Machine Learning Microserver for Neuromodulation Device Training

Gerard O'Leary, Asish O. Abraham, Akshay K. Kamath, David Groppe†, Taufik A. Valiante†* and Roman Genov
Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada
*Division of Neurosurgery, Department of Surgery, University of Toronto, Toronto, ON M5T 2S8, Canada
†Krembil Neuroscience Center, Toronto, M5T 2S8, Canada

***Abstract* - Determining a treatment for those with refractory epilepsy often requires an observation period in an Epilepsy Monitoring Unit (EMU) where neural recordings are analyzed to localize a seizure onset zone. This region can be targeted by an implanted neuromodulation device to detect and inhibit ictal activity. Due to the patient-specific nature of epilepsy, on-device machine learning has been demonstrated to improve seizure detection accuracy. However, chronic implants experience considerable recording signal variability over time, leading to a degradation in treatment efficacy. This suggests the need for post-implantation learning which is impractical to perform on a power-constrained device. Presented here is a patient-localized microserver which enables continuous model adaptation aided by unsupervised machine learning. The system employs a one-class support vector machine (OC-SVM) to identify irregular neural activity for remote clinical assessment and device re-training. The system performance is demonstrated using 500 hours of human intracranial EEG (iEEG) where a clinical seizure detection rate of 97.05% is achieved.**

I. Introduction

Epilepsy is one of the most common neurological disorders, affecting approximately 50 million people worldwide. A third of these individuals are not successfully treated with current anti-seizure medications [1]. Determining a treatment for such patients often requires inpatient monitoring to localize the origin of seizure onset. During this observation period, time segments of electroencephalographic interest are identified by technicians to assist epileptologists in patient assessment. These recordings are also monitored in real-time to alert clinical staff early in the development of a seizure. This is critical to enhance patient safety, and for the timely clinical assessment necessary for semiological classification of seizures.

Distinguishing seizure activity from normal brain activity can be a difficult task because of potentially great variation in seizure morphology [2]. Machine learning enables the utilization of large volumes of patient recordings to accurately distinguish pathological from physiological activity. This approach has led to the introduction of responsive closed-loop neuromodulation devices which proactively detect and inhibit the onset of seizures [3-5]. Such systems generally utilize supervised learning models to maintain low false-detection rates for improved power efficiency and reduced side-effects. However, the use of supervised learning classifiers for seizure detection exposes a class imbalance problem which arises from a lack of ictal recordings compared to large volumes of inter-ictal data. Furthermore, supervised classification systems require accurate data labeling, and so are vulnerable to the human error in annotating complex EEG recordings [6].

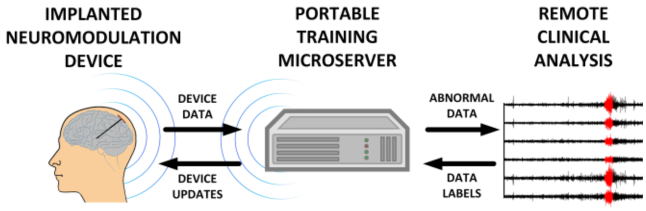

Fig. 1. Training microserver concept: System sends irregular EEG events for clinical assessment and re-trains implanted device.

Activity in EEG signal frequency bands are typically used by epileptologists to categorize irregular neural recordings. Such events include electrographic seizure onsets and inter-ictal discharges (IIDs). This assessment is generally based on temporo-spectral changes such as low-voltage fast activity in iEEG seizure onset [7]. To capture these changes, Exponentially Decaying Memory (EDM) is presented as a hardware efficient mechanism to represent temporal feature characteristics for machine learning.

The proposed system employs an unsupervised learning-based One-class Support Vector Machine (OC-SVM) [8]. This approach navigates the class imbalance and data labeling challenges by learning to distinguish normal neural activity from segments of clinical interest. The irregular recording periods indicated by the OC-SVM can be reviewed remotely by an epileptologist, enabling ictal data to be labelled and accumulated over time. With increasing volumes of data, specialized supervised learning classifiers can be trained more effectively for closed-loop applications.

Chronic neural recording implants experience considerable signal variability over time [9], leading to a gradual degradation of classifier performance. Thus, continuous model re-training is necessary to adapt to changing physiological recording conditions and maximize the treatment efficacy. This is impractical to perform on an implantable device as power-consumption is a primary consideration to reduce both heat dissipation and the risks associated with battery replacement surgery.

