 4 internal nodes are evaluated using consecutive windows. In Fig. 15(b), we evaluate two layers of the tree (3 nodes) per window, requiring 6 node evaluations in total. Finally, Fig. 15(c) evaluates all 15 nodes in parallel.

To demonstrate the trade-off between power consumption and latency, we built an oblique tree on seizure detection task. Different numbers of nodes were evaluated in parallel and the corresponding hardware cost is reported in Fig. 16. Single-path inference (Fig. 15(a)) obtained the lowest power and highest latency (power cost = 305nW, latency = 1.04s). On the other hand, concurrently evaluating all nodes in parallel reduced the latency by 2.1× but increased the power by 11.6× (Fig. 15(c)). The case of 3 nodes (Fig. 15(b)) achieved a better trade-off between power and latency, leading to 1.9× reduction in latency and 3× increase in power cost. This scheme can potentially be useful in latency-constrained applications.

C. Interpretable DTs for Neural Prostheses

Closed-loop stimulation is a safety-critical application, favoring an interpretable decision process. Another distinct advantage of tree-based models is their interpretability, in contrast to most classical machine learning and deep learning methods that lack transparency and interpretability. This is critical to understanding a specific therapeutic strategy for a particular neurological symptom or behavior. We can simply visualize the decision process of DTs and the informative biomarkers used in making predictions. Therefore, tree-based models are widely used in clinical applications that require high interpretability [133], [134].

For example, Fig.17(a) shows the contributions from time- and spectral-domain features in tremor detection task, using shapley additive explanations [135]. The feature values at the visualized window are shown on the left, and the red/blue colors represent features that indicate a high/low risk of tremor, respectively. The power over low beta and tremor bands are the most predictive features. The model predicts a tremor state according to the weighted contribution of all features.

Figure 17(b) visualizes the decision process of an oblique tree on seizure detection task. We used pie charts at internal and leaf nodes to represent the class distribution. Both seizure and non-seizure samples are mixed at the internal nodes, while each leaf node is dominated by either seizure or non-seizure samples. The decision process follows an explainable rule

list structure, with the left branch leading samples directly to a leaf node. The percentage of samples that travel through a node (internal or leaf) is shown next to that node. We also show the approximate power and latency to process each internal node. For comparison, Fig. 17(c) shows the decision process of a cost-aware oblique tree trained on the same patient. As shown in this figure, the power cost to evaluate the internal nodes is significantly reduced in the cost-aware approach. Particularly, the most notable reduction in power (i.e., node complexity) is observed at the root node, as it is the most frequently visited node in the tree. Moreover, the overall latency along the root-leaf path in Fig. 17(c) is shorter than that of Fig. 17(b), indicating a reduction of processing latency.

VII. Conclusion

In this paper, we reviewed the latest developments in closed-loop neural interface design, with a particular focus on system-on-chips that integrate machine learning for symptom detection. The current commercial and research-based closed-loop devices, advances in electrode and circuit design, and clinical applications were discussed. We reviewed various hardware approaches used to implement machine learning on neural prostheses, design trade-offs and hardware/performance comparisons. We further proposed a novel tree-based neural decoder, Depth-Variant Tree Ensemble, to reduce latency in neurological symptom detection. A cost-aware learning approach was applied to DVTE to further reduce power and latency. We also integrated various techniques, including cost-aware learning and model compression, to construct resource-efficient oblique trees. Testing on epileptic seizure and PD tremor detection tasks, the proposed model improved both power and latency, and reduced the memory requirement, while maintaining a high performance. We also discussed the interpretability of tree-based models, as an essential component for next-generation intelligent neural prostheses.

Acknowledgment

This work was partially supported by the National Institute of Mental Health Grant R01-MH-123634.

Biography



Bingzhao Zhu received the B.Sc. degree in Opto-Electronics Science and Engineering from Zhejiang University, Hangzhou, China, in 2017. He is currently pursuing the Ph.D. degree in Applied and Engineering Physics and a minor in Computer Science, at Cornell University, Ithaca, New York, USA. Since 2020, he is a visiting PhD student at Swiss Federal Institute of Technology (EPFL), Geneva, Switzerland. His research interests include brain-computer interfaces (BCI), low-power machine learning, neural signal processing, and computational imaging.



Uisub Shin received the B.E. and M.S. degrees in Electrical Engineering from Chungnam National University and Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea in 2015 and 2017, respectively. Since 2018, he has been pursuing the Ph.D. degree in Electrical and Computer Engineering at Cornell University, Ithaca, NY, USA. He is currently a visiting Ph.D. student at Swiss Federal Institute of Technology (EPFL), Geneva, Switzerland. His research interests include low-power mixed-signal IC design for biomedical applications and embedded machine learning.



Mahsa Shoaran received the B.Sc. and M.Sc. degrees from Sharif University of Technology in 2008 and 2010, respectively, and the Ph.D. degree in Electrical Engineering from Swiss Federal Institute of Technology (EPFL) in 2015. She was a Postdoctoral Fellow in Electrical Engineering and Medical Engineering at the California Institute of Technology from 2015 to 2017. She is currently an Assistant Professor in the Center for Neuroprosthetics and Electrical Engineering Institute of EPFL and director of the Integrated Neurotechnologies Laboratory. From 2017 to 2019, she was an Assistant Professor at the School of Electrical and Computer Engineering at Cornell University, Ithaca, NY. Mahsa is a recipient of the 2018 Google Faculty Research Award in Machine Learning, the Early and Advanced Swiss NSF Postdoctoral Fellowships, and the NSF Award for Young Professionals Contributing to Smart and Connected Health. She was named a Rising Star in EE/CS by MIT in 2015. Her research interests include low-power circuit and system design for biomedical applications, neural and brain-machine interfaces, machine learning hardware, and neuromodulation therapies for neurological disorders. She serves on the technical program committee of IEEE CICC, as track chair in Bio and Medical Electronics for IEEE ICECS, and review committee member for IEEE BioCAS.

References

- [1]. Little S, Pogosyan A, Neal S, Zavala B, Zrinzo L, Hariz M, Foltynie T, Limousin P, Ashkan K, FitzGerald J et al. , “Adaptive deep brain stimulation in advanced parkinson disease,” *Annals of neurology*, vol. 74, no. 3, pp. 449–457, 2013. [PubMed: 23852650]
- [2]. Yao L, Brown P, and Shoaran M, “Improved detection of Parkinsonian resting tremor with feature engineering and Kalman filtering,” *Clinical Neurophysiology*, vol. 131, no. 1, pp. 274–284, 2020. [PubMed: 31744673]
- [3]. Lo M-C and Widge AS, “Closed-loop neuromodulation systems: next-generation treatments for psychiatric illness,” *International review of psychiatry*, vol. 29, no. 2, pp. 191–204, 2017. [PubMed: 28523978]

- [4]. Ezzyat Y, Wanda PA, Levy DF, Kadel A, Aka A, Pedisich I, Sperling MR, Sharan AD, Lega BC, Burks A et al. , “Closed-loop stimulation of temporal cortex rescues functional networks and improves memory,” *Nature Communications*, vol. 9, no. 1, pp. 1–8, 2018.
- [5]. Iturrate I, Pereira M, and Millán J. d. R., “Closed-loop electrical neurostimulation: challenges and opportunities,” *Current Opinion in Biomedical Engineering*, vol. 8, pp. 28–37, 2018.
- [6]. Huang M, Harvey RL, Ellen Stoykov M, Ruland S, Weinand M, Lowry D, and Levy R, “Cortical stimulation for upper limb recovery following ischemic stroke: a small phase ii pilot study of a fully implanted stimulator,” *Topics in stroke rehabilitation*, vol. 15, no. 2, pp. 160–172, 2008. [PubMed: 18430685]
- [7]. Lozano AM, Lipsman N, Bergman H, Brown P, Chabardes S, Chang JW, Matthews K, McIntyre CC, Schlaepfer TE, Schulder M et al. , “Deep brain stimulation: current challenges and future directions,” *Nature Reviews Neurology*, vol. 15, no. 3, pp. 148–160, 2019. [PubMed: 30683913]
- [8]. Krauss JK, Lipsman N, Aziz T, Boutet A, Brown P, Chang JW, Davidson B, Grill WM, Hariz MI, Horn A et al. , “Technology of deep brain stimulation: current status and future directions,” *Nature Reviews Neurology*, pp. 1–13, 2020. [PubMed: 31827266]
- [9]. Skarpaas TL and Morrell MJ, “Intracranial stimulation therapy for epilepsy,” *Neurotherapeutics*, vol. 6, no. 2, pp. 238–243, 2009. [PubMed: 19332315]
- [10]. Meidahl AC, Tinkhauser G, Herz DM, Cagnan H, Debarros J, and Brown P, “Adaptive deep brain stimulation for movement disorders: the long road to clinical therapy,” *Movement disorders*, vol. 32, no. 6, pp. 810–819, 2017. [PubMed: 28597557]
- [11]. Yoo J, Yan L, El-Damak D, Altaf MAB, Shoeb AH, and Chandrakasan AP, “An 8-channel scalable EEG acquisition SoC with patient-specific seizure classification and recording processor,” *IEEE journal of solid-state circuits*, vol. 48, no. 1, pp. 214–228, 2013.
- [12]. Shoaran M, Haghgi BA, Taghavi M, Farivar M, and Emami-Neyestanak A, “Energy-efficient classification for resource-constrained biomedical applications,” *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 8, no. 4, pp. 693–707, 2018.
- [13]. Altaf MAB, Zhang C, and Yoo J, “A 16-channel patient-specific seizure onset and termination detection SoC with impedance-adaptive transcranial electrical stimulator,” *IEEE Journal of Solid-State Circuits*, vol. 50, no. 11, pp. 2728–2740, 2015.
- [14]. Altaf MAB and Yoo J, “A 1.83 uJ/classification, 8-channel, patient-specific epileptic seizure classification soc using a non-linear support vector machine,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 10, no. 1, pp. 49–60, 2015. [PubMed: 25700471]
- [15]. Kassiri H, Tonekaboni S, Salam MT, Soltani N, Abdelhalim K, Velazquez JLP, and Genov R, “Closed-loop neurostimulators: A survey and a seizure-predicting design example for intractable epilepsy treatment,” *IEEE transactions on biomedical circuits and systems*, vol. 11, no. 5, pp. 1026–1040, 2017. [PubMed: 28715338]
- [16]. O’Leary G, Groppe DM, Valiante TA, Verma N, and Genov R, “Nurip: Neural interface processor for brain-state classification and programmable-waveform neurostimulation,” *IEEE Journal of Solid-State Circuits*, vol. 53, no. 11, pp. 3150–3162, 2018.
- [17]. Chen W-M, Chiueh H, Chen T-J, Ho C-L, Jeng C, Ker M-D, Lin C-Y, Huang Y-C, Chou C-W, Fan T-Y et al. , “A fully integrated 8-channel closed-loop neural-prosthetic CMOS SoC for real-time epileptic seizure control,” *IEEE Journal of Solid-State Circuits*, vol. 49, no. 1, pp. 232–247, 2014.
- [18]. Taufique Z, Zhu B, Coppola G, Shoaran M, and Altaf MAB, “A low power multi-class migraine detection processor based on somatosensory evoked potentials,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 68, no. 5, pp. 1720–1724, 2021.
- [19]. Verma N, Shoeb A, Bohorquez J, Dawson J, Gutttag J, and Chandrakasan AP, “A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system,” *IEEE journal of solid-state circuits*, vol. 45, no. 4, pp. 804–816, 2010.
- [20]. Zhang F, Aghagolzadeh M, and Oweiss K, “An implantable neuroprocessor for multichannel compressive neural recording and on-the-fly spike sorting with wireless telemetry,” in *2010 Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2010, pp. 1–4.
- [21]. Zhan T, Fatmi SZ, Guraya S, and Kassiri H, “A resource-optimized VLSI implementation of a patient-specific seizure detection algorithm on a custom-made 2.2 cm² wireless device for

