Edge-centric Optimization of Multi-modal ML-driven eHealth Applications

Anil Kanduri[‡], Sina Shahhosseini[†], Emad Kasaeyan Naeini[†], Hamidreza Alikhani[†], Pasi Liljeberg[‡], Nikil Dutt[†], and Amir M. Rahmani[†,§]

[†] Department of Computer Science, University of California, Irvine
Mailing address: Bren Hall 3086, Irvine, CA 92697-3435, USA
[§] School of Nursing, University of California, Irvine
Mailing address: 106D Berk Hall, Irvine, CA 92697, USA
[‡] Department of Computing, University of Turku, Finland
Mailing address: Vesilinnantie 3, Turku 20540, Finland
Email: spakan@utu.fi, sshahos@uci.edu, ekasaeya@uci.edu, hamidra@uci.edu, pasi.liljeberg@utu.fi, dutt@uci.edu, a.rahmani@uci.edu

Abstract. Smart eHealth applications deliver personalized and preventive digital healthcare services to clients through remote sensing, continuous monitoring, and data analytics. Smart eHealth applications sense input data from multiple modalities, transmit the data to edge and/or cloud nodes, and process the data with compute intensive machine learning (ML) algorithms. Run-time variations with continuous stream of noisy input data, unreliable network connection, computational requirements of ML algorithms, and choice of compute placement among sensor-edge-cloud layers affect the efficiency of ML-driven eHealth applications. In this chapter, we present edge-centric techniques for optimized compute placement, exploration of accuracy-performance trade-offs, and cross-layered sense-compute co-optimization for ML-driven eHealth applications. We demonstrate the practical use cases of smart eHealth applications in everyday settings, through a sensor-edge-cloud framework for an objective pain assessment case study.

Keywords: Multi-modal Machine Learning, Edge Computing, eHealth Applications, Reinforcement Learning

1 Introduction

Smart eHealth applications deliver critical digital healthcare services such as disease diagnostics, clinical decision support, forecasting health status, pro-active and preventive healthcare decisions, and alerts for emergency intervention, etc [19]. eHealth applications improve the reach and quality of healthcare services, timeliness and accuracy of clinicians decisions, and reduce the burden on healthcare professionals and overall medical expenditure [72]. Smart eHealth systems integrate remote sensing, continuous monitoring, wireless transmission, data analytics, and machine learning to deliver intelligent patient-centric digital healthcare and wellbeing services [35]. eHealth applications are particularly effective for
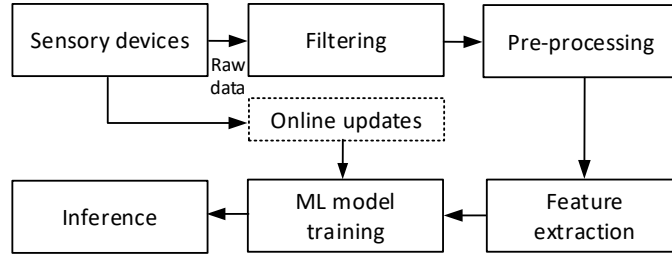
Fig. 1: ML-driven eHealth application pipeline

managing chronic patients through continuous monitoring, extracting clinically relevant data with minimal intrusion [17].

1.1 ML in Smart eHealth Applications

eHealth systems continuously monitor patients using wearable sensors for acquiring physiological parameters [16]. In addition to the bio-signals, eHealth applications also track behavioral, and environmental parameters to contextualize the patients' current situation [38]. Thus, smart eHealth applications generate huge volumes of heterogeneous input data, combining multiple streams of inputs from physiological, behavioral, and environmental parameters [19]. Analyzing such continuous stream of heterogeneous multi-modal raw data for predicting potential threats, accurate clinical decisions, and diagnostics requires broader support from the AI domain [24]. Smart eHealth systems are increasingly using ML algorithms for analyzing multi-modal input sensory data, to provide intelligent digital healthcare and wellbeing services [35]. State-of-the-art eHealth applications have applied different ML algorithms for analyzing input data, and predicting results on diagnostics, potential and health status [15]. ML-driven eHealth systems work in a pipeline of data acquisition, filtering and pre-processing, data analysis, training, inference for predictive results, followed by notification to the clients [19]. Figure 1 shows the workflow of typical ML-driven eHealth applications, where raw data acquired by sensory devices is filtered and pre-processed to remove noisy components, motion artifacts, and anomalies. This input data is then used for extracting relevant features, and training suitable ML models. Predictive results are achieved by inferencing the trained ML model, while the trained model is periodically updated with evolving input data.

1.2 Collaborative Edge Computing for Smart eHealth Applications

eHealth applications rely on traditional cloud infrastructure for training, storage, and updating of ML models, and inference tasks. However, rapidly increasing volumes of sensory input data, and uncertain network conditions imposes limitations on the efficacy of running eHealth applications on the cloud layer. Edge
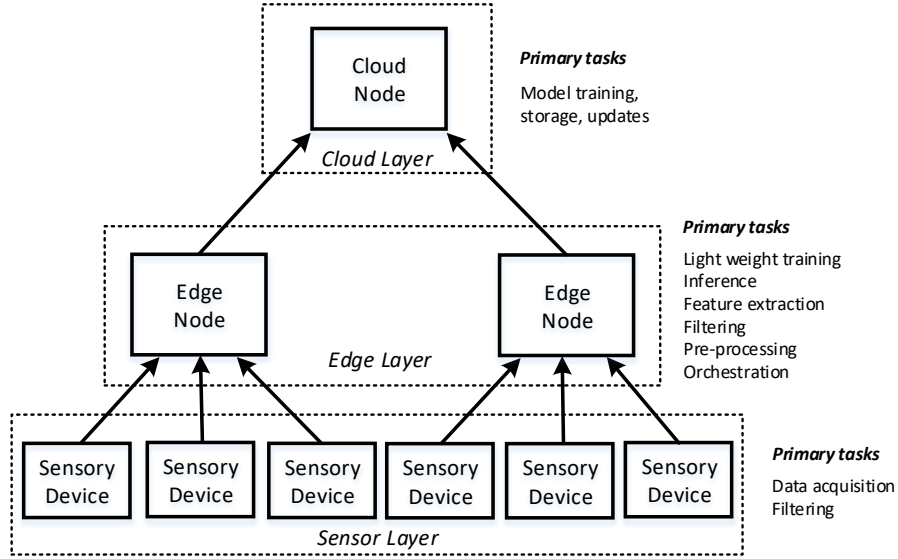
Fig. 2: Sensor-edge-cloud architecture

computing paradigm brings computational intelligence closer to the sources of input data, minimizing the reliance of smart eHealth applications on cloud infrastructure [74]. Edge computing architectures have been widely adopted for deploying ML-driven smart eHealth applications, simultaneously handling input data, compute intensity, and network constraints [54]. Figure 2 shows an overview of the hierarchical multi-layered sensor-edge-cloud architecture [69]. The *sensor layer* comprises of sensory devices such as wearable sensors, smart bio-sensors, and sensors deployed on mobile and IoT devices. The *sensor layer* primarily acquires raw data from different devices, performs lightweight tasks such as filtering, and transmits relevant inputs to the resourceful layers in the hierarchy. The *edge layer* receives input data from the sensory devices, and executes intensive tasks such as data pre-processing, feature extraction, lightweight ML model training, and inference tasks. More importantly, the *edge layer* also handles orchestration functionalities such as application-level and system-level monitoring, application partitioning, compute placement, and resource allocation. The *cloud layer* handles heavy computational tasks such as ML model training, updating, and storage, and notifications to edge nodes on model updates.

Example Scenario: We demonstrate the pipeline of a quintessential edge-centric ML-driven eHealth application through the example of arrhythmia detection application [19]. Figure 3 shows the data and control flows of the arrhythmia detection application [19] across the device-edge-cloud layers. Initially, raw electroencephalogram (EEG) signals are acquired by wearable sensory de-
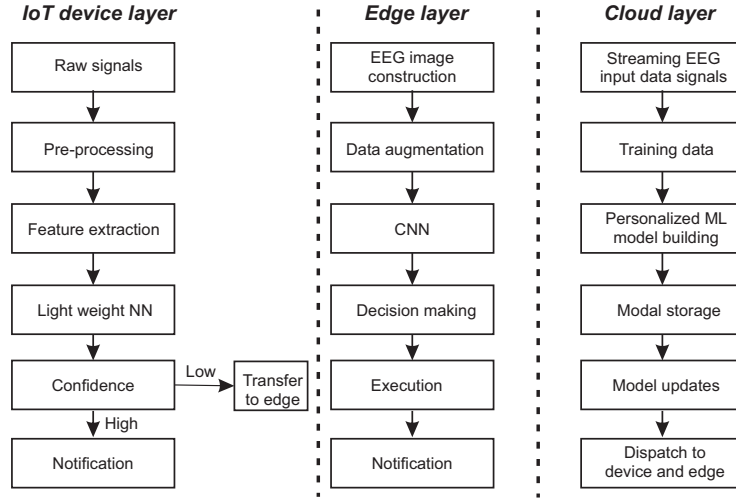
Fig. 3: Pipeline of arrhythmia detection

vices. The raw signals are pre-processed for noise removal to extract relevant features from the raw data for training an ML model. Considering the limited compute resources of the IoT device layer, a lightweight neural network (NN) model is employed to predict/detect arrhythmia. The predictions are notified to the client if the NN model has a higher confidence on prediction accuracy, while forwarding the input data to edge layer when the model confidence is lower. In this application, the edge layer uses the reconstructed EEG images for data aggregation. With relatively higher compute resources, the edge layer consists of a moderately intensive convolutional neural network (CNN) model to train on input data. The CNN model is inferred in the execution phase for arrhythmia detection/prediction, and the client is notified with the result. The cloud layer collects streaming inputs from the device and edge layers to train appropriate ML model, store the model, and update the model periodically with evolving input data. The cloud layer transmits the updated model parameters to the edge and device layers for running local inference tasks. Further the cloud layer performs advanced data analytics to generate personalized decisions for each client. It should be noted that the compute capabilities and thus tasks vary at each of the device-edge-cloud layers. Optimizing ML-driven eHealth systems requires such understanding on input data flows, computational requirements, and accuracy and performance of ML models.

1.3 Summary

Implementing ML-driven smart eHealth applications presents different challenges on sensory data acquisition, understanding application-level requirements, handling compute intensities of ML models, and energy and network constraints.