The proposed solution is a patient-localized microserver which communicates with an implanted device to enable incremental training (Fig. 1). Data recorded by the device is sent to the server and processed by an FPGA-accelerated OC-SVM. iEEG segments which are considered irregular are archived and sent to a remote epileptologist for review. Once an assessment is made, the microserver re-trains the model to be uploaded to the implanted device.

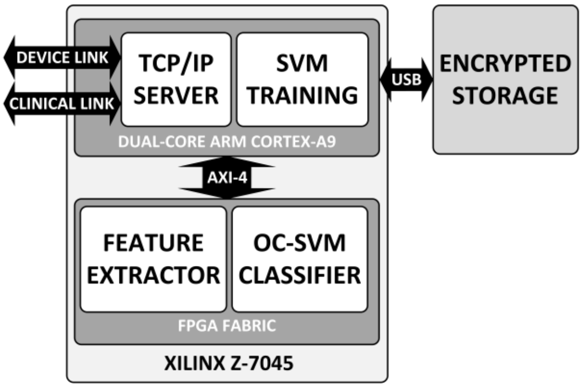

Fig. 2. Training microserver implementation overview. A Linux OS enables external communication, SVM training and access to the FPGA fabric's feature extraction and classification accelerators.

II. System Overview

The system shown in Fig. 2 is implemented using a Xilinx Zynq SoC. A dedicated dual-core CPU hosts an on-chip Linux operating system (OS) which runs in parallel with the FPGA fabric. SVM training is performed on the microserver with encrypted patient data maintained on USB storage. A TCP/IP implementation allows data to be streamed wirelessly from a compatible neuromodulation device. To reduce the implanted device's power requirements for data transmission, raw samples are sent rather than higher dimensional features. This requires feature extraction to be replicated on the microserver. Communication with a remote EEG analyst is supported via an ethernet network interface. Feature extraction and machine learning accelerators are implemented on the FPGA fabric.

III. Feature Extraction

To detect anomalous activity in neural signals, spectral energy in conventional physiological signal bands are used by epileptologists in a clinical setting to label electrographic events [10]. Signal bands of interest are extracted by passing recorded samples through parallel bandpass filters for Delta (<4 Hz), Theta (4-8 Hz), Alpha (8-13 Hz), Beta (13-30 Hz) and Gamma (30-60 Hz) bands. A 256-tap Type-1 FIR filter is implemented for each band with a symmetric impulse response, allowing coefficient multiplications to be shared. Each iEEG channel is processed sequentially and filter states are stored in BRAM between sample processing.

For each band, the absolute value of each output sample is taken as a measure of signal energy. This approximation of instantaneous energy is accumulated over a time window to generate a temporo-spectral measure of the signal.

IV. Exponentially Decaying Memory

To capture the temporal evolution of machine learning features such as signal energy, conventional methods use a windowing approach where contiguous time epochs are concatenated to form a feature vector to be classified. Using this approach, it is possible to learn temporal differences between windows for events such as seizure onset [5].

Window-based approaches have several limitations in performance and hardware efficiency. Firstly, processing larger windows requires proportionally large accumulation logic. Secondly, if classification is performed at every epoch, test vector re-ordering logic may be necessary to remove old windows and add new windows. Thirdly, the minimum detection latency is the time required to generate a window (typically multiple seconds). Lastly, as EEG recordings are patient specific, one window size may give sufficient temporal resolution in one case, but may not be optimal for another. This suggests the need to learn feature timescales in a patient-specific manner to maximize classification performance.

Exponentially Decaying Memory (EDM) is an approach which addresses the challenges outlined. Rather than accumulating and concatenating fixed windows, a continuous sampling recursive window can be defined by:

$$EDM_{(t)} = EDM_{(t-1)} - \alpha\left[EDM_{(t-1)} - x_{(t)}\right] \quad (1)$$

This approach incorporates new inputs, or degrades existing memory of a feature according to the decay rate, α. Where:

$$\alpha = \frac{1}{N}, N = 2^i, 1 < i < 16 \quad (2)$$

The EDM minimizes latency as the output is continuous and can be classified at every sample, rather than every window. Furthermore, temporal resolution is maximized as accumulation over an epoch is not required. The EDM can be implemented efficiently in hardware using shift and add operations if N is limited to powers-of-two. This efficiency allows multiple EDMs to be used in parallel, enabling multiple timescales to be processed simultaneously at a low computational cost.

In this application, after the signal energy is extracted for a given EEG band, its value is passed to a corresponding bank of EDMs. Each EDM implements a different decay rate, α, complementing one another by offering a different temporal perspective of the input feature to be used for classification. Small values of α result in longer-term memory, while larger values capture finer time resolutions.

