EA^2: Energy Efficient Adaptive Active Learning for Smart Wearables

Hamidreza Alikhani
University of California Irvine
Irvine, USA

Ziyu Wang
University of California Irvine
Irvine, USA

Anil Kanduri
University of Turku
Turku, Finland

Pasi Liljeberg
University of Turku
Turku, Finland

Amir.M. Rahmani
University of California Irvine
Irvine, USA

Nikil Dutt
University of California Irvine
Irvine, USA

ABSTRACT

Mobile Health (mHealth) applications rely on supervised Machine Learning (ML) algorithms, requiring end-user-labeled data for the training phase. The gold standard for obtaining such labeled data is by sending queries to users and gathering responses for the corresponding label, which was conventionally done through triggering questions sent at random. Active Learning (AL) methods use intelligent query-sending policies by incorporating users' contextual information to maximize the response rate and informativeness of the collected labeled data. However, wearable devices' substantial battery drainage associated with the sensing of physiological signals underscores the need for developing an efficient sensing policy in addition to a query-sending policy. In this work, we present a co-optimization framework for both sensing and querying strategies within wearable devices, leveraging contextual information and ML model's prediction confidence. We designed a Reinforcement Learning (RL) agent to quantify different contextual parameters combined with model confidence to determine sensing and querying decisions. Our evaluation of an exemplar stress monitoring application showed a 76% reduction in sensing and data transmission energy consumption, with only a 6% drop in user-labeled data.

KEYWORDS

Efficient Machine Learning, Data Efficiency, Active Learning, Context-Aware Sensing, Reinforcement Learning.

ACM Reference Format:
Hamidreza Alikhani, Ziyu Wang, Anil Kanduri, Pasi Liljeberg, Amir.M. Rahmani, and Nikil Dutt. 2024. EA^2: Energy Efficient Adaptive Active Learning for Smart Wearables. In *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED '24), August 5–7, 2024, Newport Beach, CA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3665314.3670840

1 INTRODUCTION

Affective computing applications such as stress, emotion, and pain monitoring use supervised Machine Learning (ML) algorithms to provide smart mHealth services [4]. Supervised models for mHealth

ISLPED '24, August 5–7, 2024, Newport Beach, CA, USA

ACM ISBN 979-8-4007-0688-2/24/08.
https://doi.org/10.1145/3665314.3670840

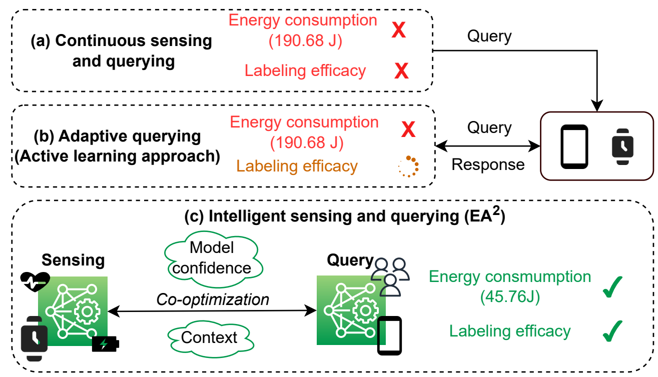

Figure 1: Energy consumption and labeling efficacy with different sensing and querying strategies.

require labeled physiological data (e.g., Photoplethysmography (PPG), Electrodermal activity (EDA), etc.) for training feasibility. Typically, raw data is labeled by sending an ecological momentary assessment (EMA) notification to the users, who in turn provide self-reports (e.g., current stress level) that can be used as ground-truth labels [11]. Deploying data-intensive ML models on resource-constrained wearable devices primarily faces the challenge of stringent energy budgets [4]. Conventional supervised learning methods spend a significant portion of wearable devices' energy budget in continuously sensing physiological signals to collect input data, leading to early battery depletion [1, 15]. However, data samples labeled by the users in response to random EMAs are often as low as 5% of total sensed samples [14], owing to diverse factors such as users' burden, activity, and proximity with a mobile device, time of the day, and frequency of EMA querying, etc. Thus, 95% of the samples that are sensed remain unlabeled, exposing the fundamental bottleneck of inefficient sensing, leading to energy drain.

