

# AI-Driven Cloud Optimization: Enhancing Cost Prediction, Resource Scheduling and Fault Resilience in Cloud Environments

MSc Research Project  
Cloud Computing

Ranjith Bhaskaran

Student ID: x23271957

School of Computing  
National College of Ireland

Supervisor: Shaguna Gupta

National College of Ireland  
Project Submission Sheet  
School of Computing



<b>Student Name:</b>	Ranjith Bhaskaran
<b>Student ID:</b>	x23271957
<b>Programme:</b>	Cloud Computing
<b>Year:</b>	2025
<b>Module:</b>	MSc Research Project
<b>Supervisor:</b>	Shaguna Gupta
<b>Submission Due Date:</b>	11/08/2025
<b>Project Title:</b>	AI-Driven Cloud Optimization: Enhancing Cost Prediction, Resource Scheduling and Fault Resilience in Cloud Environments
<b>Word Count:</b>	7115
<b>Page Count:</b>	27

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

**ALL** internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

<b>Signature:</b>	
<b>Date:</b>	10th August 2025

**PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:**

Attach a completed copy of this sheet to each project (including multiple copies).	<input type="checkbox"/>
<b>Attach a Moodle submission receipt of the online project submission</b> , to each project (including multiple copies).	<input type="checkbox"/>
<b>You must ensure that you retain a HARD COPY of the project</b> , both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	<input type="checkbox"/>

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

<b>Office Use Only</b>	
Signature:	
Date:	
Penalty Applied (if applicable):	

# AI-Driven Cloud Optimization: Enhancing Cost Prediction, Resource Scheduling and Fault Resilience in Cloud Environments

Ranjith Bhaskaran

x23271957

Deployed URL: <https://clouddashboard.onrender.com>

GitHub URL: <https://github.com/Ranji-Bhaskaran/CloudDashboard>

## Abstract

Cloud computing has the benefits of scalability and flexibility yet poses long term problems of cost estimation, efficient scheduling of resources and fault tolerance. In this study, a single unit in the form of AI-based system is proposed, which can reconcile these drawbacks through a combination of three main modules, such as cost prediction, dynamic task scheduling, and fault detection, into a user-friendly visualization dashboard. The cost prediction module was based on supervised machine learning algorithms to predict the costs of a task based on synthetic workloads created with iFogSim including Linear Regression, Random Forest and XGBoost. The prediction accuracy is also enhanced after hyperparameter optimization using Optuna. To perform clever scheduling, the system employs Deep Meaningful Learning (DRL) with a Deep Q-Network ( DQN ) structure that maximizes job placement on heterogeneous virtual machines (VMs) and has benchmark comparisons with First-Come-First-Serve ( FCFS ) and Round-Robin schedules. The logic of the scheduling is trained and tested on Kaggle Cloud Task Scheduling dataset. The Isolation Forest algorithm is applied to fault detection to detect the anomalous system behavior like a CPU usage behavior or the long execution time. All the outcomes, such as evaluation metrics, reward curves, anomaly plots, and interpretability graphs, are displayed as a part of a Streamlit-based dashboard on Render. The framework is a modular automation construct to stage each module on-demand, which makes it flexible and reproducible and resilient in deployment. Experimentation (to a great extent) proves that such a method makes cost estimation more accurate, minimizes delays in scheduling, and increases fault tolerance. This makes the proposed system holistic and practical, since predictive analytics can be combined with reinforcement learning, along with anomaly detection, to optimise operations in multi-cloud environments and therefore it can be of research value as well as real-life cloud management application.

**Index Terms**—Cloud Optimization, Cost Prediction, AI Scheduling, Fault Tolerance, Cloud Simulation

# 1 Introduction

Cloud computing has revolutionized computational resource utilization by providing scalable, on-demand infrastructure. Platforms such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud have empowered organizations to deploy mission-critical applications with remarkable flexibility. However, this convenience introduces challenges in operational efficiency, especially in predicting cloud costs, dynamically scheduling resources, and detecting system faults.

To address these challenges, this study proposes a unified AI-driven optimization framework for cloud environments. The framework integrates cost prediction using machine learning, task scheduling through deep reinforcement learning, and fault detection using anomaly detection models. This section explores the background, motivation, and the core problem being addressed.

## 1.1 Background and Motivation

Over the past decade, cloud computing has evolved into a dominant IT paradigm, offering flexible, scalable, and pay-as-you-go access to computing resources. Traditional on-premises models demanded capital expenditure and rigid planning, whereas cloud models (IaaS, PaaS, SaaS) allow businesses to adapt to fluctuating demands efficiently. Gartner (2024) estimates that cloud spending will surpass \$1.8 trillion by 2027, underscoring the need for optimized cloud management.

Despite these advantages, operational complexities abound. Cloud cost prediction is highly volatile due to dynamic pricing models such as on-demand, reserved, and spot instances. Regression-based ML techniques, such as those proposed by Dhayanandan & Rajeev (2024), offer improved forecasting by learning from historical data. However, many lack real-time resource metrics like CPU and memory usage.