V. One Class Support Vector Machine

The Support vector machine (SVM) is most commonly used as supervised learning model for classification tasks of two or more classes. The original algorithm requires a similar number of examples in each class to prevent classifier bias. In the case of seizure detection, ictal activity is rare and accurate classification is critical to prevent the onset of a seizure.

The one-class SVM [8] has been proposed as a method for datasets with such class imbalances. It can be viewed as a regular two-class SVM where the training data is taken as one class, and the origin is taken as the only member of the second

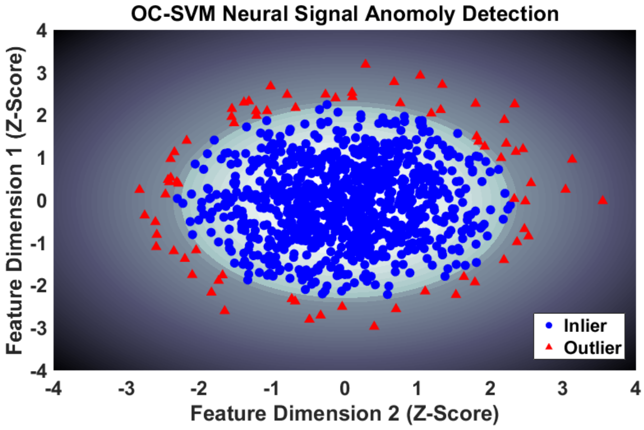

Fig. 3. One-class support vector machine enables the detection of outlier activity (red) following training using inter-ictal data (blue).

class. Training is performed without labels, where data is mapped to the kernel space and separated from the origin by a hyperplane with maximum margin. To classify an input feature vector, a decision function is evaluated to distinguish an inlier (f(x) > 0), from an outlier (f(x) < 0):

$$f(x) = \; sgn\left(\sum_{i=1}^{N} a_i \mathrm{K}(\overrightarrow{sv_i}, \vec{x}) - b\right) \qquad (3)$$

Where sv_i are the support vectors used to construct the hyperplane, a_i are the corresponding weights, b is the classifier bias term, and K is implemented here as the Radial Basis Function (RBF) kernel, defined as:

$$\mathrm{K}(\vec{x}, \overrightarrow{sv}) = \; e^{-\gamma \|\overrightarrow{sv} - \vec{x}\|^2} \qquad (4)$$

This concept is illustrated in Fig. 3, where the model is trained using normal physiological activity (blue). The extracted test vectors classified as outliers (red) could indicate anomalous activity such as inter-ictal discharges (IIDs) or subclinical seizures.

VI. Hardware Implementation

The microserver shown in Fig. 4 is implemented on a Xilinx Zynq XC7Z045, which contains dual-core ARM Cortex-A9 CPUs alongside programmable logic (PL) FPGA fabric. A Petalinux OS is used to host a TCP/IP server for remote communication with a clinician, and local communication with the implanted device for data retrieval and model uploading. The OS is also used to host a LibSVM implementation [11] to perform SVM training, to encrypt neural recordings, and for external data storage.

Feature extraction, real-time feature normalization, and RBF SVM classification are implemented on the FPGA fabric using custom accelerators which are accessible from the OS via an AXI-4 interface. Dual-port Block RAMs (BRAMs) are shared between the OS and the PL to enable online access to FPGA memories for SVM model coefficients, feature extraction coefficients, and test vectors.

Samples are streamed from the TCP/IP server to a PL BRAM which are processed in parallel by the feature extractors and classified by the SVM accelerator. If a sample is classified as an outlier, an interrupt is sent to the OS to encrypt and store the surrounding 10 minutes of data. This enables a remote clinician to access and annotate the irregular segments and re-train the implant SVM. The OC-SVM can be retrained and uploaded to the PL to refine the outlier detection segments and reduce the volume of data to be sent to the analyst for review.

