ISSN: 2395-1680 (ONLINE)
DOI: 10.21917/ijme.2026.0393

ADAPTIVE DEEP REINFORCEMENT LEARNING ENABLED HARDWARE SOFTWARE CO OPTIMIZATION FRAMEWORK FOR ENERGY EFFICIENT FPGA BASED WEARABLE HEALTHCARE

Mariam Safar Mohammed Alshahrani[1], M.K. Jayanthi Kannan[2] and Shree Nee Thirumalai Ramesh[3]

[1]Digital Government Authority (DGA), Digital Government Authority of KSA, Riyadh Province, Kingdom of Saudi Arabia
[2]School of Computing Science Engineering and Artificial Intelligence), VIT Bhopal, India
[3]Department of Medicine, Manipal University Medical College Malaysia, Malaysia

Abstract

Wearable healthcare systems have emerged as an essential component of modern medical monitoring by enabling continuous acquisition of physiological signals and real-time health assessment. Field Programmable Gate Arrays (FPGAs) have become attractive computing platforms for wearable devices because they offer high computational capability, reconfigurability, and lower energy consumption than conventional processors. However, achieving an optimal balance among hardware acceleration, software execution, latency, and power consumption remains a challenging task under dynamic healthcare workloads. Existing hardware-software partitioning approaches primarily employ static optimization strategies that cannot effectively adapt to changing workloads, patient-specific data streams, and resource availability. Consequently, these methods suffer from increased energy consumption, reduced resource utilization, and degraded real-time performance. This research proposes a novel Deep Reinforcement Learning based Adaptive Hardware Software Co Optimization Framework (DRL-AHSCOF) for FPGA-enabled wearable healthcare systems. The proposed framework integrates a Deep Q-Network (DQN) based adaptive decision engine with dynamic hardware-software partitioning, runtime FPGA resource allocation, workload-aware scheduling, and energy-aware optimization. The reinforcement learning agent continuously observes workload characteristics, hardware utilization, memory occupancy, communication overhead, and latency to determine optimal execution policies that maximize overall system efficiency while minimizing energy consumption. Experimental evaluation demonstrates that the proposed DRL-AHSCOF framework significantly improves the performance of FPGA-based wearable healthcare systems. The proposed method reduces energy consumption to 92.74 J, achieving a 26.67% reduction compared with GA-HSO. It decreases execution latency to 30.12 ms, improving response time by 40.66% compared with existing optimization approaches. Furthermore, DRL-AHSCOF achieves 95.34% FPGA resource utilization, 1846 tasks/s throughput, and 99.18% scheduling accuracy, outperforming GA-HSO, PSO-RS, and DL-AS methods. These results confirm that the proposed deep reinforcement learning-based optimization strategy effectively enhances energy efficiency, adaptive resource management, and real-time processing capability for FPGA-enabled wearable healthcare systems.

Keywords:

Deep Reinforcement Learning, FPGA Wearable Healthcare, Hardware Software Co-Optimization, Energy Efficiency, Dynamic Resource Allocation

1. INTRODUCTION

Recent advances in wearable healthcare technologies have transformed the delivery of personalized medical services by enabling continuous monitoring of physiological parameters such as electrocardiograms, blood oxygen saturation, heart rate, blood pressure, and body temperature [1]. These wearable devices facilitate early disease diagnosis, remote patient monitoring, and emergency healthcare management while reducing hospitalization costs and improving patient quality of life. The increasing adoption of artificial intelligence and Internet of Medical Things (IoMT) technologies has further enhanced the capability of wearable healthcare systems to provide intelligent decision support and predictive analytics [2]. FPGAs have become promising hardware platforms because they provide configurable parallel architectures, high processing performance, low latency, and superior energy efficiency compared with conventional embedded processors, making them particularly suitable for real-time healthcare applications [3].

Despite these advantages, wearable healthcare systems continue to encounter several technical challenges. The computational requirements of modern deep learning algorithms continue to increase, placing significant pressure on the limited energy budgets of battery-powered wearable devices. Simultaneously, wearable healthcare applications require strict real-time execution, minimal processing latency, and reliable performance under continuously varying physiological data streams [4]. Static hardware-software partitioning techniques often allocate computational tasks before system deployment without considering runtime workload variations. Such approaches result in inefficient FPGA resource utilization, excessive communication overhead, increased execution latency, and unnecessary energy consumption. Furthermore, heterogeneous computing architectures containing processors, FPGA accelerators, and memory subsystems require intelligent coordination to maximize overall system performance while preserving battery lifetime [5].

Another significant limitation arises from the inability of existing optimization frameworks to dynamically adapt to changing healthcare workloads. Wearable devices continuously experience fluctuations in sensing frequency, patient activity, environmental conditions, and computational complexity. Conventional optimization algorithms including heuristic methods, evolutionary optimization, and rule-based scheduling generally require predefined optimization objectives and cannot efficiently learn optimal execution strategies from runtime observations. Consequently, these approaches frequently produce suboptimal scheduling decisions, leading to increased hardware idle time, software bottlenecks, resource contention, and degraded Quality of Service (QoS). These limitations restrict the deployment of intelligent wearable healthcare systems in practical large-scale medical environments [6].

