STINT: Selective Transmission for Low-Energy Physiological Monitoring

Tao-Yi Lee
Khuong Vo
taoyil@uci.edu
khuongav@uci.edu
University of California, Irvine

Wongi Baek
qordnjsrl13@cau.ac.kr
Chung-Ang University
Seoul, Korea

Michelle Khine
Nikil Dutt
mkhine@uci.edu
dutt@ics.uci.edu
University of California, Irvine

ABSTRACT

Noninvasive, and continuous physiological sensing enabled by novel wearable sensors is generating unprecedented diagnostic insights in many medical practices. However, the limited battery capacity of these wearable sensors poses a critical challenge in extending device lifetime in order to prevent omission of informative events. In this work, we exploit the inherent *sparsity* of physiological signals to intelligently enable selective transmission of these signals and thereby improve the energy efficiency of wearable sensors. We propose STINT, a selective transmission framework that generates a sparse representation of the raw signal based on domain-specific knowledge, and which can be integrated into a wide range of resource-constrained embedded sensing IoT platforms. STINT employs a neural network (NN) for selective transmission: the NN identifies, and transmits only the informative parts of the raw signal, thereby achieving low power operation. We validate STINT and establish its efficacy in the domain of IoT for energy-efficient physiological monitoring, by testing our framework on EcoBP – a novel miniaturized, and wireless continuous blood pressure sensor. Early experimental results on the EcoBP device demonstrate that the STINT-enabled EcoBP sensor outperforms the native platform by 14% of sensor energy consumption, with room for additional energy savings via complementary bluetooth and wireless optimizations.

CCS CONCEPTS

• **Hardware** → **Power estimation and optimization**; • **Networks** → **Link-layer protocols**; • **Computing methodologies** → **Regularization**; **Neural networks**.

KEYWORDS

computation-communication trade-off, continuous monitoring, edge deep learning, Bluetooth, IoT healthcare

ACM Reference Format:

Tao-Yi Lee, Khuong Vo, Wongi Baek, Michelle Khine, and Nikil Dutt. 2020. STINT: Selective Transmission for Low-Energy Physiological Monitoring. In *ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED '20)*, August 10–12, 2020, Boston, MA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3370748.3406563>

This work is licensed under a Creative Commons Attribution International 4.0 License.
ISLPED '20, August 10–12, 2020, Boston, MA, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7053-0/20/08.
<https://doi.org/10.1145/3370748.3406563>

1 INTRODUCTION

Continuous monitoring of various physiological signals have enabled new insights for clinical diagnosis and post-treatment medical follow-ups, especially for chronic diseases. For instance, electrocardiogram (ECG) can highlight events that warrant medical intervention for cardiovascular diseases such as atrial fibrillation. Traditionally, continuously-monitored ECG is recorded by a Holter monitor which is both very cumbersome, and also does not facilitate real-time detection of critical events in everyday settings.

Recently, Koole et al. [1] reported a smart and wearable ECG monitor contributing to timely intervention of a progressing congenital heart disease. The key enabler of such diagnostic advantage is the real-time machine learning model that identifies and classifies each single heartbeat into categories significant in the medical context. Generally, not all categories are equally useful in diagnosis. This presents an opportunity to perform selective data transmission, where data with less contextual meaning can be acquired, processed, and transmitted with much *lower precision*, if not *altogether dropped*, to achieve low power operation. To illustrate this opportunity, this paper investigates selective data transmission for low power operation in the context of a wearable continuous noninvasive blood pressure (CNBP) sensor. We present techniques to selectively transmit data samples of radial pulse waves (*i.e.*, a type of CNBP) to achieve lower power operation and minimum diagnostic sensitivity loss, on a novel custom hardware platform named EcoBP.