- ambulatory epilepsy diagnostics,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 13, no. 6, pp. 1175–1185, 2019. [PubMed: 31634843]
- [22]. Zhu B, Shin U, and Shoaran M, “Closed-loop neural interfaces with embedded machine learning,” in *2020 27th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*. IEEE, 2020, pp. 1–4.
- [23]. He S, Baig F, Mostofi A, Pogosyan A, Debarros J, Green AL, Aziz TZ, Pereira E, Brown P, and Tan H, “Closed-loop deep brain stimulation for essential tremor based on thalamic local field potentials,” *Movement Disorders*, vol. 36, no. 4, pp. 863–873, 2021. [PubMed: 33547859]
- [24]. Kuhlmann L, Karoly P, Freestone DR, Brinkmann BH, Temko A, Barachant A, Li F, Titericz G Jr, Lang BW, Lavery D et al. , “[Epilepsyecosystem.org](https://epilepsyecosystem.org): crowd-sourcing reproducible seizure prediction with long-term human intracranial EEG,” *Brain*, vol. 141, no. 9, pp. 2619–2630, 2018. [PubMed: 30101347]
- [25]. Truong ND, Zhou L, and Kavehei O, “Semi-supervised seizure prediction with generative adversarial networks,” in *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2019, pp. 2369–2372.
- [26]. DiLorenzo DJ, Leyde KW, and Kaplan D, “Neural state monitoring in the treatment of epilepsy: Seizure prediction—conceptualization to first-in-man study,” *Brain sciences*, vol. 9, no. 7, p. 156, 2019.
- [27]. Maiwald T, Winterhalder M, Aschenbrenner-Scheibe R, Voss HU, Schulze-Bonhage A, and Timmer J, “Comparison of three nonlinear seizure prediction methods by means of the seizure prediction characteristic,” *Physica D: nonlinear phenomena*, vol. 194, no. 3–4, pp. 357–368, 2004.
- [28]. Jouny CC, Franaszczuk PJ, and Bergey GK, “Improving early seizure detection,” *Epilepsy & Behavior*, vol. 22, pp. S44–S48, 2011. [PubMed: 22078518]
- [29]. O’Leary G, Xu J, Long L, Sales Filho J, Tejeiro C, ElAnsary M, Tang C, Moradi H, Shah P, Valiante TA et al., “A neuromorphic multiplier-less bit-serial weight-memory-optimized 1024-tree brain-state classifier and neuromodulation SoC with an 8-channel noise-shaping SAR ADC array,” in *2020 ISSCC*. IEEE, 2020, pp. 402–404.
- [30]. Yang Y, Boling S, and Mason AJ, “A hardware-efficient scalable spike sorting neural signal processor module for implantable high-channel-count brain machine interfaces,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 11, no. 4, pp. 743–754, 2017. [PubMed: 28541908]
- [31]. Zhu B, Farivar M, and Shoaran M, “Resot: Resource-efficient oblique trees for neural signal classification,” *IEEE Transactions on Biomedical Circuits and Systems*, 2020.
- [32]. Chen Y, Yao E, and Basu A, “A 128-channel extreme learning machine-based neural decoder for brain machine interfaces,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 10, no. 3, pp. 679–692, 2016. [PubMed: 26672048]
- [33]. Shaikh S, So R, Sibindi T, Libedinsky C, and Basu A, “Towards intelligent intracortical BMI (i^2 BMI): Low-power neuromorphic decoders that outperform Kalman filters,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 13, no. 6, pp. 1615–1624, 2019. [PubMed: 31581098]
- [34]. Do AT, Zeinolabedin SMA, Jeon D, Sylvester D, and Kim TT, “An area-efficient 128-channel spike sorting processor for real-time neural recording with $0.175 \mu\text{W}/\text{channel}$ in 65-nm CMOS,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 1, pp. 126–137, 2019.
- [35]. Chae MS, Yang Z, Yuce MR, Hoang L, and Liu W, “A 128-channel 6 mw wireless neural recording ic with spike feature extraction and uwb transmitter,” *IEEE transactions on neural systems and rehabilitation engineering*, vol. 17, no. 4, pp. 312–321, 2009. [PubMed: 19435684]
- [36]. Navarro X, Krueger TB, Lago N, Micera S, Stieglitz T, and Dario P, “A critical review of interfaces with the peripheral nervous system for the control of neuroprostheses and hybrid bionic systems,” *Journal of the Peripheral Nervous System*, vol. 10, no. 3, pp. 229–258, 2005. [PubMed: 16221284]
- [37]. Courtine G and Sofroniew MV, “Spinal cord repair: advances in biology and technology,” *Nature medicine*, vol. 25, no. 6, pp. 898–908, 2019.

- [38]. Rahimiazghadi M, Lammie C, Eshraghian JK, Payvand M, Donati E, Linares-Barranco B, and Indiveri G, "Hardware implementation of deep network accelerators towards healthcare and biomedical applications," *IEEE Transactions on Biomedical Circuits and Systems*, 2020.
- [39]. Duun-Henriksen J, Baud M, Richardson MP, Cook M, Kouvas G, Heasman JM, Friedman D, Peltola J, Zibrandtsen IC, and Kjaer TW, "A new era in electroencephalographic monitoring? subscalp devices for ultra-long-term recordings," *Epilepsia*, vol. 61, no. 9, pp. 1805–1817, 2020. [PubMed: 32852091]
- [40]. Parvizi J and Kastner S, "Promises and limitations of human intracranial electroencephalography," *Nature neuroscience*, vol. 21, no. 4, pp. 474–483, 2018. [PubMed: 29507407]
- [41]. Herff C, Krusienski DJ, and Kubben P, "The potential of stereotactic-EEG for brain-computer interfaces: current progress and future directions," *Frontiers in neuroscience*, vol. 14, p. 123, 2020. [PubMed: 32174810]
- [42]. Anderson DN, Osting B, Vorwerk J, Dorval AD, and Butson CR, "Optimized programming algorithm for cylindrical and directional deep brain stimulation electrodes," *Journal of neural engineering*, vol. 15, no. 2, p. 026005, 2018. [PubMed: 29235446]
- [43]. Morrell MJ, "Responsive cortical stimulation for the treatment of medically intractable partial epilepsy," *Neurology*, vol. 77, no. 13, pp. 1295–1304, 2011. [PubMed: 21917777]
- [44]. Rosin B, Slovik M, Mitelman R, Rivlin-Etzion M, Haber SN, Israel Z, Vaadia E, and Bergman H, "Closed-loop deep brain stimulation is superior in ameliorating parkinsonism," *Neuron*, vol. 72, no. 2, pp. 370–384, 2011. [PubMed: 22017994]
- [45]. Arlotti M, Rosa M, Marceglia S, Barbieri S, and Priori A, "The adaptive deep brain stimulation challenge," *Parkinsonism & related disorders*, vol. 28, pp. 12–17, 2016. [PubMed: 27079257]
- [46]. Mullin JP, Shriver M, Alomar S, Najm I, Bulacio J, Chauvel P, and Gonzalez-Martinez J, "Is SEEG safe? a systematic review and meta-analysis of stereo-electroencephalography-related complications," *Epilepsia*, vol. 57, no. 3, pp. 386–401, 2016. [PubMed: 26899389]
- [47]. Viventi J, Kim D-H, Vigeland L, Frechette ES, Blanco JA, Kim Y-S, Avrin AE, Tiruvadi VR, Hwang S-W, Vanleer AC et al. , "Flexible, foldable, actively multiplexed, high-density electrode array for mapping brain activity in vivo," *Nature neuroscience*, vol. 14, no. 12, p. 1599, 2011. [PubMed: 22081157]
- [48]. Hermiz J, Rogers N, Kaestner E, Ganji M, Cleary DR, Carter BS, Barba D, Dayeh SA, Halgren E, and Gilja V, "Sub-millimeter ecog pitch in human enables higher fidelity cognitive neural state estimation," *NeuroImage*, vol. 176, pp. 454–464, 2018. [PubMed: 29678760]
- [49]. Stead M, Bower M, Brinkmann BH, Lee K, Marsh WR, Meyer FB, Litt B, Van Gompel J, and Worrell GA, "Microseizures and the spatiotemporal scales of human partial epilepsy," *Brain*, vol. 133, no. 9, pp. 2789–2797, 2010. [PubMed: 20685804]
- [50]. Shoaran M, Kamal MH, Pollo C, Vanderghenst P, and Schmid A, "Compact low-power cortical recording architecture for compressive multichannel data acquisition," *IEEE transactions on biomedical circuits and systems*, vol. 8, no. 6, pp. 857–870, 2014. [PubMed: 24723633]
- [51]. Shoaran M, Pollo C, Leblebici Y, and Schmid A, "Design techniques and analysis of high-resolution neural recording systems targeting epilepsy focus localization," in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Ieee*, 2012, pp. 5150–5153.
- [52]. Baker JL, Ryou J-W, Wei XF, Butson CR, Schiff ND, and Purpura KP, "Robust modulation of arousal regulation, performance, and frontostriatal activity through central thalamic deep brain stimulation in healthy nonhuman primates," *Journal of neurophysiology*, vol. 116, no. 5, pp. 2383–2404, 2016. [PubMed: 27582298]
- [53]. Guggenmos DJ, Azin M, Barbay S, Mahnken JD, Dunham C, Mohseni P, and Nudo RJ, "Restoration of function after brain damage using a neural prosthesis," *Proceedings of the National Academy of Sciences*, vol. 110, no. 52, pp. 21 177–21 182, 2013.
- [54]. DeMichele G and Troyk P, "Stimulus-resistant neural recording amplifier," in *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 4. IEEE, 2003, pp. 3329–3332.