At the same time, deploying such applications on multi-layered sensor-edge-cloud platforms exposes opportunities for selective processing through input data quality awareness, choice of compute placement among edge-cloud nodes exploring energy-performance trade-offs, and understanding algorithmic nature of applications to explore accuracy-performance-energy trade-offs. Collaborative sensor-edge-cloud platforms enable layer-wise partitioning of smart eHealth application pipeline, to synergistically improve the quality of the services. In the subsequent sections, we present different edge-centric optimizations for ML-driven smart eHealth applications on collaborative sensor-edge-cloud platforms.

Organization: Section 2 presents an exemplar case study of pain assessment application, describing a sensor-edge-cloud framework integrating different edge-centric optimizations. This case study is then used in the following chapters to demonstrate some of the optimization techniques. Section 3 presents techniques for improving performance metrics of edge-centric ML workloads through efficient compute placement, and exploration of accuracy-latency trade-offs. Section 4 presents techniques for improving resilience of ML-driven eHealth applications, through sense-compute co-optimization. Section 5 concludes with key insights, and open research directions.

2 Exemplar Case Study of Edge-ML-driven Pain Assessment Application

In this section, we present an exemplar case study of pain assessment application, describing application characteristics, nature of input data, and challenges in typical edge-centric ML-driven smart eHealth application. We also present a modular framework, iHurt, for deploying ML-driven eHealth applications (pain assessment in this case study) on collaborative sensor-edge-cloud platforms. The iHurt framework serves as a generic platform for prototyping smart eHealth applications for processing sensory data using ML models. Further, iHurt platform provides a test-bed for evaluating the edge orchestration, compute placement, RL-agent based offloading, and sense-compute co-optimization techniques presented in Sections 3 and 4.

2.1 Pain Assessment

Pain is a complex phenomenon, associated with several illnesses [70]. Pain is defined as "an unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage" [46]. The pain assessment "gold standard" relies on a patient's self-report of their pain intensity on a scale of 0 to 10, where 0 refers to no pain, and 10 represents the most severe pain. Pain assessment is done through tools such as Numerical Rating Scale (NRS), Visual Analogue Scale (VAS), and Verbal Rating Scale (VRS). There is a high demand for objective tools to assess patients' pain in the
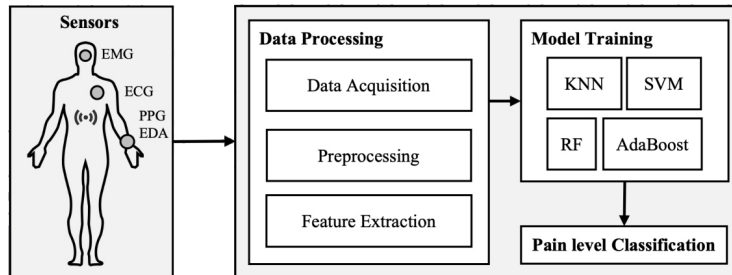
Fig. 4: Overview of pain assessment application

clinical context. Tools are needed especially when the patient's own opinion is difficult to obtain. Assessment of pain is particularly difficult when the ability of a patient to communicate is limited (e.g., during critical illness, in infants and preverbal toddlers, or in patients under sedation or anesthesia, with intellectual disabilities, and at the end of life) [9]. Inadequately treated pain has major physiological, psychological, economic, and social ramifications for patients, their families, and society [4]. Undertreatment of pain could result in many adverse effects and other complications and may evolve into chronic pain syndromes. It could also cause delayed discharge or prolonged recovery, which may incur higher health care costs and more patient suffering [65]. Overtreatment of pain, on the other hand, may result in unintended adverse consequences such as acute respiratory complications or in long-term complications such as opioid addiction. These issues are particularly pronounced for noncommunicative patients who are unable to articulate their experience of pain [7].

We demonstrate an abstract overview of the pain assessment application [30] in Figure 4. The pain assessment application is implemented as a pipeline of sensing, data processing, model training for predictive results, and inference for pain level classification. In the context of this case study, we use input data from EMG, ECG, PPG, and EDA sensors for estimation of pain. Sensory data is pre-processed for filtering qualitative inputs, followed by feature extraction. Relevant learning models are trained with the multi-modal input data sets for classifying pain level. Different phases of the pain assessment application are detailed in the following.

2.2 Sensory Data Acquisition

Objective pain assessment application collects raw data through continuous monitoring in both clinical and every day settings [19]. Raw data collected from the sensory data acquisition phase is used for training the ML models for accurate prediction of pain level. In this sub-section, we describe the nature and characteristics of sensory input data used in the pain assessment application.

Types of Signals: There are various types of signals that influence the accuracy of monitoring and assessing the affective states in pain assessment. These indicators are extracted through different forms of behavioral, physiological, and contextual/environmental sensor modalities via facial expression, speech, full-body motion, text, and physiological signals. Both behavioral and physiological manifestations of pain can be measured objectively. Behavioral pain indicators include facial expressions, body movements such as rubbing, restlessness, and head movements, and paralinguistic vocalizations such as crying and moaning. Physiological pain indicators are acquired from brain, cardiovascular, and electrodermal activities. Monitoring of these physiological, behavioral, and contextual sensor inputs can also be used other prominent eHealth applications including emotion recognition and stress monitoring. For instance, stress activates the autonomic nervous system (ANS) which can be detected through monitoring the changes in physiological signals including cardiovascular activity, and electrodermal activity, respiration rate, and blood pressure [22]. Further, physiological signals of cardiac function, temperature, muscle electrical activity, respiration, skin conductance, and brain electrical activity can be used to detect human emotions. The multitude of physiological, behavioral, and contextual signals fused together can provide valuable insights for training ML models, specifically in the domain of smart affective computing applications.

Commonly Used Sensors: Recording physiological signals require people to connect with biosensors. There are contact-based sensors (such as adhesive electrodes and wristbands) or contact-free sensors (such as cameras and microphones) to gather information from patients and analyze them. Widely used contact sensors in eHealth applications record electroencephalogram (EEG, electrical activity of the brain), electrocardiogram (ECG, heart activity (heart rate (HR) and heart rate variability (HRV))), electrodermal activity (EDA often measured using skin conductance level (SCL), and sometimes the old term "galvanic skin response" (GSR)), surface electromyogram (sEMG, muscle activity), photoplethysmogram (PPG, blood perfusion of the skin for pulse and other measures, also called blood volume pulse or BVP), respiration (RSP), or acceleration (ACC, movement). For pain monitoring, the uni-dimensional assessment tools have been questioned and debated for their oversimplification and limited applicability in non-communicative patients, since they require interactive communication between patient and caregiver [26]. As a result, physiological sources of data comprising of heart rate (HR), heart rate variability (HRV), SpO2, skin temperature, electrodermal activity (EDA), and facial expression and frontal muscle activity using computer vision or facial electromyography (EMG) and electroencephalogram (EEG) are prioritized for pain assessment. Accurate calculation of HRV parameters depends on detecting the position of peaks within ECG or PPG signals. Root Mean Square of Successive Differences (RMSSD) is an HRV parameter that is correlated to the short-term variation in the PPG signal. Figure 5 shows a minute filtered PPG signal illustrating an example, in which less than 5 seconds of the PPG signal (highlighted in red) are distorted
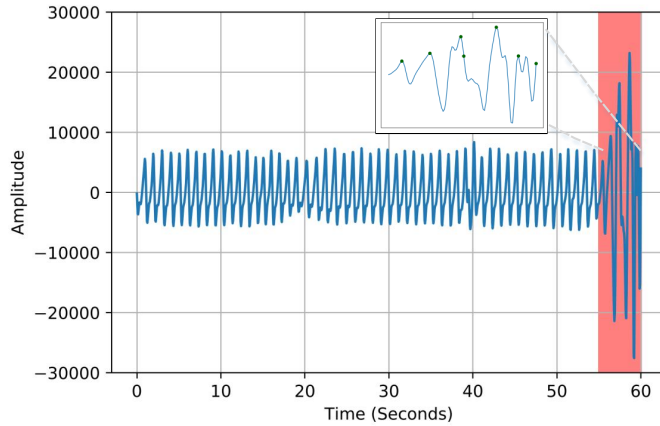
Fig. 5: One-minute windows of filtered PPG signals carrying noise

due to hand movements. Such a minor window of corrupted input data in the signal could still affect the eventual accuracy significantly. For instance, in Figure 5, few peaks are detected incorrectly within the noisy signal part, and thus the RMSSD is not reliable anymore during this window of data. The pain assessment application uses ML models to detect such abnormalities, and enable accurately predictions.

	EDA	ECG	RR	MM
BL vs. PL1	63.36	72.04	71.79	77.13
BL vs. PL2	79.24	81.13	82.14	85.64
BL vs. PL3	69.59	69.8	76.64	86.9
BL vs. PL4	63.7	63.41	66.67	74.73
Mean	68.97	71.6	74.31	81.1

Table 1: Summary of accuracy results (BL: baseline, PL: pain level, MM: multi-modal)

Multi-modal Inputs: Objective pain assessment relies on input data from multiple modalities and combinations of physiological, behavioral, and contextual parameters. All these modalities differ in terms of data, noise characteristics, comfort and ease of use, privacy concerns, and energy consumption. Using a single modality versus a combination of multiple modalities effects the computational workloads of the ML models, and eventual prediction accuracy. This is

demonstrated from a pain monitoring case study on a five levels of pain data collection [2, 10, 31, 33]. The accuracy of using each individual input modality of sensory data (EDA, ECG, RR), and multi-modal (MM) inputs for binary classification between no pain/baseline and various pain levels is shown in Table 1. It should be noted that different sensor modalities result in a range of prediction accuracies across different levels of pain. Predictions based on multi-modal input data set has the highest accuracy among each of these cases.