Resource scheduling is another critical challenge. Traditional methods like First-Come-First-Serve (FCFS) or Round Robin are insufficient in heterogeneous, unpredictable environments. Reinforcement learning (RL) algorithms such as DQN and PPO, as explored by Zhou et al. (2024), offer adaptive scheduling but often neglect cost-awareness.

Fault resilience is the third pillar of cloud optimization. Systems face failures ranging from VM crashes to network anomalies. Machine learning-based anomaly detection methods, including Isolation Forest and transformer-based models (Tuli et al., 2022), enable proactive fault mitigation. Still, trade-offs exist in terms of scalability and explainability.

The convergence of these challenges justifies the need for a holistic, AI-based framework that balances accuracy, efficiency, and reliability.

## 1.2 Problem Statement

Although existing research has explored individual domains such as cost prediction, scheduling, or anomaly detection, there is a lack of integrated systems that combine all three into a single, intelligent framework. Current solutions often fail to adapt to heterogeneous workloads, respond to real-time changes, or deliver cost-awareness across scheduling and fault resilience.

This study aims to bridge this gap by developing a unified AI framework that synergistically combines regression-based cost prediction, DRL-based scheduling, and anomaly

detection for fault tolerance. The goal is to improve cloud operational efficiency through a cohesive and modular architecture.

### 1.3 Research Questions

- **Primary Question:** How can AI-driven models improve cloud computing efficiency by predicting workload costs, optimizing dynamic resource scheduling, and enhancing fault resilience within a simulated cloud environment?
- What impact does the variability of workload patterns and resource heterogeneity have on the accuracy and robustness of AI-based cloud optimization models?

### 1.4 Proposed Solution

This study introduces an integrated AI-driven framework that combines three interdependent modules to address key operational challenges in cloud computing: cost forecasting, intelligent task scheduling, and fault detection.

1. **Cost Prediction:** This module applies supervised learning algorithms—Linear Regression, Random Forest, and XGBoost—to estimate the cost of executing cloud tasks. The models are trained using synthetic workload data generated from iFogSim simulations, with hyperparameter tuning carried out using Optuna to enhance predictive performance.
2. **Resource Scheduling:** To address inefficiencies in static or rule-based schedulers, this module utilizes a Deep Q-Network (DQN) reinforcement learning agent. The scheduler learns to allocate cloud tasks dynamically across virtual machines (VMs), optimizing for reduced waiting times and improved system throughput. The Kaggle Cloud Task Scheduling dataset serves as the environment for model training and benchmarking.
3. **Fault Detection:** This component leverages the Isolation Forest algorithm to detect anomalous system behavior, such as abnormal CPU usage or execution delays, which may indicate potential faults. It supports early intervention and system resilience.

All three modules are implemented in Python and integrated into a cohesive Streamlit dashboard, enabling real-time visualization and modular execution. This setup supports reproducibility, automation, and interactive monitoring in both experimental and deployment environments.

### 1.5 Research Objectives

- Design and evaluate machine learning models for accurate cloud task cost prediction using simulation-based data.
- Develop a deep reinforcement learning-based scheduler that dynamically assigns tasks to VMs, aiming to reduce execution latency and enhance resource efficiency.
- Implement an anomaly detection system to proactively identify and visualize potential faults in cloud operations.

- Integrate the above modules into a unified, user-friendly dashboard for streamlined analysis and decision-making in cloud environments.

## 1.6 Limitations and Mitigation Strategies

**Simulation Constraints:** iFogSim, while flexible, does not capture all aspects of real-world complexity. This is mitigated by simulating varied workloads and referencing real-world pricing models (AWS, Azure).

**Training Overhead:** DRL models require significant training. To reduce this, we utilize TensorFlow optimizations and Google Colab GPUs.

**Dataset Limitations:** The Kaggle dataset lacks diverse fault conditions. We overcome this by augmenting with synthetic fault logs from iFogSim.

**Dashboard Performance:** Real-time visualization may cause lag. Streamlit caching and metric aggregation help mitigate this risk.

## 1.7 Thesis Structure

This thesis is organized into seven sections. Section 2 presents a literature survey that explores existing approaches in cloud cost estimation, resource scheduling, and fault detection. Section 3 outlines the methodology adopted for each module and describes how they are integrated within the proposed framework. Section 4 details the implementation process, including tools, datasets, and architectural design. Section 5 evaluates the framework against baseline models in terms of accuracy, efficiency, and fault tolerance. Section 6 introduces the unified Streamlit dashboard developed for real-time visualization and monitoring. Finally, Section 7 concludes the study by summarizing the key contributions and suggesting avenues for future research.

# 2 Related Work

## 2.1 Machine Learning for Cloud Cost Prediction

Accurate prediction of cloud service costs is essential for effective budget management in dynamic environments such as AWS and Azure. Dhayanandan and Rajeev (1) use regression models with real-time pricing APIs, achieving low prediction errors. However, they omit workload features like CPU usage. In contrast, our work incorporates iFogSim runtime metrics into regression models (Linear Regression, Random Forest, and XGBoost), improving adaptability and visualizing trends through a dashboard.