- [55]. Blum RA, Ross JD, Brown EA, and DeWeerth SP, "An integrated system for simultaneous, multichannel neuronal stimulation and recording," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 54, no. 12, pp. 2608–2618, 2007.
- [56]. Johnson BC, Gambini S, Izyumin I, Moin A, Zhou A, Alexandrov G, Santacruz SR, Rabaey JM, Carmena JM, and Muller R, "An implantable 700 μ w 64-channel neuromodulation ic for simultaneous recording and stimulation with rapid artifact recovery," in *2017 Symposium on VLSI Circuits*. IEEE, 2017, pp. C48–C49.
- [57]. Rozgi D, Hokyhikyan V, Jiang W, Akita I, Basir-Kazeruni S, Chandrakumar H, and Markovi D, "A 0.338 cm³, artifact-free, 64-contact neuromodulation platform for simultaneous stimulation and sensing," *IEEE transactions on biomedical circuits and systems*, vol. 13, no. 1, pp. 38–55, 2018. [PubMed: 30582551]
- [58]. Chandrakumar H and Markovi D, "An 80-mvpp linear-input range, 1.6-g-ohm input impedance, low-power chopper amplifier for closed-loop neural recording that is tolerant to 650-mvpp common-mode interference," *IEEE journal of solid-state circuits*, vol. 52, no. 11, pp. 2811–2828, 2017.
- [59]. Zhou A, Santacruz SR, Johnson BC, Alexandrov G, Moin A, Burghardt FL, Rabaey JM, Carmena JM, and Muller R, "A wireless and artefact-free 128-channel neuromodulation device for closed-loop stimulation and recording in non-human primates," *Nature biomedical engineering*, vol. 3, no. 1, pp. 15–26, 2019.
- [60]. Mendrela AE, Cho J, Fredenburg JA, Nagaraj V, Netoff TI, Flynn MP, and Yoon E, "A bidirectional neural interface circuit with active stimulation artifact cancellation and cross-channel common-mode noise suppression," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 4, pp. 955–965, 2016.
- [61]. Shoeb AH and Gutttag JV, "Application of machine learning to epileptic seizure detection," in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 975–982.
- [62]. Zhang Z and Parhi KK, "Seizure detection using wavelet decomposition of the prediction error signal from a single channel of intracranial EEG," in *2014 36th annual international conference of the IEEE engineering in medicine and biology society*. IEEE, 2014, pp. 4443–4446.
- [63]. Zhu B, Coppola G, and Shoaran M, "Migraine classification using somatosensory evoked potentials," *Cephalalgia*, vol. 39, no. 9, pp. 1143–1155, 2019. [PubMed: 30913908]
- [64]. Yao L, Baker JL, Schiff ND, Purpura KP, and Shoaran M, "Predicting task performance from biomarkers of mental fatigue in global brain activity," *Journal of neural engineering*, vol. 18, no. 3, p. 036001, 2021.
- [65]. Esteller R, Echaz J, Tchong T, Litt B, and Pless B, "Line length: an efficient feature for seizure onset detection," in *2001 Conference Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 2. IEEE, 2001, pp. 1707–1710.
- [66]. Hjorth B, "EEG analysis based on time domain properties," *Electroencephalography and clinical neurophysiology*, vol. 29, no. 3, pp. 306–310, 1970. [PubMed: 4195653]
- [67]. Abdelhalim K, Jafari HM, Kokarovtseva L, Velazquez JLP, and Genov R, "64-channel uwb wireless neural vector analyzer soc with a closed-loop phase synchrony-triggered neurostimulator," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 10, pp. 2494–2510, 2013.
- [68]. Yao L, Brown P, and Shoaran M, "Resting tremor detection in Parkinson's disease with machine learning and Kalman filtering," in *2018 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2018, pp. 1–4.
- [69]. Yao L, Baker JL, Ryou J-W, Schiff ND, Purpura KP, and Shoaran M, "Mental fatigue prediction from multi-channel ecog signal," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 1259–1263.
- [70]. Schindler K, Leung H, Elger CE, and Lehnertz K, "Assessing seizure dynamics by analysing the correlation structure of multichannel intracranial eeg," *Brain*, vol. 130, no. 1, pp. 65–77, 2007. [PubMed: 17082199]
- [71]. Stanslaski S, Herron J, Chouinard T, Bourget D, Isaacson B, Kremen V, Opri E, Drew W, Brinkmann BH, Gunduz A et al. , "A chronically implantable neural coprocessor for

- investigating the treatment of neurological disorders,” *IEEE transactions on biomedical circuits and systems*, vol. 12, no. 6, pp. 1230–1245, 2018. [PubMed: 30418885]
- [72]. Shoaran M, Shahshahani M, Farivar M, Almajano J, Shahshahani A, Schmid A, Bragin A, Leblebici Y, and Emami A, “A 16-channel 1.1 mm² implantable seizure control soc with sub- μ w/channel consumption and closed-loop stimulation in 0.18 μ m cmos,” in 2016 IEEE Symposium on VLSI Circuits (VLSI-Circuits). Ieee, 2016, pp. 1–2.
- [73]. Altaf MAB, Tillak J, Kifle Y, and Yoo J, “A 1.83 μ J/classification nonlinear support-vector-machine-based patient-specific seizure classification SoC,” in 2013 ISSCC. IEEE, 2013, pp. 100–101.
- [74]. Zhu B and Shoaran M, “Unsupervised domain adaptation for cross-subject few-shot neurological symptom detection,” arXiv preprint arXiv:2103.00606, 2021.
- [75]. Page A, Sagedy C, Smith E, Attaran N, Oates T, and Mohsenin T, “A flexible multichannel EEG feature extractor and classifier for seizure detection,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 62, no. 2, pp. 109–113, 2014.
- [76]. Shoaran M, Pollo C, Schindler K, and Schmid A, “A fully integrated IC with 0.85- μ w/channel consumption for epileptic iEEG detection,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 62, no. 2, pp. 114–118, 2015.
- [77]. Wang T, Shoaran M, and Emami A, “Towards adaptive deep brain stimulation in parkinson’s disease: Lfp-based feature analysis and classification,” in 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2018, pp. 2536–2540.
- [78]. Bergey GK, Morrell MJ, Mizrahi EM, Goldman A, King-Stephens D, Nair D, Srinivasan S, Jobst B, Gross RE, Shields DC et al. , “Long-term treatment with responsive brain stimulation in adults with refractory partial seizures,” *Neurology*, vol. 84, no. 8, pp. 810–817, 2015. [PubMed: 25616485]
- [79]. Nair DR, Laxer KD, Weber PB, Murro AM, Park YD, Barkley GL, Smith BJ, Gwinn RP, Doherty MJ, Noe KH et al. , “Nine-year prospective efficacy and safety of brain-responsive neurostimulation for focal epilepsy,” *Neurology*, vol. 95, no. 9, pp. e1244–e1256, 2020. [PubMed: 32690786]
- [80]. Jarosiewicz B and Morrell M, “The rns system: brain-responsive neurostimulation for the treatment of epilepsy,” *Expert Review of Medical Devices*, pp. 1–10, 2020.
- [81]. Hamilton P, Soryal I, Dhahri P, Wimalachandra W, Leat A, Hughes D, Toghil N, Hodson J, Sawlani V, Hayton T et al. , “Clinical outcomes of vns therapy with aspiresr®(including cardiac-based seizure detection) at a large complex epilepsy and surgery centre,” *Seizure*, vol. 58, pp. 120–126, 2018. [PubMed: 29702409]
- [82]. Englot DJ, Rolston JD, Wright CW, Hassnain KH, and Chang EF, “Rates and predictors of seizure freedom with vagus nerve stimulation for intractable epilepsy,” *Neurosurgery*, vol. 79, no. 3, pp. 345–353, 2016. [PubMed: 26645965]
- [83]. Kawaji H, Yamamoto T, Fujimoto A, Uchida D, Ichikawa N, Yamazoe T, Okanishi T, Sato K, Nishimura M, Tanaka T et al. , “Additional seizure reduction by replacement with vagus nerve stimulation model 106 (aspresr),” *Neuroscience letters*, vol. 716, p. 134636, 2020.
- [84]. “Newronika alphas closed-loop adaptive deep brain stimulation system,” [Online]. Available: <http://www.newronika.com/>, 2021.
- [85]. Zamora M, Toth R, Ottaway J, Gillbe T, Martin S, Benjaber M, Lamb G, Noone T, Nairac Z, Constantinou TG et al. , “Dyneumo mk-1: A fully-implantable, motion-adaptive neurostimulator with configurable response algorithms,” *bioRxiv*, 2020.
- [86]. Provenza NR, Paulk AC, Peled N, Restrepo MI, Cash SS, Dougherty DD, Eskandar EN, Borton DA, and Widge AS, “Decoding task engagement from distributed network electrophysiology in humans,” *Journal of neural engineering*, vol. 16, no. 5, p. 056015, 2019. [PubMed: 31419211]
- [87]. Kremen V, Brinkmann BH, Kim I, Guragain H, Nasseri M, Magee AL, Attia TP, Nejedly P, Sladky V, Nelson N et al. , “Integrating brain implants with local and distributed computing devices: a next generation epilepsy management system,” *IEEE journal of translational engineering in health and medicine*, vol. 6, pp. 1–12, 2018.
- [88]. Cheng C-H, Tsai P-Y, Yang T-Y, Cheng W-H, Yen T-Y, Luo Z, Qian X-H, Chen Z-X, Lin T-H, Chen W-H et al. , “A fully integrated 16-channel closed-loop neural-prosthetic cmos soc with