To overcome these challenges, this research develops a novel Deep Reinforcement Learning based Adaptive Hardware

Software Co Optimization Framework (DRL-AHSCOF). Unlike traditional static optimization approaches, the proposed framework continuously learns optimal execution policies through interactions with the computing environment. A Deep Q-Network serves as the intelligent decision-making engine that dynamically determines hardware-software task allocation according to workload characteristics, FPGA resource utilization, processing latency, memory occupancy, communication cost, and energy consumption. Runtime adaptation enables continuous optimization of heterogeneous computing resources without requiring manual intervention or predefined scheduling rules.

The primary objectives of this research are to develop an adaptive hardware-software co-optimization framework capable of minimizing energy consumption while maintaining real-time processing performance for wearable healthcare applications. The framework seeks to optimize FPGA resource allocation, reduce execution latency, improve computational throughput, maximize hardware utilization, and enhance scheduling accuracy using reinforcement learning. Furthermore, the proposed system aims to achieve scalable runtime adaptation suitable for diverse healthcare monitoring scenarios involving heterogeneous sensor workloads.

The novelty of this research lies in integrating deep reinforcement learning with adaptive hardware-software co-optimization specifically designed for FPGA-enabled wearable healthcare systems. Unlike conventional optimization algorithms that rely on static heuristics or predefined optimization models, the proposed DRL-AHSCOF continuously updates its scheduling policy according to environmental feedback. The framework jointly optimizes workload partitioning, runtime resource allocation, memory management, and energy-aware scheduling using a unified reinforcement learning architecture. This adaptive capability enables efficient management of dynamic healthcare workloads while significantly improving energy efficiency and computational performance.

The proposed research makes two major contributions. A novel Deep Reinforcement Learning based Adaptive Hardware Software Co Optimization Framework (DRL-AHSCOF) is proposed to perform intelligent runtime hardware-software partitioning, adaptive FPGA resource allocation, and workload-aware scheduling for wearable healthcare systems. An integrated energy-aware optimization strategy is developed that simultaneously minimizes power consumption, execution latency, and communication overhead while maximizing FPGA utilization, throughput, and scheduling accuracy through continuous reinforcement learning, thereby improving the efficiency and scalability of next-generation wearable healthcare platforms.

2. RELATED WORKS

Wearable healthcare systems have attracted significant research attention due to their capability to provide continuous physiological monitoring with low latency and high reliability. Researchers have increasingly explored FPGA-based architectures to accelerate healthcare signal processing while reducing computational complexity and energy consumption. Recent developments have focused on integrating artificial intelligence, heterogeneous computing, and adaptive optimization techniques to improve wearable device performance [7].

Several studies have investigated hardware-software co-design methodologies for FPGA-based embedded healthcare systems. These approaches partition computational tasks between software processors and FPGA hardware accelerators according to execution complexity and resource availability. Although hardware acceleration significantly improves computational performance, most existing partitioning techniques employ static optimization strategies established during design time. Consequently, runtime workload variations cannot be efficiently accommodated, leading to increased resource underutilization and reduced energy efficiency [8].

Energy-aware scheduling has also become an important research direction for wearable healthcare platforms. Numerous optimization algorithms including Genetic Algorithms, Particle Swarm Optimization, Ant Colony Optimization, and Simulated Annealing have been introduced to minimize power consumption while maintaining acceptable computational performance. These metaheuristic algorithms demonstrate satisfactory optimization capability for offline scheduling problems. However, their convergence time, parameter sensitivity, and limited adaptability reduce their effectiveness in highly dynamic wearable healthcare environments where computational workloads continuously fluctuate [9].

Machine learning techniques have recently been incorporated into hardware optimization frameworks to improve adaptive scheduling decisions. Supervised learning algorithms including Support Vector Machines, Random Forests, Artificial Neural Networks, and Gradient Boosting have been utilized to predict workload characteristics and estimate resource requirements. Although prediction accuracy improves scheduling efficiency, supervised learning models depend heavily on labeled training datasets and generally fail to adapt to previously unseen runtime conditions. Their offline learning paradigm restricts continuous optimization during real-time healthcare monitoring [10].

Deep learning models have further enhanced intelligent optimization by learning complex relationships among workload characteristics, system resources, and execution performance. Convolutional Neural Networks, Long Short-Term Memory networks, and Transformer-based architectures have demonstrated remarkable success in healthcare data analysis and workload prediction. Nevertheless, these approaches primarily focus on predictive modeling rather than sequential decision making. Their inability to continuously optimize runtime scheduling decisions limits their applicability for dynamic FPGA resource management [11].

Reinforcement learning has recently emerged as an effective optimization strategy for heterogeneous computing environments. Deep Reinforcement Learning enables autonomous decision making by continuously interacting with the environment and maximizing long-term cumulative rewards. Several studies have applied Deep Q-Networks, Policy Gradient methods, and Actor-Critic architectures for cloud computing, edge computing, and resource scheduling applications. These methods demonstrate superior adaptability compared with conventional optimization algorithms because they continuously update scheduling policies according to environmental feedback. However, their application to FPGA-enabled wearable healthcare systems remains relatively limited [12].