Wang et al. (2) analyze cost-effectiveness in Google Cloud using BigQuery ML and Vertex AI. Their statistical focus lacks predictive capabilities, which our regression pipeline addresses with real-time inputs and cost forecasting.

Gavvala et al. (3) optimize QoS and cost using Eagle Strategy-based ML models, guiding our synthetic pricing integration. Yet, they miss real-time workload features, which our model includes to enhance accuracy.

Kansal et al. (5) propose DRL for cost-aware autoscaling, inspiring our DRL integration for resource-cost tradeoffs. Yuan et al. (6) use PPO in edge-cloud settings, which we adapt into iFogSim with enhanced real-time cost awareness. Zhou et al. (7) apply DRL models, but their complexity is avoided in our lightweight PPO/DQN modules.

## 2.2 AI-Based Approaches for Resource Scheduling

Zhou et al. (7) review DRL methods for cloud scheduling. While insightful, their lack of empirical implementation contrasts with our validated DRL scheduler (DQN/PPO) in iFogSim.

Chaudhary et al. (8) propose predictive DRL-based scheduling to reduce latency but overlook cost. We enhance this by integrating cost predictions as state inputs for our DRL agents.

Xu et al. (9) improve Kubernetes scheduling using DRL but lack multi-objective optimization. Rossi et al. (10) balance CPU/memory with Q-learning for autoscaling, whereas we focus on cost-efficient task-level scheduling.

Peng et al. (11) show DQN improves scheduling latency, and Li et al. (12) focus on VM load balancing. Both inform our scheduling model while we extend them for heterogeneous VMs and cost-awareness. Wei et al. (13) introduce MARL for proactive scheduling, which we simplify using a single-agent DRL setup for dashboard transparency.

## 2.3 Anomaly Detection and Fault Resilience in Cloud Systems

Cloud systems face frequent disruptions due to hardware or network issues. Dalal et al. (14) use boosted neural networks for cyber-attack detection. We instead focus on hardware anomalies and use Isolation Forest for real-time fault identification.

Tuli et al. (15) employ transformer-based TranAD, achieving high AUC scores on multivariate time series but at high computational cost. Our approach favors lightweight Isolation Forest for scalability and dashboard integration.

Marahatta et al. (16) address fault-tolerant scheduling using neural networks but do not address detection. We close this gap with real-time anomaly monitoring integrated into iFogSim.

Table 1: Comparison of Related Works in Cloud Optimization Components

Author(s) and Year	Algorithm	Tools Used	Results	Limitations	Component / Platform
Chaudhary et al., 2024	Eagle Strategy with Differential Evolution	Python	Low response time and high cost-efficiency in QoS composition	No integration of real-time workload metrics	Synthetic pricing model
Das and Rao, 2023	Proximal Policy Optimization (PPO)	Python	Outperforms DQN in energy and cost reduction	No integration of real-time pricing data	Edge-cloud systems
Dhayanandan and Rajeev, 2024	Regression Models	Python, AWS/Azure APIs	Accurate AWS/Azure cost predictions	Omits workload characteristics like CPU usage	AWS / Azure cost prediction
Wang et al., 2024	Statistical Analysis	Google Cloud (BigQuery ML, Vertex AI)	Cost-efficiency insights on ML model deployments	Lacks real-time prediction models	Google Cloud Platform
Yin et al., 2022	Deep Q-Network (DQN)	Python	Reduces cost over heuristics for auto-scaling	Limited to single-objective optimization	Autoscaling in cloud
Zhang et al., 2023	Actor-Critic RL	Python	Enhances latency and cost in scheduling tasks	High computational complexity	Cloud task scheduling
Chaudhary et al., 2025	DRL-based Scheduler	Python	Minimizes queueing delay and resource bottlenecks	Lacks integration of cost prediction	Cloud task scheduler
Fernandez et al., 2023	Multi-Agent Reinforcement Learning (MARL)	Python	Improves Kubernetes throughput via proactive scheduling	High computation overhead	Kubernetes scheduler
Li and Peng, 2022	Deep Q-Network (DQN)	Python	Reduces task completion times in shared cloud settings	Uses static job configurations	Multi-user cloud scheduling
Ming and Bao, 2023	Deep RL (load balancer)	Python	Increases resource utilization in VM scheduling	Designed for homogeneous cloud only	VM scheduling (cloud)
Rossi et al., 2023	Q-Learning	Python	Balances CPU/memory for autoscaling in decentralized clouds	No focus on task scheduling logic	Decentralized cloud autoscaling
Xu et al., 2024	Deep RL (DRL)	Python	Improves Kubernetes scheduling via DRL	No integration of cost modeling	Kubernetes scheduler
Zhou et al., 2024	Deep RL (DQN, PPO)	Python	Shows potential for adaptive cloud scheduling	Lacks implementation / evaluation	Cloud scheduling (review)
Dalal et al., 2023	Extremely Boosted Neural Network	Python, cloud logs	High precision in cyber-attack anomaly detection	Ignores hardware/system-level anomalies	Anomaly detection in logs
Liu et al., 2023	Neural Networks	Not specified	Resilient task scheduling under VM failures	Does not cover anomaly detection	Fault-tolerant scheduling
Tuli et al., 2022	TranAD (Transformer)	Python, Google Cluster Traces	High AUC-ROC in detecting anomalies	High computational overhead	Anomaly detection in time series