The key design constraints for the microserver include the completion of feature extraction and anomaly classification

TABLE I. FPGA Resource Utilization

XC7Z045 Resource	Feature Extractor	RBF SVM	Used*	Total	%
BRAMs	42	18	60	1090	5.5
DSPs	16	46	62	900	6.8
Flip-Flops	2262	3863	6805	437200	1.6
LUTs	2960	7400	10360	218600	4.7

** for 16 Channels, 5 frequency bands and 5 EDM decay coefficients.*

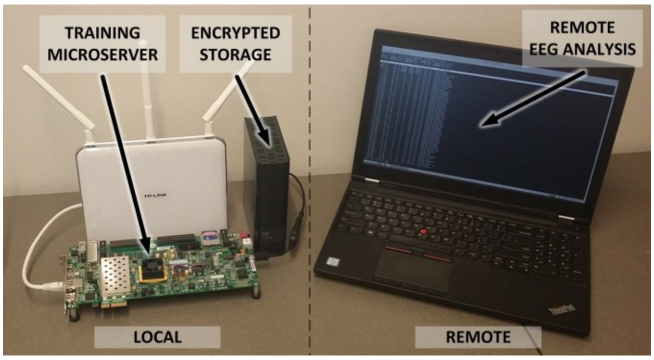

Fig. 4. Xilinx ZC-706 Prototype Training Microserver.

with a speed at least equal to the devices recording sample rate, while minimizing power consumption. The processing latency for feature extraction and classification is 1.259 ms and 0.189 ms respectively, enabling a maximum sample rate of 690 Hz for 16 recording channels. The dynamic power consumption for the system is 0.698 W at a sample rate of 256 Hz. However, due to the low FPGA fabric utilization in the system prototype (Table I), a smaller, lower-power device could be utilized for a portable system deployment.

VII. Training Method

System functionality is demonstrated using the EU intracranial EEG epilepsy database with expert-annotated clinical and subclinical seizure events [12]. Patients were selected based on a postoperative outcome of Engel class I, indicating that intracranial electrodes were positioned at an informative location. After the first 24 hours of neural recordings are accumulated, feature extraction is performed to generate the initial training set. Expert-labelled subclinical and clinical seizure events are then removed along with the surrounding 10 minutes of recordings. The OC-SVM model is trained and stored on the FPGA fabric along with feature normalization coefficients used for the training data.

Minimizing the misclassification of normal physiological neural activity while ensuring that all pathological activity is captured is a key consideration. To enable this tradeoff, the classifier output is smoothed using a moving average window which can be increased at the expense of detection latency.

Once highlighted activity has been annotated by a remote clinician, a refined supervised model can be trained on the microserver to be uploaded to the implanted device. SVM training is performed on the Zynq SoC's dual-core CPUs using a LibSVM implementation. The required training time scales linearly with the number of features used on the

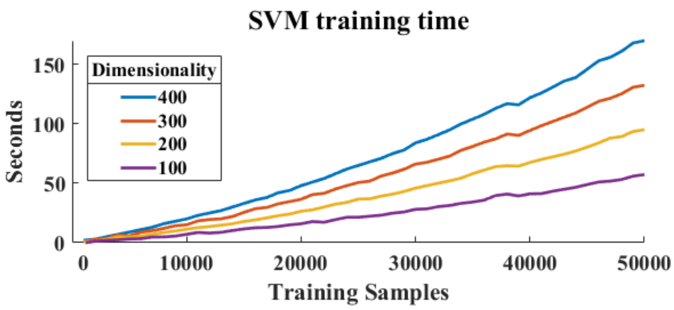

Fig. 5. SVM training time on dual-core ARM Cortex-A9 CPUs.

implanted device, and the FPGA fabric (Fig. 5). The number of training vectors is constrained by external memory to 50,000 with a dimensionality of 400. Incremental training can be performed to enable the use of larger volumes of data.

VIII. Results

The performance of the system is validated using 500 hours of iEEG data across four subjects in the expert-labelled EU epilepsy database. A combination of 16 depth and surface electrodes was determined on a per patient basis based on proximity to the seizure onset zone. The feature extraction implementation uses five spectral bands per channel, each with α decay coefficients, of 4, 6, 8, 10, 12, 14 and 16. The resulting feature vector has a dimensionality of 560. An illustration of the feature space for an electrode placed in the seizure onset zone is shown in Fig. 6 (b). The resulting OC-SVM output is shown in Fig. 6 (c) where the detected outlier activity corresponds to an expert labelled clinical seizure onset time segment.

The system performance is outlined in Table II, where a seizure detection rate of 97.05% is achieved. In the context of iEEG anomaly detection, the system performance has been optimized to capture seizure-like events at the expense of a high detection rate. The alarms per hour is reported using forward chaining validation with a moving average smoothing window length determined on a per-patient basis. These alarms could indicate unlabeled ictal activity and so are not considered false positives as in the case of supervised learning performance analysis.