Recent investigations have explored dynamic partial reconfiguration capabilities of FPGAs to improve runtime flexibility. Dynamic reconfiguration enables hardware modules to be modified during execution without interrupting overall system operation, thereby improving resource utilization and reducing hardware redundancy. Although dynamic reconfiguration enhances flexibility, existing frameworks generally rely on predefined scheduling heuristics instead of intelligent adaptive optimization, limiting their effectiveness under continuously changing healthcare workloads [13].

Several researchers have combined edge computing with wearable healthcare systems to reduce communication latency and improve data privacy. Edge-based architectures distribute computational tasks among wearable devices, edge servers, and cloud platforms according to resource availability. Intelligent workload distribution significantly reduces response time while improving system scalability. Nevertheless, existing workload allocation strategies typically overlook FPGA-specific hardware optimization and runtime hardware-software coordination, resulting in suboptimal system performance [14].

3. PROPOSED METHOD

The proposed DRL-AHSCOF is designed to achieve energy-efficient execution of wearable healthcare applications on FPGA-enabled heterogeneous computing platforms. The framework intelligently partitions computational tasks between software processors and FPGA hardware accelerators while continuously adapting to dynamic workload variations through DRL. Unlike conventional static partitioning approaches, the proposed framework continuously monitors system states, evaluates resource utilization, predicts execution efficiency, and determines optimal scheduling actions that minimize energy consumption while maximizing throughput and computational performance. The reinforcement learning agent learns the optimal execution policy by interacting with the runtime environment and updating its decision model according to observed rewards. Consequently, the framework simultaneously optimizes execution latency, hardware utilization, communication overhead, and power consumption under continuously changing healthcare workloads.

The proposed framework consists of two major operational sections:

- Adaptive Runtime State Analysis and Deep Reinforcement Learning Decision Engine
- Dynamic Hardware–Software Co-Optimization and Energy-Aware FPGA Resource Scheduling

These two sections operate iteratively throughout system execution to ensure continuous optimization of wearable healthcare applications.

3.1 HEALTHCARE WORKLOAD ACQUISITION AND RUNTIME STATE REPRESENTATION

The first stage begins with continuous acquisition of physiological data generated from wearable healthcare sensors. Multiple biomedical sensors produce heterogeneous streams including electrocardiogram (ECG), blood oxygen saturation (SpO_2), blood pressure, respiratory rate, body temperature, and motion signals. Since these physiological streams exhibit different sampling frequencies and computational complexities, the framework first transforms incoming workloads into a unified runtime state representation. This representation serves as the observation space of the reinforcement learning environment. The runtime state contains processor utilization, FPGA logic occupancy, available memory resources, communication bandwidth, queue length, estimated execution latency, workload priority, and remaining battery energy. Every arriving healthcare task is characterized according to computational complexity, deadline constraints, required hardware resources, and communication requirements. The complete runtime system state at decision instant t is represented as $S_t=[W_t,U_{CPU,t},U_{FPGA,t},M_t,B_t,L_t,E_t,Q_t]$, where, W_t denotes workload complexity, $U_{CPU,t}$ represents processor utilization, $U_{FPGA,t}$ denotes FPGA utilization, M_t is available memory, B_t is communication bandwidth, L_t is execution latency, E_t denotes remaining battery energy, and Q_t represents waiting task queue. The normalized system state is obtained by

$$\mathbf{S}_t=\left[\frac{W_t}{W_{\max}},\frac{U_{CPU,t}}{100},\frac{U_{FPGA,t}}{100},\frac{M_t}{M_{\max}},\frac{B_t}{B_{\max}},\frac{L_t}{L_{\max}},\frac{E_t}{E_{\max}},\frac{Q_t}{Q_{\max}}\right] \quad (1)$$

The parameter is normalized into the interval [0,1], enabling stable reinforcement learning convergence. The normalized state vector provides a comprehensive description of current hardware conditions and computational requirements. Since wearable healthcare workloads vary continuously according to patient activities and sensing frequencies, this adaptive representation allows the learning agent to observe instantaneous operating conditions before making scheduling decisions.

3.2 DEEP REINFORCEMENT LEARNING-BASED ADAPTIVE DECISION ENGINE

The normalized runtime state is forwarded to the Deep Reinforcement Learning decision engine. A Deep Q-Network (DQN) estimates the expected cumulative reward associated with every possible scheduling action. The objective is to identify the hardware-software execution policy that minimizes energy consumption while satisfying latency constraints. The state-action value function is expressed as

$$Q^{\pi}(s,a)=\mathrm{E}\left[\sum_{k=0}^{\infty}\gamma^k R_{t+k}\,\middle|\,S_t=s,A_t=a\right] \quad (2)$$

The DQN approximates the optimal value function through $Q(s,a;\theta)=f\left(S_t;\theta\right)$. The immediate reward simultaneously considers energy, latency, utilization, and throughput.

$$R_t=\alpha\left(\frac{T_t}{T_{\max}}\right)+\beta\left(\frac{U_{FPGA}}{100}\right)-\delta\left(\frac{E_t}{E_{\max}}\right)-\eta\left(\frac{L_t}{L_{\max}}\right) \quad (3)$$

The neural network parameters are updated using

$$\theta_{\text{new}}=\theta_{\text{old}}-\lambda\nabla_{\theta}\begin{pmatrix}R_t+\gamma\max_{a'}\\ Q\left(S_{t+1},a';\theta^{-}\right)-Q\left(S_t,A_t;\theta\right)\end{pmatrix}^2 \quad (4)$$

Through continuous interaction, the agent gradually discovers the optimal scheduling policy without requiring predefined optimization rules.

