Multi-Task Learning Mixed-Signal Classifier for In-situ Detection of Atrial Fibrillation and Sepsis

Sudarsan Sadasivuni†1, Sumukh Prashant Bhanushali†1, Sai Srinivasa Singamsetti†, Imon Banerjee*, and Arindam Sanyal† (1equal contributing first-authors)
*Department of Biomedical Informatics, Emory University, Atlanta GA 30322, USA.
†Department of Electrical Engineering, University at Buffalo, Buffalo, NY 14260, USA. Email: ssadasiv@buffalo.edu

2021 IEEE Biomedical Circuits and Systems Conference (BioCAS) | 978-1-7281-7204-0/21/$31.00 ©2021 IEEE | DOI: 10.1109/BIOCAS49922.2021.9644994

***Abstract*—This paper presents an on-chip analog machine learning (ML) classifier IC for detecting atrial fibrillation (AFib) and sepsis from electrocardiogram (ECG) signal. The proposed technique allows continuous in-situ health surveillance using wearables with embedded AI for early detection of underlying health issues. The analog classifier uses custom activation function and performs in-memory computation (IMC) with switched-capacitor circuits for reduced data movement. Designed in 65nm, the test chip achieves average accuracy of 98.2% for AFib detection, and 90.7% for predicting sepsis 4 hours before onset. The energy efficiency of the test-chip is 12.9nJ/classification which is $4\times$ better than state-of-the-art.**

***Index Terms*—Machine learning, atrial fibrillation, sepsis, mixed-signal classifier and in-memory computation**

I. Introduction

Atrial fibrillation (AFib) results in more than 150,000 underlying cause of deaths in the USA annually. However, most AFib patients are asymptomatic, leading to reduced awareness and less chances of managing stroke risks. Sepsis is another significant cause of death in the USA with close to 40% mortality rate after onset and with 80% of the patients having onset outside hospital. Abnormalities in electrocardiogram (ECG) signal of patients can be used as an early indicator for both AFib and sepsis onset. Continuous health surveillance using wearables with built-in artificial intelligence (AI) is a potential solution for risk management while securing patient data privacy. However, AI analysis is typically computationally intensive and it is difficult to embed AI model within resource constrained wearables. Approaches to reduce energy consumption for AI analysis involve low-precision computation and in-memory computation (IMC) to reduce data movement. State-of-the-art mixed-signal AI circuits have typically only demonstrated on-chip vector matrix multiplication (VMM) or the first hidden/convolutional layer, with rest of the AI model implemented in software [1], [2].

This work presents a mixed-signal, multi-task learning (MTL) classifier for detecting AFib and sepsis from temporal ECG signal (Fig. 1(a)), with the 3-layer artificial neural network (ANN) classifier implemented on-chip. The MTL ANN model gives high accuracy for both AFib and sepsis prediction tasks, since both tasks are fundamentally identifying abnormalities in ECG signal. The key contributions of this work are - 1) demonstration of fully integrated analog MTL ANN with switched-capacitor IMC, 2) custom activation functions that leverage intrinsic analog nonlinearity, and 3) an error-aware AI training methodology that trains the ANN with circuit models to ensure good match between software model and ANN circuit. Compared to SRAM based IMC, switched-capacitor IMC has better linearity for VMMs (see Fig. 1(b)). Linearity of VMM using SRAM cells is fundamentally limited by nonlinear relationship between discharge current (I_{ds}) and voltage on BL/BLB, and the VMM results are not linear over the full dynamic range [1]. In addition, matching capacitors is easier than matching transistors and I_{ds} across large SRAM array. The proposed classifier is demonstrated on single lead ECG data from Physionet 2017 dataset and on data from a patient cohort admitted to Emory University Hospital (EUH) between 2014 to 2018.