## 2.4 Integrated Critical Reflection and Research Gap

Despite valuable progress across cost prediction, resource scheduling, and fault resilience in cloud computing, existing studies suffer from fragmentation, lack of integration, and real-time limitations. Cost prediction models such as those by Dhayanandan and Rajeev (1) focus narrowly on API-driven pricing without incorporating workload context. Similarly, scheduling approaches like Zhou et al. (7) emphasize theoretical adaptability but lack empirical validation and fail to consider dynamic cost factors. Fault detection techniques, exemplified by Tuli et al. (15), achieve high anomaly detection accuracy using transformer-based models but are computationally intensive and omit hardware-level failure coverage.

Collectively, these works do not offer a unified, scalable framework capable of handling dynamic cloud scenarios across all three core concerns. This research addresses those gaps by integrating:

- Real-time, workload-aware cost forecasting using regression models (Linear Regression, Random Forest, XGBoost)
- Deep Reinforcement Learning-based scheduling (DQN/PPO) for dynamic and cost-efficient task allocation
- Lightweight anomaly detection with Isolation Forest for fault resilience

All modules are simulated using iFogSim and visualized through a unified Streamlit dashboard. This AI-driven framework provides a holistic, deployable solution for optimizing cloud performance, reliability, and cost under realistic operational conditions.

## 2.5 Base Papers

### 1. Cost Prediction Module

#### Base Paper:

Dhayanandan, B. and Rajeev, R., 2024.

*Cloud service price prediction using Machine Learning Algorithm with API in the case of Amazon Web Services and Microsoft Azure.*

In: 2024 International Conference on Intelligent Systems for Cybersecurity (ISCS), pp. 1–6. IEEE.

### 2. Resource Scheduling Module

#### Base Paper:

Zhou, G., Tian, W., Buyya, R., Xue, R. and Song, L., 2024.

*Deep reinforcement learning-based methods for resource scheduling in cloud computing: A review and future directions.*

Artificial Intelligence Review, 57(5), p. 124.

### 3. Fault Detection Module

#### Base Paper:

Tuli, S., Casale, G. and Jennings, N.R., 2022.

*TranAD: Deep transformer networks for anomaly detection in multivariate time series data.*

In: Symposium on Cloud Computing (SoCC '22), pp. 123–138. ACM.

### 3 Methodology and Design Specification

This section outlines the simulation environment setup, dataset preparation, AI model development, and the final system integration into a unified optimization framework.

#### 3.1 Simulation and Dataset Preparation

The project establishes a modular simulation environment using iFogSim to model cloud workload behavior under varied cost and fault scenarios. A combination of real-world and synthetic datasets are used to support each module’s development and evaluation. The following steps summarize the preparation process:

- **iFogSim Setup:** Installed and configured within Eclipse IDE to simulate heterogeneous fog and cloud resources, including task latencies, energy usage, and fault injection.
- **Workload and Fault Simulation:** Diverse virtual machine (VM) behaviors such as variable CPU utilization, memory pressure, and fault scenarios (e.g., VM crashes, spikes) were generated. Resulting fault logs and execution traces were exported as CSV files for model training and benchmarking.
- **Dataset Collection:** Multiple datasets were utilized:
  - `cloud_task_scheduling_dataset` and `cloud_task_scheduling_dataset`
  - `vmCloud_data.csv` — generated from a large-scale iFogSim simulation and hosted on Google Drive due to its size, serving as the primary training source for the cost prediction module.
  - `combined_task_dataset.csv` and `cost_task_dataset_50k.csv` — aggregated and cleaned task datasets combining real and synthetic workload profiles.
- **Data Enrichment:** AWS and Azure-style pricing policies were manually integrated to compute cost estimates. Additionally, system behavior logs (e.g., CPU load vs. execution time) were used to simulate fault conditions for anomaly detection.