3.3 ADAPTIVE HARDWARE–SOFTWARE TASK PARTITIONING

Following the optimal action predicted by the reinforcement learning agent, computational tasks are dynamically assigned either to software processors or FPGA hardware accelerators. Computationally intensive signal-processing operations such as filtering, feature extraction, and neural inference are mapped onto FPGA hardware, whereas lightweight control operations remain on embedded processors. The task assignment objective minimizes total execution cost

$$C_{\text{Total}} = \sum_{i=1}^{N} \left[x_i C_{HW,i} + (1-x_i) C_{SW,i} \right] \tag{5}$$

where x_i=1 indicates FPGA execution, and x_i=0 software execution. Execution cost consists of

$$C_{HW} = \omega_1 E_{HW} + \omega_2 L_{HW} + \omega_3 M_{HW} \tag{6}$$

where, E_{HW} is the hardware energy, L_{HW} is latency, and M_{HW} memory usage. Similarly, $C_{SW} = \omega_4 E_{SW} + \omega_5 L_{SW} + \omega_6 M_{SW}$. The framework compares both execution costs before assigning tasks dynamically. Hardware assignment continuously changes according to available FPGA resources and runtime workload intensity. This adaptive strategy prevents processor congestion while maximizing hardware parallelism.

3.4 ENERGY-AWARE FPGA RESOURCE SCHEDULING AND RUNTIME OPTIMIZATION

The final stage performs runtime FPGA scheduling based on the hardware-software partition generated previously. The scheduler allocates configurable logic blocks, DSP slices, memory blocks, and routing resources while simultaneously minimizing energy consumption. Overall energy consumption is represented by

$$E_{\text{Total}} = E_{\text{Dynamic}} + E_{\text{Static}} + E_{\text{Communication}} + E_{\text{Memory}}$$

The optimization objective becomes

$$\min J = \alpha E_{\text{Total}} + \beta L_{\text{Total}} + \delta C_{\text{Comm}} - \mu U_{FPGA} - \nu T$$

where L_{Total} is the execution latency, C_{Comm} is the communication cost, U_{FPGA} is the FPGA utilization, $T_{throughput}$, $\alpha,\beta,\delta,\mu,v$ optimization coefficients. Resource utilization is computed as:

$$U_{FPGA} = \frac{CLB_{\text{used}} + DSP_{\text{used}} + BRAM_{\text{used}}}{CLB_{\text{total}} + DSP_{\text{total}} + BRAM_{\text{total}}} \times 100$$

The throughput is $T = \frac{N_{\text{Completed}}}{\Delta t}$, where, $N_{\text{Completed}}$ completed healthcare tasks, and Δt execution interval. The scheduler repeatedly evaluates runtime utilization, execution latency, communication overhead, and remaining battery energy before reallocating FPGA resources. Whenever workload characteristics change, the updated runtime state is fed back into the Deep Reinforcement Learning agent, creating a closed-loop optimization cycle. Continuous feedback enables the framework to maintain high processing efficiency even under highly dynamic healthcare monitoring scenarios. The proposed DRL-AHSCOF therefore establishes an intelligent adaptive optimization architecture in which reinforcement learning continuously improves hardware-software partitioning and FPGA resource scheduling. By jointly optimizing energy consumption, latency, hardware utilization, throughput, and runtime adaptability, the framework provides an effective solution for next-generation energy-efficient FPGA-based wearable healthcare systems.

4. RESULTS AND DISCUSSION

The proposed DRL-AHSCOF framework is evaluated using a simulation-based experimental environment developed for FPGA-enabled wearable healthcare systems. The framework is implemented using Python 3.10 with TensorFlow 2.12 for Deep Reinforcement Learning model development and Xilinx Vivado Design Suite 2023.1 for FPGA resource analysis and hardware performance evaluation. The experiments are executed on a computer system equipped with an Intel Core i7 processor, 32 GB RAM, and NVIDIA RTX 3060 GPU.

4.1 EXPERIMENTAL SETUP AND SIMULATION PARAMETERS

The experimental parameters used for evaluating the proposed DRL-AHSCOF framework are presented in Table.1. The parameters represent wearable healthcare workload characteristics, FPGA hardware configuration, reinforcement learning settings, and simulation constraints. These configurations ensure a realistic evaluation environment for adaptive hardware-software co-optimization.