Recently, mHealth applications have adopted active learning methods, where EMAs are sent to the users selectively – to collect labels for data with higher utility [14]. Selective user querying through active learning enables training on a curated subset of the most insightful labeled data to improve prediction accuracy with minimal training overhead. This approach targets improving the quality of labeling and users' experience with a minimal number of queries. However, both conventional random EMA methods and active learning methods still rely on continuous sensing, largely being agnostic to (i) the disparity between the number of sensed and/or queried samples to the number of labeled samples by the user and (ii) the efficacy of labeled samples in contributing towards

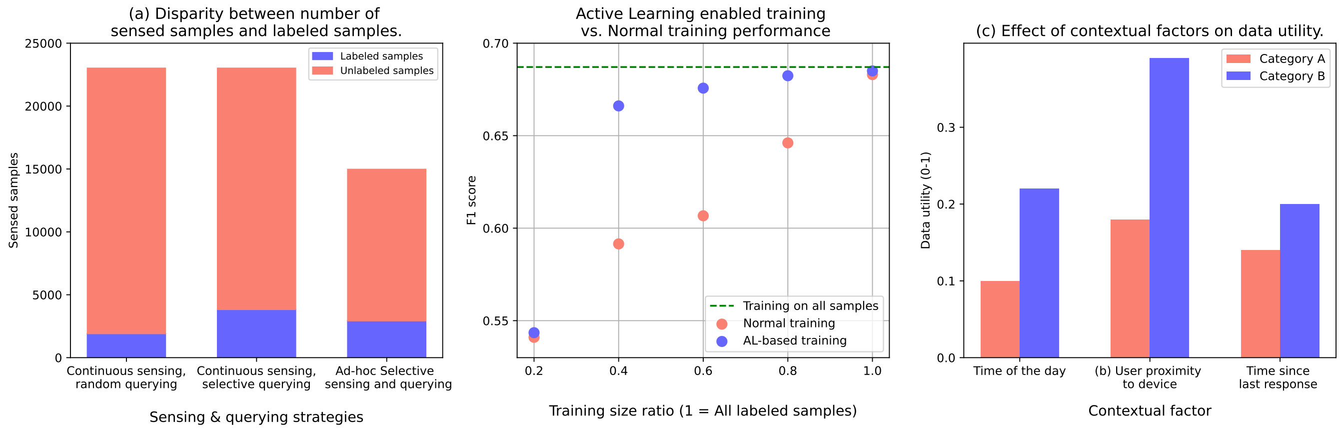

Figure 2: (a) Disparity between sensed samples and labeled data amount with different sensing and querying strategies. (b) Active learning enabled training vs normal training and corresponding F1 score gain with different training sizes. (c) Effect of different contextual factor on data utility.

prediction accuracy. Existing active learning approaches thus result in unnecessary sensing of data which is potentially un-insightful and/or unlabeled, also draining energy resources.

In this work, we address the inefficient sensing and energy drain of active learning methods on wearable devices through holistic context-aware selective sensing and querying. We design a comprehensive context-awareness model that combines user activity (e.g., call and screen touch status, notifications, etc.), interactive user labeling feedback (query response properties), and ML model confidence on the sensed data. Our context-awareness model prunes the design space to select rewarding instances where sensing has higher data utility and subsequent querying has a likelihood of user labeling response. We design a Reinforcement Learning (RL) [9, 12] agent that models context awareness to guide selective sensing and querying with the objectives of improving sensing efficiency (implicitly minimizing energy consumption) and maximizing labeling efficacy. Our sense-query co-optimization approach provides (i) significant reduction in sensing unnecessary samples, (ii) higher labeling efficacy, and (iii) retention of prediction accuracy and training feasibility even with lower number of sensed samples. Figure 1 summarizes different querying strategies, highlighting our contributions and limitations of existing methods. Conventional (Figure 1 (a)) methods continuously sense physiological data and query randomly, resulting in sensing higher volume of unnecessary data (leading to higher energy consumption) and poor labeling efficacy with fewer responses from users. Active learning methods (Figure 1 (b)) improve the relative labeling efficacy with selective querying; although fail to preserve sensing efficiency due to continuous sensing. Our framework combines contextual parameters from different entities across sensing and querying phases viz., user interaction with the device, users' response on labeling, and ML model's confidence in predicting the outcomes of a specific sample. This enables effective selection of instances to sense – improving sensing efficiency and limiting energy drain, followed by querying – retaining labeling efficacy and training feasibility. **Our contributions:**

- Design of comprehensive context-awareness model to guide selective sensing and querying, by combining user behavior (through interactions with the mobile device), users' labeling response feedback, and ML model confidence.
- Design and implementation of an RL agent to guide co-optimization of selective sensing and querying for sensing efficient active learning on wearables.
- Evaluation of our selective sensing framework's efficacy over real-world application of stress monitoring using physiological data from 35 individuals, in comparison with other relevant methods.