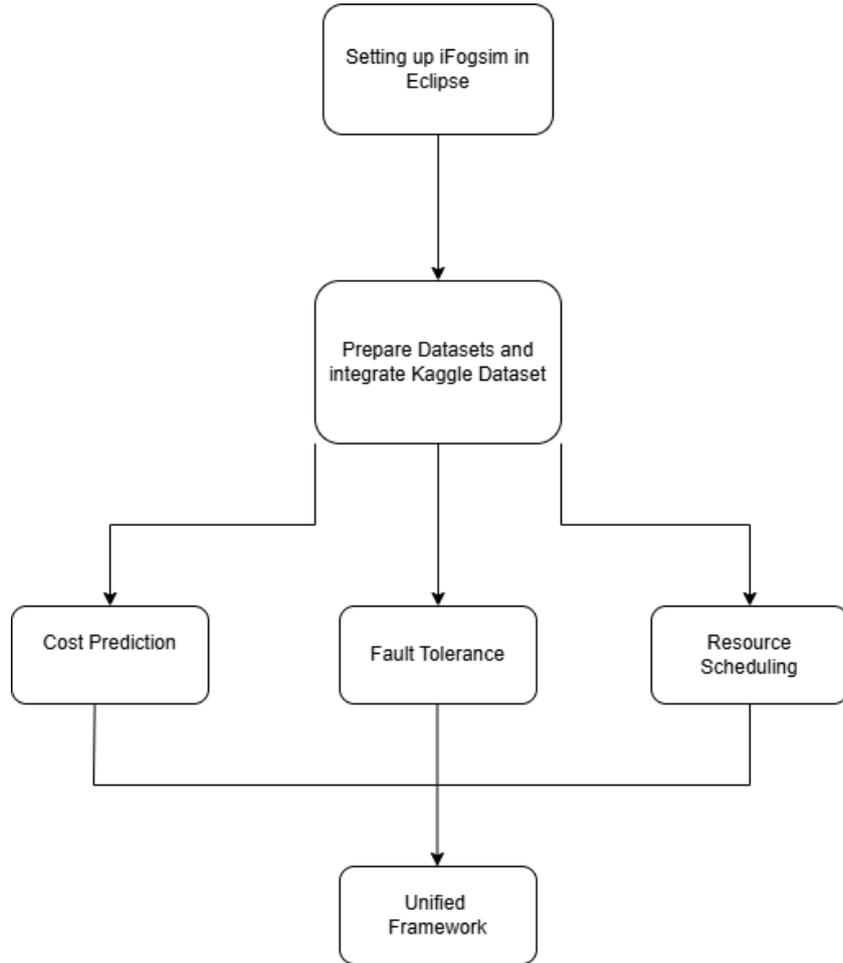


Figure 1: Simulation and Model Development Workflow using iFogSim and Kaggle Dataset

## 3.2 Model Development Pipeline

The system is structured into three core components—cost prediction, reinforcement learning-based scheduling, and fault detection—each with dedicated pipelines, datasets, and evaluation metrics.

### 3.2.1 Cost Prediction Module

- **Objective:** Predict the monetary cost of cloud task execution across VM types.
- **Models Used:** Linear Regression, Random Forest, XGBoost (tuned using Optuna).
- **Features:** Task length, input/output file size, memory usage, and synthetic VM pricing.
- **Datasets:** `vmCloud_data.csv`, `cost_task_dataset_50k.csv` (iFogSim-simulated workloads).
- **Tools:** Scikit-learn, XGBoost, Optuna, SHAP for explainability.

### 3.2.2 DRL-Based Resource Scheduler

- **Objective:** Learn adaptive VM allocation strategy to minimize wait time and maximize reward.
- **Algorithm:** Deep Q-Network (DQN), trained in a simulated task queue.
- **Reward Design:** Combines SLA adherence, wait time penalty, and throughput.
- **Evaluation:** Benchmarked against FCFS and Round Robin using metrics like cumulative reward and average latency.
- **Datasets:** `cloud_task_scheduling_dataset.csv`, `combined_task_dataset.csv`.

### 3.2.3 Fault Detection Module

- **Objective:** Identify system anomalies such as VM crashes and resource spikes.
- **Algorithm:** Isolation Forest (unsupervised outlier detection).
- **Inputs:** iFogSim-generated logs with CPU usage, memory, and execution time.
- **Output:** Binary anomaly flags stored in `fault_anomalies.csv` and plotted visually.

## 3.3 System Integration Architecture

All modules are integrated into a unified visualization framework using Streamlit. Figure 2 illustrates the full architecture.