Table.1. Experimental and Simulation Parameters

Parameter	Value
Simulation Platform	Xilinx Vivado 2023.1
Programming Framework	TensorFlow 2.12
Processor	Intel Core i7-12700
Main Memory	32 GB
GPU Accelerator	NVIDIA RTX 3060
FPGA Device	Xilinx Zynq UltraScale+MPSoC
FPGA Logic Resources	274K LUTs
DSP Units	2520
BRAM Capacity	38 Mb
Number of Healthcare Tasks	10,000
Training Episodes	500
Learning Rate	0.001
Discount Factor	0.95
Exploration Rate	1.0
Minimum Exploration Rate	0.01
Batch Size	64
Simulation Duration	1000 s

4.2 PERFORMANCE EVALUATION METRICS

The performance of the proposed DRL-AHSCOF framework is evaluated using five important metrics: energy consumption, execution latency, FPGA resource utilization, throughput, and scheduling accuracy.

4.2.1 *Energy Consumption:*

Energy consumption represents the total power required for executing healthcare workloads on the FPGA-based platform. Lower energy consumption indicates improved battery efficiency and extended operational lifetime of wearable devices. The metric is calculated by considering dynamic switching power, static leakage power, memory access energy, and communication overhead.

$$E_{\text{total}} = P_{\text{dynamic}} \times T_{\text{execution}} + P_{\text{static}} \times T_{\text{execution}} + E_{\text{memory}} + E_{\text{communication}}$$

4.2.2 *Execution Latency:*

Execution latency measures the time required to complete healthcare processing tasks from input acquisition to final output generation. Lower latency is essential for real-time medical monitoring applications where rapid response is required.

$$L_{\text{total}} = T_{\text{processing}} + T_{\text{communication}} + T_{\text{scheduling}} + T_{\text{memory}}$$

4.2.3 *FPGA Resource Utilization:*

FPGA resource utilization evaluates how effectively available hardware resources are employed during healthcare workload execution. Higher utilization indicates efficient hardware allocation and reduced resource wastage.

$$U_{FPGA} = \frac{R_{\text{used}}}{R_{\text{available}}} \times 100$$

4.2.4 *Throughput:*

Throughput represents the number of healthcare processing tasks completed within a specific time interval. Higher throughput demonstrates improved computational efficiency and system scalability.

$$\text{Throughput} = \frac{N_{\text{completed}}}{T_{\text{simulation}}}$$

4.2.5 *Scheduling Accuracy:*

Scheduling accuracy indicates the ability of the DRL agent to select the optimal hardware-software execution strategy. It measures the percentage of correctly optimized task allocation decisions compared with the ideal scheduling solution.

$$\text{Accuracy} = \frac{N_{\text{optimal decisions}}}{N_{\text{total decisions}}} \times 100$$

4.3 DATASET DESCRIPTION

The proposed framework is evaluated using a synthesized wearable healthcare workload dataset generated from multiple physiological monitoring applications.

Table.2. Dataset Description for FPGA Wearable Healthcare Evaluation

Dataset Parameter	Value
Dataset Type	Synthetic IoMT Benchmark Dataset
Total Samples	50,000
Healthcare Applications	5
Features per Sample	18
Data Sources	ECG, SpO_2, Heart Rate, Temperature, Motion
Task Categories	Low, Medium, High Complexity
Sampling Frequency	250 Hz
FPGA Resource Features	LUT, DSP, BRAM Utilization
Energy Features	Dynamic and Static Energy
Latency Features	Processing and Communication Delay
Dataset Split	80% / 20%

4.4 EXISTING METHODS FOR COMPARISON

The existing optimization approaches are selected from previous research studies for comparison with the proposed DRL-AHSCOF framework. These methods represent conventional optimization, machine learning-based optimization, and adaptive scheduling strategies used in FPGA and wearable healthcare systems. The selected approaches include Genetic Algorithm based Hardware Software Optimization (GA-HSO), Particle Swarm Optimization based Resource Scheduling (PSO-RS), and Deep Learning based Adaptive Scheduling (DL-AS).

4.5 ENERGY CONSUMPTION REDUCTION COMPARISON OVER TRAINING EPOCHS

Energy consumption is a critical factor in FPGA-based wearable healthcare systems because these devices operate with limited battery capacity. The proposed DRL-AHSCOF framework continuously optimizes task allocation and FPGA resource scheduling to reduce unnecessary power consumption. The Table.3 presents the energy consumption comparison among GA-HSO, PSO-RS, DL-AS, and the proposed DRL-AHSCOF framework over different training epochs.

Table.3. Energy Consumption Comparison over Training Epochs

Training Epochs	GA-HSO (J)	PSO-RS (J)	DL-AS (J)	Proposed DRL-AHSCOF (J)
5	184.62	178.45	171.83	158.72
10	176.54	169.82	162.47	147.38
15	168.73	161.24	153.92	138.56
20	161.45	152.86	145.38	130.14
25	154.62	145.73	137.84	122.36
30	148.25	138.94	130.62	115.47
35	142.18	132.56	123.91	108.82
40	136.54	126.37	117.46	102.65
45	131.12	120.83	111.72	97.18
50	126.46	115.62	106.34	92.74

The Table.3 demonstrates that the proposed DRL-AHSCOF framework achieves the lowest energy consumption throughout the complete evaluation period. At the initial 5th epoch, the proposed framework consumes 158.72 J, whereas GA-HSO, PSO-RS, and DL-AS consume 184.62 J, 178.45 J, and 171.83 J, respectively. This indicates an immediate improvement of 14.03%, 11.05%, and 7.63% compared with the existing approaches. As training progresses, the reinforcement learning agent continuously improves its decision policy through environmental feedback, resulting in further energy reduction. At the 50th epoch, DRL-AHSCOF reduces energy consumption to 92.74 J, while GA-HSO, PSO-RS, and DL-AS require 126.46 J, 115.62 J, and 106.34 J, respectively. The proposed method achieves reductions of 26.67%, 19.79%, and 12.80% compared with GA-HSO, PSO-RS, and DL-AS.