2.3 ML-driven Objective Pain Assessment

It is imperative to design and develop an objective monitoring tool to improve the well-being and care processes of patients with a more accurate assessment and more timely treatment. While this raises significant technical challenges, requiring a combination of sensing, signal processing, and machine-learning skill sets, it also has a tremendous potential to pave the way for next-generation human-modeling methods. Machine learning and deep learning techniques have become immensely popular for classification tasks, as well as for other recognition and pattern matching tasks. ML models can be used for accurate objective pain assessment, using input data from different modalities. The combination of vast amount of multi-modal input data, increased computing power and more intelligent methods enables fast and automated production of machine-learning algorithms able to analyze complex data with accurate results.
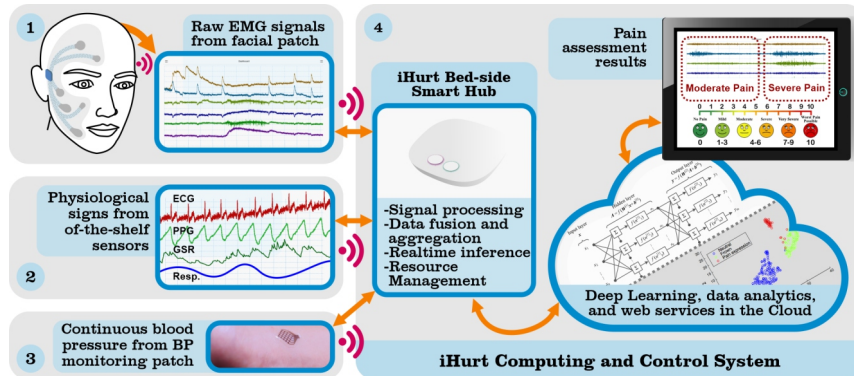

Fig. 6: The Objective Pain Assessment Technology

iHurt Platform Figure 6 presents a complete objective pain assessment technology developed jointly by researchers from University of California, Irvine (UCI) and University of Turku (UTU), Finland. This is an end-to-end system for multi-modal data acquisition, data processing and analyzing at the gateway. The first step in building a multimodal system is to process the raw signals collected

during trials. While this varies according to the deployed sensor data, a typical pre-processing pipeline consists of the following: Filtering (data cleansing, noise reduction, and artifact removal), segmentation (partitioning into time intervals) and normalization (w.r.t. to a baseline signal). These steps are followed by feature extraction to obtain features within various domains. Finally, the processed data can be used to build prediction models using machine learning. Different types of ML models can be used to train over multi-modal input data sets. Choice of the learning model determines the integration and fusion of multi-modal data at feature, decision, or intermediate levels (early, late, hybrid fusion). Most approaches classify pain using support vector machines (SVM) [23, 77], random forests (RF) [29, 76], and nearest neighbors (NN) [1]. Other widely used models include Adaboost and XGBoost [55], which are ensemble methods to reduce bias and variance in predictive data analysis. The ML models can be computationally fine-tuned differently to pursue various objectives. Sense-making knobs such as early exit, model selection, and input modality selection presented in Section 4 can be explored in this context. Figure 7 shows five different classification methods using respiratory signals including ADABoost, XGBoost, RF, SVM, and K-NN classifiers in comparison with a state-of-the-art method, RESP [10]. This case study down sampled pain levels into three levels of pain (pain levels 1-3), besides the baseline of no pain. Note that they achieved higher accuracy compared with the state-of-the-art while using only 88% less features. We can maintain accuracy in presence of noise, or using useful features from the reliable modalities, while also meeting the requirements set on the computational performance. The orchestration functionalities presented in Section 3 can guide these decisions on exploring accuracy-performance trade-offs.
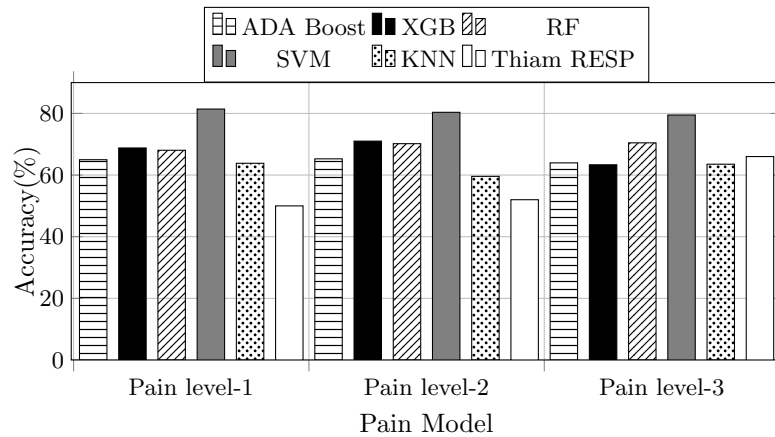

Fig. 7: Validation accuracy of classifiers on top-8 features for different pain levels compared against the baseline

Other Exemplar eHealth Applications: Emotion recognition and stress monitoring are two other widely used eHealth applications that rely on input data similar to that of the pain assessment application. Emotion recognition uses verbal inputs, and cues such as tone of voice, facial expressions, postures, gestures and also through physiological signals [37]. Stress monitoring application detects the existence of stress in each period of time using physiological signals [25]. The Electrodermal Activity (EDA) or Galvanic Skin Response (GSR) is one of the physiological signals generated by the human body, which can be used to detect stimuli in individuals. Fall detection is another exemplar application that uses 3-d accelerometer data to detect falls by using classification models [14]. The input data quality assessment, using multi-modal ML models, configuration of ML models, and edge orchestration techniques used for pain assessment an be analogously applied to other similar applications of emotion recognition, stress monitoring, and fall detection.

3 Edge-centric Optimization of ML-based eHealth Workloads

Machine learning (ML) is advancing real-time and interactive user services in healthcare domain [56]. ML applications are primarily deployed on cloud infrastructure to meet the compute intensity and storage requirements of ML algorithms, and address the resource constraints of user-end wearable sensory devices [6]. However, unpredictable network constraints including variable signal strength and availability of the network affect real-time delivery of cloud services [36]. The edge computing paradigm allows deployment of ML applications closer to the user-end devices, minimizing the latency of service delivery, reducing the total network load, and alleviating privacy concerns.

Systems Perspective Collaborative sensor-edge-cloud architecture presents multiple execution choices for workload partitioning and compute placement including execution on a single sensor, edge, cloud nodes, and any possible combinations of these devices. Considering the variable accuracy nature of ML algorithms, these execution choices expose a wide range of energy-performance-accuracy trade-off space. Choosing an optimal execution option under varying system dynamics, available energy budget of devices, network constraints, and error resilience of ML workloads is a complex run-time challenge.

Application Perspective Depending on the ML model employed, different applications feature varying compute, data, and communication intensities. At an application level, there is diversity in terms of sensitivities to latency, throughput, infrastructure availability, and accuracy. Further, different application execution choices result in different energy consumption patterns [32]. Considering these application level variations, the choice of execution of an application is a subject of multiple factors varying at run-time [21].

Edge Orchestration Edge orchestration techniques handle workload partitioning, distribution, and scheduling of ML workloads, considering both applications' and systems' perspectives. Understanding the varying compute intensi-

ties of applications, latency and throughput requirements, user-interaction and responsiveness expectations, quality of interconnection network including signal strength, availability, load balancing, compute and storage capacities of underlying hardware elements put together makes application orchestration a stochastic process [59, 62]. To maximize the efficiency of edge-enabled health care services, run-time solutions that holistically consider both requirements and opportunities vertically across the user, device, application, network, edge node, and platform layers are necessary [58]. In this Section, we present state-of-the-art edge orchestration techniques for efficient compute placement with rule-based heuristics, and optimized orchestration through reinforcement learning.

3.1   Dynamics of Compute Placement

Computation offloading techniques transfer the execution an application, or a task within an application, to a resourceful device for improving performance [44]. Compute placement determines the choices on partitioning an application, and selection of the external resourceful edge/cloud nodes onto which the partitioned task is to be offloaded [6]. Some of the existing compute placement and offloading strategies do not consider the diversity in applications' compute and communication requirements, eventual performance gain with offloading and compute placement choices, and potential latency penalties incurred with those choices. For optimal performance gains, compute placement techniques have to consider the dynamically varying application and network characteristics, energy budgets, and accuracy-performance tradeoffs [18, 60, 68].

We explore the intricacies of running smart eHealth applications on multi-layer sensor-edge-cloud platforms to demonstrate the dynamics of compute placement. We choose stress monitoring [47], fall detection [14], and pain assessment [25] as representative workloads from ML-driven digital healthcare systems. The stress monitoring application uses predictive models to extract statistical features from physiological parameters of Electro-dermal Activity (EDA) and Galvanic Skin Response (GSR), and predict stress levels [3]. The fall detection application uses de-noising, feature extraction, and decision-tree training for classification of fall and no-fall events [58]. The pain assessment application uses pre-processing, feature extraction, and SVM classification for determining level of pain [40]. In summary, each application typically includes the pipeline of data pre-processing, feature extraction, and classification ML tasks. We execute these workloads on real hardware testbed that emulates a baseline sensor-edge-cloud platform. We consider an edge platform with configurable on-board sensors for sensing at variable sampling rates, and connectivity to cloud infrastructure. We consider three compute placement policies with the execution choices of running the ML workloads: i) fully on the edge device, ii) fully on the cloud node, iii) partially on the edge and partially on the cloud. For evaluation, we define *latency* metric as the time taken to respond to a request including the communication and execution costs.

Figure 8 shows the latency of stress monitoring, fall detection, and pain monitoring applications with different compute placement choices, over different
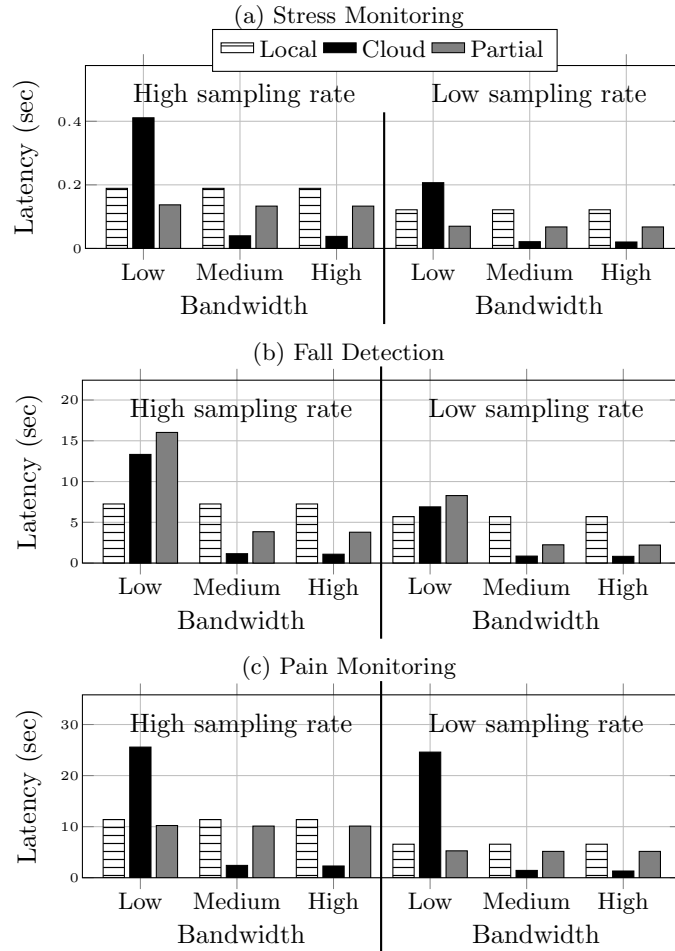
Fig. 8: Latency with different compute placement strategies for eHealth applications under different network bandwidth (low, medium, high) and sampling rates (high and low). (a) Stress monitoring, (b). Fall detection, (c). Pain monitoring.

network bandwidths, and sampling rates of input sensory devices. The compute placement choices Local, Cloud, Partial represent execution of applications on edge node, cloud node, and collaborative edge-cloud nodes. We used three levels of available bandwidths – low, medium, and high. Available bandwidth influences the latency of execution, specifically incurred in transmitting data from edge to cloud nodes. We used two levels of sampling rates – high and low, for input data sensory devices. Sampling rate determines the total volume of data being transmitted from edge to cloud nodes, effecting the latency.

