

AMSER: Adaptive Multimodal Sensing for Energy Efficient and Resilient eHealth Systems

Emad Kasaeyan Naeini¹, Sina Shahhosseini¹, Anil Kanduri², Pasi Liljeberg², Amir M. Rahmani¹, Nikil Dutt¹

¹ Dept. of CS, University of California, Irvine, USA, ² Dept. of Computing, University of Turku, Finland

ekasaeya@uci.edu, sshahos@uci.edu, spakan@utu.fi, pakrili@utu.fi, a.rahmani@uci.edu, dutt@uci.edu

Abstract—eHealth systems deliver critical digital healthcare and wellness services for users by continuously monitoring physiological and contextual data. eHealth applications use multi-modal machine learning kernels to analyze data from different sensor modalities and automate decision-making. Noisy inputs and motion artifacts during sensory data acquisition affect the i) prediction accuracy and resilience of eHealth services and ii) energy efficiency in processing garbage data. Monitoring raw sensory inputs to identify and drop data and features from noisy modalities can improve prediction accuracy and energy efficiency. We propose a closed-loop monitoring and control framework for multi-modal eHealth applications, AMSER, that can mitigate garbage-in garbage-out by i) monitoring input modalities, ii) analyzing raw input to selectively drop noisy data and features, and iii) choosing appropriate machine learning models that fit the configured data and feature vector - to improve prediction accuracy and energy efficiency. We evaluate our AMSER approach using multi-modal eHealth applications of pain assessment and stress monitoring over different levels and types of noisy components incurred via different sensor modalities. Our approach achieves up to 22% improvement in prediction accuracy and 5.6× energy consumption reduction in the sensing phase against the state-of-the-art multi-modal monitoring application.

I. INTRODUCTION

eHealth applications deliver digital healthcare and wellness services for users by continuously monitoring and analyzing physiological and contextual sensory data [1]. eHealth applications feature: i) high volumes of input data from multi-modal sensory devices, ii) intense machine learning algorithmic computations for data analysis, and iii) limited energy and compute resources, particularly at the sensing layer (e.g., wearable sensors) [2]. A large category of eHealth applications are designed to monitor users in *everyday settings* through wearable sensors, acquiring often noisy continuous longitudinal data, compared to that of reliable clinical-grade sensors used in clinical settings [3]. Input sensory data perturbations including noisy and unreliable signal components, motion artifacts, and physical failure of sensors challenges the resilience of eHealth applications [4]. Processing perturbed sensory inputs affects prediction accuracy of machine learning algorithms, while incurring significant computational and energy expense [5].

Multi-modal eHealth applications are executed in two phases viz., *sensing* - to acquire complementary data from multi-modal sensors for an application, and *sense-making* - to analyze the acquired data for predictive results [6]. Input data perturbations at the multi-modal sensing phase affects the computational workload and energy consumption in the subsequent sense-making phase. Selective sensing through data filtering and pre-processing reduces the total data volume to be sensed and processed at the sensing phase itself

which consequently leads to energy savings. Selective sensing extracts insightful data to alleviate compute pressure, reduce energy consumption and improve prediction accuracy in the sense-making phase. State-of-the-art multi-modal applications use selective sensing techniques such as data filtering and compression [7], context-aware sensing [8], and heterogeneity-aware sensing [9] to reduce the total input data volume. Other techniques target model optimization to reduce the compute intensity of sense-making algorithms [10]. Existing optimizations for multi-modal machine learning (MMML) based eHealth applications address multi-modal sensing and sense-making independently [11]–[13]. Improving end-to-end system metrics of performance, energy consumption, and prediction accuracy necessitates *joint optimization of sensing and sense-making phases*. This exposes opportunities for i) adaptive sensing: selectively extracting quality-related input data and features - to alleviate compute pressure and reduce energy consumption incurred in processing garbage and/or less insightful data, and ii) adaptive sense-making: choosing suitable machine learning models by considering selective features and modalities from the sensing phase - to improve prediction accuracy. Joint optimization of multi-modal sensing and sense-making requires continuous monitoring of sensor modalities to identify input data perturbations, selective feature aggregation to isolate quality-related inputs, and choice of learning algorithms suitable for the given sensor modalities and feature vectors. To this end, we propose an adaptive multi-modal sensing and sense-making approach, AMSER, for improving performance, energy consumption, and prediction accuracy of eHealth applications. To demonstrate the potential of our AMSER approach in improving both sensing and sense-making, we design a proof-of-concept rule-based controller for adaptive feature selection, machine learning model selection, and sensor configuration, based on the monitored signal qualities. We integrate the controller into our framework running real eHealth applications to guide the sensing and sense-making joint optimization decisions in real-time for resilient, low-latency, and energy efficient eHealth services. We summarize our contributions as follows:

- Design of sensor-edge framework for multi-modal machine learning based eHealth applications, capable of monitoring input signal quality, and detection of discrepancies to guide sensing and sense-making optimization decisions.
- Design of a proof-of-concept rule-based controller for adaptive feature selection, modality selection, machine learning model selection, and sensor configuration based on sensor modality monitoring - to improve performance, energy, and

prediction accuracy.

- Demonstration of our framework’s efficiency on two eHealth case studies: pain assessment and stress monitoring.

II. MOTIVATION

We illustrate the resilience challenges in MMML-based eHealth services and the significance of adaptive sensing and control in addressing those challenges through a motivational example of a pain monitoring application [14]. The pain monitoring application acquires physiological data from different modalities viz., Electrocardiography (ECG), Electrodermal Activity (EDA), and Photoplethysmography (PPG) sensors to capture the autonomic nervous system activity against pain. Figure 1 shows the pipeline composed of multi-modal sensor data acquisition, feature aggregation, and machine learning model selection - for training and prediction. In this example, sensors from each of the ECG, EDA, and PPG modalities are sampled at 500Hz (2 channels), 4Hz, and 64 Hz, respectively. We extract relevant features from the raw input data and aggregate the features from all the modalities into an early-fused feature vector. We select a suitable machine learning model and train using the aggregated features to predict the pain levels. The model yields an average accuracy of 81% considering no noise in the input modalities. In practical scenarios, sensory data from one or more modalities can be noisy, and/or have data unavailability with perturbations such as motion artifacts, physical damages and battery shutdown [2]. In this example, we describe the implications of such sensory data perturbations on prediction accuracy and energy consumption, using different approaches of application orchestration.

(a) Processing noisy data. Figure 1 (a) shows the baseline scenario of processing sensory data without considering the noisy components. The ECG, EDA, and PPG modalities generate data of 8KB, 32B, and 512B, with 52, 42, and 42 features, respectively. In this case, the data from the ECG modality has a significant noisy component (e.g., due to improper patch contact to the chest), while the EDA and PPG modalities (e.g., collected from the wrist) generate quality inputs. Ignoring the input quality, features from all the modalities are fused into a vector comprising 136 features. The model pool comprises a set of machine learning models suitable for different feature vectors and input modalities. From the model pool, a machine learning model that suits the feature vector fused from 3 modalities (3-modal, 136 feature) is chosen for training and prediction. This model yields an average accuracy of 51%, with an end-to-end latency of 776 ms, and energy consumption of 5.14 J and 3.85 J, for the edge and the sensor devices, respectively. The loss in prediction accuracy can be attributed to the noisy data from the ECG modality (garbage-in garbage-out effect). Further, this levies an unnecessary energy consumption incurred in sensing and processing the noisy data.

(b) Processing with selective feature aggregation. Figure 1 (b) shows an optimized scenario with selective feature aggregation for handling modalities with noisy data. By considering the noisy components and aggregating only a selected set of features from ECG modality, we improve prediction accuracy, while lowering the computational effort and energy consumption. In this scenario, the *Selective feature aggregation* selects