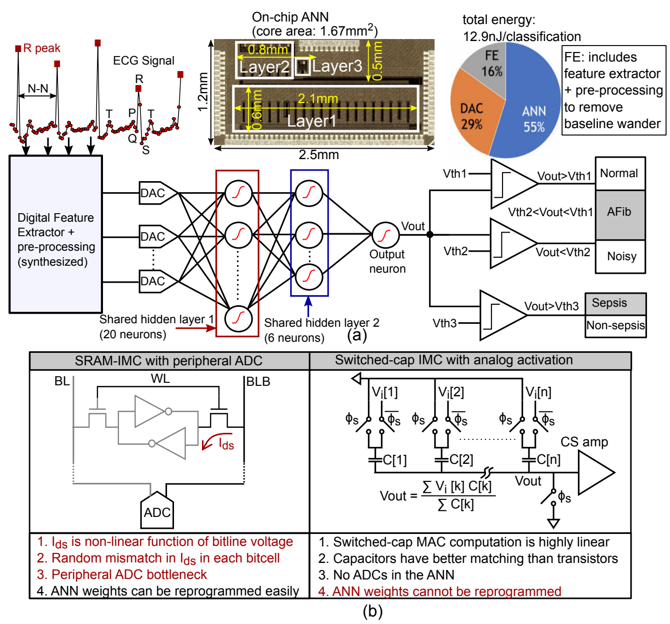

Fig. 1: a) Multi-task learning ANN for atrial fibrillation and sepsis prediction from ECG signal, b) comparison of IMC with SRAM vs switched-capacitor

II. Multi-task learning ANN Training

A. Description of dataset

The 2017 PhysioNet dataset comprises of ECG recordings lasting from 9 seconds to over 60 seconds. The ECG recordings are sampled at 300Hz, and contain normal sinus rhythm, AFib and noisy data of 5971 patients. The sepsis dataset is obtained from EUH with approval from Emory Institutional Review Board,. The cohort consisted of 800 patients admitted to the ICUs at two hospitals within the Emory Healthcare system from 2014 to 2018. For each patient, there is at least 8 hours of ECG signal recordings, sampled at 300Hz, from the time of admission in the ICU. 400 patients in the cohort had sepsis, with onset time assigned using Sepsis-3 criterion. The goal of this work is to detect sepsis 4 hours before onset to allow adequate time for the 3-hour recommended sepsis treatments that have been shown to significantly improve sepsis outcomes [3].

B. Feature extraction

The input features are calculated on 30 second windows of the ECG signal, and only time-domain features are used for low-cost hardware implementation. The time-domain features are calculated from first-order statistical measures of shape, dispersion, location and distribution of R peaks, QRS complexes, PR intervals, ST intervals, QT intervals, and NN intervals (see Fig. 1) in the single-lead ECG signal, and involves calculation of standard deviation, mean, median and maximum/minimum values. 63 time-domain features are used for AFib detection, while a subset of 14 features, computed on NN intervals and R peaks, are used for sepsis prediction. The feature extractor (FE) removes baseline wander from ECG signal by subtracting the median value from each segment.

C. ANN model training and circuit design

A three-layer MTL ANN is trained for detecting AFib and predicting sepsis onset. The first two hidden layers have 20 and 6 neurons respectively, and use a custom tanh activation function, while the output layer uses custom softmax activation. The result of softmax activation at the output neuron is compared with threshold voltages for classification into 'normal/AFib/noisy' and 'sepsis/non-sepsis' categories as shown in Fig. 1. The threshold voltages for the two prediction tasks are calculated during the ANN training phase to optimize loss function for each task. Fig. 2 shows circuit schematic of neurons in the hidden and output layers. Switched-capacitor circuits are used to perform in-memory, charge-domain multiply-and-accumulate (MAC) operations. Bottom-plate sampling technique is used to suppress charge injection.

The custom activation functions are realized using common-source differential amplifiers as shown in Fig. 2. The custom tanh activation circuit uses fully differential common-source amplifier with output offset cancellation as shown in Fig. 2(a). During sampling phase, ϕ_s, the differential inputs of the amplifier are shorted together, and the offset is stored in the capacitor, C_{off}, which is subtracted from the amplifier output during evaluation phase. The custom softmax activation circuit uses a common-source differential amplifier with single ended output as shown in Fig. 2(b).