2 BACKGROUND AND MOTIVATION

Energy efficiency. We demonstrate the potential energy gain with selective sensing for active learning through a real-world exemplar stress monitoring application. This application uses Photoplethysmography (PPG) signals collected from 35 participants using their smart watches [14] (application detailed in Section 4). Figure 2 (a) shows the number of samples sensed by the PPG sensor, with each sensing window set to 15 minutes. We present the data sensed using three different strategies viz., (a) conventional approach with continuous sensing and random querying, (b) active learning approach with continuous sensing and selective querying, and (c) adaptive active learning approach with selective sensing and selective querying. Figure 2 (a) demonstrates the disparity between the number of sensed samples and the subsequent number of answered queries (labeled samples). Both conventional and active learning approaches sense continuously, collecting an equal amount of PPG data and consuming the same amount of energy. However, the advantage of active learning is in improving the quality of labels collected for raw PPG data. In this example, we use an ad-hoc selective sensing policy that skips sensing every third sample. This selective sensing approach senses 35% fewer PPG samples than the continuous sensing approaches, which reflects proportionally on energy savings. The key insight here is that a 35% reduction in the total number of samples eventually resulted in a reduction of only 24% of answered queries. This can be attributed to joint decision-making on (i.e., co-optimization of) sensing and querying, exploiting the opportunities presented by contextual parameter variations. Thus, a context-aware selective sensing and querying policy can provide significant energy gains with minimal/no loss in labeling efficacy.

Labeling efficacy. The primary challenge with selective sensing is potential degradation of prediction accuracy due to fewer input data samples. We present the dynamics of prediction accuracy with the increasing number of selected labeled samples. Figure 2 (b) compares the F1 score of a Random Forest (RF) model for stress level prediction on different volumes of labeled data using selective subsets (active learning) versus random subsets (normal training), based on PPG data from the study in [14]. Trivially, active learning method achieves a higher F1 score by prioritizing uncertain instances. Moreover, the accuracy difference between successive subsets in this method with increasing labeled samples is minimal, which is attributed to the fact that active learning prioritizes the subset of data samples over which the model's confidence is uncertain. For instance, reducing labeled samples from 80% to 60% results in only a 1.4% F1 score drop in the AL-based method. This presents opportunities for exploiting selective sensing to reduce the total number of sensed and labeled samples without enduring significant prediction accuracy loss.

Leveraging contextual parameters for selective sensing. Contextual parameters including users' proximity to the mobile device, activity, response to EMAs for labeling data, and model's confidence expose diverse opportunities for appropriate selective sensing decisions. Identifying user's proximity and willingness to interact with their wearables can simultaneously improve querying response rates and energy efficiency through selective sensing. We demonstrate the influence of different contextual factors on *sensing utility* i.e., the ratio of sensed data that has a corresponding label from users in a defined time window. Figure 2 (c) shows the sensing utility ratio for different contextual parameters.

Time of the day: We split the time of the day into two categories viz., category A– 8pm-8am and category B – 8am-8pm. It can be observed that users' response rate is higher for category B (i.e., during daytime), since users are more likely to interact with their phones during the day as opposed to night/sleep time.

User proximity to device: We define *proximity to device* as combination of factors including screen touch status, number of notifications popped up on the phone, number of received/sent messages and calls in the past T minutes etc (T is a configurable time window parameter). These metrics are defined by the AWARE features [2] in the context of estimating quality of experience. In this case, Category A again shows less interaction of users' with the mobile device (e.g., having 5 or less than 5 calls/messages/notifications on the phone in the past X minutes), while Category B represents a higher user interaction with the device.

Time since last response: Finally, we capture the users' likelihood of labeling a specific sample based on the time elapsed since the last response from the user. Typically, as the time since the last response increases, the probability of observing another response decreases. Here, Category A represents the instances where the time elapsed since the last query is more than 2 hours while Category B shows the instances where the time elapsed from previous response is less than 2 hours. In summary, across these different contextual factors, the sensing efficiency rate increases from 0.10 to 0.22 (based on time of the day), 0.18 to 0.39 (based on user proximity to device), and 0.14 to 0.20 (based on time since last response). We combine the aforementioned factors to develop context awareness that guides co-optimization of sensing and querying. We design an RL agent to quantify the impact of different contextual parameters for sensing efficient active learning on wearables in everyday settings. Details of our framework architecture and RL agent are described in the next Section.