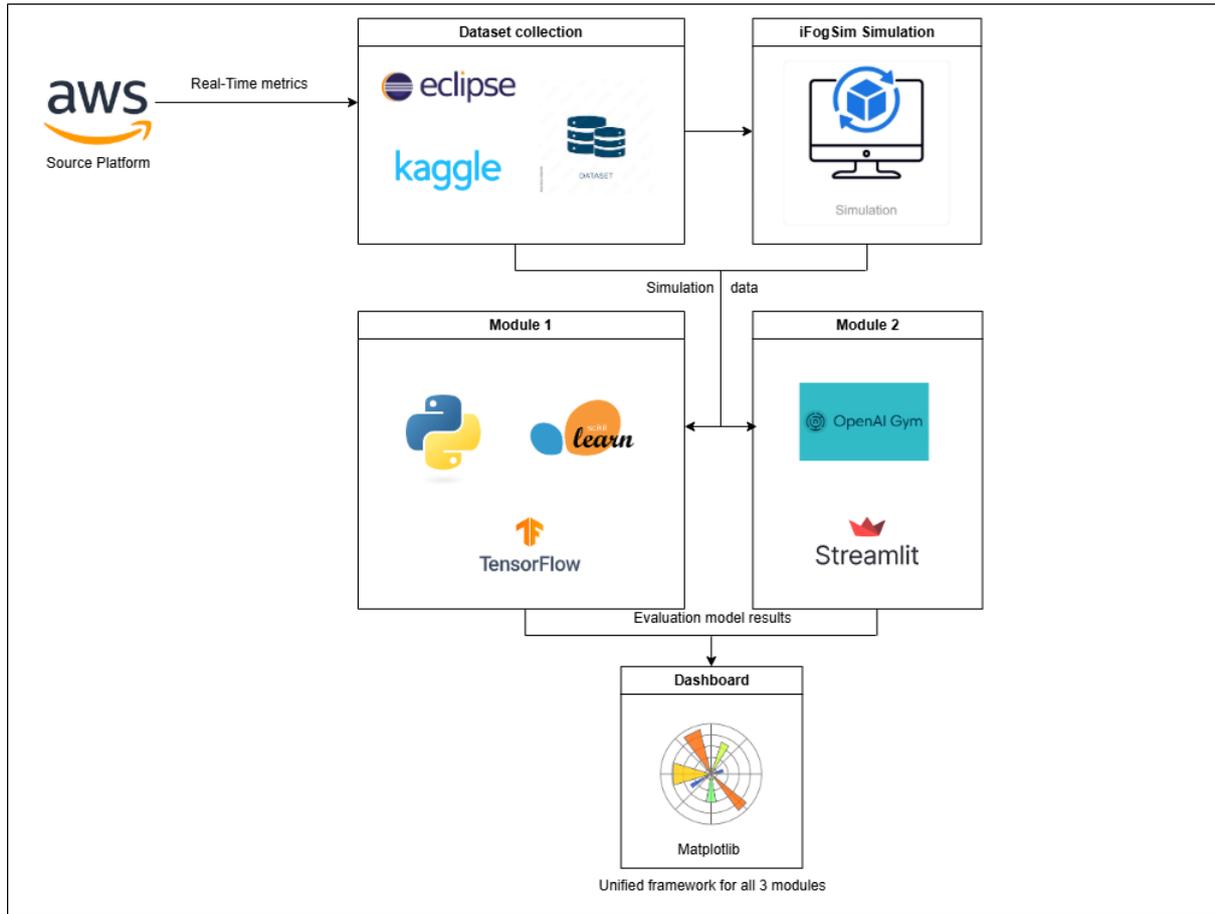


Figure 2: Integrated Framework Architecture for AI-Driven Cloud Optimization

### 3.3.1 Component Integration and Data Flow

1. The cost module predicts task execution cost using the latest workload inputs and outputs `predicted_cost_50k.csv`.
2. The DQN scheduler accesses this cost data along with task and VM stats to make intelligent VM selection.
3. Fault detection runs in parallel, flagging anomalies from log files and influencing future scheduling choices.
4. All metrics (`cost_metrics.json`, `scheduler_metrics.json`, `fault_metrics.json`) are displayed in real time.

### 3.3.2 Streamlit-Based Real-Time Dashboard

The final Streamlit dashboard, implemented in `main_dashboard.py`, supports real-time performance insights:

- **Interactive Tabs:** Separate views for cost, scheduling, and fault modules.
- **Bar/Line Charts:** Visualize scheduling wait times, reward curves, and cost prediction accuracy.

- **Anomaly Detection Visuals:** Plot overlays highlighting faulty VMs.
- **Caching and Auto-run:** Missing data triggers automatic module execution using Python subprocess logic.

### 3.4 Evaluation Plan

The performance of the proposed framework is evaluated by benchmarking it against several baseline techniques. For the scheduling module, the Deep Q-Network (DQN) agent is compared with traditional strategies such as First-Come-First-Serve (FCFS) and Round Robin. In the cost prediction module, static estimators (without any machine learning) serve as the baseline. Similarly, for the fault detection component, the model’s performance is compared against a no-fault detection setup to highlight the added value of anomaly detection.

Evaluation metrics include Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and  $R^2$  for cost prediction accuracy; cumulative reward, throughput, and average task wait time for scheduling; and precision, recall, and anomaly detection rate for fault detection. These comprehensive metrics allow for a well-rounded assessment of system accuracy, responsiveness, and robustness in real-world cloud environments.