- wireless power and bidirectional data telemetry for real-time efficient human epileptic seizure control,” *IEEE Journal of Solid-State Circuits*, vol. 53, no. 11, pp. 3314–3326, 2018.
- [89]. Wang Y, Sun Q, Luo H, Chen X, Wang X, and Zhang H, “26.3 a closed-loop neuromodulation chipset with 2-level classification achieving 1.5 v pp cm interference tolerance, 35db stimulation artifact rejection in 0.5 ms and 97.8% sensitivity seizure detection,” in *2020 IEEE International Solid-State Circuits Conference-(ISSCC)*. IEEE, 2020, pp. 406–408.
- [90]. Lee KH and Verma N, “A low-power processor with configurable embedded machine-learning accelerators for high-order and adaptive analysis of medical-sensor signals,” *IEEE Journal of Solid-State Circuits*, vol. 48, no. 7, pp. 1625–1637, 2013.
- [91]. Huang S-A, Chang K-C, Liou H-H, and Yang C-H, “A 1.9-mw svm processor with on-chip active learning for epileptic seizure control,” *IEEE Journal of Solid-State Circuits*, vol. 55, no. 2, pp. 452–464, 2019.
- [92]. Shoaran M, Haghi BA, Farivar M, and Emami A, “Efficient feature extraction and classification methods in neural interfaces,” in *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2017 Symposium*, vol. 47, no. 4, 2017, pp. 31–35.
- [93]. Yang J and Sawan M, “From seizure detection to smart and fully embedded seizure prediction engine: A review,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 14, no. 5, pp. 1008–1023, 2020. [PubMed: 32822304]
- [94]. Shoaib M, Lee KH, Jha NK, and Verma N, “A 0.6–107 μ w energy-scalable processor for directly analyzing compressively-sensed eeg,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, no. 4, pp. 1105–1118, 2014.
- [95]. Taghavi M, Haghi BA, Farivar M, Shoaran M, and Emami A, “A 41.2 nj/class, 32-channel on-chip classifier for epileptic seizure detection,” in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2018, pp. 3693–3696.
- [96]. Shoaran M, Farivar M, and Emami A, “Hardware-friendly seizure detection with a boosted ensemble of shallow decision trees,” in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2016, pp. 1826–1829.
- [97]. Breiman L, “Random forests,” *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [98]. Chen T and Guestrin C, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. ACM, 2016, pp. 785–794.
- [99]. Freund Y and Schapire RE, “A decision-theoretic generalization of on-line learning and an application to boosting,” *Journal of computer and system sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [100]. Aslam AR, Iqbal T, Aftab M, Saadeh W, and Altaf MAB, “A 10.13 μ J/classification 2-channel deep neural network-based SoC for emotion detection of autistic children,” in *2020 IEEE Custom Integrated Circuits Conference (CICC)*. IEEE, 2020, pp. 1–4.
- [101]. Fang W-C, Wang K-Y, Fahier N, Ho Y-L, and Huang Y-D, “Development and validation of an EEG-based real-time emotion recognition system using edge AI computing platform with convolutional neural network system-on-chip design,” *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 4, pp. 645–57, 2019.
- [102]. Wagenaar JB, Worrell GA, Ives Z, Dimpelmann M, Litt B, and Schulze-Bonhage A, “Collaborating and sharing data in epilepsy research,” *Journal of clinical neurophysiology: official publication of the American Electroencephalographic Society*, vol. 32, no. 3, p. 235, 2015. [PubMed: 26035676]
- [103]. Swann NC, de Hemptinne C, Thompson MC, Miocinovic S, Miller AM, Ostrem JL, Chizeck HJ, Starr PA et al. , “Adaptive deep brain stimulation for parkinson’s disease using motor cortex sensing,” *Journal of neural engineering*, vol. 15, no. 4, p. 046006, 2018. [PubMed: 29741160]
- [104]. Shute JB, Okun MS, Opri E, Molina R, Rossi PJ, Martinez-Ramirez D, Foote KD, and Gunduz A, “Thalamocortical network activity enables chronic tic detection in humans with tourette syndrome,” *NeuroImage: Clinical*, vol. 12, pp. 165–172, 2016. [PubMed: 27419067]
- [105]. Zhu B, Taghavi M, and Shoaran M, “Cost-efficient classification for neurological disease detection,” in *2019 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2019, pp. 1–4.

- [106]. Watts J, Khojandi A, Shylo O, and Ramdhani RA, "Machine learning's application in deep brain stimulation for parkinson's disease: A review," *Brain Sciences*, vol. 10, no. 11, p. 809, 2020.
- [107]. Lozano AM, Giacobbe P, Hamani C, Rizvi SJ, Kennedy SH, Kolivakis TT, Debonnel G, Sadikot AF, Lam RW, Howard AK et al. , "A multicenter pilot study of subcallosal cingulate area deep brain stimulation for treatment-resistant depression," *Journal of neurosurgery*, vol. 116, no. 2, pp. 315–322, 2012. [PubMed: 22098195]
- [108]. Chang S-Y, Wu B-C, Liou Y-L, Zheng R-X, Lee P-L, Chiueh T-D, and Liu T-T, "An ultra-low-power dual-mode automatic sleep staging processor using neural-network-based decision tree," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 66, no. 9, pp. 3504–3516, 2019.
- [109]. Restuccia D, Vollono C, Piero I. d., Martucci L, and Zanini S, "Different levels of cortical excitability reflect clinical fluctuations in migraine," *Cephalalgia*, vol. 33, no. 12, pp. 1035–1047, 2013. [PubMed: 23575822]
- [110]. Plow EB and Machado A, "Invasive neurostimulation in stroke rehabilitation," *Neurotherapeutics*, vol. 11, no. 3, pp. 572–582, 2014. [PubMed: 24353109]
- [111]. Schiff ND, Giacino JT, Kalmar K, Victor J, Baker K, Gerber M, Fritz B, Eisenberg B, O'connor J, Kobylarz E et al. , "Behavioural improvements with thalamic stimulation after severe traumatic brain injury," *Nature*, vol. 448, no. 7153, pp. 600–603, 2007. [PubMed: 17671503]
- [112]. Park S-Y, Cho J, Na K, and Yoon E, "Modular 128-channel δ - δ σ analog front-end architecture using spectrum equalization scheme for 1024-channel 3-d neural recording microsystems," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 2, pp. 501–514, 2017.
- [113]. Taghavi M and Shoaran M, "Hardware complexity analysis of deep neural networks and decision tree ensembles for real-time neural data classification," in 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER). IEEE, 2019, pp. 407–410.
- [114]. Mandal A, Peña D, Pamula R, Khateeb K, Murphy L, Yazdan-Shahmorad A, Perlmutter S, Pape F, Rudell JC, and Sathe V, "A 46-channel vector stimulator with 50mv worst-case common-mode artifact for low-latency adaptive closed-loop neuromodulation," in 2021 IEEE Custom Integrated Circuits Conference (CICC). IEEE, 2021, pp. 1–2.
- [115]. Musk E et al. , "An integrated brain-machine interface platform with thousands of channels," *Journal of medical Internet research*, vol. 21, no. 10, p. e16194, 2019. [PubMed: 31642810]
- [116]. Yoon D-Y, Pinto S, Chung S, Merolla P, Koh T-W, and Seo D, "A 1024-channel simultaneous recording neural soc with stimulation and real-time spike detection," in 2021 Symposium on VLSI Circuits. IEEE, 2021, pp. 1–2.
- [117]. Muller R, Le H-P, Li W, Ledochowitsch P, Gambini S, Bjorninen T, Koralek A, Carmena JM, Maharbiz MM, Alon E et al. , "A minimally invasive 64-channel wireless μ ecog implant," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 1, pp. 344–359, 2014.
- [118]. Sharma M, Gardner AT, Strathman HJ, Warren DJ, Silver J, and Walker RM, "Acquisition of neural action potentials using rapid multiplexing directly at the electrodes," *Micromachines*, vol. 9, no. 10, p. 477, 2018.
- [119]. Uehlin JP, Smith WA, Pamula VR, Perlmutter SI, Rudell JC, and Sathe VS, "A 0.0023 mm²/ch. delta-encoded, time-division multiplexed mixed-signal ecog recording architecture with stimulus artifact suppression," *IEEE transactions on biomedical circuits and systems*, vol. 14, no. 2, pp. 319–331, 2019. [PubMed: 31902767]
- [120]. Wang S, Chaovalitwongse WA, and Wong S, "Online seizure prediction using an adaptive learning approach," *IEEE transactions on knowledge and data engineering*, vol. 25, no. 12, pp. 2854–2866, 2013.
- [121]. Feng L, Li Z, and Wang Y, "Vlsi design of svm-based seizure detection system with on-chip learning capability," *IEEE transactions on biomedical circuits and systems*, vol. 12, no. 1, pp. 171–181, 2017.
- [122]. Murmann B, Bankman D, Chai E, Miyashita D, and Yang L, "Mixed-signal circuits for embedded machine-learning applications," in 2015 49th Asilomar Conference on Signals, Systems and Computers. IEEE, 2015, pp. 1341–1345.