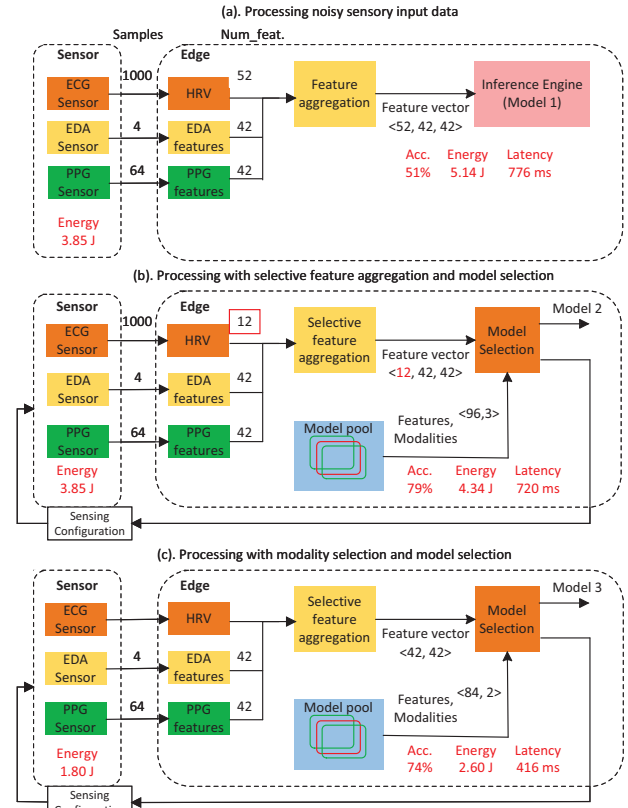


Fig. 1: Example experimental scenario demonstrating adaptive sensing configuration for feature and modality selection.

12 features from ECG modality (shown in red box) among the original 52 full scale features, and aggregates them with EDA and PPG features. The updated feature vector is sent to the *Model Selection* to choose a machine learning model that suits the feature vector for training and prediction. This model with selective feature aggregation yields an average accuracy of 79%, which is 35% higher than that of the baseline scenario with no input data considerations and only 2% less than the baseline with the no-noisy modality condition (ideal condition). Reduced features from the noisy ECG modality lowers the computational effort and energy consumption. In this case, selective feature aggregation improved the latency by 7% and energy consumption of the edge device by 16%. The sensing energy consumption does not change in this scenario since sensing configuration is unchanged.

(c) Processing with modality selection: Figure 1 (c) shows the optimized scenario with modality selection for handling modalities with unavailable data. In this scenario, data from the ECG modality is completely unavailable. With data from the EDA and PPG modalities being available, features from these modalities can be aggregated to build an updated feature vector. The *Selective feature aggregation* aggregates 42 features from EDA and 42 features from PPG modalities. From the *Models Pool*, the *Model Selection* chooses a machine learning model that suits the updated feature vector with reduced number of modalities and features (2-modal, 84 feature) for training and prediction. This model yields an average accuracy

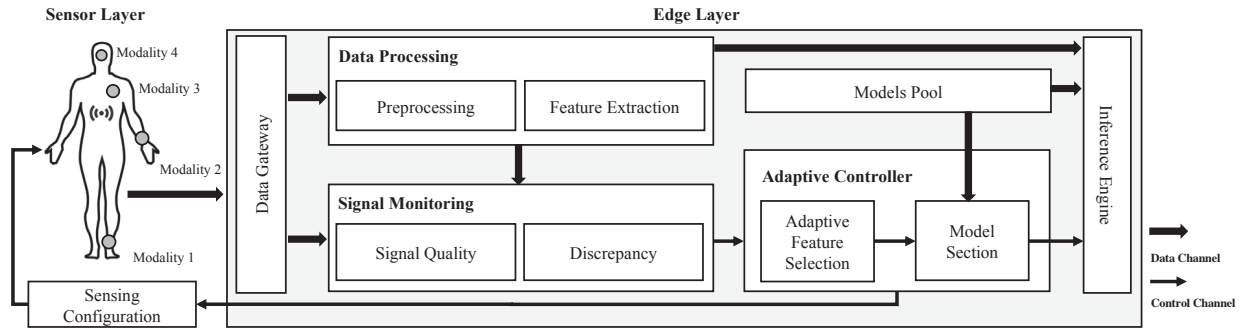


Fig. 2: System Architecture Overview.

of 74%, which is 31% higher than that of the baseline scenario with no input data considerations. Note that reducing the number of modalities to 2 (from 3 in Scenario (b)) lowers the accuracy only by 7% (compared to the ideal scenario). Dropping the unavailable ECG modality reduces the total data volume to be processed significantly. If the information on the sense-making decision is fed back to the sensing layer to adaptively change the sensing configuration, further energy saving at the sensing layer is also possible, which is critical in battery-powered wearable sensors. Overall, this leads to 46% improved latency, 50% lower energy consumption at the edge device, and 53% lower energy consumption at the sensors as compared to the baseline scenario of processing noisy data.