4.6 EXECUTION LATENCY REDUCTION COMPARISON OVER TRAINING EPOCHS

Execution latency directly affects the responsiveness of wearable healthcare systems, especially for emergency monitoring applications. The proposed DRL-AHSCOF framework minimizes latency by dynamically assigning computationally intensive tasks to FPGA accelerators and reducing communication delays. The latency comparison results are presented in Table.4.

Table.4. Execution Latency Comparison Over Training Epochs

Training Epochs	GA-HSO (ms)	PSO-RS (ms)	DL-AS (ms)	Proposed DRL-AHSCOF (ms)
5	82.64	78.92	74.56	67.43
10	78.35	73.84	69.72	61.58
15	73.92	69.25	64.83	56.47
20	69.86	64.72	60.14	51.36
25	66.24	60.53	55.76	46.82
30	62.85	56.84	51.62	42.73
35	59.47	53.26	47.83	39.12
40	56.38	50.14	44.57	35.84
45	53.42	47.36	41.72	32.91
50	50.76	44.83	38.96	30.12

The latency performance presented in Table.4 confirms the effectiveness of DRL-AHSCOF in achieving faster healthcare task execution. At the 5th epoch, the proposed method records a latency of 67.43 ms, which is lower than GA-HSO (82.64 ms), PSO-RS (78.92 ms), and DL-AS (74.56 ms). The reduction is achieved because the reinforcement learning agent identifies computationally intensive operations and transfers them to FPGA accelerators with higher parallel processing capability.

At the 50th epoch, DRL-AHSCOF achieves 30.12 ms latency, while GA-HSO, PSO-RS, and DL-AS obtain 50.76 ms, 44.83 ms, and 38.96 ms, respectively. The proposed framework improves latency by 40.66%, 32.81%, and 22.69% compared with the three existing approaches. The improvement results from adaptive scheduling decisions, reduced task migration overhead, and optimized FPGA resource utilization.

4.7 RESOURCE UTILIZATION OVER TRAINING EPOCHS

Efficient utilization of FPGA resources is essential for maximizing computational performance while avoiding hardware underutilization. The Table.5 presents FPGA resource utilization achieved by different optimization approaches.

Table.5. FPGA Resource Utilization Comparison Over Training Epochs

Training Epochs	GA-HSO (%)	PSO-RS (%)	DL-AS (%)	Proposed DRL-AHSCOF (%)
5	61.42	64.18	67.36	72.84
10	63.56	66.74	69.82	76.25
15	65.83	69.15	72.46	79.68
20	68.24	71.62	74.83	82.47
25	70.56	74.38	77.21	84.96
30	72.84	76.94	79.63	87.32
35	75.26	79.35	81.92	89.48
40	77.63	81.72	84.26	91.24
45	79.85	83.96	86.43	93.16
50	82.14	86.25	88.72	95.34

As shown in Table.5, the proposed DRL-AHSCOF framework achieves superior FPGA utilization compared with existing methods. At the 5th epoch, the proposed method achieves 72.84% utilization, outperforming GA-HSO, PSO-RS, and DL-AS by 11.42%, 8.66%, and 5.48%, respectively. The improvement indicates that the reinforcement learning model effectively identifies suitable hardware execution opportunities. At the 50th epoch, the proposed framework reaches 95.34% FPGA utilization, whereas GA-HSO, PSO-RS, and DL-AS achieve 82.14%, 86.25%, and 88.72%, respectively. The proposed method provides utilization improvements of 13.20%, 9.09%, and 6.62%. The improvement is attributed to adaptive resource allocation and continuous optimization of FPGA configuration according to workload requirements.

4.8 THROUGHPUT COMPARISON OVER TRAINING EPOCHS

Throughput represents the processing capability of the FPGA-based wearable healthcare system and indicates the number of healthcare tasks completed within a given time period. Higher throughput is essential for handling continuous physiological data streams generated by multiple wearable sensors. The proposed DRL-AHSCOF framework improves throughput by intelligently distributing computational workloads between software processors and FPGA accelerators. The throughput comparison among GA-HSO, PSO-RS, DL-AS, and DRL-AHSCOF is presented in Table.6.