- [123]. Zhang J, Wang Z, and Verma N, "In-memory computation of a machine-learning classifier in a standard 6t sram array," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 4, pp. 915–924, 2017.
- [124]. Tanno R, Arulkumaran K, Alexander DC, Criminisi A, and Nori A, "Adaptive neural trees," *arXiv preprint arXiv:1807.06699*, 2018.
- [125]. Hazimeh H, Ponomareva N, Mol P, Tan Z, and Mazumder R, "The tree ensemble layer: Differentiability meets conditional computation," in *International Conference on Machine Learning*. PMLR, 2020, pp. 4138–4148.
- [126]. Burrello A, Benatti S, Schindler KA, Benini L, and Rahimi A, "An ensemble of hyperdimensional classifiers: Hardware-friendly short-latency seizure detection with automatic ieeg electrode selection," *IEEE journal of biomedical and health informatics*, 2020.
- [127]. Bandarabadi M, Rasekhi J, Teixeira CA, Netoff TI, Parhi KK, and Dourado A, "Early seizure detection using neuronal potential similarity: A generalized low-complexity and robust measure," *International journal of neural systems*, vol. 25, no. 05, p. 1550019, 2015. [PubMed: 25997912]
- [128]. Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q, and Liu T-Y, "Lightgbm: A highly efficient gradient boosting decision tree," in *Advances in neural information processing systems*, 2017, pp. 3146–54.
- [129]. Shi K and Howard D, "Sleep transistor design and implementation-simple concepts yet challenges to be optimum," in *2006 International Symposium on VLSI Design, Automation and Test*. IEEE, 2006, pp. 1–4.
- [130]. Yao L and Shoaran M, "Enhanced classification of individual finger movements with ECoG," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2019, pp. 2063–2066.
- [131]. Lin D, Talathi S, and Annapureddy S, "Fixed point quantization of deep convolutional networks," in *International Conference on Machine Learning*, 2016, pp. 2849–2858.
- [132]. Han S, Mao H, and Dally WJ, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," in *International conference on learning representation*, 2016.
- [133]. Artzi NS, Shilo S, Hadar E, Rossman H, Barbash-Hazan S, Ben-Haroush A, Balicer RD, Feldman B, Wiznitzer A, and Segal E, "Prediction of gestational diabetes based on nationwide electronic health records," *Nature medicine*, vol. 26, no. 1, pp. 71–76, 2020.
- [134]. Lundberg SM, Nair B, Vavilala MS, Horibe M, Eisses MJ, Adams T, Liston DE, Low DK-W, Newman S-F, Kim J et al. , "Explainable machine-learning predictions for the prevention of hypoxaemia during surgery," *Nature Biomedical Engineering*, vol. 2, no. 10, p. 749, 2018.
- [135]. Lundberg SM and Lee S-I, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30*, Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, and Garnett R, Eds. Curran Associates, Inc., 2017, pp. 4765–4774. [Online]. Available: <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>

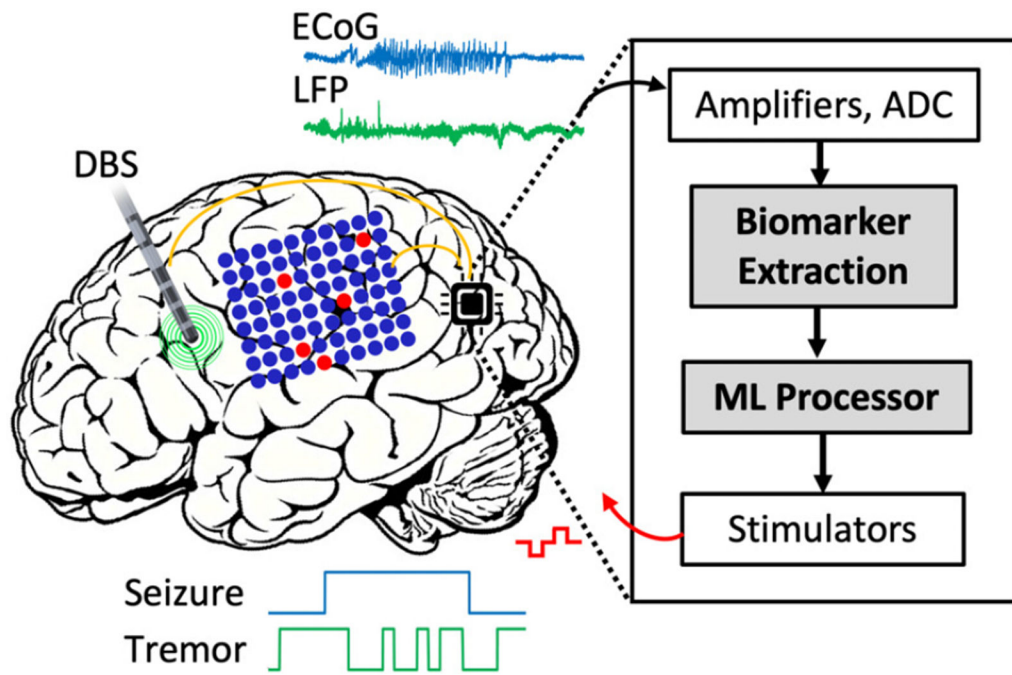


Fig. 1: Symbolic view of a closed-loop neural prosthesis. Multi-channel neural signals such as ECoG and LFP are recorded by cortical and deep-brain electrodes and sent to the implantable microchip. The on-chip biomarker extraction and ML processor detect the onset of symptoms and trigger a therapeutic neurostimulator.

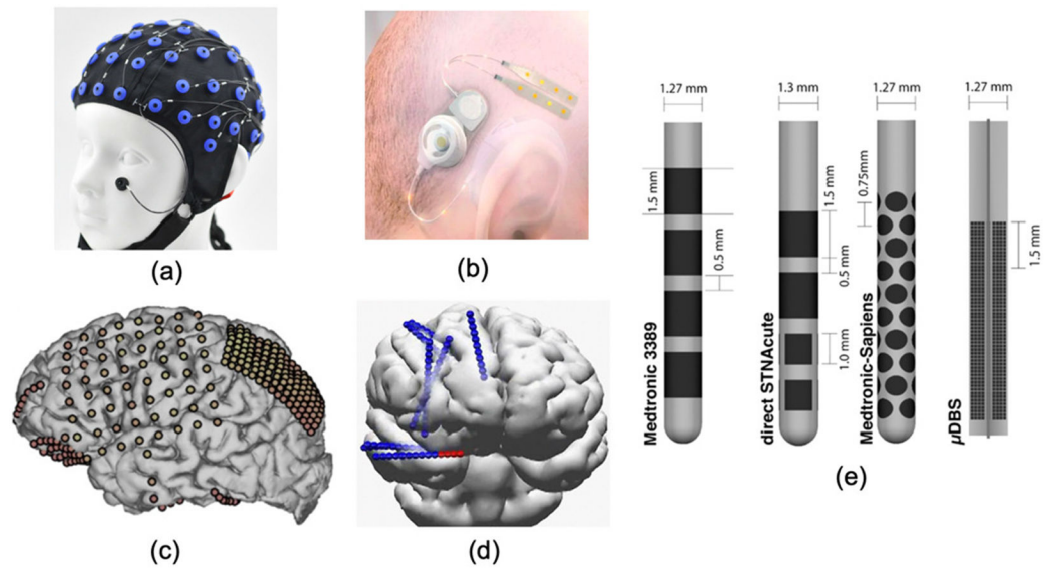


Fig. 2: Standard and emerging electrodes for neural recording and stimulation via noninvasive, minimally-invasive, and invasive technologies; (a) Standard scalp-EEG electrodes. (b) The Epios subscalp EEG device for chronic epilepsy monitoring [39]. (c) Standard and high-density ECoG [40]. (d) Stereo-EEG leads [41]. (e) Clinical DBS (Medtronic’s FDA-approved 3389, left), emerging directional DBS leads (8-channel direct STNacute and 40-channel Medtronic-Sapiens, middle) and the Willsie and Dorval 1760-contact micro-DBS lead (right) [42].

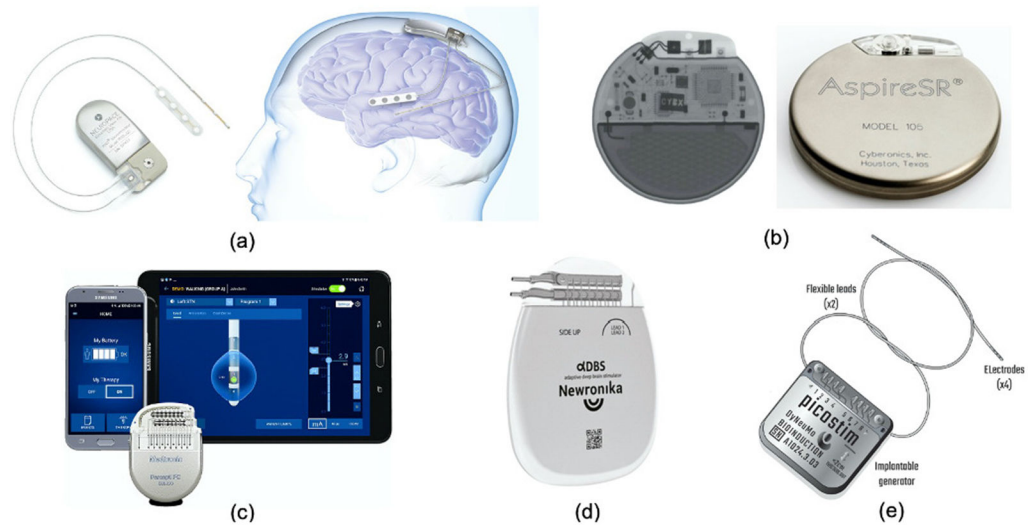


Fig. 3: Existing clinical or research-based closed-loop neuromodulation devices (with or without on-device ML); (a) The NeuroPace RNS device for epilepsy. (b) The AspireSR (Cyberonics, now known as LivaNova) device for epilepsy. (c) The Medtronic Percept PC device for movement disorders. (d) The Newronika AlphaDBS system for Parkinson’s disease. (e) The DyNeuMo Mk-1 system for movement disorders.