Our paper makes the following contributions:

- We present STINT, a selective transmission framework that generates a sparse representation of a raw signal based on domain-specific knowledge, and which can be integrated into a wide range of resource-constrained embedded sensing IoT platforms.
- We develop a multilayer perceptron (MLP) neural network that efficiently approximates the sequence-to-sequence relationship between successive input signals and labels, and which can be deployed on resource constrained IoT microcontroller units (MCUs).
- We present a real case study of applying STINT on EcoBP, a novel wearable piezocapacitive continuous blood pressure sensor platform.
- Early experimental results demonstrate energy savings of 14% on the native EcoBP platform, with room for additional energy savings via communication protocol and hardware optimizations.

The subsequent sections are organized as follows. Section 2 presents related work on low-power sensing techniques. Section 3 defines the selective transmission problem on a resource-constrained device, and presents our neural network formulation. Section 4 presents our case study applying STINT on EcoBP, a novel wearable CNBP platform. Section 5 presents our experimental results and analysis of energy savings using STINT on the EcoBP platform. Section 6 concludes the paper.

2 BACKGROUND AND RELATED WORK

Many algorithmic schemes have been proposed to optimize the computation-communication trade-offs in sensor networks. The availability of efficient embedded computation can help the communication system to minimize usage of either energy or bandwidth which are scarce resources in many IoT applications.

Samie et al. [2] introduced a two-level architecture in which a lightweight yet accurate classifier (LYAC) decides to off-load samples to the fog layer for computation and the other LYAC classifies samples that are chosen to be processed locally. Lee et al. [3] proposed a joint cross-layer optimization to optimize the data transmission strategy based on the objectives guided by the application.

Compressive sensing (CS) [4] is also a promising technique that projects the raw signal from sensors to a smaller subspace before transmission, thereby reducing communication energy. Dixon et al. [5] and Roy et al. [6] both adopted CS to reduce acquisition and communication energy in ECG systems.

Hardware techniques such as the Bluetooth Low Energy (BLE) platform proposed by Wang et al. [7] specifically optimizes for wearable medical devices by employing various best practices in power management, and a low energy communication protocol. Recent efforts have investigated extending the battery-lifetime of a BLE platform to several days by exploiting a photovoltaic cell, and maximum power point tracking for energy harvesting [8]. These techniques relax the overall system energy budget.

However, thanks to the compact form factor, most wearable IoT sensors are still energy sensitive and have limited computational resources. We address the challenges arose in computation-energy trade-off (instead of general computation-communication trade-off that covers bandwidth limits) on the edge devices by complementing prior efforts in following two dimensions:

- 1) we present STINT, a novel *selective transmission* strategy to realize wearable low power IoT sensors, and
- 2) we apply this scheme on EcoBP, a novel CNBP platform that integrates a piezocapacitive sensor, and power management circuits to balance noise performance and energy saving for continuous blood pressure monitoring.

3 OUR SOLUTION: STINT

Any physiological signal of length T can be represented as a time-series $\mathbf{x} = [x_T, x_{T-1}, \dots, x_1] \in \mathbb{R}^{1 \times T}$. Suppose an oracle $O(\cdot)$ generates truthful low-dimension representation $\mathbf{y}$ of $\mathbf{x}$ for a given task $\mathcal{T}$ (e.g., deciding whether $\mathbf{x}$ is indicative of a disease. In this case $\mathbf{y} \in \mathbb{R}^1$).

3.1 Selective Transmission

We can now define the problem of *selective transmission*.

Figure 1: (a) Native system transmission (b) STINT-enabled selective transmission. SIPO: serial-input-parallel-output; PISO: parallel-input-serial-output.

A selector $s(\mathbf{x}) \in \mathbb{R}^{1 \times T} \rightarrow [\text{false}, \text{true}]^{1 \times T}$

$$s(\mathbf{x}) \triangleq \tilde{s}(\mathbf{x}; \theta) \geq \tau \quad (1)$$

is a function that maps the raw signal into a Boolean time-series indicating whether each sample x_t is important or not, based on a decision threshold τ . Determining optimum selector is equivalent to finding