Table.6. Throughput Comparison Over Training Epochs

Training Epochs	GA-HSO (Tasks/s)	PSO-RS (Tasks/s)	DL-AS (Tasks/s)	Proposed DRL-AHSCOF (Tasks/s)
5	842	876	914	982
10	895	934	976	1086
15	948	996	1042	1174
20	1004	1058	1116	1268
25	1057	1124	1193	1365
30	1112	1196	1274	1458
35	1168	1265	1356	1552

40	1224	1338	1438	1647
45	1281	1412	1523	1742
50	1338	1487	1608	1846

The Table.6 demonstrates that the proposed DRL-AHSCOF framework consistently achieves higher throughput compared with existing optimization methods. At the initial 5th epoch, the proposed framework achieves 982 tasks/s, whereas GA-HSO, PSO-RS, and DL-AS provide 842 tasks/s, 876 tasks/s, and 914 tasks/s, respectively. The proposed method improves throughput by 16.63%, 12.11%, and 7.44% compared with these approaches. This improvement is obtained through adaptive workload distribution and efficient utilization of FPGA parallel processing capabilities. As the training progresses, the reinforcement learning model improves its scheduling decisions by learning from previous execution experiences. At the 50th epoch, DRL-AHSCOF reaches 1846 tasks/s, while GA-HSO, PSO-RS, and DL-AS achieve 1338 tasks/s, 1487 tasks/s, and 1608 tasks/s, respectively. The proposed framework improves throughput by 27.50%, 24.01%, and 14.80% compared with the existing approaches. The increased throughput results from optimized task mapping, reduced processing bottlenecks, and effective exploitation of FPGA acceleration resources. The DQN-based scheduling mechanism enables continuous adaptation according to workload complexity, thereby maintaining high computational efficiency under dynamic healthcare scenarios.

4.9 SCHEDULING ACCURACY COMPARISON OVER TRAINING EPOCHS

Scheduling accuracy evaluates the capability of the optimization framework to select appropriate hardware-software execution strategies for incoming healthcare workloads. Higher scheduling accuracy indicates that the system can efficiently identify optimal resource allocation decisions while satisfying energy and latency constraints. The Table.7 presents the scheduling accuracy comparison between existing methods and the proposed DRL-AHSCOF framework.

Table.7. Scheduling Accuracy Over Training Epochs

Training Epochs	GA-HSO (%)	PSO-RS (%)	DL-AS (%)	Proposed DRL-AHSCOF (%)
5	82.36	84.25	86.74	90.82
10	84.58	86.72	89.36	93.15
15	86.73	88.96	91.54	94.86
20	88.42	90.84	93.25	96.18
25	90.16	92.73	94.82	97.12
30	91.74	94.18	95.96	97.86
35	93.28	95.42	96.85	98.24
40	94.62	96.58	97.42	98.56
45	95.76	97.36	98.06	98.78
50	96.84	98.12	98.42	99.18

The Table.7 confirms that the proposed DRL-AHSCOF framework provides superior scheduling accuracy throughout the evaluation process. At the 5th epoch, the proposed method achieves 90.82% accuracy, whereas GA-HSO, PSO-RS, and DL-AS achieve 82.36%, 84.25%, and 86.74%, respectively. This demonstrates that the reinforcement learning agent quickly identifies effective execution policies even during the early training stage. At the 50th epoch, the proposed framework achieves a maximum scheduling accuracy of 99.18%, exceeding GA-HSO (96.84%), PSO-RS (98.12%), and DL-AS (98.42%). The accuracy improvement reaches 2.34%, 1.06%, and 0.76%, respectively. Although existing deep learning approaches achieve reasonable prediction performance, they lack continuous adaptation during runtime. The proposed DRL-based strategy improves decision accuracy by continuously updating the policy model according to observed system feedback. This adaptive learning capability enables accurate hardware-software allocation decisions for diverse wearable healthcare workloads.

The comprehensive evaluation demonstrates that the proposed DRL-AHSCOF framework significantly improves the performance of FPGA-based wearable healthcare systems compared with GA-HSO, PSO-RS, and DL-AS approaches. The energy consumption analysis presented in Table.3 shows that the proposed framework reduces energy usage from 126.46 J to 92.74 J at the 50th epoch, achieving a reduction of 26.67% compared with GA-HSO. Similarly, the execution latency results in Table.4 indicate that DRL-AHSCOF reduces processing delay to 30.12 ms, providing improvements of 40.66%, 32.81%, and 22.69% compared with GA-HSO, PSO-RS, and DL-AS. The FPGA utilization results in Table.5 demonstrate that the proposed framework achieves 95.34% resource utilization, which is significantly higher than GA-HSO (82.14%), PSO-RS (86.25%), and DL-AS (88.72%). Improved resource utilization contributes directly to enhanced throughput performance. As shown in Table.6, the proposed method achieves 1846 tasks/s, exceeding DL-AS by 14.80% and GA-HSO by 27.50%. Furthermore, the scheduling accuracy results in Table.7 confirm that DRL-AHSCOF reaches 99.18% accuracy, demonstrating effective decision-making capability.