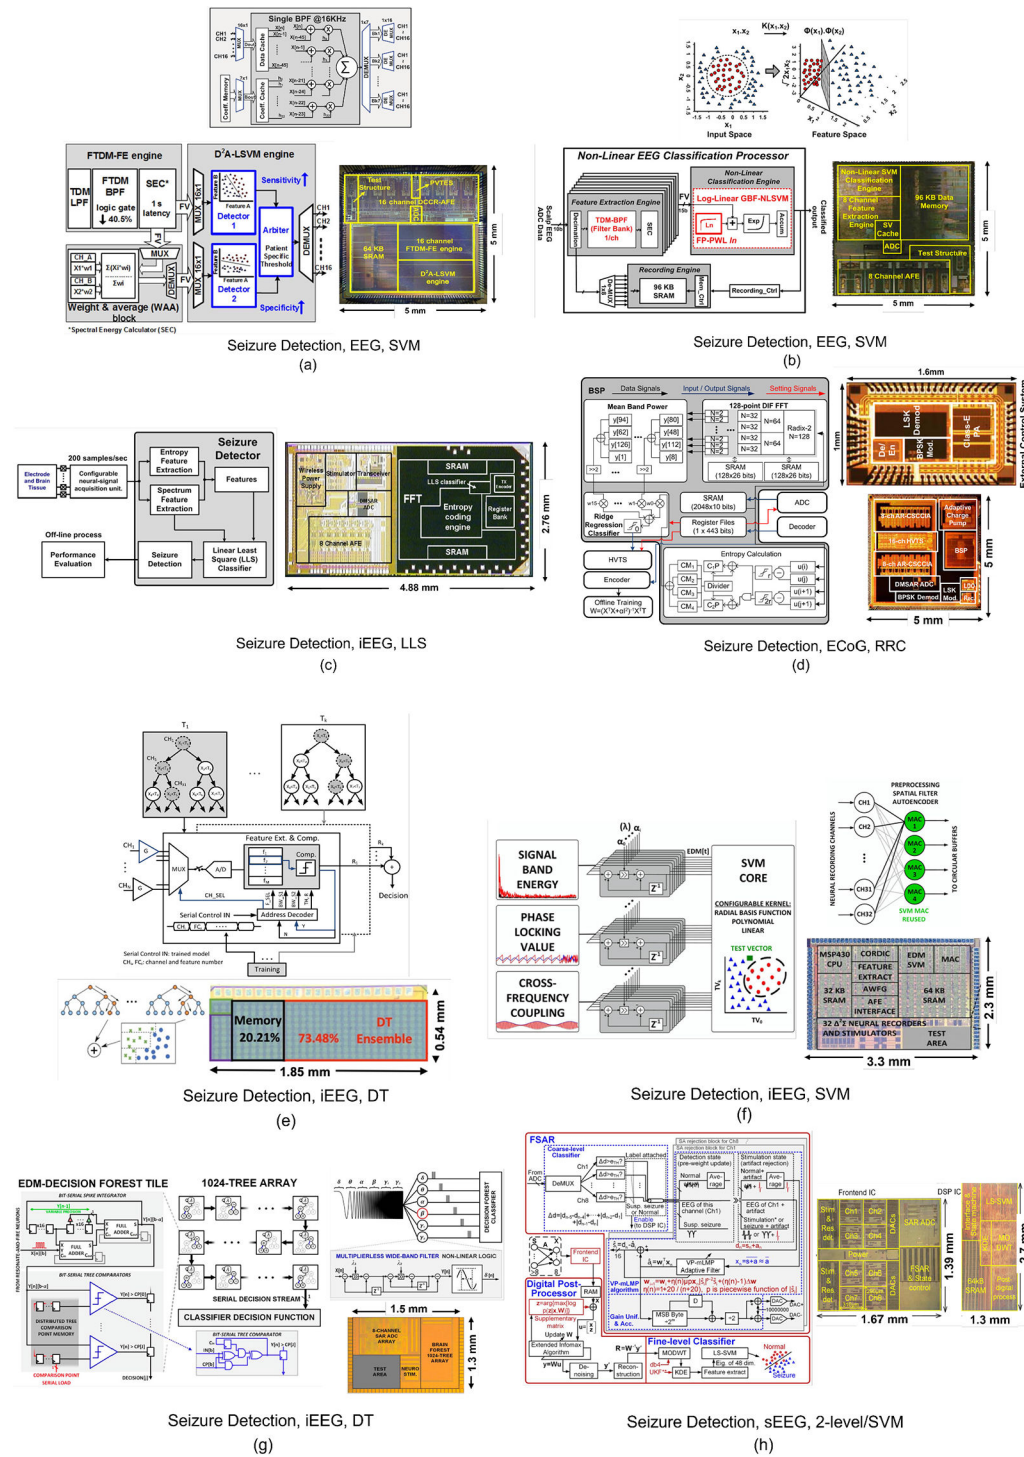


Fig. 4: Hardware architectures and chip micrographs of ML-embedded neural prostheses for epilepsy: (a) Linear dual-detector SVM classifier and closed-loop transcranial neurostimulator [13], (b) non-linear SVM-based seizure detector [14], (c) linear least square (LLS) classifier and closed-loop stimulator [17], (d) ridge regression classifier (RRC) and

closed-loop stimulator [88]. (e) Gradient-boosted tree ensemble for seizure detection [12], (f) exponentially decaying-memory SVM and closed-loop stimulator [16], (g) AdaBoost decision tree classifier and closed-loop stimulator [29], (h) two-level coarse/fine classifier and closed-loop stimulator [89].

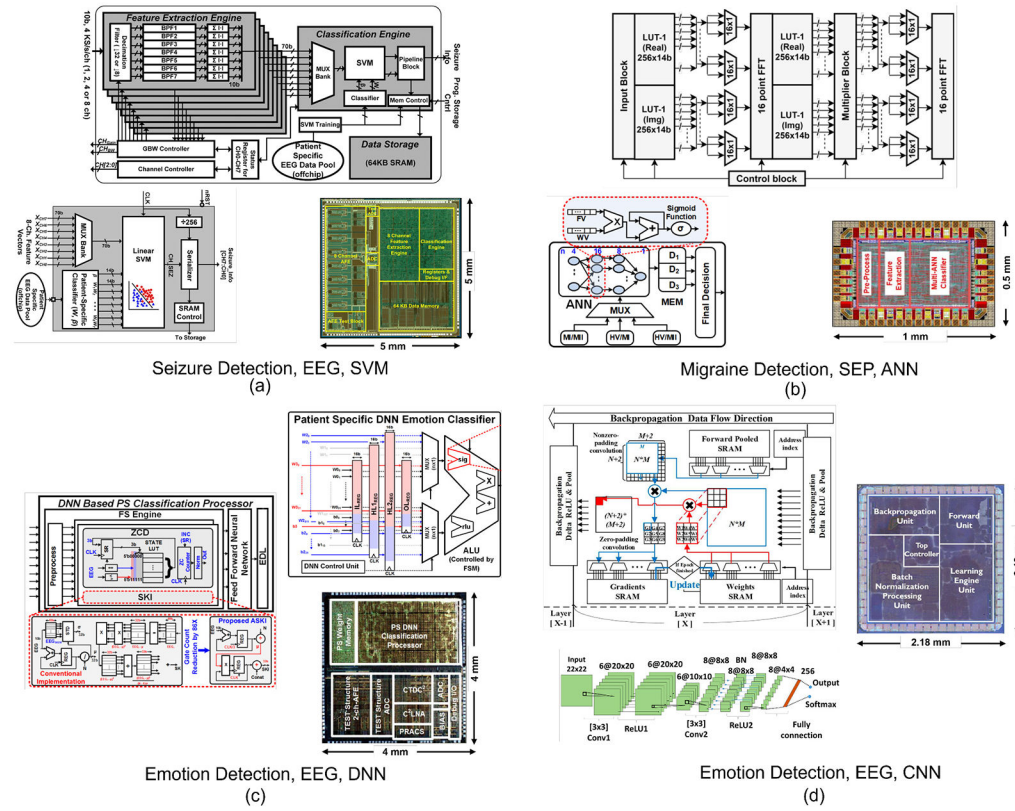


Fig. 5: Hardware architectures and chip micrographs of ML-embedded neural prostheses for various applications: (a) Linear SVM for epilepsy [11], (b) ANN for migraine state detection [18], (c) DNN for Autism emotion detection [100], (d) CNN for emotion detection [101].




Fig. 6:

A DVTE with eight decision trees. Unlike conventional tree ensembles that uniformly set the maximum depth on all trees, the maximum depths in a DVTE are different (1–4). The internal and leaf nodes are shown in blue and black, respectively.

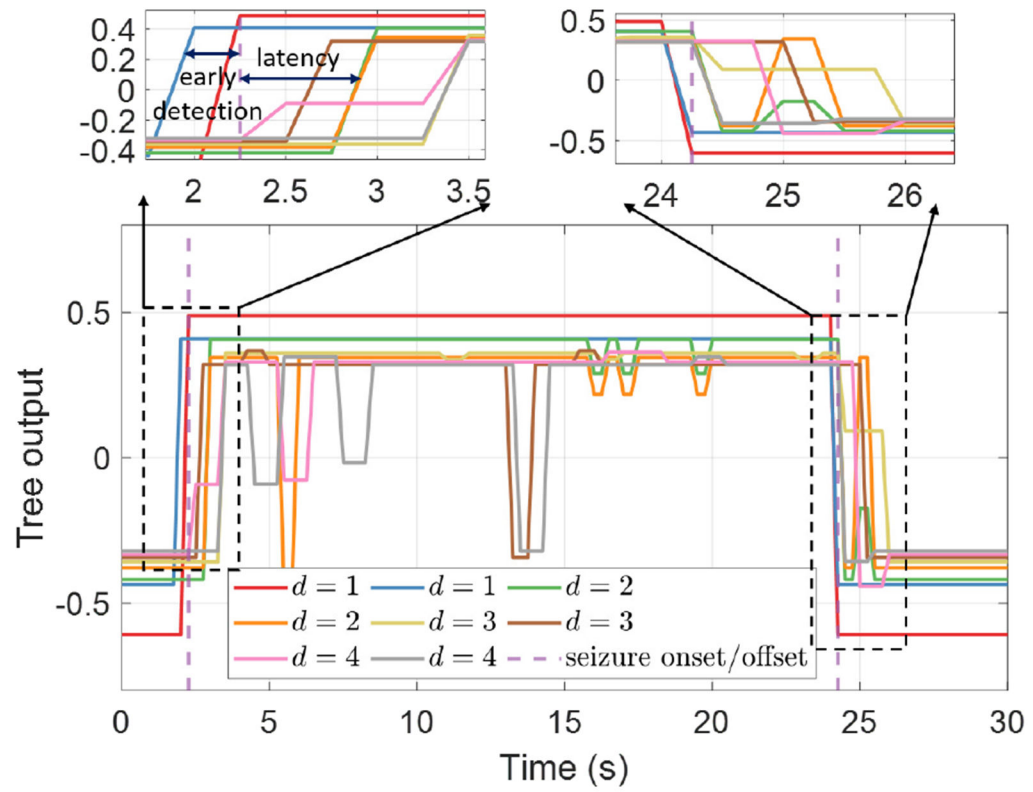


Fig. 7:
 The outputs of decision trees in a DVTE. Latency is defined as the time difference between the expert-marked seizure onset and the state change of each tree's output. d is the maximum depth of each tree.