The evaluation shows that network variation in terms of available bandwidth substantially effects the latency and choice of compute placement. For example, consider the scenario shown in Figure 8 (a), under high sampling rate and high

bandwidth. In this case, compute placement on the cloud is the optimal choice in comparison with the edge, and partial choices. Since there is sufficient network bandwidth, the penalty of transmitting data from edge to cloud is minimal, and the performance gain of executing the application on the cloud is also significant. In contrast, for the same scenario with low bandwidth (LBW), compute placement on the cloud has the highest latency, owing to the higher penalty of transmitting data from edge to cloud under low bandwidth. In this case, the partial (edge-cloud) execution has a better latency.

In addition to the bandwidth availability, the application's nature can substantially influence the choice of compute placement. For instance, consider the scenario in Figure 8 (b), under high sampling rate and low bandwidth. The fall detection application's compute-communication ratio, and lower bandwidth makes the local execution optimal, as opposed to offloading to the cloud node.

The sensing configuration of an application alters the volume of data generated for transmission and processing. For instance, consider the scenario in Figure 8 (c), under low sampling rate and medium bandwidth. As the sampling rate is lower, the penalty incurred in transmitting data to the cloud node is minimized. With lower data volume and availability of either medium (and/or high) bandwidth results in significant improvement in latency.

It should be noted that lowering the sampling rate potentially sacrifices the accuracy, although the latency is improved. While the accuracy loss is subjective to the error resilience of the application, missing insightful input data samples could lead to mis-predictions and critical errors. Mis-predictions can be minimized by setting an upper bound on accuracy requirements and/or bounds on lowering the sampling rates. In such scenarios, the error resilience of an application also influences the compute placement decisions.

3.2 Using RL for Optimization

Finding the optimal orchestration policy for an unknown and dynamic system is critical since dynamicity of environment (e.g., network condition, workload arrival at computing nodes, user traffic, and application characteristics) changes over time. Most current solutions are based on design time optimization, without considerations on varying system dynamics at run-time [5, 11–13, 34, 41, 45, 51, 64, 79, 80]. A complex system that runs a variety of applications in uncertain environmental conditions requires dynamic control to offer high-performance or low-power guarantees [32, 58, 60, 61, 63]. Considering the run-time variation of system dynamics, and making an optimal orchestration choice requires intelligent monitoring, analysis, and decision making. Existing heuristic and rule-based orchestration methods require an extensive design space exploration to make optimal compute placement decisions. Further, such a solution based on exhaustive search at run-time becomes practically infeasible for latency critical services. In this context, different offline and online machine learning models have been adopted for run-time resource management of distributed systems, to handle the complexity of orchestration choices. Among these models, reinforcement learning approach (RL) is effective in developing an understanding and
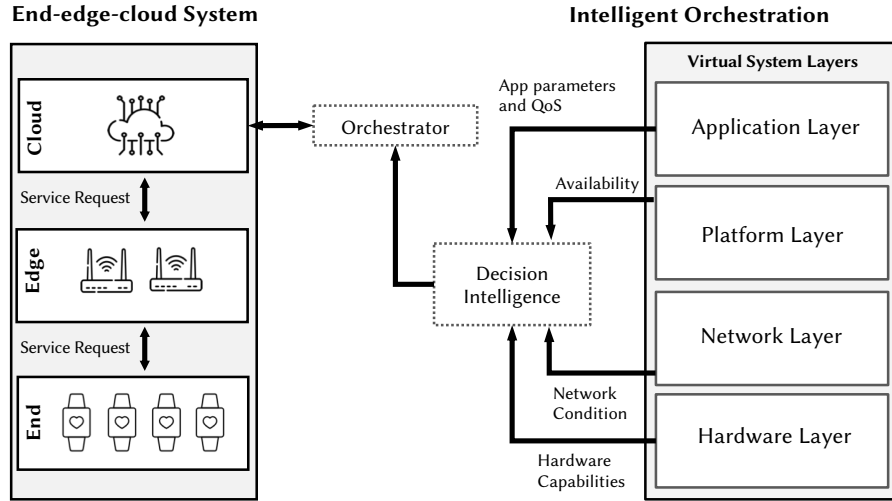
Fig. 9: Overview of intelligent orchestration in end-edge-cloud architectures

interpreting varying system dynamics [48, 53]. Reinforcement learning enables identification of complex dynamics between influential system parameters, and online decision making to optimize objectives such as response time, energy consumption, and quality of service [66]. RL approach allows formulating policies at run-time using the input data collected over time. Specifically, RL approach uses a reward function to quantify the effect of an action on the system state. This allows optimizing orchestration choices over time, considering the system wide context and objectives. In this sub-section, we present the design of reinforcement learning agent for orchestrating ML workloads on sensor-edge-cloud platforms.

Orchestration Framework: Figure 9 shows an overview of generic intelligent orchestration framework for multi-layered end-edge-cloud architectures [63]. This framework uses system-wide information for intelligent orchestration through virtual system layers that include application, platform, network and hardware layers. Each of the virtual system layers provide inputs for monitoring system and application dynamics such as application adjustment parameters, accuracy requirements, availability of devices for execution, network characteristics, and hardware capabilities. Each execution choice affects the performance and energy consumption of the user end-device, based on the system parameters such as hardware capabilities, network conditions, and workload characteristics. Each layer exhibits a diverse set of requirements, constraints, and opportunities to trade-off performance and efficiency that vary over time. For example, the application layer focuses on the user's perception of algorithmic correctness of services, while the platform layer focuses on improving system parameters such as energy drain and data volume migrated across nodes. Both application and

platform layers have different measurable metrics and controllable parameters to expose different opportunities that can be exploited for meeting overall objectives. The network layer provides connectivity for data and control transfer between different physical devices. In addition, the hardware layer provides hardware capabilities for computing nodes in the system. Run-time system dynamics affect orchestration strategies significantly in addition to requirements and opportunities. Sources of run-time variation across the system stack include workload of a specific computing node, connectivity and signal strength of the network, mobility and interaction of a given user, etc. Information on run-time cross-layer requirements and run-time variations provide necessary feedback to make appropriate decisions on system configurations such as offloading policies.

RL Agent for Orchestration: Making the optimal orchestration choice considering these varying dynamics is an NP-hard problem, while brute force search of a large configuration space is impractical for real-time applications. Understanding the requirements at each level of the system stack and translating them into measurable metrics enable appropriate orchestration decision making. On the other hand, heuristic, rule-based, and closed-loop feedback control solutions are slow in convergence due to the large state space [66]. To address these limitations, reinforcement learning (RL) approaches have been adapted for the computation offloading problem [57]. RL builds specific models based on data collected over initial epochs, and dramatically improves the prediction accuracy [66]. We design an RL agent that monitors system-wide parameters and chooses a suitable action that maximizes the efficiency of orchestration decisions.
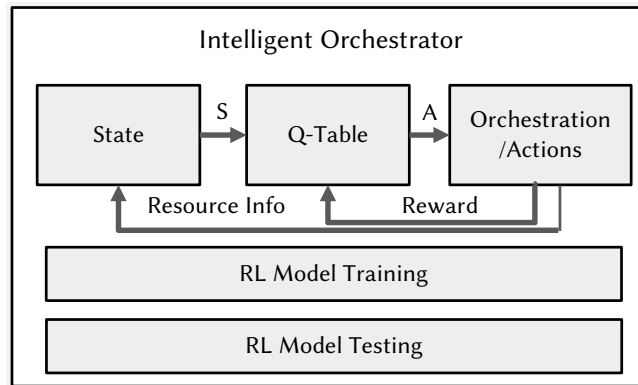

Fig. 10: RL agent for intelligent orchestration.

Figure 10 shows the workflow of the RL agent for making orchestration decisions. The RL agent component is deployed within the decision intelligence and orchestration blocks in orchestration framework (Figure 2). The RL agent

receives resource information (e.g., processor utilization, available memory, available bandwidth) from the virtual system layers of the orchestration framework (Figure 2). The RL agent also collects the reward information (response time in this case) from the environment to learn an optimal action that maximizes the reward. The agent builds the Q-Table for Q-Learning algorithm, based on cumulative reward obtained from the environment over time. This signifies the efficacy of a specific orchestration decision (action) in achieving the target of minimizing latency, and enables subsequent optimal orchestration decisions.

We demonstrate the efficacy of using the RL agent for orchestrating ML workloads on sensor-edge-cloud platforms. While the focus of this Chapter is optimizing smart eHealth applications, we use image classification task in the experimentation for the purpose of demonstrating multi-user ML workloads. We implement a scenario where a device-edge-cloud architecture serves up to five end-users to execute ML services simultaneously. In this scenario, users sends requests for image classification to an agent located at the cloud. In addition, end-users share their resource availability to the agent. The agent decides to orchestrate the ML tasks based on three static (i.e., device-only, edge-only, and cloud-only) and one RL-optimized strategies. In the device-only strategy, each end-device executes the inference service on a local device. Thus, varying number of users has no effect on the average response time in this case. In the edge and cloud only strategies, simultaneous requests compete for edge and cloud resources. This increases the average response time significantly, as the number of users increase. In the **RL-optimized** strategy, the resource availability is continuously observed and the ML tasks are orchestrated accordingly.
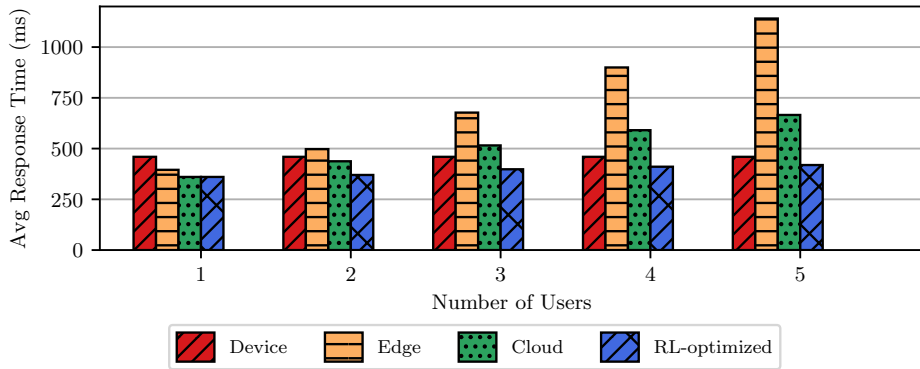

Fig. 11: DL inference orchestration in an end-edge-cloud system which runs an image classification application.