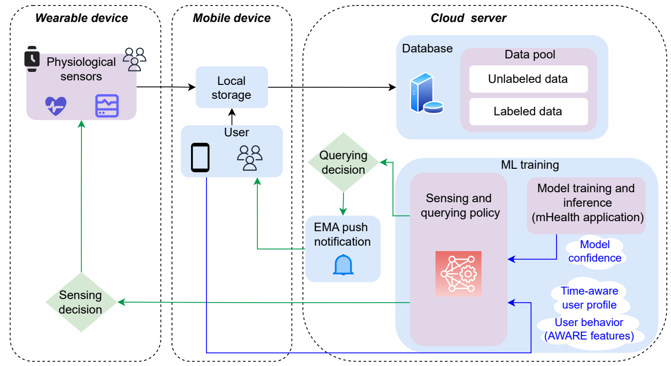

Figure 3: EA^2 framework architecture and interaction of different component (wearable, phone, and cloud server).

3 SENSING EFFICIENT ACTIVE LEARNING FRAMEWORK

3.1 System architecture

Figure 3 shows the proposed EA^2 framework, providing infrastructure for collecting and labeling physiological signals for mHealth applications in everyday settings through wearable devices, mobile phones, and cloud server computing and database storage resources. The wearable device is equipped with various physiological sensors capable of continuous or selective sensing. Based on the sensing decision, the physiological signal is collected by sensors, and if needed, a personalized query (EMA push notification) will be sent to the user's phone device from the cloud server. Upon receiving a response from the user for labeling the sensed signal, the labeled and unlabeled data are gathered and stored in local storage on the phone device and will be periodically dispatched to the cloud server database. Based on the application at hand (e.g., stress monitoring), the desired machine learning model will be trained on the received data through supervised learning methods. The trained ML model will be utilized for performing inference on the sensed data, while considering model confidence for active learning enabled data labeling as one of the inputs influencing the querying decisions. Moreover, time-aware user profile properties (based on query response features as shown in Table 1) and user behavior captured through AWARE [2] features on the mobile device are other additional components that form the overall contextual information. This is fed as input to the sensing and querying module for taking intelligent actions with the objective of maximizing sensing and labeling efficiency and efficacy. Details of the intelligent sensing and querying module implementation are explained as follows.

3.2 RL-based context-awareness model for selective sensing and querying

The choice on sense-query configuration depends on conflicting objectives i.e., selective sensing for lower energy consumption versus continuous sensing for higher labeling efficacy. Understanding contextual parameters such as user activity and likelihood of a user response can guide the sense-query configuration settings. To this end, we formulate the sense-query co-optimization as a reinforcement learning problem to balance the conflicting objectives of reducing energy consumption while improving labeling efficacy. We design an RL agent to determine sense-query co-optimization

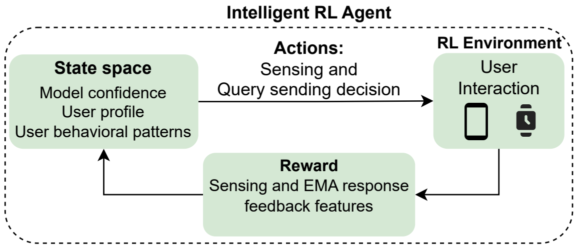

Figure 4: Overview of the RL agent for intelligent sensing and querying.

decisions, by understanding diverse contextual parameters. Figure 4 shows an overview of the RL agent, with environment, state, and action spaces. The RL agent integrates real-time user behavior patterns from the AWARE framework [2], and time-aware user profiles to form a comprehensive state space for decision-making. Using the combined state representation, the RL agent makes informed decisions that optimize both energy consumption and labeling efficacy simultaneously. The following sections provide a detailed breakdown of the RL agent's architecture and algorithms.

State Space. The state vector s represents the current status of essential parameters that enable the RL agent's sense-query optimization decisions. The state vector has three principle components viz., model confidence, time-aware user profile, and real-time user behavior patterns, as shown in Table 1.