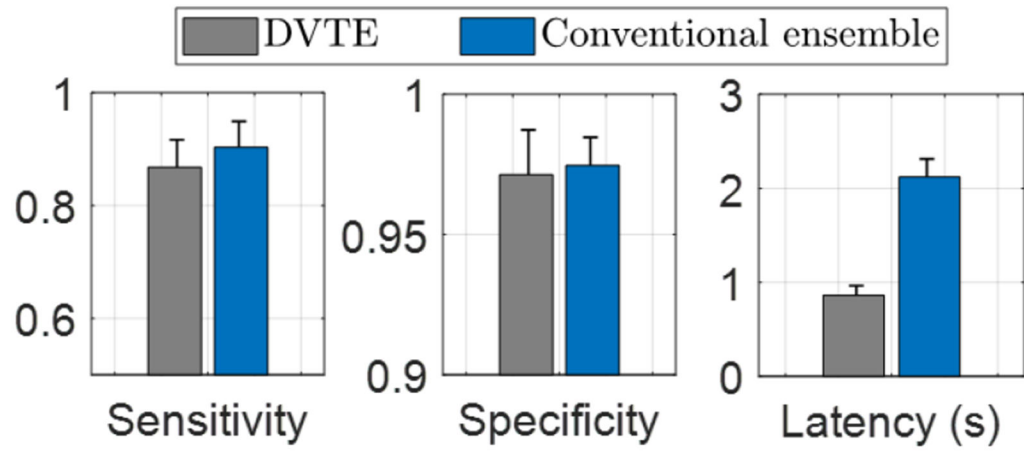


Fig. 8: Performance comparison of DVTE and conventional tree ensemble with a maximum depth of 4. DVTE reduced the latency by 2.5× with a marginal performance reduction (<3% in sensitivity and <1% in specificity). Error bars indicate the standard errors across patients.




Fig. 9: Hardware cost as a function of the regularization coefficient C in DVTE. Large C imposes strong regularization and reduces the power/latency cost. The power cost was calculated as the average power consumption to extract features along the decision path. Latency was estimated as the average time to traverse a root-to-leaf decision path in the tree.

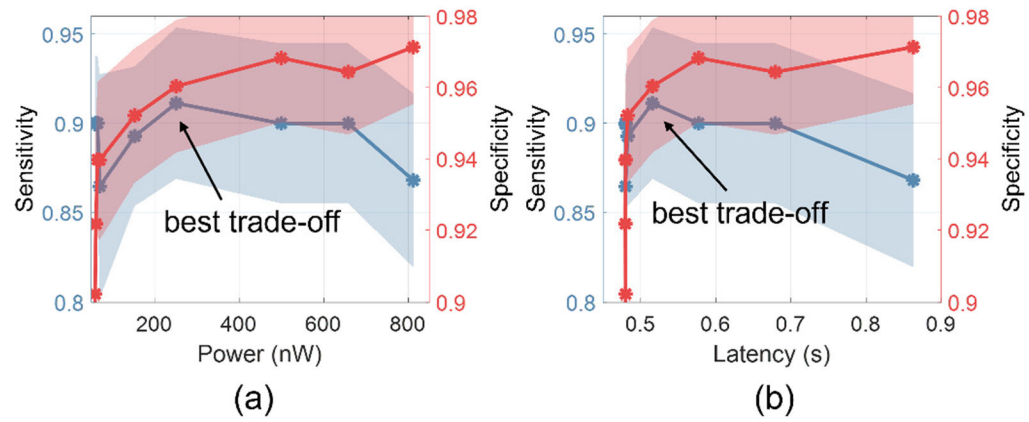


Fig. 10: Seizure detection performance as a function of (a) power consumption and (b) latency. Shaded area indicates the standard errors across patients. The experiment was performed using DVTE and the following setting: 8 trees, depths varying from 1 to 4.

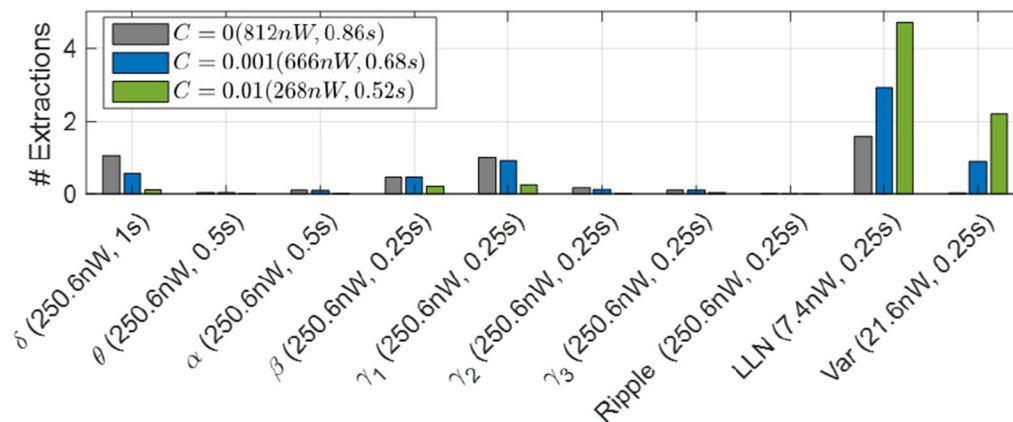


Fig. 11:

The number of extracted features in DVTE for different regularization coefficients. With greater C , the cost-aware model tends to use hardware-friendly features (e.g., LLN, Var). Features with longer windows (δ , θ , α) are also penalized in the cost-aware model. The power cost and latency for each C are shown in the legend, while the X-axis shows individual feature costs. For $C = 0.01$, we achieved an average power cost of 268nW and latency of 0.52s.

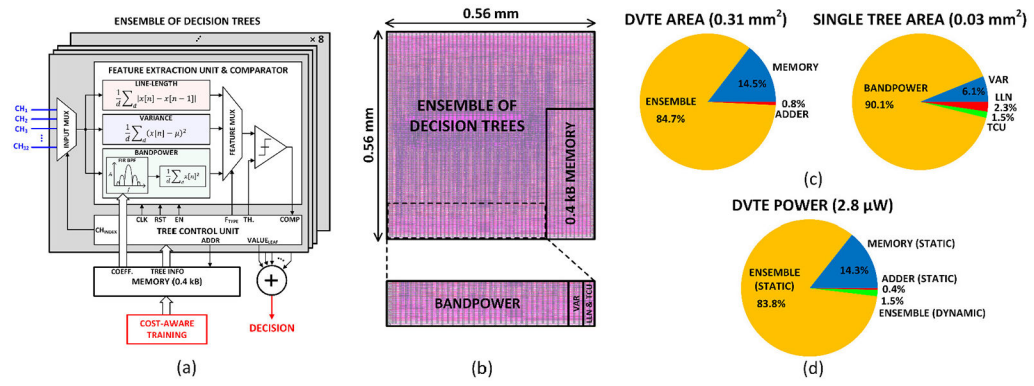


Fig. 12: Hardware implementation of the proposed DVTE classifier: (a) system architecture, (b) layout, (c) area breakdown of the DVTE processor and a single decision tree, and (d) system power breakdown.

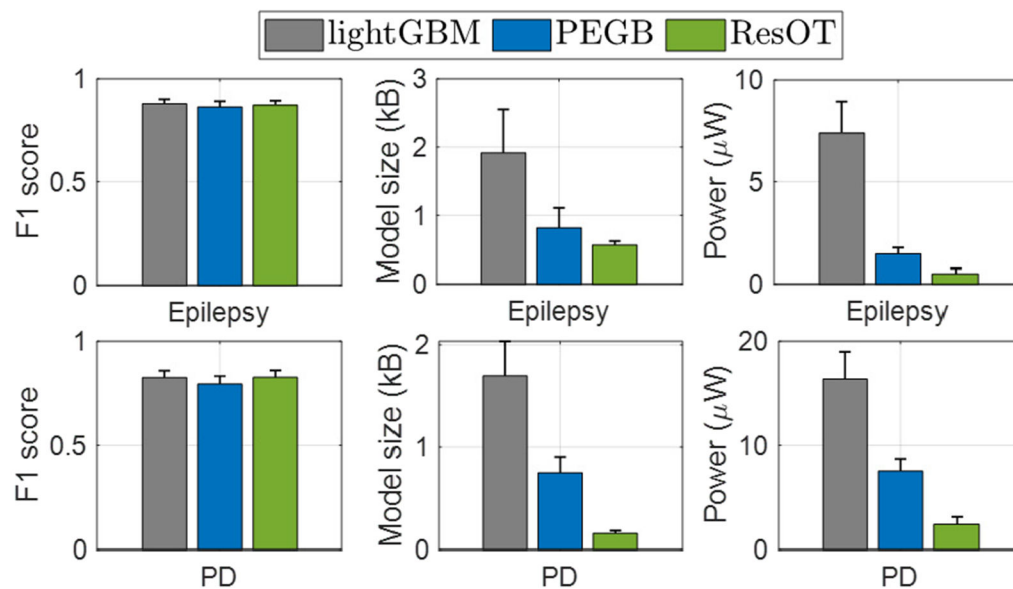


Fig. 13: Comparison of ResOT and axis-aligned tree ensemble on seizure and tremor detection tasks. The conventional gradient boosted ensemble (lightGBM [128]) and gradient boosting with power-efficient training (PEGB) were included. For PEGB, We used fixed-point thresholds and leaf weights, in contrast to floating points in lightGBM [105]. Cross-subject standard errors are shown by error bars.





Fig. 14:

The number of extracted features with different regularization coefficients in ResOT. With greater power- and latency-aware regularization terms, oblique trees prioritize low-power and low-latency features.

**Fig. 15:**

Parallel node evaluation scheme. (a) One internal node is evaluated per window. (b) Two layers (maximum 3 nodes) are concurrently evaluated per window. (c) All nodes are evaluated in parallel. The evaluated nodes are shown in color and bold lines represent the decision path. Node colors represent successive windows.