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \sum_{\mathbf{x} \in \mathcal{X}} \left(\|O(\mathbf{x}) - O(\tilde{s}(\mathbf{x}; \theta) \odot \mathbf{x})\|_2 + \lambda \sum_i \mathcal{R}(s_i) \right) \quad (2)$$

$$= \arg \min_{\theta \in \Theta} \sum_{\mathbf{x} \in \mathcal{X}} \left(\|\mathbf{y} - O(\tilde{s} \odot \mathbf{x})\|_2 + \lambda \mathcal{R}(\mathbf{s}) \right) \quad (3)$$

where $\tilde{s} = \tilde{s}(\mathbf{x}; \theta) = [s_T, s_{T-1}, \dots, s_1]$, $\odot$ denotes Hadamard (elementwise) product of two vectors, and λ controls how important it is to minimize the number of selected samples with the regularization term $\mathcal{R}$.

Figure 1(a) shows the native transmission scheme, and 1(b) illustrates STINT's selective transmission scheme. Note that we assume a transmitter only transmits on *nonzero* inputs. The term $\mathbf{s} \odot \mathbf{x}$ in Eq. 3 can be viewed as a *sparse* representation of the raw signal, induced by selector $s(\mathbf{x}; \theta)$. A specific implementation of the selective transmission scheme will be detailed in the following subsection.

3.2 Evaluation Metrics

To evaluate the advantage of STINT, we propose a figure of merit (FoM) η that measures the ratio of Matthews correlation coefficient (MCC) [9] and total energy consumption $\mathcal{E}$ in micro-Joules. Formally, FOM η is defined as

$$\eta = \frac{\text{MCC}}{\mathcal{E}} \quad (4)$$

$$= \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\mathcal{E} \sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}} \quad (5)$$

where TP, TN, FP, and FN denote the count of the four possible outcomes of a binary classifier. MCC is chosen over other metrics such as accuracy and F_1 score because of its insensitivity to class imbalance and symmetry between true positives (TP) and true negatives (TN). To further analyze STINT, we need to define two additional quantities, namely recall and precision. The definitions are adapted from binary classification to evaluate the decision quality of selective transmission. Recall, or sensitivity R is defined as the count of transmitted abnormal samples to the raw count of

abnormal samples, and precision P is defined as the ratio between transmitted abnormal samples to the total number of transmitted samples. They are calculated as

$$R = \frac{|\{\hat{y}_i | \hat{y}_i \in \mathcal{A}\}|}{|\{x_i | x_i \in \mathcal{A}\}|} = \frac{TP}{TP + FN} = \frac{TP}{TP + (n^+ - TP)} \quad (6)$$

$$P = \frac{|\{\hat{y}_i | \hat{y}_i \in \mathcal{A}\}|}{|\{\hat{y}_i\}|} = \frac{TP}{TP + FP} = \frac{TP}{TP + (n^- - TN)} \quad (7)$$

where $\mathcal{A}$ is the set of samples that are considered abnormal (or suspicious). n^+ and n^- denote the numbers of condition positive and the numbers of condition negative, respectively. Higher recall is desirable in STINT because less abnormal samples are missed. On the other hand, high precision is also desired because less normal, or unnecessary samples are transmitted. Note that MCC has the difference between the product of TP and TN and the product of FP and FN in the numerator. Maximizing MCC is equivalent to jointly maximizing precision and recall — a desirable state of STINT. One may argue that maximizing accuracy (defined as the ratio of TP+TN and TP+TN+FP+FN) seems to be another valid objective. This is not true, because accuracy is sensitive to class imbalance. For example, if the abnormality event is as rare as 3%, a classifier that always returns normal will achieve an accuracy of $970/1000 = 97\%$ for 1000 transmitted samples. This useless classifier has MCC of 0 which reflects its deficiency.

Both energy $\mathcal{E}$ and MCC will be determined empirically on the EcoBP platform which is described in Section 4. Figure of merit between the native system, and the STINT selective transmission system depicted in Fig. 1a, and Fig. 1b, respectively, will then be compared in Section 5.

3.3 Multilayer Perceptron Realization of STINT

While in principle many candidate artificial neural networks (NNs) could be used to realize the selector function, popular DNN models suffer from excessive resource and memory usage that limit their deployment on resource-constrained MCUs. Consequently, we chose a *multilayer perceptron* (MLP) that enables efficient realization on the MCU thanks to its smaller memory footprint and modest resource usage. The selector $s(\mathbf{x}; \theta)$ that captures the sequence-to-sequence relation between raw signal $\mathbf{x} = x_T, x_{T-1}, \dots, x_1$, and the labeling $\mathbf{y} = y_T, y_{T-1}, \dots, y_1$ can be approximated with an MLP defined by following equations:

$$\mathbf{a}^{(1)} = \text{ReLU}(\mathbf{W}^{(1)} \mathbf{x} + \mathbf{b}^{(1)}) \quad (8)$$

$$\hat{\mathbf{y}} = \sigma(\mathbf{W}^{(2)} \mathbf{a}^{(1)} + \mathbf{b}^{(2)}) \quad (9)$$

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (10)$$

Model parameters θ can be expressed as a 4-tuple $(\mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \mathbf{b}^{(1)}, \mathbf{b}^{(2)})$ where the weight matrices $\mathbf{W}^{(1)}$, and $\mathbf{W}^{(2)}$ are masked such that the activations $\mathbf{a}_i^{(1)}$, and $\hat{\mathbf{y}}_i$ at timestamp i do not depend on input samples in the future (i.e., $x_{i+1}, \dots, x_T$). Equivalently, function $s(\mathbf{x}; \theta)$ is a concatenation of T time-series regression models, each of which estimates the abnormality of a sample i based on i samples of history.

Figure 2: Hardware block diagram of EcoBP

Figure 3: The EcoBP system

Training of the selective transmitter on dataset $\mathcal{D}$ can be achieved by minimizing categorical cross entropy loss between the ground truth targets $\mathbf{y}$ and the estimates $\hat{\mathbf{y}}$. That is,

$$\hat{\theta} = \arg \min_{\theta \in \Theta} -\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^T \mathbf{y}_{ij} \log \hat{\mathbf{y}}_{ij} + \lambda (\hat{\mathbf{y}}_{ij} - \hat{\mathbf{y}}_{ij} \log \hat{\mathbf{y}}_{ij}) \quad (11)$$

where M denotes the number of training samples in $\mathcal{D}$, and $\mathbf{y}_{ij}$ signifies the abnormality at timestamp j of i -th training sample. L_1 regularization is imposed on the output as well as negative entropy is added to penalize confident output distributions [10]. This regularization technique is invariant to the parameterization of the underlying neural network.

4 CASE STUDY: ECOBP

4.1 System Overview

EcoBP, is our novel wearable piezocapacitive continuous blood pressure sensor platform. The block diagram of EcoBP is illustrated in Fig. 2. The first prototype of the hardware is based on Mbitent Lab MMR board which is designed with Nordic Semiconductor BLE SoC solution nRF52832. We extended the board with additional capacitance-to-digital (CDC) converter and a proprietary capacitive sensor to read pulse waves from radial arteries. Fig. 3a shows a prototype EcoBP.

4.1.1 MCU and Bluetooth Radio. The main SoC nRF52832 integrates an ARM-M4 micro-controller unit (MCU) and Bluetooth LE radio. The MCU has access to 512kB flash, 64kB RAM and a floating

point unit (FPU) running at 64 MHz. On SoC level, it is integrated with 32 GPIO pins, a SPI controller, and a I²C controller which we use to communicate with capacitance-to-digital converter (CDC).

4.1.2 Firmware. The firmware of the EcoBP is implemented in C along with nRF SDK version 15. Core application code includes a custom Bluetooth service, and a clone of MLP which is trained offline, using Keras [11]. The Bluetooth service offers a characteristic which *notifies* raw reading of the CDC, the probability of abnormality reported by MLP, and the timestamp associated with. Notification is only issued when the probability is greater than a prescribed threshold τ which can be adjusted through another writable Bluetooth characteristic.

4.1.3 Piezocapacitive Sensor. We deployed a soft flexible capacitive pressure sensor that utilizes nano-structured wrinkled Au thin films (similar to Kim et al. [12] [13]). The wrinkled Au structures provide robust integration with the underlying polymeric substrate allowing the thin film to withstand mechanical deformation, while maintaining electrical performance. More importantly, the soft mechanical characteristics of the sensor allowed for improved interfacing to the human body, enabling accurate and precise measurements of physiological signals, including beat to beat blood pressure (BP). The pressure sensor conforms to the body and measures the pressure from the skin exerted by arterial wall expansion. The measured pressure correlates with physiological beat to beat BP.

The sensor prototype in Fig. 3b is constructed as a parallel plate variable capacitor which responds to pressure by exhibiting a higher nominal capacitance. Pressure sensitivity of EcoBP was improved by incorporating micro-ridged structures within the capacitive based sensor to create a hybrid air-silicone dielectric layer. When compressed, the effective dielectric constant increases due to the reduced air gap volume.

4.1.4 Capacitor to Digital Converter. AD7746 CDC is the device on-board to digitize the capacitance continuously. The 24-bit Δ - Σ analog-to-digital converter is capable of sampling either single or dual channel capacitance at 90Hz output data rate. The capacitance dynamic range is $\pm 4\text{pF}$ which can be further extended to $\pm 25\text{pF}$ using additional range extension circuitry. The CDC communicates with the host MCU on a I²C bus.

5 EVALUATION

5.1 Synthetic Blood Pressure Waveforms

Synthetic blood pressure waveforms are generated based on the principle of pulse decomposition analysis (PDA) [14] for system evaluation. The rationales for using artificial waveforms include: 1) a well-defined, and parameterized generative source enables reproducible system evaluation of the system, and 2) it obviates the need for Institutional Review Board (IRB) approval, thereby further enhancing reproducibility.

We implemented, and open-sourced a beat-to-beat waveform synthesizer in Python whose output can be directly fed into an arbitrary waveform generator. Based on PDA, each pulse $W_j(t; \Delta)$ associated with j -th heartbeat can be represented as a superposition

Figure 4: Hidden Markov model that generates synthetic blood pressure waveform. $s_i \in \{\text{ABN}, \text{NOR}\}$ are binary variables denoting the abnormality. $d_i \in \{0, 1, 2\}$ models the amount of deviation from the mean value given s_i . Δ_A and Δ_N are two randomized data generator sampling from concentric multi-dimensional Gaussian spheres. The range set K of the mixtures is $\{0, 1, 2\}$. For brevity, the outputs are denoted by W_j which represents a time-series. $W_j(t)$

of 5 triangle-Gaussian wavelets that can be expressed as

$$w_i(t; \delta_i) = \mathcal{T}(t; v_t) * \mathcal{G}(t; v_g, \sigma) \quad (12)$$

where $\mathcal{T}(t; v_t)$ denotes an m_t -sample, and unity amplitude triangular function while $\mathcal{G}(t; v_g, \sigma)$ represents the v_g -sample, and unity amplitude Gaussian function with standard deviation of σ . Operator $*$ denotes convolution. For brevity, the parameterization $\delta_i = \{v_t, v_g, \sigma\}$ will be omitted in the following. Therefore, a pulse can be represented as

$$W_j(t; \Delta_j) = \sum_{i=1}^5 \alpha_i w_i(t - \tau_i) \quad (13)$$

where Δ_j is the set of parameters, α_i is the relative amplitude, and τ_i is the time delay between i -th wavelet and the first wavelet. Δ_j is a set which includes all lower case Greek alphabets presented above, that fully describes a pulse.

A realistic recording fragment can then be generated through cascading multiple pulses $W_j(t)$, $\forall j \in [1, n]$. To model normal and abnormal pulses, we introduced events similar to arrhythmia (by varying μ , and τ), and hyper-/hypo-tension (by varying α). These two sets of parameters can then parameterize a dynamic hidden Markov model (HMM) which generates pulses with randomness. Fig. 4 illustrates the HMM which produces pulse wave fragments. The *transition* probabilities between normal and abnormal (e.g. $P(s_t = \text{NOR} | s_{t-1} = \text{ABN}, d_{t-1} = 0)$) are modeled in a fixed conditional probability table.