Figure 11 shows the average response time for different numbers of active users for regular network conditions, using different orchestration strategies. The x-axis represents the number of active users. Each bar represents a different orchestration decision made by using the corresponding orchestration

strategy [61, 63]. In RL-optimized approach the average response time remains constant while the number of users is less than three. This is due to the orchestration decision of distributing the services across edge and cloud layers. As the number of users increase to three, the services start competing for resources, leading to an increase in the average response time. With the number of users increasing from three to five, the average response time increases, but at a relatively lower rate, exhibiting efficient utilization of the edge and cloud resources. As the number of users increase, the efficiency of the RL-optimized approach over the static strategies is more prominent.

4 Sense-Compute Co-optimization of ML-driven eHealth Applications

In this section, we describe sense-compute co-optimization approaches for improving resiliency of smart eHealth applications. Common use cases of ML models often handle complete and clean input data, with no specific sensing challenges. However, using ML methods in smart eHealth applications on edge devices requires considerations on challenges from the sensory data acquisition phase [73]. With different types of implanted, wearable, on-body, and remote sensors, there is a higher probability of noise, motion artifacts, and missing input data from sensors [43]. For critical healthcare IoT applications, input data perturbation from motion artifacts, physical failure of sensors, network anomalies and other factors can affect the prediction accuracy of ML models significantly [49]). On the other hand, the sense-making (computation) phase of eHealth applications faces the challenge of limited computational capacity of the edge devices for running ML models [73]. Additionally, the ML models should be resilient to probable sensing anomalies like noisy or missing input data, and maintain higher prediction accuracy even with potential garbage input signals [43]. Moreover, some applications (e.g., pain assessment in clinical healthcare monitoring systems) require near real-time response time, emphasizing the need for ML inference performance [32]. Addressing these multitude of challenges necessitates a co-optimization approach that jointly handles sensing and sense-making phases for system-wide exploration of suitable optimizations. A simple schematic for the interaction between sensing and sense-making modules is shown in Figure 12. Sensing-awareness can be developed by monitoring, and analyzing continuous streams of input data. This intelligence can be used to control the sensing and sense-making configurations simultaneously, by fine-tuning ML models to fit input data characteristics.

4.1 Handling Input Data Perturbations

Sensing Phase Knobs: Monitoring input data originating from the sensory devices for anomalies, discrepancies, and noisy components etc., provides insights into quality of input data [42]. Understanding the quality of input data can be exploited for selective sensing, such that unreliable sensors can be downsampled,
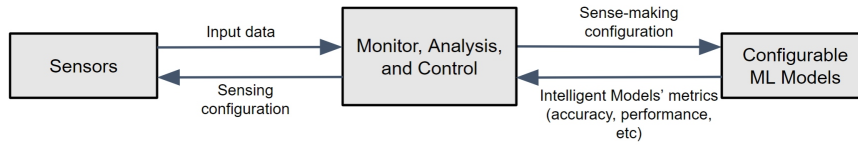
Fig. 12: Sensing and Sense-making (compute) modules' interaction with their action knobs in sensor level and intelligent models.

to dampen the effect of unreliable input data. When the input data from a specific sensor is noisy, the sampling rate of the sensory device can be decreased so that the garbage data would not waste the computational resources in the edge layer. In acute scenarios, input data from a specific device might be completely un-reliable (e.g., when a heartbeat sensor is detached from the human body). In this case, detection of the sensor detachment can guide the computation phase to ignore the input from that specific device, and rely on the data from other input modalities [47]. It should be noted that both selective sensing and disabling a sensor modality minimizes the network latency penalty with reduced input data volume. In specific scenarios of unreliable network connection, sensing phase knobs can be opportunistically triggered to complement the network delay with reduced input data volume [78].

Sense-making Phase Knobs: Addressing input data perturbation from multiple modalities requires appropriate and proportional actions in the ML algorithmic phase. For instance, selective or greedy feature selection from the pre-processed input data can reduce the noisy data components fed into the ML models [39]. This approach is effective in reducing un-necessary computation over noisy input data, although potentially affecting the accuracy of the learning models [50]. Selecting an appropriate ML model from a pool of pre-trained models is another strategy for handling noisy and missing input data. This requires implementation of a model pool, consisting of different ML models that are trained to handle specific combinations of input modalities [50]. Algorithmic optimizations such as meta learning [20], fast reinforcement learning [8], and few-shot learning [75] can also be used for providing efficient ML optimizations particularly for edge devices.

Co-optimization Knobs: Besides the presented methods for optimizing sensing and sense-making phases in the edge devices, it is quite important to use co-optimization techniques so that each of these phases complements the other to form a holistic system with efficiency and robustness. To achieve this goal, meaningful interaction between these phases is required. For example, the sensing information from input data can be used as a trigger to adjust specific sense-making knobs so that the computation models can adapt to the recent changes

in the input sources. One example of this approach is using the early-exit technique [67] [27] in the neural network when the input data has good quality with negligible noise. In this way, an acceptable confidence threshold can be achieved in less time by skipping deeper layers in neural network models. Moreover, the sensing module in the edge layer can send information about input sources with high or low reliability of their data to the sense-making module, and then the machine learning models in the edge layer can adjust importance weights to those input sources by using attention mechanisms inside their architecture [52] [71].

Example Scenarios: We demonstrate the advantages of sense-compute co-optimization through an example of multi-modal pain assessment application [50]. The pain assessment application uses inputs from three modalities – Electrocardiography (ECG), Electrodermal Activity (EDA), and Photoplethysmography (PPG). The ECG, EDA, and PPG modalities have sampling rates of 100, 4, and 64, and generate 52, 42, and 42 features. Figure 13 outlines sensor data acquisition, feature extraction, feature aggregation, and inference task for predicting pain levels under different scenarios. Figure 13 (a) shows the scenario, where application is executed without modality-awareness. In this scenario, data from the ECG sensor is noisy, yet feature vectors from the noisy modality are fed into the ML model, yielding a baseline prediction accuracy of 51%. Figure 13 (b) shows the scenario where modality-awareness is considered while executing the application. In this scenario, the application is executed with selective feature aggregation, by selecting fewer features from the noisy ECG modality. This reduces the total number of features from the ECG modality to 12. An appropriate ML model to suit the updated feature vector is selected from a model pool, which comprises of pre-trained models. Minimizing the features from noisy ECG modality improves the prediction accuracy to 79%, while also reducing the energy consumption and improving the performance, in comparison with the baseline scenario (a). Figure 13 (c) shows the scenario where modality-awareness is used to select specific modalities with quality input data. In this scenario, the noisy ECG modality is completely dropped, and data from the EDA and PPG modalities is processed. Similar to the scenario (b), an appropriate ML model that suits EDA and PPG inputs is selected from the model pool. Dropping an entire modality of data significantly reduces the computational effort and energy consumption, in comparison with scenarios (a) and (b). It should be noted that the prediction accuracy with only two modalities is 74%, which is higher than the baseline from scenario (a), while being marginally lower than the prediction accuracy from selective feature aggregation from scenario (b).

4.2 Sense-Compute Co-optimization Framework

We present the conceptual design of sense-compute co-optimization for multi-modal eHealth applications through the AMSER framework [50]. Figure 14 shows an overview of the AMSER framework for sense-compute co-optimization in multi-layered sensor-edge-cloud platforms. The AMSER framework uses run-
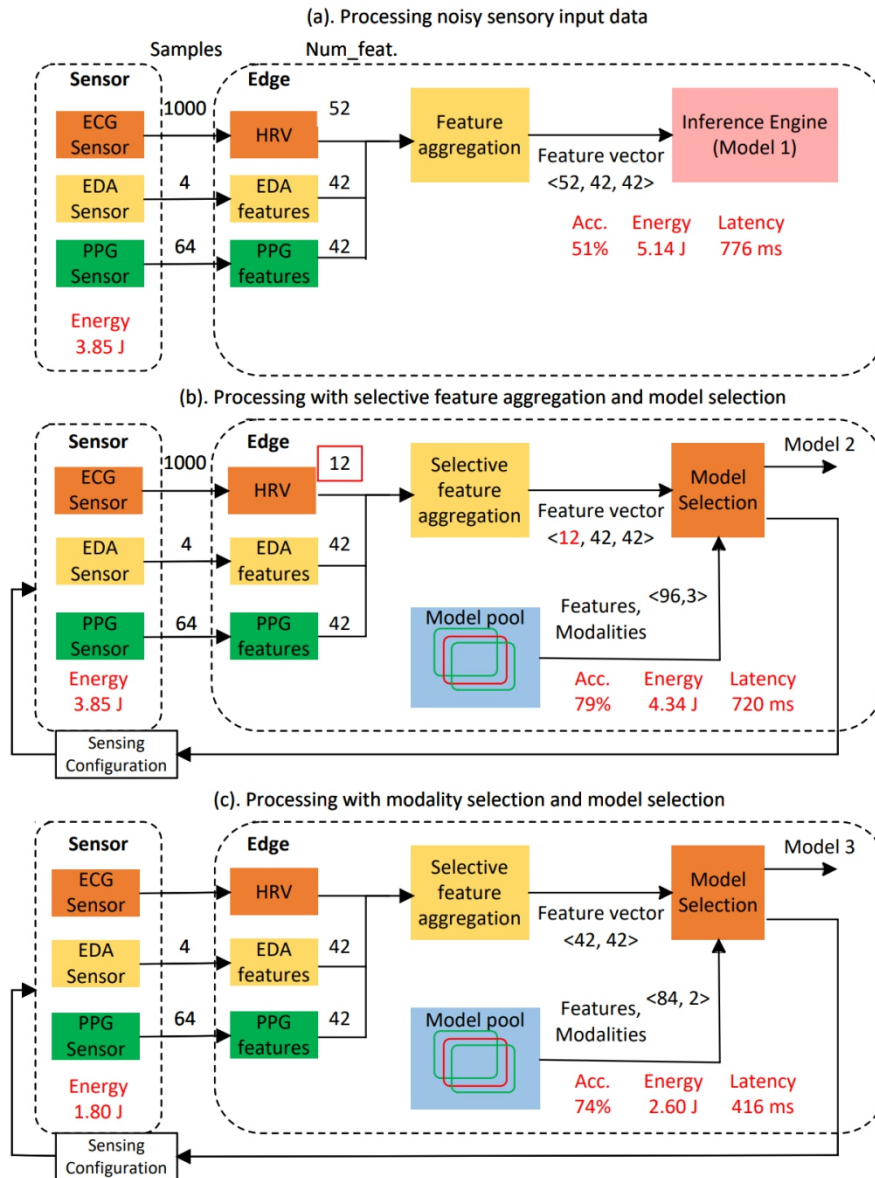
Fig. 13: Motivational scenarios for sense-compute co-optimization