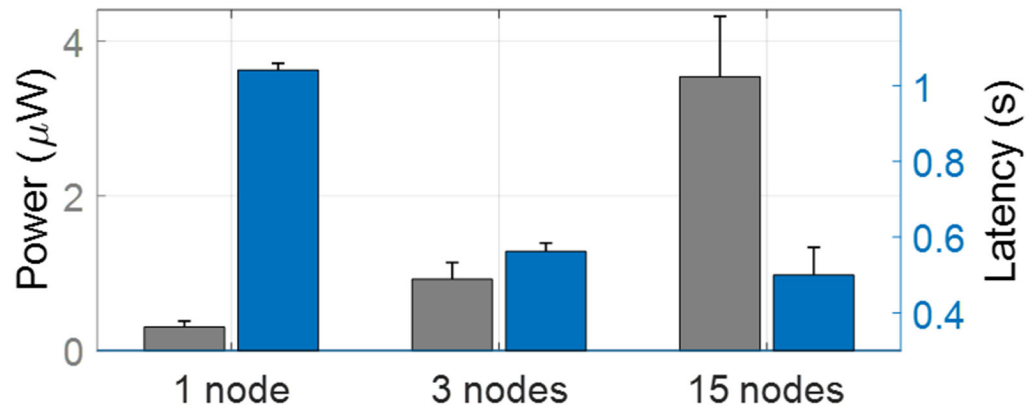


Fig. 16: The power-latency trade-off with parallel node evaluation. With more nodes evaluated in parallel, latency is reduced at the cost of increased power consumption. Experiments were conducted with ResOT on epilepsy task.

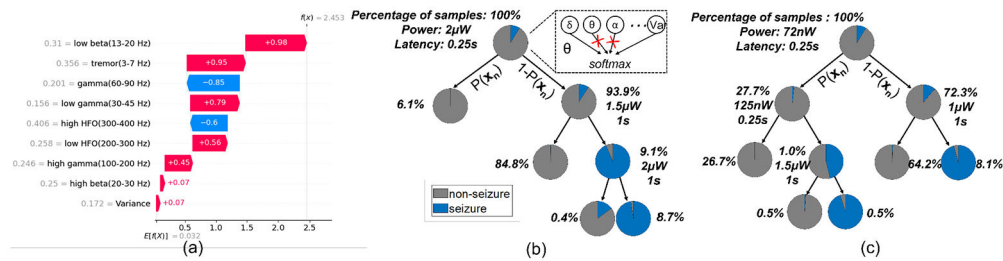


Fig. 17:

(a) Interpretation of the tremor detection process using a tree ensemble and shapley additive explanations. Features plotted in red predicted an increased risk of tremor, while those in blue were associated with a low tremor risk. (b) Visualization of the seizure detection process in an oblique tree in one arbitrary patient. We show the percentage of samples visiting each node, and the required power and latency to evaluate each internal node. There are multiple “short paths” which allow dynamic early exiting. (c) Visualization of a cost-aware oblique tree, showing a significant reduction in the power cost of the root node.

TABLE I:

Performance Summary of Machine Learning SoCs for Epilepsy

Parameter	JETCAS'18 [12]	JSSC'13 [17]	ISSCC'20 [29]	JSSC'18 [88]	ISSCC'18 [16]	ISSCC'20 [89]	TBCAS'16 [14]	JSSC'13 [11]	JSSC'15 [13]	This Work
Process	65 nm	180 nm	65 nm	180 nm	130 nm	180 nm	180 nm	180 nm	180 nm	65 nm
Classifier	XGB DT	LLS	AdaBoost DT	RRC	EDM-SVM	coarse/fine SVM	Non-Lin SVM	Lin-SVM	Dual-LSVM	DVTE [‡]
Features	LLN, Pow, Var, BPF	Ent., Spec.	RAF-BPF	FFT, ApEn	PLV, CFC, BPF	MODWT-KDE	TDM-BPF	BPF	FTDM-BPF	LLN, Var, BPF
Signal Modality	iEEG	iEEG	iEEG	ECoG	iEEG	Stereo-EEG	EEG	EEG	EEG	iEEG
Closed-loop	N	Y	Y	Y	Y	Y	N	N	Y	N
# of Sensing Channels	32	8	8	16	32 [‡]	8	8	8	16	32
ML Energy Efficiency	41.2 nJ/class.	77.9 μ J/class.	36 nJ/class.	62.5 μ J/class.	168.6 μ J/class.	14.2 μ J/class.	1.31 μ J/class. ^{‡‡}	1.49 μ J/class. ^{‡‡}	1.85 μ J/class.	5.6 nJ/class.
ML Power	206.4 μ W	882 μ W [‡]	9.6 μ W ^{‡*}	2.5 mW [‡]	674.4 μ W	1.16 μ W	156.6 μ W [‡]	193.8 μ W [‡]	216.7 μ W [‡]	2.8 μW
Total Area (ML Area)	1 (1) mm ²	13.47 (4.85 [*]) mm ²	1.95 (0.42) mm ²	25 (2.52 [*]) mm ²	7.59 (3.32) mm ²	5.83 (3.51) mm ²	25 (5.55 [*]) mm ²	25 (7.37 [*]) mm ²	25 (7.47 [*]) mm ²	0.31 (0.31) mm²
Sampling Rate/Ch.	5 kS/s	62.5 kS/s	256 S/s	2 kS/s	256 S/s	1 kS/s ^{**}	128 S/s ^{‡‡}	128 S/s ^{‡‡}	128 S/s ^{‡‡}	500 S/s
Sensitivity	83.7%	92% [‡]	96.7%	97.8% [‡]	97.7%	97.8%	95.1%	82.7% ^{‡‡}	95.7%	91.1%
Specificity	88.1%	N.A.	0.8 FAR ^{**}	N.A.	0.185 FAR ^{**}	99.7%	0.27 FAR [§]	4.5% FPR	98% (0.27 FAR ^{**})	96%
Dataset (# patients)	iEEG.org (26)	Rats	EU-iEEG	ECoG (5)	EU-iEEG (4)	CHB-MIT [§] (23)	CHB-MIT (24)	CHB-MIT (24)	CHB-MIT (23)	iEEG.org (11)
Latency	1.1 s	0.8 s	N.A.	0.76 s	<0.1 s ^{§§}	<0.3 s ^{§§}	2 s	<2 s ^{‡‡}	1 s ^{§§}	0.52 s ^{§§}
ML Energy/Ch.	1.29 nJ/S	1.76 nJ/S	4.69 nJ/S	78.1 nJ/S	82.3 nJ/S	0.145 nJ/S	153 nJ/S	189 nJ/S	106 nJ/S	0.175 nJ/S
ML Area/Ch.	0.031 mm ²	0.606 mm ²	0.053 mm ²	0.157 mm ²	0.104 mm ²	0.439 mm ²	0.694 mm ²	0.921 mm ²	0.467 mm ²	0.01 mm²
ML E-A FoM	40.3 pJ-mm ² /S	1.07 nJ-mm ² /S	248 pJ-mm ² /S	12.3 nJ-mm ² /S	8.5 nJ-mm ² /S	63.6 pJ-mm ² /S	106.1 nJ-mm ² /S	174.3 nJ-mm ² /S	49.4 nJ-mm ² /S	1.7 pJ-mm²/S

[‡]ML (feature extractor and classifier) power consumption estimated from power breakdown

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

^{†*} ML dynamic power (static power not reported)

[†] Also applicable to Parkinson tremor detection. Post place-and-route results.

^{**} Variable (256, 1k, 125kS/s)

^{††} Seizure detection rate

[§] With 2000 seizure samples generated by synthetic minority oversampling technique

^{††} Rapid eye blink detection

[†] 4-channel post dimensionality reduction

^{††} As reported in [13]

^{*} ML (feature extractor and classifier) area estimated from chip micrograph

^{††} Accuracy metric

^{*†} Number of false alarms per hour

^{§§} Processing (system) latency

^{††} After on-chip decimation

TABLE II:

Comparison of Machine Learning SoCs

Parameter	TCAS-H'21 [18]	JETCAS'19 [101]	CICC'20 [100]
Process	180 nm	28 nm	180 nm
Classifier	Multi-ANN [†]	CNN	DNN
Application	Migraine Detection	Emotion Detection	Emotion Detection
Features	HFO, BPF, Peak latency	Off-chip	ZCD, SK
Signal Modality	SEP	EEG	EEG
Closed-loop	N	N	N
# of Sensing Channels		6	2
ML Energy Efficiency	N.A.	N.A.	10.13 $\mu\text{J}/\text{class}$.
ML Power	249 μW	76.61 mW	N.A.
Total Area (ML Area)	0.5 (0.5) mm^2	3.35 (3.35) mm^2	16 (6.02 [*]) mm^2
Sampling Rate/Ch.	5 kS/s	250 S/s	N.A.
Accuracy	76%	83.4% ^{**}	85.2%
Dataset (# patients)	MI, MII (42), HV (15)	DEAP (32)	DEAP (32), SEED
Latency	50 ms [‡]	0.45 s [‡]	<1min [‡]
ML Energy/Ch.	49.8 nJ/S	51.1 $\mu\text{J}/\text{S}$	N.A.
ML Area/Ch.	0.5 mm^2	0.558 mm^2	3.01 mm^2
ML E-A FoM	24.9 nJ· mm^2/S	29 $\mu\text{J}\cdot\text{mm}^2/\text{S}$	N.A.

[†] Post place-and-route results.

^{*} ML (feature extractor and classifier) area estimated from chip micrograph

^{**} Accuracy metric

[‡] Processing (system) latency

TABLE III:

Epilepsy features, their power and latency costs

Features and description		Power (nW)	Latency (s)
Delta (δ):	Bandpower over 1-4Hz	250.6	1
Theta (θ):	Bandpower over 4-8Hz	250.6	0.5
Alpha (α):	Bandpower over 8-13Hz	250.6	0.5
Beta (β):	Bandpower over 13-30Hz	250.6	0.25
Low-Gamma (γ_1):	Bandpower over 30-50Hz	250.6	0.25
Gamma (γ_2):	Bandpower over 50-80Hz	250.6	0.25
High-Gamma (γ_3):	Bandpower over 80-150Hz	250.6	0.25
Ripple:	Bandpower over 150-250Hz	250.6	0.25
Line-Length (LLN):	$\frac{1}{d} \sum_d x[n] - x[n-1] $ $d =$ window size	7.4	0.25
Variance (Var):	$\frac{1}{d} \sum_d (x[n] - \mu)^2$, $\mu = \frac{1}{d} \sum_d (x[n])$	21.6	0.25

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript