Sensor-Classifier Co-Optimization for Wearable Human Activity Recognition Applications

Anish NK[1], Ganapati Bhat[1], Jaehyun Park[2], Hyung Gyu Lee[3], Umit Y. Ogras[1]
[1]Arizona State University, USA [2]University of Ulsan, Korea [3]Daegu University, Korea

***Abstract*— Advances in integrated sensors and low-power electronics have led to an increase in the use of wearable devices for health and activity monitoring applications. These devices have severe limitations on weight, form-factor, and battery size since they have to be comfortable to wear. Therefore, they must minimize the total platform energy consumption while satisfying functionality (e.g., accuracy) and performance requirements. Optimizing the platform-level energy efficiency requires considering both the sensor and processing subsystems. To this end, this paper presents a sensor-classifier co-optimization technique with human activity recognition as a driver application. The proposed technique dynamically powers down the accelerometer sensors and controls their sampling rate as a function of the user activity. It leads to a 49% reduction in total platform energy consumption with less than 1% decrease in activity recognition accuracy.**

***Index Terms*—Wearable computing, human activity recognition, IoT, flexible hybrid electronics (FHE), health monitoring.**

I. Introduction

Low-power wearable devices have increasing popularity since they enable a wide range of applications, such as health monitoring, activity tracking, and activity recognition. In particular, health and activity monitoring applications proliferate through wearable devices that integrate motion sensors and local processing capabilities. For example, monitoring the activities of patients with chronic disorders, such as heart disease and diabetes, helps in tracking the progress of the disease over time [2]. Similarly, human activity recognition (HAR) using smartphones and wearable devices enables tracking daily activities, including walk, stand, and exercise [13]. This information helps movement disorder specialists in analyzing specific motor functions outside the clinic. Indeed, the use of smartwatches has sharply risen as they integrate activity and health monitoring features, such as heart rate monitor, sedentary alert, and fitness tracking. Hence, wearable devices facilitate advanced round-the-clock human monitoring, tracking, and activity recognition.

Wearable devices are severely limited in size, weight, and form-factor because they are attached to the body. Consequently, large and heavy batteries are prohibitive. Due to these constraints, the success of wearable devices depends critically on meeting performance requirements using minimum energy consumption. For instance, wearable devices used for HAR must meet minimum recognition accuracy constraint while minimizing the total platform energy. Hence, designers need to co-optimize all major platform components that affect accuracy and energy consumption.

978-1-7281-2437-7/19/$31.00

Optimizing the platform-level energy efficiency requires considering both the sensor and processing subsystems. For instance, human activity recognition consists of two phases - (1) data collection from sensors for an activity period ranging from one to a few tens of seconds, (2) processing the raw data to recognize the target activity. Data collection offers significant energy savings potential since sensors stay on whenever there is user activity. Unlike the sensors, processing resources can enter sleep states during the data collection phase. However, the energy consumption can still be significant despite short active time since the power consumption during the processing is relatively large compared to the power consumption of the sensors.

In this paper, we propose two sensor-classifier co-optimization techniques for wearable IoT devices targeting human activity recognition. The first technique monitors the user activity and operates the accelerometer sensors in a low-power mode (with 50× lower power) until there is a significant change in the user's activity. The accelerometer is triggered to sample the user movements by comparing the change in the user activity against a parameterized threshold. Decreasing this threshold leads to higher sensitivity (*always-on in the limit*), while increasing it reduces energy consumption. Since different thresholds affect the quality of feature data, we also optimize the HAR classifier to match the sensor behavior. The second technique dynamically controls the accelerometer sampling rate to optimize the activity in the sensor and processing requirements in the classifier. Similar to the first technique, we match the classifier to the sensor data rate to minimize the loss in the recognition accuracy. Finally, we combine these two techniques and evaluate them in a custom hardware prototype.

The major contributions of this work include:

- A framework that co-optimizes sensors and classifiers to enhance the energy consumption of wearable devices,
- Extensive evaluations using data from user-subject studies and a custom hardware prototype.

II. Related Work

Human activity recognition has received significant research attention due to its applications in patient monitoring, fitness tracking, and patient rehabilitation [4, 10]. Recent approaches use wearable devices with integrated motion sensors to infer activities, such as walk and exercise [3, 6]. Wearable devices continuously collect and process the sensor data to identify user activity. Due to their small form-factor and limited battery capacity, they must minimize energy consumption to extend the life-time [5, 14]. At the same time, user satisfaction requires high accuracy [11]. Therefore, wearable devices need to co-optimize the recognition accuracy and energy consumption.

Optimization of wearable devices has been an active research area. For example, Williamson et al. characterize energy consumption of sensing, processing and wireless data communication in wearable devices [15]. Then, they present potential solutions such as compressive sensing, FIFO and power gating to reduce energy consumption in wearable devices. Liu et al. selectively identify the best sampling points to maintain high accuracy while reducing sensing and analysis energy overheads [9]. Similarly, Krause et al. propose variable sampling rates and selective sampling to utilize the accuracy-power consumption tradeoff in HAR [8]. However, they do not put the sensors into low-power states during low-intensity activities to minimize energy consumption. In contrast to these approaches, we propose a framework to co-optimize the sensors and classifiers to reduce energy consumption in HAR applications. We evaluate the proposed approach on a custom prototype using data from 22 user studies.

III. DYNAMIC SENSOR-CLASSIFIER CO-OPTIMIZATION

Recognition accuracy is one of the most important design considerations in HAR applications. Higher sampling frequencies for sensors and classifiers leads to higher recognition accuracy. However, this comes at the cost of higher energy consumption that can reduce the operating lifetime of the device. As a result, wearable devices have to co-optimize accuracy and energy consumption per activity.

Energy consumption in HAR applications consists of two major components: the sensor and processor energy consumption. The sensor energy is typically a significant part of the total energy consumption as the sensors stay powered on continuously. However, this may not be necessary when the user is not active (e.g., during sitting or standing). Operating the sensors in low-power mode during such activities can lead to significant energy savings. Similarly, the sampling frequency of the sensors can be reduced as a function of the activity to enable further energy savings. Furthermore, significant processor energy savings can be achieved by tailoring the classifiers to the behavior of the sensor. However, the HAR accuracy can reduce significantly unless these techniques are co-optimized with the activity classifier. In what follows, we utilize these insights to enable two sensor-classifier co-optimization techniques that minimize the energy consumption of HAR while maintaining high accuracy.

A. Dynamic Sensor Power Gating

We can achieve significant energy savings by operating the sensors in low-power mode during low-intensity activities. For instance, sensor outputs do not change significantly when the user is not active, e.g., during sitting or standing. One can detect significant changes in user activity by monitoring the sensor output before transmitting it to the processor. We use this behavior and a threshold for changes in sensor data to reduce the energy consumption of the device.

Figure 1 illustrates the proposed dynamic power gating approach. The HAR engine uses the accelerometer data to determine the current activity. Our specific implementation uses a 3-axis accelerometer and passive stretch sensor [6]. The inputs to the proposed dynamic sensor control technique are the current activity and accelerometer readings. The dynamic sensor control block uses these inputs to set the power state of the accelerometer sensor [7]. If the change in the sensor values is small, it puts the accelerometer into the ultra-low-power mode. This power state measures the 3-axis acceleration at a low frequency. The accelerometer continues to operate in this mode until the dynamic sensor control block detects that the change in the sensor value is greater than a threshold δ_{th}. Specifically, for each new sample we calculate the absolute changes in the x-, y-, and z-axis of the accelerometer δ_x, δ_y, and δ_z as:

$$\delta_x = |a_x[n] - a_x[n-1]|, \qquad \delta_y = |a_y[n] - a_y[n-1]|$$
$$\delta_z = |a_z[n] - a_z[n-1]|, \qquad \delta = \max(\delta_x, \delta_y, \delta_z) \qquad (1)$$

where $a_x[n], a_y[n]$, and $a_z[n]$ denote the accelerometer readings normalized to the unit of gravity at time instance n for the x-,y- and z-axis, respectively. Similarly, $a_x[n-1]$, $a_y[n-1]$, and $a_z[n-1]$ denote the previous values of the accelerometer. The accelerometer is turned on only if the maximum of δ_x, δ_y, and δ_z exceeds the set threshold δ_{th}, i.e. $\delta > \delta_{th}$. This optimization reduces the sensor energy consumption significantly, as we demonstrate in Section IV.

Altering the power-states of the accelerometer affects the sensor values input to the HAR engine since there are fewer samples from the accelerometer. This results in a change in the inputs features fed to the classifier in the HAR engine. The change can cause significant degradation in accuracy of activity classification, which is an undesirable outcome. Hence, it is necessary to co-optimize sensor power management techniques and classifiers. To this end, we propose to retrain classifiers with the sensor data for each threshold. At run-time, we can choose the appropriate classifier as a function of the threshold. With this approach, we can achieve significant reduction in energy consumption with little to no effect on classification accuracy.

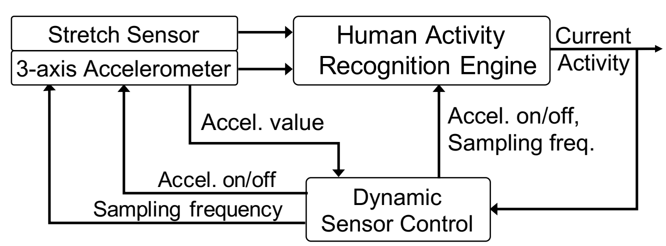

Fig. 1: Overview of dynamic sensor-classifier co-optimization

B. Dynamic Sampling Rate Control

Dynamic sampling rate control of sensors is an efficient technique to minimize energy consumption. Let P_{LP} and P_{Max} be the power consumption of the accelerometer in the low-power state and while is it operating at the highest sampling frequency f_{Max}, respectively. We can model the power consumption of the accelerometer as a function of sampling frequency f_s as:

$$P(f_s) = P_{LP} + P_{Max}(f_s/f_{Max})^{\alpha} \quad (2)$$

The parameter α models the non-linearity of power consumption as a function of the frequency. In our experiments, we use $\alpha = 1$ and employ the analytical sensor power models [7, 12].

In low-intensity activities, such as sitting or lying down, the outputs of the sensor do not change significantly. Since the power consumption decreases with frequency as shown in Equation 2, we employ lower sampling frequencies can be used to reduce energy consumption. However, using a smaller sampling frequency also changes the input features of the classifier and reduces the classification accuracy. Therefore, we retrain the classifier for each possible sampling frequency to retain the accuracy. The baseline classifier used for activity classification is trained with datasets obtained by running the accelerometer at 250 Hz. Then, we analyze the energy-accuracy trade-off to choose the right sampling frequency during low-intensity activities.

The results of energy-accuracy analysis are utilized by the dynamic sensor control block shown in Figure 1. The HAR engine classifies activity based on inputs from the accelerometer and feeds the classification output to the dynamic sensor control block. The sensor control block scales down the frequencies during low-intensity activities supervised by the energy-accuracy analysis with varying sampling frequencies.

IV. Experimental Results

Experimental setup: We evaluate the proposed dynamic sensor power management techniques using our custom prototype shown in Figure 2. The prototype integrates a Texas Instruments (TI) CC2650 MCU [1] for computation and Invensense MPU-9250 accelerometer (motion sensor). In addition, the prototype includes power measurement points. We use the wake-on-motion feature in the accelerometer to set a threshold to wake up the accelerometer, as described in Section III-A. Specifically, the accelerometer operates in low-power mode until it detects motion that exceeds the set threshold. We also use a passive stretch sensor (Figure 2) attached to a knee sleeve. The data from this sensor is used to divide the activities into distinct windows (1–3 s duration) for feature extraction and classification using deep neural network (DNN) with two hidden layers [6]. Using this setup, we validate the proposed techniques on an extensive dataset comprising of activities from 22 users, 7 activities and the transitions between them.

A. Evaluation and Analysis of Sensor Power Gating Technique

We first analyze the effect of the accelerometer power gating on recognition accuracy and power consumption. Figure 3 shows the accuracy of classification as we increase the threshold for waking up the accelerometer. The x-axis of the figure is normalized to the unit of gravity. Therefore, a value of 1 indicates that the change in accelerometer value is at least 9.8 m/s^2. We compare the accuracy under two scenarios. In the first scenario, there is a single classifier trained with a

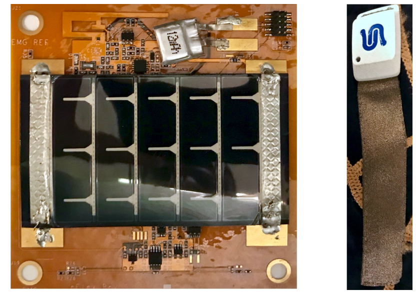

Fig. 2: Custom hardware prototype and the stretch sensor

threshold of zero, i.e. the accelerometer is assumed to be always on. The single classifier stores only a single neural network and does not have to switch between classifiers as the value of the threshold changes. In the second scenario, the DNN classifier is trained separately for each threshold. Then, these DNNs are stored on the device such that the appropriate classifier can be chosen at runtime.

We observe that the accuracy of activity recognition reduces as the value of the threshold is increased when we use a single classifier. Specifically, as the threshold increases from zero to about 0.2, the accuracy decreases rapidly to about 40%. Then, the decrease in accuracy is slower until it reaches an accuracy of 10%. The reduction in accuracy is expected since the feature vectors seen at the time of training may differ from the features observed after using the threshold value. In contrast, the drop in accuracy is significantly lower when we train a classifier for each threshold. The accuracy drops gradually to 81% until the threshold increases to 0.2. Beyond this point, the loss in accuracy is negligible. Furthermore, when the threshold is 1 (i.e., the accelerometer is always off), the classifier uses data from only the stretch sensor. This leads to an accuracy of about 76%. The accelerometer data exhibits similar patterns for multiple activities. The accuracy decreases with increasing threshold because the features are unable to differentiate between activities easily.

Next, we compare the energy consumption per activity window as a function of the threshold. The left y-axis of Figure 4 shows the energy per activity when we change the threshold for waking up the accelerometer. We observe that the energy per activity rapidly reduces by 46% for a threshold of 0.025 and very gradually reduces thereafter. Furthermore, we compare a combined metric of energy per accuracy to analyze energy consumption and accuracy together. The right y-axis in Figure 4 shows the energy per accuracy metric as a function of the threshold. We observe that the energy per accuracy decreases up to a threshold of 0.025 and increases

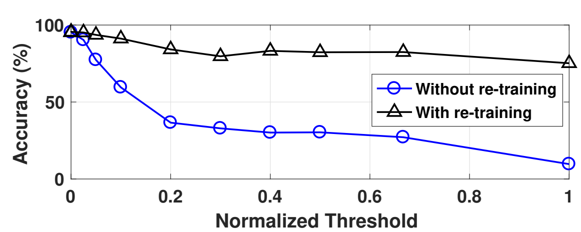

Fig. 3: Classification accuracies with sensor power gating

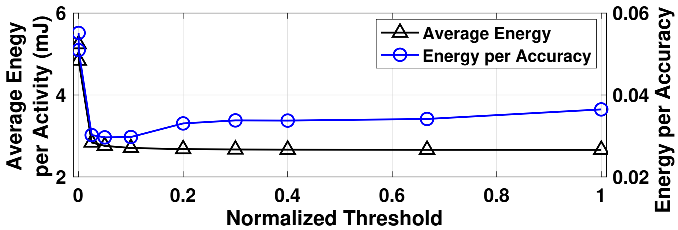

Fig. 4: Comparison of energy per activity and energy per accuracy with sensor power gating

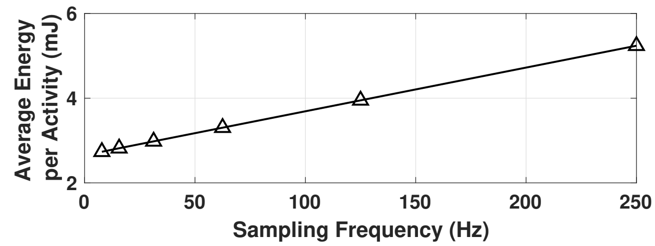

Fig. 6: Energy per activity with sampling rate control

thereafter. This shows that a threshold of 0.025 achieves the best trade-off between accuracy and energy consumption. In summary, sensor-classifier co-optimization enables 46% lower energy per activity with less than 1% reduction in recognition accuracy at a power gating threshold of 0.025.

B. Evaluation and Analysis of Sampling Rate Control

This section analyzes the accuracy and energy consumption of the HAR classifier as a function of the sampling frequency of the accelerometer. We focus on controlling the sampling frequency of the accelerometer since it accounts for a large portion of the energy consumption in the HAR application. We first sub-sample our HAR data set to obtain data sets with lower sample frequencies. Then, we analyze the effect of lower sampling frequencies on accuracy and energy consumption. Similar to the case of power gating, the classification accuracy is analyzed with and without retraining the classifier. There is a significant drop in accuracy when we use a single classifier for the full range of sampling frequencies, as shown in Figure 5. In contrast, retraining the classifier for each sampling frequency leads to lower loss in accuracy even when the sampling frequency is lowered. In particular, we observe that the classifier is able to achieve an accuracy greater than 90% for sampling frequencies as low as 8 Hz. Unlike the case of dynamic sensor power gating, the energy consumption is linearly proportional to the sampling frequency and the energy consumption reduces by 48% when the sampling frequency is 8 Hz, as we see in Figure 6. Therefore, the optimal decision is to choose a sampling frequency that meets the accuracy requirements.