Action Space. Upon receiving the time-aware user profiles and real-time user behavior patterns at the cloud layer, the RL agent periodically (every 15 minutes) evaluates actions across two primary dimensions: sensing and querying. These actions are (i) not sensing, (ii) sensing but not querying the user for the label, and (iii) sensing and querying for collecting the label. If the decision is to sense but not send the query, the sample will be stored as unlabeled data for future use for semi-supervised learning . If the decision is sensing and also querying for the current instance, based on the presence of user response, it will be stored in an unlabeled or labeled data pool. It is important to highlight the fact that state space components related to model confidence are available after sensing the signal, hence it is reflected on the action of either querying or not querying if the decision in the previous step is to sense the signal.

Reward Function. The reward function is precisely designed to prioritize instances best suited for both querying and sensing. Actions taken at opportune moments, such as timely sensing or querying, receive a higher reward. On the other hand, the RL agent faces penalties for actions during user inactivity or excessive sensing and querying. The reward function's criteria are fine-tuned to encapsulate the desired characteristics for both actions. Factors such as the classifier's model confidence, the time elapsed since the last query, the interval since the previous sensing, and user response rate are pivotal metrics that shape our reward function. Given a state space matrix S_t at time t and a weight vector W, the immediate reward R_t can be computed as:

$$R_t(a) = S_t(a) \cdot W(a) \tag{1}$$

• $S_t(a)$ is the state space matrix at time t for a specific action a. The elements of S_t represent metrics such as model confidence, and context information, as shown in Table 1. The values in this matrix are updated based on the actions and the environment's response.

Table 1: State Space Components

State	Description
Model Confidence	
Class Probability	Prediction label probability
Uncertainty Window	Distance from the decision boundary
Time-Aware User Profile	
Response Rate	Response rate during different times.
Time since Last Query	Query frequency.
Time since Last Sensing	Sensing frequency.
Time of Day	Current hour.
Real-time User Behavior (AWARE features)	
Communication	Calls and messages.
Screen & Touch	User screen interactions.
Application	Application usage and notifications.

• $W(a)$ is the weight vector for action a, which assigns importance to each metric in S_t based on action's significance.

• The reward R_t varies depending on the action taken and the updated states in S_t. The weights in $W(a)$ are determined during training to prioritize between sensing and querying, while minimizing unnecessary actions during user inactivity.

Deep Q-Learning Network and Policy Function. Deep Q Learning is a model-free, online, off-policy reinforcement learning method. In the context of our problem, it is used to determine the optimal action for the agent in a given state to maximize the expected cumulative reward. The agent's decision-making process is modeled using a neural network, which approximates the Q-value function.

Q-value Function: The Q-value function, denoted as $Q(s, a)$, represents the expected cumulative reward of taking action a in state s and following the optimal policy thereafter. The Q-value function can be recursively defined by the Bellman equation:

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a') \tag{2}$$

where r is the immediate reward after taking action a in state s, γ is the discount factor, modeling the agent's consideration for future rewards in the range $[0, 1]$, and s', a' are the next state and action.

Neural Network Approximation: The Q-value function is approximated using a neural network, referred to as the Deep Q-Network (DQN). Given a state s as input, the DQN outputs the Q-values for all possible actions.

Policy Function: The policy function, denoted as $\pi(s)$, determines the action the agent should take in a given state. In our implementation, an epsilon-greedy strategy is employed:

$$\pi(s) = \begin{cases} \text{random action} & \text{with probability } \epsilon \\ \arg\max_a Q(s, a) & \text{with probability } 1 - \epsilon \end{cases} \tag{3}$$

This strategy ensures a balance between exploration (random actions) and exploitation (taking the action with the highest predicted Q-value). Over time, ϵ is decayed to reduce the likelihood of taking random actions and increase the reliance on the learned Q-values.

Training: The DQN is trained using experiences collected during the agent's interaction with the environment. We train a personalized RL agent for each experience as a tuple (s, a, r, s'). The loss function used for training is the Mean Squared Error (MSE) between the predicted Q-values and the target Q-values as follows:

$$\text{Loss} = \left(Q(s, a) - (r + \gamma \max_{a'} Q(s', a'))\right)^2 \tag{4}$$