time monitoring functionalities of signal quality monitoring, discrepancy detection, abnormality detection, and confidence level monitor to understand the quality of input data modalities, and confidence of the ML model in predicting results. Insights from the run-time monitoring are used to configure the sam-

pling rates of different sensory devices through the sensing controller. Based on the run-time monitoring, the sense-making (compute) optimizer uses ML configuring knobs – adaptive feature selection, model selection, neural network attention, and early exit mechanisms to configure ML models for current input data sets. For example, data from a specific modality is labeled as uncertain when the signal quality is below a specific Signal to Noise Ratio (SNR). Under such scenarios, the edge level sensing optimizer feeds the learning models with reliable input modalities, while dropping the noisy modality. The sense-making optimizer then selects an ML model from the model pool, that is suitable for the available input data modalities. The model pool contains different pre-trained ML models suitable for different combinations of reliable input modalities.
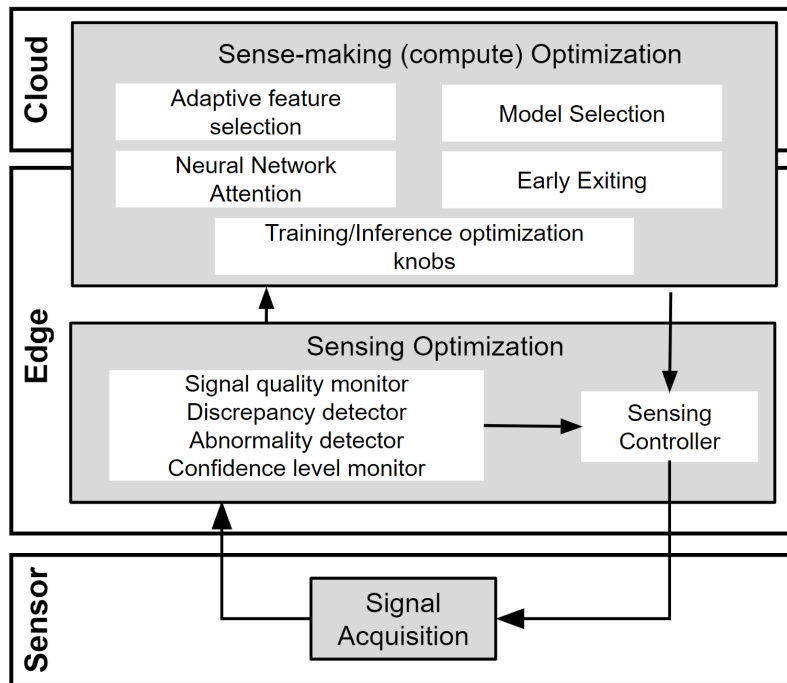

Fig. 14: An overview of sensing and sense-making (compute) optimization in a sensor-edge-cloud architecture

The results for accuracy and performance (speedup) gain with sense-compute co-optimization in comparison with the baseline [28] is shown in Figures 15 and 16. ML-driven eHealth applications of pain assessment and stress monitoring applications were used as input workloads. For a comprehensive evaluation, four different scenarios (S1-S4) with different noise components in input data of modalities are used. Scenario 1 (S1) is the baseline with no noise components.

Scenario 2 (S2), has a wandering noise added to the original data. Scenario 3 has an additional motion artifact added to one input modality on top of the wandering noise, which makes that modality completely unreliable. In scenario 4 (S4), two modalities suffer from severe motion artifact noise. In scenario 2 (S2), the AMSER framework activates the selective feature selection to handle the noisy input modalities, resulting in higher prediction accuracy in comparison with the baseline. In scenario 3 with unreliable input (S3), the AMSER approach drops the entire unreliable modality, yielding a better prediction accuracy and performance gain. Similarly, the AMSER framework drops two noisy modalities in scenario 4 (S4), resulting in better prediction accuracy, and significantly higgher performance. The sensing awareness and synergistic compute knob actuation of the AMSER platform provides significant improvements in accuracy, efficiency, and performance, in comparison with existing dis-joint sensing and compute optimization approaches.

(a) **Pain**   (b) **Stress**
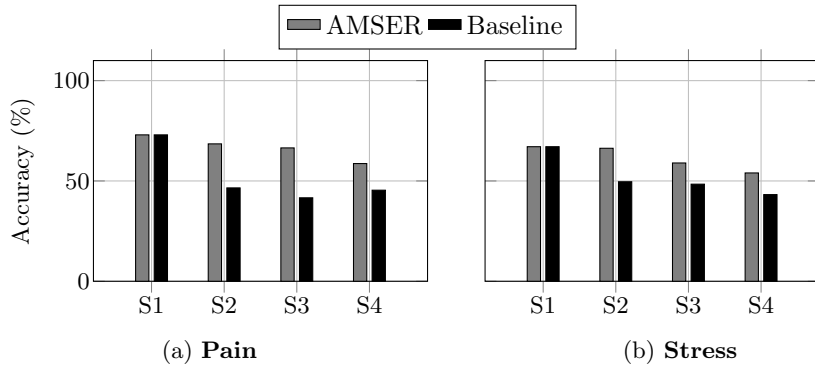
Fig. 15: Accuracy analysis in comparison with baseline [28] [50])

(a) Performance gain   (b) Data volume reduction
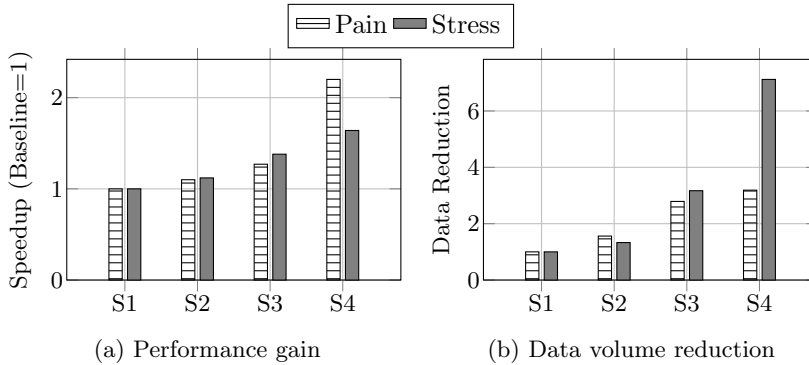
Fig. 16: (a) Performance analysis of the edge-layer device for AMSER framework vs. Baseline study. (b) Data transfer volume between the sensors and edge-layer device for AMSER framework vs. Baseline study.

5 Conclusions

In this chapter, we presented edge-centric optimizations for ML-driven eHealth applications through compute placement, improving the compute placement decisions through reinforcement learning agent, and cross-layered sense-compute co-optimizations. We also presented an exemplar case study of objective pain assessment to demonstrate the use cases of edge-ML based smart eHealth applications, common data flows, computation challenges, and frameworks for deploying smart eHealth applications.

5.1 Key Insights

ML-driven smart healthcare applications have different input data characteristics, computational requirements, and quality metrics. Continuous stream of input data, varying network conditions, and computational requirements of different ML models create dynamic workload scenarios. At an application-level, requirements include higher prediction accuracy of ML models, latency of inferencing results from ML models, resilience, and an overall higher quality of service. At a system-level, requirements include availability of compute nodes in edge and cloud layers, compute capabilities of edge nodes to meet performance requirements of ML models, network utilization, and overall energy efficiency. Considering both application and system-level parameters simultaneously is necessary for optimizing edge-centric ML-driven smart eHealth applications. Optimized compute placement has higher efficacy in meeting both application and system requirements simultaneously. Accuracy-performance trade-offs can also be explored within the compute placement phase, by configuring the choice of ML models. Both model-based and model-free reinforcement learning agents can guide the compute placement decisions on choice of execution node and tuning accuracy-performance trade-offs with high degree of convergence. Further, multi-modal eHealth applications are prone to input data perturbations, which also presents an opportunity to exploit the inherent resilience to selectively process input data. This brings sensing-awareness into computation, and compute-awareness to sensing through bi-directional feedback. Cross-layered sense-compute co-optimization improves sensing, computation, and communication aspects of edge-ML based eHealth applications holistically.

5.2 Open Research Directions

Data Quality Management: Input data quality is an essential component for improving prediction accuracy of ML-driven smart eHealth applications. Processing exclusively quality input data also improves the bandwidth utilization and latency of tasks run on the edge nodes. Some of the techniques presented in this chapter address the sensing aspects through continuous monitoring and analysis of input data quality. Qualitative assessment of sensory data can be improved significantly beyond the rule-based monitors using cognitive learning models. Design of autonomous models for input data quality management

remains an open challenge. Autonomous models enable reasoning for different input perturbations to assess true quality of sensory inputs. Consequently, the garbage data that is un-necessarily processed is minimized, supporting the scalability of edge ML solutions. Quality assessment of input data is significant in other sensor-driven domains such as autonomous driving, robotics, and computer vision etc.