Finally, we combine the two techniques together using the optimal threshold value of 0.025 from energy per accuracy curve in Figure 4 with varying sampling frequencies. The combination of both techniques enables 49% reduction in energy per activity with less than 1% loss in classification accuracy. In summary, the proposed sensor-classifier optimization techniques can achieve significant energy savings with minimal loss in accuracy.

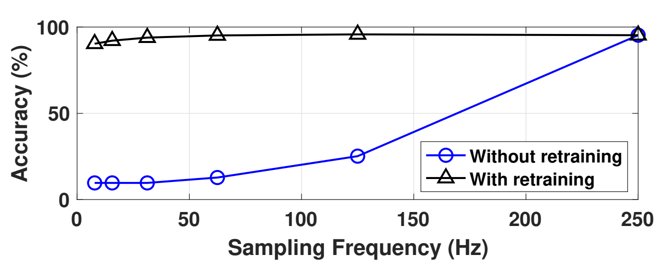

Fig. 5: Classification accuracies with sampling rate control

V. Conclusion

Advances in low-power sensors and processors have fueled an increase in the use of wearable devices for health and activity monitoring. With tight form-factor and weight constraints, these devices must operate with limited energy budgets and small batteries. To this end, this paper presented a sensor-classifier co-optimization technique for wearable devices using human activity recognition as a driver application. We dynamically power down the accelerometer and lower the sampling frequency when the user is performing low-intensity activities. Using these optimizations, the proposed approach achieves up to 49% reduction in total platform energy consumption with less than 1% decrease in the accuracy.

References

[1] TI SensorTag. https://store.ti.com/cc2650stk.aspx accessed 05/10/2019.
[2] G. Appelboom *et al.*, "Smart Wearable Body Sensors for Patient Self-Assessment and Monitoring," *Archives of Public Health*, vol. 72, no. 1, p. 28, 2014.
[3] F. Attal *et al.*, "Physical Human Activity Recognition using Wearable Sensors," *Sensors*, vol. 15, no. 12, pp. 31 314–31 338, 2015.
[4] A. Avci *et al.*, "Activity Recognition using Inertial Sensing for Healthcare, Wellbeing and Sports Applications: A Survey," in *Int. Conf. on Arch. of Comput. Syst.*, 2010, pp. 1–10.
[5] G. Bhat *et al.*, "Near-Optimal Energy Allocation for Self-Powered Wearable Systems," in *Proc. of ICCAD*, 2017, pp. 368–375.
[6] G. Bhat *et al.*, "Online Human Activity Recognition using Low-Power Wearable Devices," in *Proc. of ICCAD*, 2018, pp. 72:1–72:8.
[7] InvenSense, "Motion Processing Unit," 2016, https://www.invensense.com/products/motion-tracking/9-axis/mpu-9250, accessed 05/10/2019.
[8] A. Krause *et al.*, "Trading Off Prediction Accuracy and Power Consumption for Context-Aware Wearable Computing," in *Proc. Int. Symp. on Wearable Comput.*, 2005, pp. 20–26.
[9] Q. Liu *et al.*, "Gazelle: Energy-Efficient Wearable Analysis for Running," *IEEE Trans. Mobile Comput.*, no. 9, pp. 2531–2544, 2017.
[10] A. Mosenia *et al.*, "Wearable Medical Sensor-Based System Design: A Survey," *IEEE Trans. Multi-Scale Comput. Syst.*, vol. 3, no. 2, pp. 124–138, 2017.
[11] A. Ozanne, D. Johansson, U. Hällgren Graneheim, K. Malmgren, F. Bergquist, and M. Alt Murphy, "Wearables in Epilepsy and Parkinson's diseaseA Focus Group Study," *Acta Neurologica Scandinavica*, vol. 137, no. 2, pp. 188–194, 2018.
[12] J. Park *et al.*, "Optimizing Operations per Joule for Energy Harvesting IoT Devices," Technical Report, Arizona State University, [Online], https://elab.engineering.asu.edu/wp-content/uploads/2018/06/energy_gesture.pdf.
[13] M. Shoaib *et al.*, "A Survey of Online Activity Recognition Using Mobile Phones," *Sensors*, vol. 15, no. 1, pp. 2059–2085, 2015.
[14] C. Wang *et al.*, "Low-Power Technologies for Wearable Telecare and Telehealth Systems: A Review," *Biomed. Eng. Let.*, vol. 5, no. 1, pp. 1–9, 2015.
[15] J. Williamson *et al.*, "Data Sensing and Analysis: Challenges for Wearables," in *Proc. of ASPDAC*, 2015, pp. 136–141.