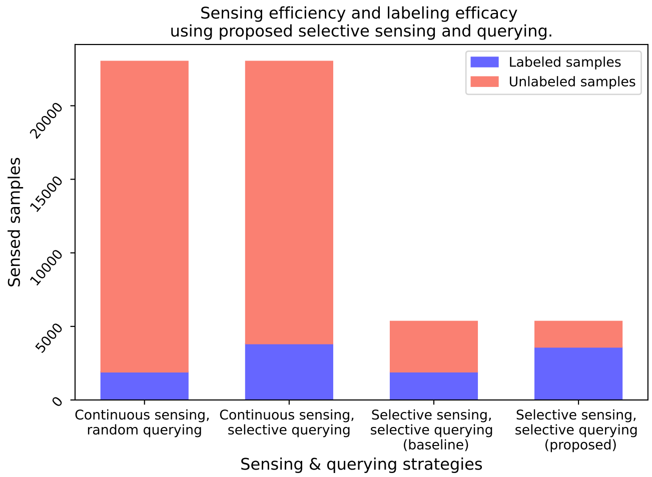

Figure 5: Comparison of sensing and labeling efficacy for different policies.

4 EVALUATION

Workload: We evaluate the efficacy of our proposed framework in terms of sensing efficiency and labeling efficacy for an exemplary application of human stress monitoring [15] using PPG signal collected from a Samsung Galaxy Gear Sport smartwatch [10]. Labels were collected from users via sending an EMA notification to their smart device asking them to rate their current stress level from 0 to 5. The user behavior patterns are collected from a mobile application installed on the user's phone named AWARE [2]. The framework is built upon ZotCare [5], an efficient mHealth service platform. The dataset used in this study is collected sporadically from March 2022 to May 2023, provided by 35 participants ranging in age from 19 to 29 years. We aggregated 23,012 samples [13]. The specific PPG features we extracted by using HeartPy [3] are: BPM, IBI, SDNN, SDSD, RMSSD, pNN20, pNN50, MAD, SD1, SD2, S, SD1/SD2, and BR. The respective Institutional Review Board (IRB) granted approval for all aspects of this investigation.

Sensing efficiency: Figure 5 shows a gain in sensing efficiency (reduction in volume of sensing potentially un-insightful data), and labeling efficacy (in terms of captured labeled samples compared to baseline with continuous sensing strategy). For a comprehensive evaluation of our framework, we compared our approach against (i) *continuous sensing and random querying*,(ii) *continuous sensing and selective querying* (based on [14] with capturing maximum available labeled data), (iii) *baseline selective sensing and querying* – which uses random skipping sensing policy (instead of continuous sensing) based on time of the day such that reduction in sensing samples would be nearly equal to that of our framework. In this context, the baseline selective sensing and querying reduces the sensing events to 1/10 during the time between 10 pm and 8 am, and reduces the sensing events frequency to 1/3 for the rest of the day). Our framework achieves a 76% reduction in sensing events in comparison with the aforementioned relevant active learning approaches. Further, our approach results in a higher number of labeled samples in comparison with our own baseline selective sensing and querying by 88%, demonstrating better labeling efficacy. In comparison with continuous sensing and selective querying, our approach has only 6% fewer labeled samples, despite sensing 76% less number of total samples. Both the sensing efficiency and labeling efficacy of our framework can be attributed to our framework's comprehensive context-awareness and sense-query co-optimization approach.

Table 2: Sensing and transmission Energy consumption with different strategies

Method	Sensing energy (PPG sensor)	Transmission energy (BLE)
Continuous sensing	190.26 J	415.38 mJ
Ad-hoc Selective sensing	125.57 J	274.16 mJ
Intelligent Sensing (proposed method)	45.66 J	99.69 mJ

Energy Efficiency: We evaluate the impact of our proposed method on reducing the total energy consumption. Since the power consumption of PPG module in Samsung smartwatch is not publicly available, we model the energy consumption using state-of-the-art LED-based PPG sensors [8]. In our framework, sensing driven energy savings are achieved through (i) reduction in total time period of actively collecting PPG signal, and (ii) reduction in amount of raw PPG data transmitted to smart phone for further computation and offloading to server database. Considering these aspects, we formulate energy consumption E_{PPG} of the PPG module as:

$$E_{PPG} = P_{PPG_{act}} * t_{act} + P_{PPG_{down}} * t_{down} \tag{5}$$

where $P_{\mathrm{PPG_{act}}}$ and $P_{\mathrm{PPG_{down}}}$ are power consumption of PPG module during active and shutdown modes, and t_{act} and t_{down} are the time periods of active data collection and sensor shutdown. We use MAX30101 SoC [7, 8] as the baseline for our energy model, with power consumption of 5.5 mW and 3.5 μW in active and shutdown modes. In our setup, a data collection event features 2 minutes of continuous sensing. The continuous sensing approach triggers data collection for every 5 minutes, while the ad-hoc strategy triggers data collection for two consecutive 5 minute windows followed by one 5 minute window of sensor shutdown. With the observation window set to 24 hours, the continuous and ad-hoc strategies consume 190.26J and 125.57J in sensing, as shown in Table 2. By reducing total time period of actively collecting PPG data, our proposed approach results in 45.66J of energy consumption, with a gain of 76% and 63% compared to continuous and ad-hoc sensing. The reduction in collected data size is calculated based on 20Hz sensor sampling rate in this study [14], and size of the data chunk in each reading. Our framework provides collected PPG data reduction of 2.10 MB and 1.16 MB compared to baseline and ad-hoc sensing, respectively. Considering the data is transmitted from watch to phone using bluetooth (BLE 5) with average transmission rate of 1Mbps and power consumption of 18.78 mW and 6.28 μW and during connection and idle time respectively [6, 8], our framework provides overall reduction in bluetooth transmission energy of 315.69 mJ and 174.46 mJ compared to baseline and ad-hoc methods.

Stress monitoring model performance: For evaluating our context aware method against the baseline approach while keeping the sensing efficiency the same for both methods, we trained a Random Forest (RF) model on the available labeled data (100 estimators, and maximum depth of 5). Due to the high reduction in the number of available labeled data in the baseline method, the F1 score of the RF model is 0.60 compared to the 0.67 for our proposed framework.

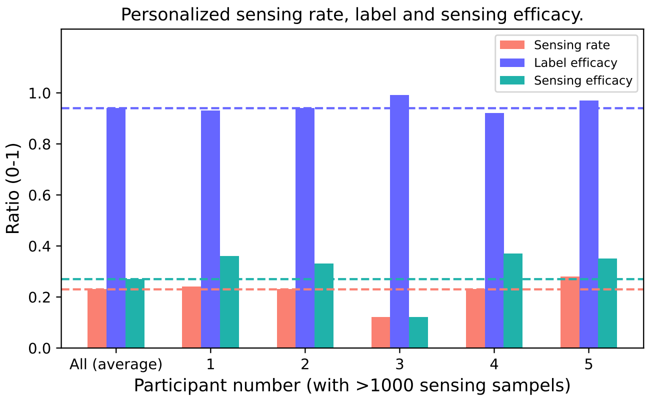

Figure 6: Personalized sensing rate and sensing and labeling efficacy for users with higher contributions in the study.

Personalization: Figure 6 shows the personalized results for sensing ratio, sensing, and labeling efficacy on participants with more than 1000 collected samples. The RL agent modifies its sense-query actions for different individuals based on their response behavior and interactions. For instance, user number 3 has the least participation in the study in terms of answered queries. Ultimately, the RL agent has a high performance in capturing almost all labels from this user with a lower sensing rate than average using context-awareness. A similar behavior can be observed for other users with variations in final sensing rate and sensing/labeling efficacy which is correlated with their engagement in the study in terms of answered EMAs and smart watch active time.

Ablation study: To demonstrate the impact of context-aware sense-query co-optimization, we conducted an ablation study by altering the RL agent's state space parameters and retraining it. Table 3 displays the scenarios with various state space components and their corresponding sensing ratios (ratio of collected instances by the RL agent over the number of sensed samples using continuous sensing), labeling efficacy, and sensing efficacy. Labeling efficacy measures the proportion of labeled samples by the user with at least one available sensing sample in the last 15 minutes, while sensing efficacy measures the proportion of sensed instances with at least one available label from the user in the next 15 minutes. These two metrics implicitly capture the temporal density of sensing and querying to mitigate overfitted decision-making on selective sensing and querying. As shown in Table 3, continuous sensing provides no sensing efficiency, which results in capturing all the labels, but the sensing efficacy is extremely low (0.07). Using only user time-aware profile reduces the sensing ratio to 0.33 but lowers labeling efficacy to 0.51 and sensing efficacy to 0.09. Training with user behavior patterns improves performance due to finer real-time characteristics, achieving better capture of labeled data with fewer sensing events. The proposed method, utilizing both parameter sets and model confidence, achieves the highest labeling and sensing efficacy (0.94 and 0.28) with an acceptable sensing ratio.