The example scenarios presented in Figure 1 demonstrates the advantages of monitoring input modalities for selective feature aggregation, modality selection, and model selection. However, this requires continuous monitoring of input modalities, and intelligent control for selective aggregation, modality, and model selection. To this end, we propose an automated framework for adaptive multi-modal sensing to improve resilience and energy efficiency of MMML-based eHealth applications.

III. ADAPTIVE MULTIMODAL SENSING-SENSEMAKING FRAMEWORK

We deploy a commonly used eHealth system architecture to implement our AMSER approach, that consists of multiple sensor devices at the sensor layer and computing resources at the edge layer. Figure 2 shows an overview of the proposed framework. The *sensor layer* includes the multi-modal signal capability needed for eHealth monitoring applications. The edge layer is a computing device providing data gateway, data processing, signal monitoring, adaptive control, and inference at the edge. We detail each level of the framework below:

Data Gateway receives raw signal data from multi-modal sensors. The physiological signals can potentially contain two main types of noises: Baseline Wander (BW) and Motion Artifacts (MA).

Data Processing performs pre-processing and feature extraction with the data collected from multi-modal sensors. *Pre-processing* module consists of data synchronization of multi-modal signals, and bio-signal processing (e.g., peak detection) from different modalities such as ECG, PPG, and EDA that are essential for feature extraction, and finally filtering sensor modality inputs to produce clean signals. *Feature Extraction*

module extracts informative and non-redundant values from the filtered signals. We extract time domain (mean, standard deviation, rms) and frequency domain (power spectral density, median frequency, central frequency) handcrafted features among the automatic features using a variational autoencoder for all modalities [15]. This process facilitates subsequent learning, leading to better interpretations for signal quality assessments.

Signal Monitoring handles quality assurance of the data and features to be used in the inference engine. This module observes the system parameters, disruptive events, data quality, and control flow to analyze the situation and context. *Signal Monitoring* assesses data and its extracted features quality by monitoring key parameters from the sensing phase to identify events and triggers for joint optimization of sensing and sense-making. For example, a disruptive event in input data of a specific modality (e.g., motion artifacts or sensor detachment) is a trigger at the sensing phase. In this regard, as shown in Algorithm 1, signal monitoring assesses data (D_M^i) and its extracted feature (F_M^i) quality and outputs the level of reliability of each modality (Rel_M^i).

Signal quality: We consider three levels of signal quality viz., reliable, noisy, and uncertain. *Signal quality* will assess the quality of sensing in terms of electrodes attachment, signal-to-noise ratio (SNR), and motion artifacts to determine the level of reliability. A signal with no noisy components is labeled as *reliable*. Signals extracted from sensor modalities whose electrodes are detached from the subject's body and signals with an SNR lower than a fixed threshold [16] are labeled as *noisy*. Signals with an SNR above the fixed threshold yet below the acceptable levels are labeled as *uncertain*.

Discrepancy Detector: The *Discrepancy Detector* acts as a second layer for assessing the quality of signals that are labeled as *reliable* in the *Signal Quality* phase by analyzing modality-specific parameters to identify any discrepancies. *Discrepancy Detector* will recognize confounding factors in each modality to identify levels of signal reliability. For instance, for ECG and PPG modalities, we use the range of HR, RR interval (RRi) length, the ratio of Max RRi to min RRi with their thresholds as the parameters to determine discrepancies [16]. As another example, for EDA modality, we use low amplitude and very low and steady tonic level for more than 1 minute as the discrepancy detecting parameter [17], [18].

Adaptive Controller analyzes inputs from the signal monitoring module to adaptively select features and modalities,

Algorithm 1: Signal Monitoring Strategy.