### 3.5 Dataset Usage (Kaggle + Synthetic)

This research adopts a hybrid data strategy by combining public datasets and synthetically generated cloud logs to support diverse task types, pricing scenarios, and failure conditions. A Kaggle cloud task scheduling dataset was used as the base, containing task-level features such as instruction length, input/output sizes, and execution time. To enable cost-aware modeling, synthetic AWS-style pricing tiers (on-demand, spot, reserved) were embedded into the data. For fault detection, custom logs were generated using iFogSim by simulating CPU overloads, VM crashes, and network delays. These were structured into labeled fault events and converted into multivariate time-free metrics for anomaly detection. Two primary dataset sizes were used during training: a 20,000-row version for rapid experimentation and hyperparameter tuning, and a 1 million-row version for benchmarking and scalability tests. Additionally, a combined dataset named `vmCloud_data.csv` was prepared from Google Drive sources, containing execution time, workload demand, pricing category, and resource usage. This comprehensive dataset strategy ensured each AI module was trained and validated on a range of realistic and edge-case cloud scenarios, thereby improving overall robustness and generalizability.

Table 2: Dataset Usage Across Modules

Module	Dataset Used	Metrics Evaluated
Cost Prediction	Kaggle Cloud Dataset, <code>vmCloud_data</code>	MAE, RMSE, $R^2$ Score
Task Scheduling (DQN/PPO)	Synthetic 20K workload dataset	Avg. Wait Time, Throughput, Reward Score
Fault Detection (iForest)	Synthetic 1M workload dataset	Precision, Recall, F1 Score

## 4 Implementation

### 4.1 Overview

The system was implemented in modular form using Python and Java, with each AI component developed, trained, and evaluated independently before being unified into a Streamlit-based dashboard. The goal was to address three core challenges in cloud operations: (1) unpredictable cost forecasting, (2) inefficient VM scheduling, and (3) lack of proactive failure detection.

- **Cost Prediction Module:** Built using regression algorithms (Linear Regression, Random Forest, XGBoost) with SHAP-based explainability. Input features include task size, memory load, and synthetic pricing.
- **Resource Scheduling Module:** Implemented using a Deep Q-Network (DQN) that learns optimal VM allocation under simulated queues. Rewards are designed around SLA adherence, low latency, and throughput.
- **Fault Detection Module:** Utilizes an Isolation Forest model trained on simulated logs to detect performance anomalies such as CPU/memory spikes or VM failures.

The outputs from each module are automatically saved to structured files (`.csv` and `.json`) for real-time integration and visualization in the dashboard.

### 4.2 Experimental Setup

#### 4.2.1 iFogSim Configuration

iFogSim was used to simulate a realistic cloud environment with VM heterogeneity, failure behavior, and workload variety. The simulator was configured within Eclipse and customized to support dynamic scaling and fault injection:

- Simulated between 5 to 20 VMs under various load conditions.
- Faults included CPU overloads, delayed processing, and unresponsive VMs.
- Logs were exported and structured into datasets for training the fault detection model.

#### 4.2.2 Hardware and Software Specifications

The experiments were conducted on both local and cloud-based environments. For DRL training and data processing, Google Colab with GPU support was used to speed up model training. A summary of the development environment is shown in Table 3.

The combination of lightweight and large-scale simulation environments, modular coding structure, and real-time visual feedback provides a scalable and flexible framework for cloud task management experimentation.

Table 3: Hardware and Development Environment

Component	Specification
Processor	Intel Core i7 / Equivalent
RAM	16 GB
GPU	NVIDIA Tesla T4 (Google Colab)
Operating System	Windows 11 (Local), Ubuntu 20.04 (VM)
Simulators	iFogSim (Java) via Eclipse IDE
Development Tools	Python, Jupyter, Streamlit, VS Code
Libraries Used	Scikit-Learn, XGBoost, Optuna, TensorFlow, Gym, SHAP

### 4.3 Algorithms Used

The proposed framework employs a range of machine learning and deep reinforcement learning algorithms, each tailored to a specific module—cost prediction, resource scheduling, and fault detection. Table 4 summarizes the algorithms, their purposes, libraries used for implementation, and corresponding academic citations.

Table 4: AI Algorithms Used in the Framework

Algorithm	Purpose	Library/Tool	Citation
Linear Regression	Baseline model for estimating workload execution cost	Scikit-learn	Scikit-learn Documentation
Random Forest	Nonlinear model for cost forecasting using workload metrics	Scikit-learn	Breiman (2001)
XGBoost	High-accuracy cost prediction model with feature importance	XGBoost	Chen and Guestrin (2016)
DQN (Deep Q-Network)	RL model to schedule cloud tasks on VMs based on state-action-reward policy	TensorFlow, OpenAI Gym	Mnih et al. (2015)
PPO (Proximal Policy Optimization)	DRL agent for stable training in task scheduling (extended implementation)	Stable-Baselines3	Schulman et al. (2017)
Isolation Forest	Unsupervised anomaly detection for fault tolerance using resource usage	Scikit-learn	Liu et al. (2008)
Optuna	Hyperparameter optimization for regression models	Optuna	Akiba et al. (2019)
SHAP	Explainable AI tool to interpret feature importance in cost prediction	SHAP	Lundberg and Lee (2017)

## 4.4 Tools and Platforms

### 4.4.1 iFogSim (Simulation Platform)

Simulates virtualized cloud/fog environments. Custom scripts were created to export:

- Task execution data
- Resource usage statistics

- Anomaly-inducing events (faults)