Contextual Edge Orchestration: Orchestration techniques for collaborative sensor-edge-cloud architectures improve a multitude of metrics in terms of performance, turn around time, energy efficiency, accuracy trade-off exploration, and network utilization. Orchestration techniques presented in this chapter enable intelligent compute placement, off-loading, and accuracy configuration decisions through rule-based heuristics. However, contextualization of system dynamics to reason for orchestration decisions, and exploration of accuracy-performance-energy trade-offs with context-awareness is an open research direction. Reinforcement learning has been widely used for optimal orchestration decisions at run-time, considering the varying system dynamics. The efforts in collecting training data, online updating, and convergence time influence the efficacy of such learning methods. In this perspective, design of model-free, and few-shot learning models for run-time edge orchestration is a promising open research direction.

Sense-Compute Co-optimization: Cross-layered sense-compute co-optimization is the most effective strategy for improving sensor dependant, edge-based ML applications. In this chapter, we presented adaptive sensing and sensing-aware computing techniques that uses system-wide monitoring and intelligence for sense-compute co-optimization. This approach focuses on selecting appropriate ML models based on quality of input modalities, exploiting the inherent resilience of multi-modal ML applications. Adaptive feature selection, and ML model selection enforce the idea of sense-compute co-optimization at a coarse-grained level, and require multiple pre-trained models. Incorporating fine-grained edge-layer ML model configurations such as early exit, greedy feature selection, neural network model attention, and saliency maps etc., can complement the model selection strategy. The feasibility of such edge-layer based ML model tuning in collaboration with cloud-layer based model selection is another open research direction.

References

1. Adibuzzaman, M., Ostberg, C., Ahamed, S., Povinelli, R., Sindhu, B., Love, R., Kawsar, F., Ahsan, G.M.T.: Assessment of pain using facial pictures taken with a smartphone. In: 2015 IEEE 39th Annual Computer Software and Applications Conference. vol. 2, pp. 726–731. IEEE (2015)
2. Aqajari, S.A.H., Cao, R., Kasaeyan Naeini, E., Calderon, M.D., Zheng, K., Dutt, N., Liljeberg, P., Salanterä, S., Nelson, A.M., Rahmani, A.M.: Pain assessment tool

with electrodermal activity for postoperative patients: Method validation study. JMIR Mhealth Uhealth 9(5), e25258 (May 2021)

3. Aqajari, S.A.H., Naeini, E.K., Mehrabadi, M.A., Labbaf, S., Rahmani, A.M., Dutt, N.: Gsr analysis for stress: Development and validation of an open source tool for noisy naturalistic gsr data. arXiv preprint arXiv:2005.01834 (2020)

4. Arif-Rahu, M., Grap, M.J.: Facial expression and pain in the critically ill non-communicative patient: state of science review. Intensive and critical care nursing 26(6), 343–352 (2010)

5. Bao, W., Li, W., Delicato, F.C., Pires, P.F., Yuan, D., Zhou, B.B., Zomaya, A.Y.: Cost-effective processing in fog-integrated internet of things ecosystems. In: Proceedings of the 20th ACM International Conference on Modelling, Analysis and Simulation of Wireless and Mobile Systems. pp. 99–108 (2017)

6. Barbera, M.V., Kosta, S., Mei, A., Stefa, J.: To offload or not to offload? the bandwidth and energy costs of mobile cloud computing. In: 2013 Proceedings Ieee Infocom. pp. 1285–1293. IEEE (2013)

7. Barr, J., Fraser, G.L., Puntillo, K., Ely, E.W., Gélinas, C., Dasta, J.F., Davidson, J.E., Devlin, J.W., Kress, J.P., Joffe, A.M., et al.: Clinical practice guidelines for the management of pain, agitation, and delirium in adult patients in the intensive care unit. Critical care medicine 41(1), 263–306 (2013)

8. Barreto, A., Hou, S., Borsa, D., Silver, D., Precup, D.: Fast reinforcement learning with generalized policy updates. Proceedings of the National Academy of Sciences 117(48), 30079–30087 (2020), `https://www.pnas.org/doi/abs/10.1073/pnas.1907370117`

9. Breivik, H., Borchgrevink, P.C., Allen, S.M., Rosseland, L.A., Romundstad, L., Breivik Hals, E., Kvarstein, G., Stubhaug, A.: Assessment of pain. BJA: British Journal of Anaesthesia 101(1), 17–24 (2008)

10. Cao, R., Aqajari, S., Kasaeyan Naeini, E., Rahmani, A.M.: Objective pain assessment using wrist-based ppg signals: A respiratory rate based method. In: 43rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE (2021), accepted for publication

11. Cao, X., Wang, F., Xu, J., Zhang, R., Cui, S.: Joint computation and communication cooperation for mobile edge computing. In: 2018 16th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt). pp. 1–6. IEEE (2018)

12. Chamola, V., Tham, C.K., Chalapathi, G.S.: Latency aware mobile task assignment and load balancing for edge cloudlets. In: 2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops). pp. 587–592. IEEE (2017)

13. Chang, Z., Zhou, Z., Ristaniemi, T., Niu, Z.: Energy efficient optimization for computation offloading in fog computing system. In: GLOBECOM 2017-2017 IEEE Global Communications Conference. pp. 1–6. IEEE (2017)

14. Chatzaki, C., Pediaditis, M., Vavoulas, G., Tsiknakis, M.: Human daily activity and fall recognition using a smartphone's acceleration sensor. In: International Conference on Information and Communication Technologies for Ageing Well and e-Health. pp. 100–118. Springer (2016)

15. Chetty, G., Yamin, M.: Intelligent human activity recognition scheme for ehealth applications. Malaysian Journal of Computer Science 28(1), 59–69 (2015)

16. Dogan, A.Y., Constantin, J., Ruggiero, M., Burg, A., Atienza, D.: Multi-core architecture design for ultra-low-power wearable health monitoring systems. In: 2012 Design, Automation & Test in Europe Conference & Exhibition (DATE). pp. 988–993. IEEE (2012)

17. Duch, L., Basu, S., Braojos, R., Ansaloni, G., Pozzi, L., Atienza, D.: Heal-wear: An ultra-low power heterogeneous system for bio-signal analysis. IEEE Transactions on Circuits and Systems I: Regular Papers 64(9), 2448–2461 (2017)
18. Eshratifar, A.E., Abrishami, M.S., Pedram, M.: Jointdnn: an efficient training and inference engine for intelligent mobile cloud computing services. IEEE Transactions on Mobile Computing (2019)
19. Farahani, B., Barzegari, M., Aliee, F.S., Shaik, K.A.: Towards collaborative intelligent iot ehealth: From device to fog, and cloud. Microprocessors and Microsystems 72, 102938 (2020)
20. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International conference on machine learning. pp. 1126–1135. PMLR (2017)
21. Gia, T.N., Jiang, M., Rahmani, A.M., Westerlund, T., Liljeberg, P., Tenhunen, H.: Fog computing in healthcare internet of things: A case study on ecg feature extraction. In: 2015 IEEE international conference on computer and information technology; ubiquitous computing and communications; dependable, autonomic and secure computing; pervasive intelligence and computing. pp. 356–363. IEEE (2015)
22. Greene, S., Thapliyal, H., Caban-Holt, A.: A survey of affective computing for stress detection: Evaluating technologies in stress detection for better health. IEEE Consumer Electronics Magazine 5(4), 44–56 (2016)
23. Gruss, S., Treister, R., Werner, P., Traue, H.C., Crawcour, S., Andrade, A., Walter, S.: Pain intensity recognition rates via biopotential feature patterns with support vector machines. PloS one 10(10), e0140330 (2015)
24. Gupta, D., Rodrigues, J.J., Peng, S.L., Nguyen, N.: Artificial intelligence for ehealth. Frontiers in Public Health 10 (2022)
25. Han, H.J., Labbaf, S., Borelli, J.L., Dutt, N., Rahmani, A.M.: Objective stress monitoring based on wearable sensors in everyday settings. Journal of Medical Engineering & Technology 44(4), 177–189 (2020)
26. Jiang, M., Mieronkoski, R., Rahmani, A.M., Hagelberg, N., Salanterä, S., Liljeberg, P.: Ultra-short-term analysis of heart rate variability for real-time acute pain monitoring with wearable electronics. In: 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). pp. 1025–1032. IEEE (2017)
27. Ju, W., Bao, W., Ge, L., Yuan, D.: Dynamic Early Exit Scheduling for Deep Neural Network Inference through Contextual Bandits, p. 823–832. Association for Computing Machinery, New York, NY, USA (2021), https://doi.org/10.1145/3459637.3482335
28. Kächele, M., Thiam, P., Amirian, M., Werner, P., Walter, S., Schwenker, F., Palm, G.: Multimodal data fusion for person-independent, continuous estimation of pain intensity. In: Iliadis, L., Jayne, C. (eds.) Engineering Applications of Neural Networks. pp. 275–285. Springer International Publishing, Cham (2015)
29. Kächele, M., Werner, P., Al-Hamadi, A., Palm, G., Walter, S., Schwenker, F.: Biovisual fusion for person-independent recognition of pain intensity. In: International Workshop on Multiple Classifier Systems. pp. 220–230. Springer (2015)
30. Kasaeyan Naeini, E., Jiang, M., Syrjälä, E., Mieronkoski, R., Calderon, M.D., Zheng, K., Dutt, N., Liljeberg, P., Salanterä, S., Nelson, A., Rahmani, A.M.: Research protocol for the smart pain assessment employing behavioral and physiologic indicators. In: JMIR Journal of Research Protocols (revision submitted) (2020)
31. Kasaeyan Naeini, E., Jiang, M., Syrjälä, E., Calderon, M.D., Mieronkoski, R., Zheng, K., Dutt, N., Liljeberg, P., Salanterä, S., Nelson, A.M., Rahmani, A.M.:

Prospective study evaluating a pain assessment tool in a postoperative environment: Protocol for algorithm testing and enhancement. JMIR Res Protoc 9(7), e17783 (Jul 2020)

32. Kasaeyan Naeini, E., Shahhosseini, S., Subramanian, A., Yin, T., Rahmani, A.M., Dutt, N.: An edge-assisted and smart system for real-time pain monitoring. In: 2019 IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE). pp. 47–52 (2019)

33. Kasaeyan Naeini, E., Subramanian, A., Calderon, M.D., Zheng, K., Dutt, N., Liljeberg, P., Salantera, S., Nelson, A.M., Rahmani, A.M.: Pain recognition with electrocardiographic features in postoperative patients: Method validation study. J Med Internet Res 23(5), e25079 (May 2021)

34. Kattepur, A., Dohare, H., Mushunuri, V., Rath, H.K., Simha, A.: Resource constrained offloading in fog computing. In: Proceedings of the 1st Workshop on Middleware for Edge Clouds & Cloudlets. pp. 1–6 (2016)