Data: $D_M^i, F_M^i, sensorStatus_M^i$
Result: Rel_M^i

```
1  $Rel_M^i \leftarrow RELIABLE$ ;  
2 if  $sensorStatus_M^i == DETACHED$  then  
3    $Rel_M^i \leftarrow NOISY$ ;  
4 else  
5   if  $SNR \leq threshold1$  then  
6      $Rel_M^i \leftarrow NOISY$ ;  
7   else  
8     if  $threshold1 \leq SNR \leq threshold2$  then  
9        $Rel_M^i \leftarrow UNCERTAIN$ ;  
10 for each modality  $D_M^i$ ;  
11 if  $Rules(D_M^i) == PASS$  then  
12    $Rel_M^i \leftarrow RELIABLE$ ;  
13 else  
14    $Rel_M^i \leftarrow UNCERTAIN$ ;
```

and configure sensing and computation models. We store pre-trained models for different combinations of signal modalities and aggregated feature vectors in a pool of models. The adaptive controller is designed to holistically and dynamically control the quality of sensing and sense-making, as shown in Algorithm 2. The adaptive controller considers the reliability of each sensor modality (Rel_M^i), the feature vector (F_M^i) extracted from *Signal Monitoring*, and the model pool for feature aggregation and model selection. The controller module adaptively selects the features that are relevant for prediction accuracy, and a learning model that is suitable for the updated feature vector using a simple rule based approach. In case a modality is labelled as *uncertain*, we process the modality by dropping some of the less prominent features. If a modality is *noisy*, we drop the input data from the modality entirely. The controller will generate a signal to update the sensing configuration, such that *noisy* sensor modalities can be turned off. Sensor modalities are turned into the idle mode when the noise cannot be mitigated and thereby eliminating the unnecessary computation and communication penalty. Finally, *Model Selection* utilizes the selected features, and the current models available in the *Model Pool* to load a preferable model for the *Inference Engine* taking into consideration the sensing status and situation. Algorithm 2 is a simple proof-of-concept rule-based policy to show the potentials of this approach. More advanced algorithms (e.g., using reinforcement learning or fuzzy control) can be devised to make intelligent decisions.

IV. EVALUATION

We evaluate the proposed adaptive sensing framework through two case studies of pain assessment [19]–[21] and stress monitoring [22]. Both case study applications require continuous monitoring of multiple modalities including ECG, EMG, PPG, and EDA signals, to detect pain and stress levels. Our evaluation platform comprises a sensory node - to collect physiological signals of ECG, EMG, PPG, and EDA from the subject, and an edge node - to execute the inference. We implement the proposed method on a ODROID-XU3 with an octa-core Exynos processor as the edge device. In

Algorithm 2: Adaptive Controller.

Data: Rel_M^i, F_M^i, MP
Result: Inference Model

```
1  $Sens_M^i \leftarrow ON$ ;  
2 if  $Rel_M^i == NOISY$  then  
3    $Sens_M^i \leftarrow OFF$ ;  
4 else  
5   if  $Rel_M^i == UNCERTAIN$  then  
6      $F_M^i \leftarrow$  Choose subset of  $F_M^i$   
7   else  
8      $F_M^i \leftarrow F_M^i$   
9 if  $Sens_M^i$  and  $F_M^i$  then  
10   $Inf_M^i \leftarrow MP \cdot M^i$ 
```

the following, we detail the two case studies and evaluate the case study applications' accuracy, energy efficiency, and performance using our AMSER proposed framework through different scenarios with different types and levels of noisy and unreliable input modalities. Further, we compare our AMSER approach against state-of-the-art multi-modal pain monitoring application that use classification of physiological signals [23] as the baseline ideal scenario.

A. Case Studies

We validate and demonstrate the efficacy of our AMSER approach through two case studies from the affective computing domain: 1) pain assessment and 2) stress monitoring.