5. CONCLUSION

This research presents a DRL-AHSCOF for improving the energy efficiency and computational performance of FPGA-enabled wearable healthcare systems. The proposed framework addresses the limitations of conventional static optimization approaches by introducing an adaptive decision-making mechanism capable of continuously learning and optimizing hardware-software execution strategies. Experimental evaluation demonstrates significant improvements across all performance metrics. The proposed framework reduces energy consumption to 92.74 J, achieving a 26.67% reduction compared with GA-HSO. It decreases execution latency to 30.12 ms, providing a 40.66% improvement compared with the conventional optimization method. The framework also achieves 95.34% FPGA resource utilization, 1846 tasks/s throughput, and 99.18% scheduling accuracy. These improvements demonstrate the effectiveness of reinforcement learning-based adaptive optimization for dynamic wearable healthcare workloads. The proposed DRL-AHSCOF framework successfully balances energy efficiency, computational performance, and resource utilization through intelligent runtime optimization. The integration of deep reinforcement learning enables autonomous adaptation to changing healthcare workloads without requiring predefined scheduling rules. Therefore, the proposed framework provides a practical solution for developing next-generation low-power, high-performance, and intelligent FPGA-based wearable healthcare systems.

REFERENCES

[1] M.I. Khan and B. Da Silva, "Harnessing FPGA Technology for Energy-Efficient Wearable Medical Devices", *Electronics*, Vol. 13, No. 20, pp. 4094-4106, 2024.

[2] M. Vaithianathan, M. Patil, S.F. Ng and S. Udkar, "Energy-Efficient FPGA Design for Wearable and Implantable Devices", *ESP International Journal of Advancements in Science and Technology*, Vol. 2, No. 2, pp. 37-51, 2024.

[3] F. Porreca, F. Frustaci and R. Gravina, "A Codesign Framework for the Development of Next Generation Wearable Computing Systems", *Sensors*, Vol. 25, No. 21, pp. 6624-6637, 2025.

[4] M.I. Frutuoso, H.C. Neto and R.P. Duarte, "Energy-Efficient and Real-Time Wearable for Wellbeing-Monitoring IoT System based on SoC-FPGA", *Algorithms*, Vol. 16, No. 3, pp. 141-156, 2023.

[5] D. Wu, S. Zhao, J. Yang and M. Sawan, "Software-Hardware Co-Design for Energy-Efficient Continuous Health Monitoring via Task-Aware Compression", *IEEE Transactions on Biomedical Circuits and Systems*, Vol. 17, No. 2, pp. 180-191, 2023.

[6] S. Lakshmikanth, C. Chethan, R. Chethan Kumar and S.H. Jeevith, "Design and Optimization of Energy Efficient FPGA Based Architecture for Wearable Medical Devices", *Proceedings of International Conference on Knowledge Engineering and Communication Systems*, pp. 1-7, 2026.

[7] S. Satpathy, M.C. Trivedi, S. Malhotra, S. Chaudhary and V. Sharma, "FPGA-Enhanced IoT Methods for Disease Pre-Screening and Prediction: An Energy Optimization Approach", *Measurement: Sensors*, Vol. 33, pp. 101213-101226, 2024.

[8] M.R. Akshayraj, V.P. Gopi, G. Lakshminarayanan, G.R. Gangadharan and J.U. Kidav, "Energy-Efficient Hardware Design for CNN-based ECG Signal Classification in Wearable Bio-Medical Devices", *Proceedings of International Symposium on VLSI Design and Test*, pp. 1-7, 2024.

[9] O. Hassan, T. Paul, M.M.H. Shuvo, D. Parvin and S.K. Islam, "Energy Efficient Deep Learning Inference Embedded on FPGA for Sleep Apnea Detection", *Journal of Signal Processing Systems*, Vol. 94, No. 6, pp. 609-619, 2022.

[10] T. Sutikno, A.Z. Jidin and L. Handayani, "Energy-Efficient Reconfigurable Architectures for Edge AI in Healthcare IoT: Trends Challenges and Future Directions", *International Journal Reconfigurable Embedded System*, Vol. 15, pp. 1-20, 2026.

[11] A. Yadav, V.K. Gupta and B. Kumar, "FPGA-based Medical Image Processing using Hardware-Software Co-design Approach", *IEEE Transactions on Biomedical Circuits and Systems*, Vol. 35, No. 1, pp. 1-24, 2025.

[12] N. Akhtar, A.R. Buzdar and M.U. Khan, "Cardio-Edge: Hardware-Software Co-design Implementation of LSTM Based ECG Classification for Continuous Cardiac Monitoring on Wearable Devices", *International Journal of Advanced Computer Science and Applications*, Vol. 16, No. 7, pp. 1-14, 2025.

[13] N.B. Gaikwad, S.K. Khare, D. Mendhe, H. Mir, S. Kosta and U.R. Acharya, "FPGA SoC Implementation of Adaptive Deep Neural Network based Multimodal Edge Intelligence for Internet of Medical Things", *IEEE Access*, Vol. 14, pp. 1-23, 2025.

[14] Y. Tuncel, A. Krishnakumar, A.L. Chithra, Y. Kim and U. Ogras, "A Domain-Specific System-on-Chip Design for Energy Efficient Wearable Edge AI Applications", *Proceedings of ACM/IEEE International Symposium on Low Power Electronics and Design*, pp. 1-6, 2022.