35. Khan, M.A., Alkaabi, N.: Rebirth of distributed ai—a review of ehealth research. Sensors 21(15), 4999 (2021)

36. Khelifi, H., Luo, S., Nour, B., Sellami, A., Moungla, H., Ahmed, S.H., Guizani, M.: Bringing deep learning at the edge of information-centric internet of things. IEEE Communications Letters 23(1), 52–55 (2018)

37. Koelstra, S., Muhl, C., Soleymani, M., Lee, J.S., Yazdani, A., Ebrahimi, T., Pun, T., Nijholt, A., Patras, I.: Deap: A database for emotion analysis; using physiological signals. IEEE transactions on affective computing 3(1), 18–31 (2011)

38. Kreps, G.L., Neuhauser, L.: New directions in ehealth communication: opportunities and challenges. Patient education and counseling 78(3), 329–336 (2010)

39. Kwak, N., Choi, C.H.: Input feature selection for classification problems. IEEE Transactions on Neural Networks 13(1), 143–159 (2002)

40. Laitala, J., Jiang, M., Syrjälä, E., Naeini, E.K., Airola, A., Rahmani, A.M., Dutt, N.D., Liljeberg, P.: Robust ecg r-peak detection using lstm. In: Proceedings of the 35th Annual ACM Symposium on Applied Computing. pp. 1104–1111 (2020)

41. Liu, J., Mao, Y., Zhang, J., Letaief, K.B.: Delay-optimal computation task scheduling for mobile-edge computing systems. In: 2016 IEEE International Symposium on Information Theory (ISIT). pp. 1451–1455. IEEE (2016)

42. Lou, P., Shi, L., Zhang, X., Xiao, Z., Yan, J.: A data-driven adaptive sampling method based on edge computing. Sensors 20(8) (2020), https://www.mdpi.com/1424-8220/20/8/2174

43. Ma, M., Ren, J., Zhao, L., Tulyakov, S., Wu, C., Peng, X.: Smil: Multimodal learning with severely missing modality. arXiv preprint arXiv:2103.05677 (2021)

44. Mach, P., Becvar, Z.: Mobile edge computing: A survey on architecture and computation offloading. IEEE Communications Surveys & Tutorials 19(3), 1628–1656 (2017)

45. Mao, Y., Zhang, J., Song, S., Letaief, K.B.: Power-delay tradeoff in multi-user mobile-edge computing systems. In: 2016 IEEE Global Communications Conference (GLOBECOM). pp. 1–6. IEEE (2016)

46. Merskey, H.: Pain terms: a list with definitions and notes on usage. recommended by the iasp subcommittee on taxonomy. Pain 6, 249–252 (1979)

47. Montesinos, V., Dell'Agnola, F., Arza, A., Aminifar, A., Atienza, D.: Multi-modal acute stress recognition using off-the-shelf wearable devices. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 2196–2201 (2019)

48. Mousavi, S.S., Schukat, M., Howley, E.: Deep reinforcement learning: an overview. In: Proceedings of SAI Intelligent Systems Conference. pp. 426–440. Springer (2016)
49. Naeini, E.K., Azimi, I., Rahmani, A.M., Liljeberg, P., Dutt, N.: A real-time ppg quality assessment approach for healthcare internet-of-things. Procedia Computer Science 151, 551–558 (2019)
50. Naeini, E.K., Shahhosseini, S., Kanduri, A., Liljeberg, P., Rahmani, A.M., Dutt, N.: Amser: Adaptive multi-modal sensing for energy efficient and resilient ehealth systems. IEEE/ACM Design, Automation and Test in Europe Conference (DATE'22) (2022)
51. Nan, Y., Li, W., Bao, W., Delicato, F.C., Pires, P.F., Zomaya, A.Y.: A dynamic tradeoff data processing framework for delay-sensitive applications in cloud of things systems. Journal of Parallel and Distributed Computing 112, 53–66 (2018)
52. Ning, H., Ye, X., Sada, A.B., Mao, L., Daneshmand, M.: An attention mechanism inspired selective sensing framework for physical-cyber mapping in internet of things. IEEE Internet of Things Journal 6(6), 9531–9544 (2019)
53. Park, J., Samarakoon, S., Bennis, M., Debbah, M.: Wireless network intelligence at the edge. Proceedings of the IEEE 107(11), 2204–2239 (2019)
54. Rahmani, A.M., Gia, T.N., Negash, B., Anzanpour, A., Azimi, I., Jiang, M., Liljeberg, P.: Exploiting smart e-health gateways at the edge of healthcare internet-of-things: A fog computing approach. Future Generation Computer Systems 78, 641–658 (2018)
55. Schapire, R.E.: Explaining adaboost. In: Empirical inference, pp. 37–52. Springer (2013)
56. Schmidhuber, J.: Deep learning in neural networks: An overview. Neural networks 61, 85–117 (2015)
57. Sen, T., Shen, H.: Machine learning based timeliness-guaranteed and energy-efficient task assignment in edge computing systems. In: 2019 IEEE 3rd International Conference on Fog and Edge Computing (ICFEC). pp. 1–10. IEEE (2019)
58. Seo, D., Shahhosseini, S., Mehrabadi, M.A., Donyanavard, B., Lim, S.S., Rahmani, A.M., Dutt, N.: Dynamic ifogsim: A framework for full-stack simulation of dynamic resource management in iot systems. In: 2020 International Conference on Omni-layer Intelligent Systems (COINS). pp. 1–6. IEEE (2020)
59. Shahhosseini, S., Anzanpour, A., Azimi, I., Labbaf, S., Seo, D., Lim, S.S., Liljeberg, P., Dutt, N., Rahmani, A.M.: Exploring computation offloading in iot systems. Information Systems p. 101860 (2021)
60. Shahhosseini, S., Azimi, I., Anzanpour, A., Jantsch, A., Liljeberg, P., Dutt, N., Rahmani, A.M.: Dynamic computation migration at the edge: Is there an optimal choice? In: Proceedings of the 2019 on Great Lakes Symposium on VLSI. pp. 519–524 (2019)
61. Shahhosseini, S., Hu, T., Seo, D., Kanduri, A., Donyanavard, B., Rahmani, A.M., Dutt, N.: Hybrid learning for orchestrating deep learning inference in multi-user edge-cloud networks. arXiv preprint arXiv:2202.11098 (2022)
62. Shahhosseini, S., Kanduri, A., Mehrabadi, M.A., Naeini, E.K., Seo, D., Lim, S.S., Rahmani, A.M., Dutt, N.: Towards smart and efficient health monitoring using edge-enabled situational-awareness. In: 2021 IEEE 3rd International Conference on Artificial Intelligence Circuits and Systems (AICAS). pp. 1–4. IEEE (2021)
63. Shahhosseini, S., Seo, D., Kanduri, A., Hu, T., Lim, S.S., Donyanavard, B., Rahmani, A.M., Dutt, N.: Online learning for orchestration of inference in multi-user end-edge-cloud networks. ACM Transactions on Embedded Computing Systems (TECS) (2022)

64. Sheng, Z., Mahapatra, C., Leung, V.C., Chen, M., Sahu, P.K.: Energy efficient cooperative computing in mobile wireless sensor networks. IEEE Transactions on Cloud Computing 6(1), 114–126 (2015)

65. Stites, M.: Observational pain scales in critically ill adults. Critical care nurse 33(3), 68–78 (2013)

66. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. MIT press (2018)

67. Teerapittayanon, S., McDanel, B., Kung, H.T.: Branchynet: Fast inference via early exiting from deep neural networks. In: 2016 23rd International Conference on Pattern Recognition (ICPR). pp. 2464–2469 (2016)

68. Teerapittayanon, S., McDanel, B., Kung, H.T.: Distributed deep neural networks over the cloud, the edge and end devices. In: 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS). pp. 328–339. IEEE (2017)

69. et al, I.A.: Hich: Hierarchical fog-assisted computing architecture for healthcare iot. ACM Transactions on Embedded Computing Systems 16(5) (2017)

70. Tompkins, D.A., Hobelmann, J.G., Compton, P.: Providing chronic pain management in the "fifth vital sign" era: Historical and treatment perspectives on a modern-day medical dilemma. Drug and Alcohol Dependence 173, S11–S21 (2017), prescription Opioids: new perspectives and research on their role in chronic pain management and addiction

71. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems 30 (2017)

72. Versluis, A., van Luenen, S., Meijer, E., Honkoop, P.J., Pinnock, H., Mohr, D.C., Neves, A.L., Chavannes, N.H., van der Kleij, R.M.: Series: ehealth in primary care. part 4: Addressing the challenges of implementation. European Journal of General Practice 26(1), 140–145 (2020)

73. Wang, X., Han, Y., Leung, V.C.M., Niyato, D., Yan, X., Chen, X.: Convergence of edge computing and deep learning: A comprehensive survey. IEEE Communications Surveys Tutorials 22(2), 869–904 (2020)

74. Wang, X., Han, Y., Leung, V.C., Niyato, D., Yan, X., Chen, X.: Convergence of edge computing and deep learning: A comprehensive survey. IEEE Communications Surveys & Tutorials 22(2), 869–904 (2020)

75. Wang, Y., Yao, Q., Kwok, J.T., Ni, L.M.: Generalizing from a few examples: A survey on few-shot learning. ACM computing surveys (csur) 53(3), 1–34 (2020)

76. Werner, P., Al-Hamadi, A., Limbrecht-Ecklundt, K., Walter, S., Gruss, S., Traue, H.C.: Automatic pain assessment with facial activity descriptors. IEEE Transactions on Affective Computing 8(3), 286–299 (2016)

77. Werner, P., Al-Hamadi, A., Niese, R., Walter, S., Gruss, S., Traue, H.C.: Towards pain monitoring: Facial expression, head pose, a new database, an automatic system and remaining challenges. In: Proceedings of the British Machine Vision Conference. pp. 1–13 (2013)

78. Yick, J., Mukherjee, B., Ghosal, D.: Wireless sensor network survey. Computer networks 52(12), 2292–2330 (2008)

79. You, C., Huang, K.: Exploiting non-causal cpu-state information for energy-efficient mobile cooperative computing. IEEE Transactions on Wireless Communications 17(6), 4104–4117 (2018)

80. Zhang, K., Mao, Y., Leng, S., Zhao, Q., Li, L., Peng, X., Pan, L., Maharjan, S., Zhang, Y.: Energy-efficient offloading for mobile edge computing in 5g heterogeneous networks. IEEE access 4, 5896–5907 (2016)