1) *Pain assessment*: State-of-the-art pain assessment tools analyze changes in physiological signals indicating different pain levels. In this work, we evaluate AMSER via the pain assessment pre-trained models. These models are trained using UCI_iHurtDB, a multi-modal dataset from postoperative patients in hospital from 20 patients [24]. Data is collected from an eight-channel biopotential acquisition system for EMG and ECG recording and Empatica E4 wristband for PPG and EDA recording. We used the sensor modalities of 500 Hz ECG and EMG, 64 Hz PPG, and 4 Hz EDA during our experiments. We pre-processed each modality and segmented them into 60 second windows. For the prediction model to handle potential noisy components in physiological signals, we augmented the raw input training data with true types of noises such as BW and MA. Noisy components induced in raw data are based on the characteristics of each modality and extraction point such as chest, face, or wrist. Then, we extract a combination of time domain, frequency domain and automatic features from this window and trained using ML algorithms.

2) *Stress Monitoring*: The state-of-the-art for stress monitoring presents objective daily stress measurements in everyday settings based on physiological signals [25], [26]. In this work, we evaluate our proposed AMSER approach via stress monitoring application based on pre-trained models using WESAD - a multi-modal dataset featuring physiological and motion data recorded from both a wrist- and chest-worn device of 15 subjects during a lab study [22]. Sensor modalities used during our experiments are a 700 Hz ECG, 64 Hz PPG, and 4 Hz EDA. We then pre-processed each modality and created 60second window segmentation. We augmented true types

TABLE I: Experimental Scenarios for Pain and Stress applications. +*BW*, and +*MA* represent the presence of baseline wandering and motion artifact noise, respectively.

		ECG	EMG	PPG	EDA
Pain	S1	-	-	-	-
	S2	+BW	+BW	+BW	+BW
	S3	+BW	+BW	+BW+MA	+BW
	S4	+BW+MA	+BW+MA	+BW	+BW
Stress	S1	-	N/A	-	-
	S2	+BW	N/A	+BW	+BW
	S3	+BW+MA	N/A	+BW+MA	+BW
	S4	+BW+MA	N/A	+BW	+BW+MA

of noises such as BW and MA based on the characteristics of each modality. Then, we extracted a combination of time domain, frequency domain and automatic features from this window and trained using ML algorithms.

B. Scenarios

We evaluated our proposed AMSER approach in 4 different scenarios with different noise component levels for each case study. Table I summarizes different scenarios (S1-S4), showing the type of noise induced in each scenario. Scenario 1 is the ideal baseline case with no noise components, Scenario 2 has *uncertain* signal components, Scenario 3 has one modality with *noisy* component, and Scenario 4 has two *noisy* modalities. In Scenario 1, there is no noise component, and thus no modality selection or feature selection is applied. Scenario 2 comprises of baseline wandering noise in addition to the original data for all the modalities. With the presence of noise, our proposed AMSER approach selects specific significant features, reducing the total number of features from the original feature vector. In Scenario 3, we add additional motion artifacts on top of one modality while all the modalities still have a baseline wander noise. In this scenario, the noisy modality is completely unreliable. Thus, our proposed method will drop the modality and uses a different model from *Model Pool* with 3 and 2 modalities for pain and stress case studies, respectively. Scenario 4 comprises 2 modalities facing severe motion artifact noises. In case of pain assessment, ECG and EMG modalities are noisy with entirely unreliable data. Our proposed adaptive sensing will drop both the ECG and EMG modalities and select a learning model with 2 modalities for pain assessment application. In case of stress monitoring, ECG and EDA modalities are noisy. Our proposed adaptive controller drops both the ECG and EDA modalities and selects a learning model with the single PPG modality for stress monitoring.

We evaluate the accuracy and energy efficiency of both the pain monitoring application (which estimates the pain intensity using 4 physiological signals ECG, EMG, EDA, and PPG), and the stress monitoring application (which calculates the stress level using 3 physiological signals ECG, PPG, and EDA) using the iHurt_DB [24]. We trained different models for each levels of pain intensities, varying in the types of modalities combined for the decision-making system. After pre-processing, we extract a set of unique features from each modality viz., 52 features of ECG, 42 features of EDA, and 42 features of PPG [19]–[21], [26].

TABLE II: Decision made by AMSER during each scenario.