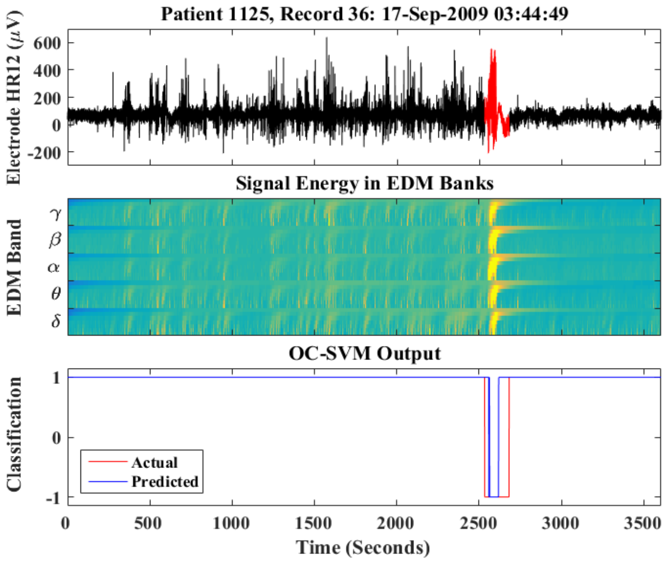

Fig. 6. (a) Recorded neural signals, (b) extracted features and (c) the classifier output, where red indicates the expert-labels and blue is the output calculated on the FPGA fabric.

TABLE II. EU Database System Performance

Patient	EU1096	EU442	EU548	EU1125
Hours Analyzed	147	118	129	108
Clinical Seizures	7	8	17	12
Seizures Detected	7	8	15	12
Alarms Per Hour	0.81	1.21	1.47	1.33
Detection Rate	100%	100%	88.2%	100%

IX. Conclusions

A machine learning microserver is presented to enable continuous post-implantation adaptation for personalized seizure-control neuromodulation devices. The system demonstrates the efficacy of OC-SVMs to assist in the labelling of complex iEEG recordings for training supervised learning models. The concept of the patient-localized microserver addresses the need for life-long learning in personalized biomedical devices. As chronic recording implants become more prevalent, the accumulation of larger volumes of ictal data using this method could enhance the performance of supervised learning classifiers, improving the treatment efficacy and the quality of life for patients.

X. References

[1] M.J. England, C.T. Liverman, A.M. Schultz, and L.M. Strawbridge. "Epilepsy across the Spectrum: Promoting Health and Understanding," Washington, DC: The National Academies Press, 2012.

[2] R. Wennberg, "Introduction to EEG for Nonepileptologists Working in Seizure Prediction and Dynamics." In Epilepsy: The Intersection of Neurosciences, Biology, Mathematics, Engineering, and Physics, pp. 23-39. CRC Press, 2011.

[3] K.H. Lee and N. Verma, "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, 2013.

[4] M. A. B. Altaf, C. Zhang, and J. Yoo, "A 16-channel patient-specific seizure onset and termination detection SoC with impedance-adaptive transcranial electrical stimulator," IEEE J. Solid-State Circuits, vol. 50, no. 11, pp. 2728–2740, Nov. 2015.

[5] A. Shoeb, "Application of Machine Learning to Epileptic Seizure Onset Detection and Treatment," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, 2009.

[6] SR Haut, et. al. "Interrater reliability among epilepsy centers: multicenter study of epilepsy surgery," Epilepsia 43, no. 11 (2002): 1396-1401.

[7] R.S. Fisher et. al., "High-frequency EEG activity at the start of seizures," J Clin Neurophysiol. 9:441–8. 1992

[8] B. Schölkopf, J.C. Platt, J. Shawe-Taylor, A.J. Smola, R.C. Williamson, "Estimating the support of a high-dimensional distribution," Technical report, Microsoft Research, MSR-TR-99-87, 1999.

[9] H. Ung et. al., "Temporal behavior of seizures and interictal bursts in prolonged intracranial recordings from epileptic canines," Epilepsia, 57(12), 1949-1957.

[10] D Groppe et. al. "Dominant frequencies of resting human brain activity as measured by the electrocorticogram," NeuroImage. 79(C), 1–11.

[11] Chih-Chung Chang and Chih-Jen Lin, "LIBSVM : a library for support vector machines," ACM Trans. on Intelligent Systems and Techn., 2:27:1--27:27, 2011.

[12] Matthias Ihle et. al. "EPILEPSIAE - A European epilepsy database," Comput. Methods Prog. Biomed. 106, 3, 2012), pp. 127-138.