5.2 Experimental Setup

Our experimental setup is shown in Fig. 5. To exploit the advantage of synthetic waveform, instead of inputting pressure signal through a piezocapacitive sensor, equivalent voltage waveforms are fed into the CDC using arbitrary waveform generator (AWG) to exercise the system in a controlled fashion. Each experimental run is invoked by a Python routine which synchronizes waveform generation, Bluetooth communication, and power monitoring between three parallel processes. Inter-process messaging is setup to align critical events such as beginning and the end of Bluetooth notification.

Figure 5: Experiment setup. Major parts are a Nordic nRF52832 breakout board, J-link programmer, a fuel gauge, 400 mAh LiPo battery, a AD7746 CDC, arbitrary waveform generator, and a I²C reader

Table 1: Performance summary of EcoBP with different decision threshold. Note that when threshold $\tau = 0$, STINT is disabled. λ that leads to maximum FoM is selected for each threshold value. Note that FoM computes energy in μJ instead of mJ. Numbers in parenthesis are relative energy to $\tau = 0$.

τ	Energy (mJ)	P	R	MCC	FoM	λ
0.0	10.19 (1.00)	0.23	1.00	0.000	0.00	0.3
0.25	8.88 (0.87)	0.85	0.94	0.861	96.99	0.7
0.5	8.81 (0.86)	0.95	0.89	0.901	102.30	0.7
0.75	8.79 (0.86)	0.98	0.86	0.896	101.98	0.3
1.0	8.46 (0.83)	1.00	0.09	0.264	31.19	0.2

Fig. 6 is the time-domain example of STINT configured with optimum settings $\tau = 0.5$ and $\lambda = 0.7$. Since the MCU has a built-in floating point unit, we did not quantize the MLP model into integer equivalent which may provide additional energy savings. The top pane of each time domain plot represents the waveform generated by AWG, and the transmitted samples selected by the MLP running on the MCU. The bottom pane shows the ground truth of the abnormality (ABN) or normality (NOR) along with the probability estimate reported by MLP. Note that for illustration purpose, the Bluetooth payload is included with the probability output which can be omitted for lower energy consumption.

5.3 Numerical Results

Table 1 summarizes the performance of STINT among different decision thresholds. The optimal figure of merit 102.30 occurs at $\tau = 0.5$ and $\lambda = 0.7$, where the STINT-enabled system can achieve 14% of relative energy saving over the native system. Overall, the system has 27% gain in terms of FoM. Besides, precision increases significantly from 0.23 at the baseline to 0.95 at the optimal setting,

Figure 6: System input / output when threshold $\tau = 0.5$, $\lambda = 0.7$.

while recall drops from 1 to 0.89. In terms of energy saving, this is desirable since the higher the precision, the lower the total system energy. Nevertheless, depending on different medical contexts, one can choose a more relaxed setting of a lower τ and λ to reduce the negative effect on the recall so that abnormal portions of the raw signal can be completely transmitted.

In our experiments, we observed that Bluetooth parameters (e.g., notification interval and payload-to-header ratio) can greatly influence the observed energy savings. The numbers we reported in Table 1 use 100-ms notification interval, and a payload size of 20 bytes (i.e. 5 samples in one packet); these settings present additional room for energy optimization if the expected rate of abnormality is lower. Due to the limitation of the Bluetooth protocol version 4.2, empty packets (PDU) will still be transmitted to keep the session alive thereby adding another energy burden. These observations indicate the possibility for even larger energy gains, prompting future investigation of wireless protocols that support adaptive notification mechanisms and variable length packets.