5 CONCLUSION

We presented a context-aware energy efficient framework for active learning on wearable devices. Our framework is based on an RL agent that co-optimizes sensing and querying policies by combining user-level contextual parameters and the ML model's prediction confidence. We demonstrated our framework's energy efficiency

Table 3: Ablation study for observing effect of using user time-aware profiles and behavior patterns

State-space components	Sensing ratio	Labeling efficacy	Sensing efficacy
None (Continuous sensing)	1.0	1.0	0.07
User time-aware profiles	0.33	0.51	0.09
Behavior patterns	0.18	0.72	0.24
All (proposed method)	0.23	0.94	0.28

and sensing/labeling efficacy on an exemplary application of stress monitoring using a PPG signal collected through a Samsung Galaxy Gear Sport watch. Our framework reduces sensing data volume by 76% with only missing 6% of the user-annotated data compared to the baseline. Our future work will extend and evaluate the EA^2 framework on more applications and studies.

6 ACKNOWLEDGEMENTS

This work was partially supported by NSF Smart and Connected Communities (S&CC) grant CNS-1831918, Nokia Foundation and Kaute Saatio, Finland.

REFERENCES

[1] S. Béres and L. Hejjel. 2021. The minimal sampling frequency of the photoplethysmogram for accurate pulse rate variability parameters in healthy volunteers. *Biomedical Signal Processing and Control* 68 (2021), 102589. https://doi.org/10.1016/j.bspc.2021.102589

[2] D. Ferreira et al. 2015. AWARE: mobile context instrumentation framework. *Frontiers in ICT* 2 (2015), 6.

[3] P. Van Gent et al. 2019. HeartPy: A novel heart rate algorithm for the analysis of noisy signals. *Transportation research part F: traffic psychology and behaviour* 66 (2019), 368–378.

[4] A. Kanduri et al. 2023. Edge-centric Optimization of Multi-modal ML-driven eHealth Applications. In *Embedded Machine Learning for Cyber-Physical, IoT, and Edge Computing: Use Cases and Emerging Challenges*. Springer, 95–125.

[5] S. Labbaf et al. 2023. ZotCare: A Flexible, Personalizable, and Affordable mHealth Service Provider. *arXiv preprint arXiv:2307.01905* (2023).

[6] Silicon Labs. 2024. Optimizing Current Consumption in Bluetooth Low Energy Devices. https://docs.silabs.com/bluetooth/5.0/general/system-and-performance/optimizing-current-consumption-in-bluetooth-low-energy-devices

[7] Jinseok Lee et al. 2018. A Low-Power Photoplethysmogram-Based Heart Rate Sensor Using Heartbeat Locked Loop. *IEEE Trans. on Biomedical Circuits and Sys.* 12, 6 (2018), 1220–1229. https://doi.org/10.1109/TBCAS.2018.2876671

[8] Tommaso Polonelli et al. 2021. H-Watch: An Open, Connected Platform for AI-Enhanced COVID19 Infection Symptoms Monitoring and Contact Tracing. In *IEEE Int. Symp. on Circuits and Systems (ISCAS)*. 1–5. https://doi.org/10.1109/ISCAS51556.2021.9401362

[9] Mohammadreza Rostami et al. 2023. Federated learning using variance reduced stochastic gradient for probabilistically activated agents. In *2023 American Control Conference (ACC)*. IEEE, 861–866.

[10] Samsung. 2023. Gear Sport Smartwatch. https://www.samsung.com/us/mobile/wearables/smartwatches/gear-sport-blue-sm-r600nzbaxar/ Available at: https://www.samsung.com/us/mobile/wearables/smartwatches/gear-sport-blue-sm-r600nzbaxar/.

[11] S. Shiffman, A. A. Stone, and M. R. Hufford. 2008. Ecological momentary assessment. *Annu. Rev. Clin. Psychol.* 4 (2008), 1–32.

[12] R. S. Sutton and A. G. Barto. 2018. *Reinforcement learning: An introduction*. MIT press.

[13] A. Tazarv et al. 2021. Data Collection and Labeling of Real-Time IoT-Enabled Bio-Signals in Everyday Settings for Mental Health Improvement *(GoodIT '21)*. Association for Computing Machinery, New York, NY, USA, 186–191. https://doi.org/10.1145/3462203.3475918

[14] A. Tazarv et al. 2023. Active Reinforcement Learning for Personalized Stress Monitoring in Everyday Settings. In *IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*. 44–55. https://doi.org/10.1145/3580252.3586979

[15] Z. Wang et al. 2024. Differential Private Federated Transfer Learning for Mental Health Monitoring in Everyday Settings: A Case Study on Stress Detection. *arXiv preprint arXiv:2402.10862* (2024).