		ECG	EMG	PPG	EDA	Feature Vol (%)
Pain	S1	✓	✓	✓	✓	100
	S2	✓	✓	✓	✓	63.8
	S3	✓	×	✓	✓	35.8
	S4	×	×	✓	✓	31.3
Stress	S1	✓	N/A	✓	✓	100
	S2	✓	N/A	✓	✓	75.4
	S3	✓	N/A	✓	✓	31.6
	S4	×	N/A	✓	✓	14




Fig. 3: Accuracy analysis for AMSER vs. Baseline [23] for Pain and Stress applications.

C. Accuracy Evaluation

We compare the accuracy achieved by the model in each scenario with noisy components (S2-S4) against the baseline ideal scenario *S1* with no adaptive modality and feature selection. Figure 3 shows the accuracy (in %) for pain assessment and stress monitoring case studies under different scenarios. Our proposed adaptive multimodal sensing technique provides an improved accuracy in each of the scenarios for both the case studies, with a maximum gain of 22% in S2 for pain assessment. In the presence of a lower level of noise i.e., *uncertain* signal, our proposed AMSER approach still utilizes the *uncertain* modality with selected features, instead of either dropping or holding the entire feature sets of the noisy modality. This approach leads to a 22% and 17% of accuracy improvement for pain and stress applications respectively, in comparison with the baseline model which uses all the features from the noisy modality. With a higher level noise i.e., *noisy* modality in scenario *S3*, our proposed AMSER approach leads to 15% and 10% accuracy improvement for pain and stress applications, respectively, compared to the baseline which uses all the modalities. Lastly, scenario *S4* shows that dropping 2 noisy modalities results in better accuracy (13% and 11% improvement for pain and stress applications, respectively), as compared to the baseline using noisy modalities.

D. Efficiency and Performance Evaluation

We evaluate the efficiency and performance of AMSER in comparison with the baseline. Figure 4 shows the energy efficiency improvement during four scenarios for edge and sensor devices. We report the results for both Pain and Stress applications. Figure 4 (a) shows the energy efficiency of the pain assessment application. In Scenario *S4*, the energy gain is $2.19\times$ and $5.63\times$ higher than the baseline for the




Fig. 4: Energy efficiency analysis of the edge and sensor device for AMSER vs. Baseline [23] for Pain and Stress applications.




Fig. 5: (a) Performance analysis of the edge device for AMSER vs. Baseline [23] for Pain and Stress applications. (b) Data volume transferred between the sensors and edge for AMSER vs. Baseline.

edge and sensors, respectively. In this case, AMSER provides 13.27% better accuracy by dropping two noisy modalities (See Figure 3 and Table II). Figure 4 (b) shows the energy efficiency for the stress monitoring application. In Scenario *S3*, our evaluation shows that AMSER improves the energy efficiency of the stress monitoring application by $1.32\times$ and $2.63\times$ for the edge and sensor devices, respectively (See Figure 4 (b)). In this case, AMSER improves the accuracy and performance by 10.56% and 27%, respectively (See Figure 3). Through the Scenario *S3* for Stress application, AMSER keeps 31.6% of the features and drops the unreliable features (See Table II). Figure 5 (a) shows the performance gains at the edge device for the pain assessment and stress monitoring applications using our proposed AMSER approach, as compared to the baseline Scenario *S1*. In Scenarios *S2-S4*, our proposed AMSER approach reduces the number of features and/or drops input data from *noisy* modalities. This lowers the total computational effort, leading to an increase in performance, as compared to the baseline. Figure 5 (b) shows the data volume transferred between the sensory devices and edge node using our proposed AMSER approach, in comparison with the baseline. Our AMSER approach adaptively selects features and/or drops *noisy* modalities, which reduces the total input data volume significantly. For instance, in Scenario *S4*, two noisy modalities are entirely dropped, leading to more than $6\times$ reduction in transmitted data volume. Such reduction in data volume to be transmitted further reduces the communication latency and energy expense incurred in transmitting noisy data.