Fig. 7-9 detail the energy consumption, MCC, and FoM over the design space defined by output regularization and threshold. The regularizer acts as a positive confidence penalty that is expected to reduce the number of false positives, which in turn cuts down unnecessary transmission energy. The parameter λ controls the strength of the penalty. In Fig. 7, it is evident that as λ increases, the energy consumption slowly declines, though the energy level transition is more rapid with respect to the threshold. As shown in Fig. 8, MCC values are at the highest flat region when the τ is around 0.4 to 0.8, and λ is around 0.2 to 0.8. Also, Fig. 9 depicts a similar contour pattern where the highest FoM values are in the middle region. The pattern proves that the increase in MCC comes with the fall in energy consumption, since the more accurate the model is, the lower the false positives.

6 CONCLUSION AND FUTURE WORK

We proposed the STINT framework for wireless biomedical sensors that minimizes energy consumption by selectively transmitting informative samples. STINT exploits the inherent sparsity of physiological signals and deploys a compact NN to identify and transmit only the informative parts of the raw signal. Experimental results demonstrated 14% of energy saving with minimum degradation of

Figure 7: Energy consumption of STINT over different (τ, λ) .

Figure 9: FoM of STINT over different (τ, λ) .

Figure 8: MCC of STINT over different (τ, λ) .

recall on a real-world EcoBP platform that processes continuous blood pressure readings. We believe STINT has the potential to generate further energy and performance gains through complementary schemes such as the wireless protocol, improved radio transmitter, adaptive/feedback mechanisms to adjust the decision threshold, and the selector function for optimal performance; these are our ongoing and future research directions.

REFERENCES

- [1] MAC Koole et al. “At last, mobile health leading to a diagnosis in a young patient with congenital heart disease”. In: *Netherlands Heart Journal* 27.3 (2019), pp. 162–163.
- [2] Farzad Samie, Lars Bauer, and Jörg Henkel. “Hierarchical Classification for Constrained IoT Devices: A Case Study on Human Activity Recognition”. In: *IEEE IoT Journal* (2020).
- [3] Tao-Yi Lee, Marco Levorato, and Nikil Dutt. “DNN-Assisted Sensor for Energy-Efficient ECG Monitoring”. In: *2019 IEEE GLOBECOM*. IEEE, 2019, pp. 1–6.
- [4] Wolfgang Härdle, Henry Horng-Shing Lu, and Xiaotong Shen. *Handbook of big data analytics*. Springer, 2018.
- [5] Anna MR Dixon et al. “Compressed sensing reconstruction: Comparative study with applications to ECG bio-signals”. In: *2011 IEEE ISCAS*. IEEE, 2011, pp. 805–808.
- [6] Rohan Basu Roy et al. “Sparse Encoding Algorithm for Real-Time ECG Compression”. In: *Recent Trends in Signal and Image Processing*. Springer, 2019, pp. 31–38.
- [7] Anping Wang et al. “EcoBT: Miniature, versatile mote platform based on Bluetooth Low Energy Technology”. In: *2014 IEEE iThings*. IEEE, 2014, pp. 148–154.
- [8] Cheng-Ting Lee et al. “EcoMicro: a miniature self-powered inertial sensor node based on Bluetooth low energy”. In: *Proceedings of the ISLPED*. 2018, pp. 1–6.
- [9] Brian W Matthews. “Comparison of the predicted and observed secondary structure of T4 phage lysozyme”. In: *BBA-Protein Structure* 405.2 (1975), pp. 442–451.
- [10] Gabriel Pereyra et al. “Regularizing neural networks by penalizing confident output distributions”. In: *arXiv preprint arXiv:1701.06548* (2017).
- [11] François Chollet et al. *Keras*. <https://keras.io>. 2015.
- [12] Joshua Kim et al. “Soft Wearable Pressure Sensors for Beat-to-Beat Blood Pressure Monitoring”. In: *Advanced healthcare materials* 8.13 (2019).
- [13] Joshua Kim et al. “Highly stretchable wrinkled gold thin film wires”. In: *Applied physics letters* 108.6 (2016).
- [14] Martin C Baruch et al. “Pulse decomposition analysis of the digital arterial pulse during hemorrhage simulation”. In: *Nonlinear biomedical physics* 5.1 (2011), p. 1.