To ensure that software training results with custom activation functions match IC measurements, we apply a hardware-

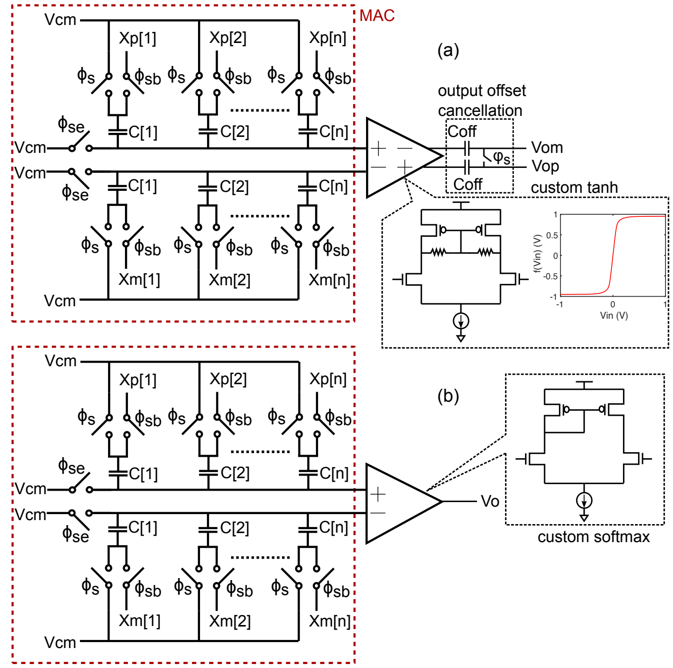

Fig. 2: Circuit schematic of a) hidden neuron with custom tanh activation b) output neuron with custom softmax neuron

software co-design methodology [4]. The amplifier transfer curve and its derivative are imported into the ANN training in Matlab which starts with random weights. Stochastic gradient descent function is used to optimize the ANN weights at each epoch by minimizing the loss function. Once the ANN is fully trained, the weights are encoded as capacitor values in the MAC circuits. The ANN weights are quantized to 4-bit in the hidden layers, and 6-bit in the output layer. The weight quantization is done during the training iterations to preserve accuracy during testing. 4fF unit capacitor, with mismatch standard deviation of 0.4%, is selected as LSB weight in the MAC circuits to ensure that classification accuracy remains close to 99% even in presence of random mismatch.

III. Measurement Results

Fig. 1(a) shows the microphotograph of the test-chip with core area of 1.67mm^2. The FE, and DACs to convert digital features into analog signals, are implemented off-chip. The on-chip ANN consumes 7.1μW for each inference while operating from 1.1V power supply at 1kHz, resulting in an energy consumption of 7.1nJ/inference. The DACs and digitally synthesized FE are estimated to consume 3.8nJ and 2nJ respectively for each inference. Thus, the test-chip has an estimated energy consumption of 12.9nJ/inference. The energy consumption will increase to 13.6nJ/inference if analog front-end amplifier and 14-bit ADC for digitizing ECG signal is integrated on-chip.

A. AFib detection results

The AFIb dataset is randomly split into 4767 training samples and 1204 test samples. Fig. 3(a) shows the measured confusion matrix on the test set. The test-chip detects AFib with 98.8% accuracy, specificity of 1 and sensitivity of 0.89. Threshold voltages for the class boundaries are calculated during foreground calibration step that applies the training samples to the test-chip and calculates the threshold voltages to maximize classification accuracy on the training samples. Fig. 3(b) shows measured histogram of accuracy and sensitivity for 1000 repeated evaluations on the test set. Small standard deviation in accuracy and sensitivity for repeated evaluations demonstrate robustness against noise.

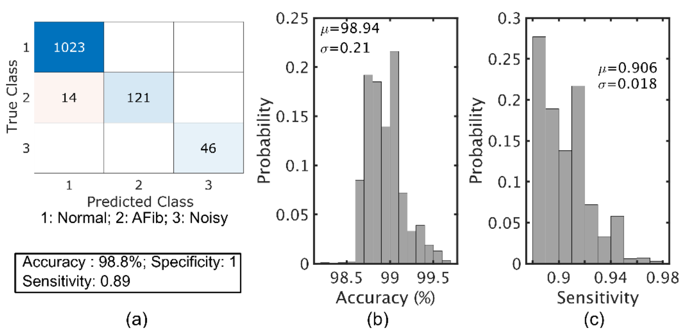

Fig. 3: a) Measured confusion matrix for AFib dataset, b) accuracy and sensitivity for 1000 evaluations

Fig. 4(a) shows the measured accuracy and sensitivity as the power supply voltage is swept from 1.2V to 0.8V. Classification accuracy reduces with supply voltage. Fig. 4(b) shows the measured accuracy and sensitivity for 4 test chips. The class boundaries are calculated for each test-chip through foreground calibration using training samples. The average accuracy and sensitivity across 4 test-chips are 98.2% and 0.89 respectively. Table I compares our prototype with state-of-the-art ASICs demonstrated on AFib detection tasks. The proposed ANN consumes the lowest energy thanks to analog ANN.

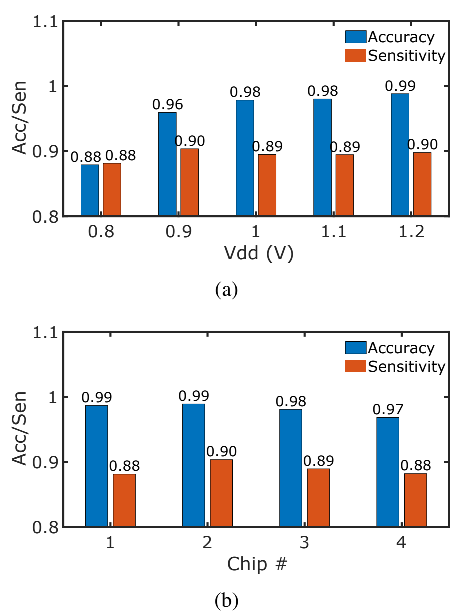

Fig. 4: Measured accuracy and sensitivity for a) as a function of supply voltage, b) multiple chips

TABLE I: Comparison with state-of-the-art ASICs

	JSSC 2019 [5]	TBioCas 2019 [6]	JSSC 2020 [7]	ISSCC 2021 [8]	This work
Process	65nm	180nm	40nm	65nm	65nm
Area (mm^2)	5.9	0.92	0.24	1.74	1.67
Accuracy	–	99.3%	96%	99.3%	98.2%[1]
Type	digital				AMS[2]
Energy	0.33μJ	3.21μJ	51.6nJ	2.25μJ	**12.9nJ**
Model	ANN	ANN	TDDL[3]	ANN	ANN
Class #	2	5	2	2	3

[1]average of 4 chips, [2]AMS: analog/mixed-signal, [3] TDDL: task-driven dictionary learning

B. *Sepsis prediction results*

The sepsis dataset is split randomly into 620 training samples and 180 test samples. Fig. 5 shows the measured accuracy and sensitivity of sepsis prediction as a function of time before onset. As expected, the prediction accuracy and sensitivity improves closer to onset. This work predicts sepsis 4 hours before onset to allow sufficient time for sepsis treatment [3].

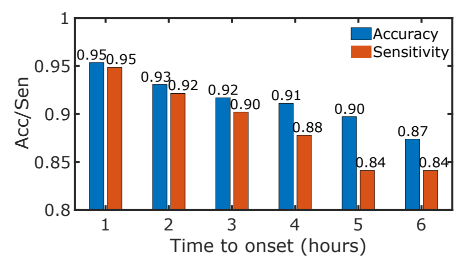

Fig. 5: Measured accuracy and sensitivity of sepsis prediction before onset

Fig. 6(a) shows the measured confusion matrix on the test set. The test-chip predicts sepsis with 91.1% accuracy, specificity of 0.94 and sensitivity of 0.88. Similar to AFib dataset, threshold voltage for the class boundaries are calculated on the training samples to maximize prediction accuracy on the train set. Fig. 6(b) shows measured histogram of accuracy and sensitivity for 1000 repeated evaluations on the test set. Small standard deviation in accuracy and sensitivity for repeated evaluations demonstrate robustness against noise.

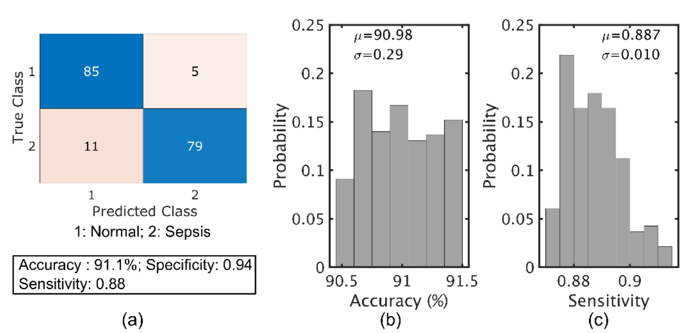

Fig. 6: a) Measured confusion matrix for sepsis dataset, b) accuracy and sensitivity for 1000 evaluations

Fig. 7(a) shows the measured accuracy and sensitivity as the power supply voltage is swept from 1.2V to 0.8V. Prediction accuracy reduces with supply voltage. Fig. 7(b) shows the measured accuracy and sensitivity for 4 test chips. The average accuracy and sensitivity across 4 test-chips are 90.7% and 0.84 respectively. Table II compares our sepsis prediction work with

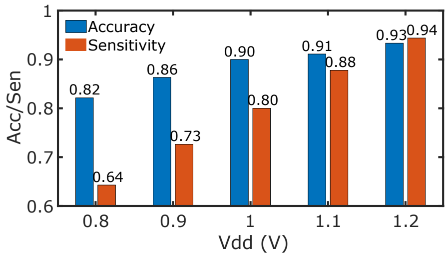

(a)

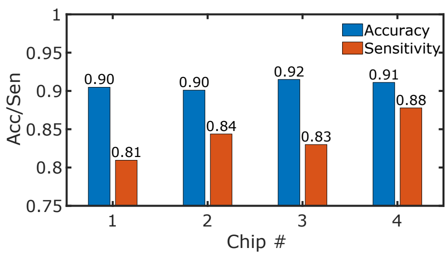

(b)

Fig. 7: Measured accuracy and sensitivity for a) as a function of supply voltage, b) multiple chips

state-of-the-art. Our work has the highest prediction accuracy 4 hours before sepsis onset. To the best of our knowledge, there are no custom ASICs in the literature that perform sepsis prediction. Fig. 8 compares our work with digital baseline synthesized in 65nm, and with state-of-the-art AI ASICs for different bio-medical applications.

TABLE II: Comparison with state-of-the-art

	CinC 2019 [9]	CCM 2019 [10]	JAMIA 2020 [11]	Nature 2021 [12]	This work
Accuracy	84.5%	67%	–	–	**90.7%**[1]
Sensitivity	0.66	0.85	0.84	0.86	0.84[1]
$\mathbf{t}_{onset}$[2]	4 hours				
Model	LSTM[3]	SM[4]	RNN[5]	RF[6]	ANN

[1]average of 4 chips, [2]hours before sepsis onset; [3]LSTM: long-short term memory; [4]SM: survival model; [5]RNN: recurrent neural network; [6]RF: random forest

IV. Conclusion

This work has presented an analog machine learning classifier IC for AFib detection, and sepsis prediction from patient ECG signal. The combination of switched-capacitor IMC and custom analog activation circuits results in $4\times$ improvement in energy efficiency without sacrificing prediction accuracy.

Acknowledgment

This material is based on research sponsored by Air Force Research Laboratory under agreement number FA8650-18-2-5402. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

Comparison with digital baseline

	MTL-ANN	Digital baseline
Accuracy	98.2%[1] (AFib)	99% (AFib)
	90.7%[1] (Sepsis)	91.2% (Sepsis)
Energy/ inference	12.9nJ[2] (13.6nJ)[4]	22.2nJ[3] (22.9nJ)[4]

[1] average of 4 test-chips; [2]energy of ANN, DAC and FE; [3]energy of FE and digital ANN; [4]after including estimated energy of analog front-end and 14-bit ADC

Comparison with AI ASICs

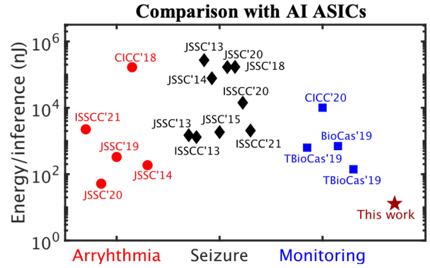

Fig. 8: Comparison with digital baseline and state-of-the-art

References

[1] Q. Dong *et al.*, " A 351TOPS/W and 372.4 GOPS compute-in-memory SRAM macro in 7nm FinFET CMOS for machine-learning applications," in *ISSCC*, 2020, pp. 242–244.

[2] H. Valavi, P. J. Ramadge, E. Nestler, and N. Verma, "A mixed-signal binarized convolutional-neural-network accelerator integrating dense weight storage and multiplication for reduced data movement," in *IEEE Symposium on VLSI Circuits*, 2018, pp. 141–142.

[3] L. Pruinelli *et al.*, "Delay within the 3-hour surviving sepsis campaign guideline on mortality for patients with severe sepsis and septic shock," *Critical care medicine*, vol. 46, no. 4, p. 500, 2018.

[4] S. T. Chandrasekaran *et al.*, "Fully Integrated Analog Machine Learning Classifier Using Custom Activation Function for Low Resolution Image Classification," *IEEE TCAS–I*, vol. 68, no. 3, pp. 1023–1033, 2021.

[5] S. Yin *et al.*, "A 1.06-μ W Smart ECG Processor in 65-nm CMOS for Real-Time Biometric Authentication and Personal Cardiac Monitoring," *IEEE JSSC*, vol. 54, no. 8, pp. 2316–2326, 2019.

[6] Y. Zhao, Z. Shang, and Y. Lian, "A 13.34 μW Event-Driven Patient-Specific ANN Cardiac Arrhythmia Classifier for Wearable ECG Sensors," *IEEE transactions on biomedical circuits and systems*, vol. 14, no. 2, pp. 186–197, 2019.

[7] K.-C. Chen, C.-Y. Chou, and A.-Y. Wu, "A Tri-Mode Compressed Analytics Engine for Low-Power AF Detection With On-Demand EKG Reconstruction," *IEEE Journal of Solid-State Circuits*, 2020.

[8] J. Liu *et al.*, "BioAIP: A Reconfigurable Biomedical AI Processor with Adaptive Learning for Versatile Intelligent Health Monitoring," in *ISSCC*, vol. 64, 2021, pp. 62–64.

[9] B. Roussel, J. Behar, and J. Oster, "A Recurrent Neural Network for the Prediction of Vital Sign Evolution and Sepsis in ICU," in *Computing in Cardiology (CinC)*. IEEE, 2019, pp. Page–1.

[10] S. Nemati *et al.*, "An interpretable machine learning model for accurate prediction of sepsis in the ICU," *Critical care medicine*, vol. 46, no. 4, p. 547, 2018.

[11] A. D. Bedoya *et al.*, "Machine learning for early detection of sepsis: an internal and temporal validation study," *JAMIA open*, vol. 3, no. 2, pp. 252–260, 2020.

[12] K. H. Goh *et al.*, "Artificial intelligence in sepsis early prediction and diagnosis using unstructured data in healthcare," *Nature communications*, vol. 12, no. 1, pp. 1–10, 2021.