V. CONCLUSION

We proposed AMSER, a novel adaptive, closed-loop multimodal sensing framework for energy efficient and resilient eHealth applications. AMSER guides sensing and sense-making by monitoring input signal quality and discrepancy detection via an intelligent control framework that reduces garbage data to achieve both energy efficiency as well as improved quality. We demonstrated the AMSER framework's efficiency with two eHealth case studies on pain assessment and stress detection, achieving up to $2.2\times$ and $5.6\times$ performance and energy gains respectively, compared to the baseline with no optimization, while achieving up to 22% better accuracy. Our future work will focus on devising more intelligent resource management policies for adaptive control and co-optimization of both sensing and sense-making.

REFERENCES

- [1] Colin O'Keeffe et al. Role of ambulance response times in the survival of patients with out-of-hospital cardiac arrest. *EMJ*, 2011.
- [2] Emad Kasaeyan Naeini et al. An edge-assisted and smart system for real-time pain monitoring. In *CHASE*, 2019.
- [3] G. Yang et al. A health-iot platform based on the integration of intelligent packaging, unobtrusive bio-sensor, and intelligent medicine box. *IEEE Transactions on Industrial Informatics*, 2014.
- [4] Emad Kasaeyan Naeini et al. A real-time ppg quality assessment approach for healthcare internet-of-things. *ANT*, 2019.
- [5] Ghulam M. et al. A comprehensive survey on multimodal medical signals fusion for smart healthcare systems. *Information Fusion*, 2021.
- [6] Jiquan Ngiam et al. Multimodal deep learning. In *ICML*, 2011.
- [7] Meenu Rani et al. A systematic review of compressive sensing: Concepts, implementations and applications. *IEEE Access*, 2018.
- [8] Charith Perera et al. Context aware computing for the internet of things: A survey. *IEEE communications surveys & tutorials*, 2013.
- [9] Sharief MA Oteafy. A framework for heterogeneous sensing in big sensed data. In *GLOBECOM*, 2016.
- [10] D. Amiri et al. Context-aware sensing via dynamic programming for edge-assisted wearable systems. *ACM HEALTH*, 2020.
- [11] Masinelli et al. Self-aware machine learning for multimodal workload monitoring during manual labor on edge wearable sensors. *Design&Test*, 2020.
- [12] Murchana Baruah and Bonny Banerjee. Modality selection for classification on time-series data. In *MileTS*, 2020.
- [13] Sicong Liu et al. On-demand deep model compression for mobile devices: A usage-driven model selection framework. In *MobiSys*, 2018.
- [14] Philipp Werner et al. Automatic recognition methods supporting pain assessment: A survey. *IEEE Transactions on Affective Computing*, 2019.
- [15] Aqajari et al. pyeda: An open-source python toolkit for pre-processing and feature extraction of electrodermal activity. *ANT*, 2021.
- [16] C. Orphanidou et al. Signal-quality indices for the ecg and ppg: Derivation and applications to wireless monitoring. *JBHI*, 2015.
- [17] Mathias Benedek et al. A continuous measure of phasic electrodermal activity. *Journal of neuroscience methods*, 2010.
- [18] Hugo F. Posada-Quintero et al. Electrodermal activity is sensitive to cognitive stress under water. *Frontiers in Physiology*, 2018.
- [19] Naeini et al. Pain recognition with electrocardiographic features in postoperative patients: Method validation study. *JMIR*, 2021.
- [20] Aqajari et al. Pain assessment tool with electrodermal activity for postoperative patients: Method validation study. *JMU*, 2021.
- [21] Rui Cao et al. Objective pain assessment using wrist-based ppg signals: A respiratory rate based method. In *EMBC*, 2021.
- [22] Philip Schmidt et al. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *ICMI*, 2018.
- [23] M. Kächele et al. Multimodal data fusion for person-independent, continuous estimation of pain intensity. In *EANN*, 2015.
- [24] Naeini et al. Prospective study evaluating a pain assessment tool in a postoperative environment: Protocol for algorithm testing and enhancement. *JRP*, 2020.
- [25] Victoriano Montesinos et al. Multi-modal acute stress recognition using off-the-shelf wearable devices. In *EMBC*, 2019.
- [26] A. Tazarv et al. Personalized stress monitoring using wearable sensors in everyday settings. *arXiv*, 2021.