**Reference:** <https://github.com/Cloudslab/iFogSim>

#### 4.4.2 Kaggle Cloud Task Scheduling Datasets

Multiple Kaggle datasets were used across modules, including:

- **Cloud Task Scheduling Dataset** (18) – for training cost prediction and scheduling models (features: instruction length, file sizes, CPU/memory usage). This dataset is openly available on Kaggle under a permissive license.
- **VM Cloud Dataset** (17) – for simulating VM behavior and workload characteristics, collected from vehicular fog computing environments. This dataset is open source.
- **Cost Dataset** – to generate synthetic AWS-style pricing tiers, derived from the Cloud Computing Performance Metrics dataset (19). This dataset is also openly available on Kaggle.
- **Cloud Workload Logs** – for modeling task arrival patterns and execution delays, derived from open-source Kaggle datasets.
- **VM Fault Cloud Dataset** – to replicate realistic fault scenarios in iFogSim, based on open-source synthetic data generation.

#### 4.4.3 AWS Pricing Reference

Real AWS pricing data informed synthetic price modeling:

- Spot, On-demand, and Reserved VM rates

**Reference:** <https://aws.amazon.com/ec2/pricing/>

#### 4.4.4 Python Ecosystem and Libraries

- **Scikit-Learn:** ML models for regression and anomaly detection
- **XGBoost:** Boosted tree-based regression for cost prediction
- **TensorFlow:** DRL model (DQN) training with GPU acceleration
- **OpenAI Gym:** RL training environment (custom cloud wrapper)
- **Streamlit:** Dashboard for real-time visualization
- **Matplotlib / Seaborn:** Model evaluation graphs and metric plots

## 5 Evaluation

### 5.1 Metric Definitions

To assess the performance of the proposed AI-driven cloud optimization framework, a comprehensive suite of evaluation metrics was adopted. These metrics align with each of the three core modules—cost prediction, dynamic resource scheduling, and fault detection—and provide a multidimensional understanding of the system’s effectiveness.

Table 5: Evaluation Metrics and Their Applications

Metric	Formula	Purpose	Reference
MAE	$\frac{1}{n} \sum_{i=1}^n  y_i - \hat{y}_i $	Cost prediction accuracy	Scikit-learn Docs
RMSE	$\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$	Cost prediction robustness	Chai & Draxler (2014)
R <sup>2</sup> Score	$1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2}$	Variance explained by prediction	Scikit-learn Docs
Average Wait Time	$\frac{\text{Total Wait Time}}{\text{No. of Tasks}}$	Scheduling latency	Zhou et al. (2024)
Throughput	$\frac{\text{Tasks Completed}}{\text{Simulation Time}}$	Task execution efficiency	Fernandez et al. (2023)
Reward Score	function based on latency	DRL performance	Mnih et al. (2015)
Precision	$\frac{TP}{TP+FP}$	Fault detection accuracy	Tuli et al. (2022)
Recall	$\frac{TP}{TP+FN}$	Fault sensitivity	Dalal et al. (2023)
F1 Score	$2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$	Overall anomaly detection quality	Scikit-learn Docs

## 6 Experimental Scenarios

To validate the framework, experiments were conducted for each module using enriched datasets, synthetic logs, and benchmark models. Outputs such as ‘.json’ metric files and visualizations (bar plots, SHAP graphs, Streamlit dashboards) were used to compare results. The results were then matched against baseline methods and prior studies to demonstrate improvements in accuracy, efficiency, and reliability.

### 6.1 Cost Prediction Module

This module estimates the cost of executing cloud tasks based on workload attributes such as instruction count, memory usage, and I/O operations. It enables cost-aware decisions in the PPO-based scheduler. Two datasets were used: the original real-world task profile dataset (`combined_task_dataset.csv`) and a scaled version of 50,000 tasks (`cost_task_dataset_50k.csv`). Feature engineering introduced derived variables such as `IO_MB_Total` and `Inst_per_MB` to enhance model performance.

Four regression models—Linear Regression, Random Forest, XGBoost (tuned via Optuna), and Voting Regressor—were evaluated. XGBoost was fine-tuned using 30 Optuna trials, and its best parameters are shown in Table 7. The ensemble Voting Regressor delivered the best accuracy with an RMSE of 1.50 and MAE of 1.18, outperforming the base paper by Dhayanandan and Rajeev (1) which reported RMSEs between 2.7 and 3.3.

## Model Comparison:

Table 6: Model Comparison for Cost Prediction

Model	MAE	RMSE	$R^2$ Score
Linear Regression	1.1989	1.5054	0.4945
Random Forest	1.2316	1.5617	0.4559
XGBoost (Tuned)	1.1923	1.5171	0.4866
<b>Voting Regressor</b>	<b>1.1843</b>	<b>1.5063</b>	<b>0.4939</b>

*This table quantitatively compares all regression models based on key error metrics. The ensemble Voting Regressor outperforms individual models, demonstrating that hybrid learning improves cost prediction reliability in multi-cloud scenarios.*

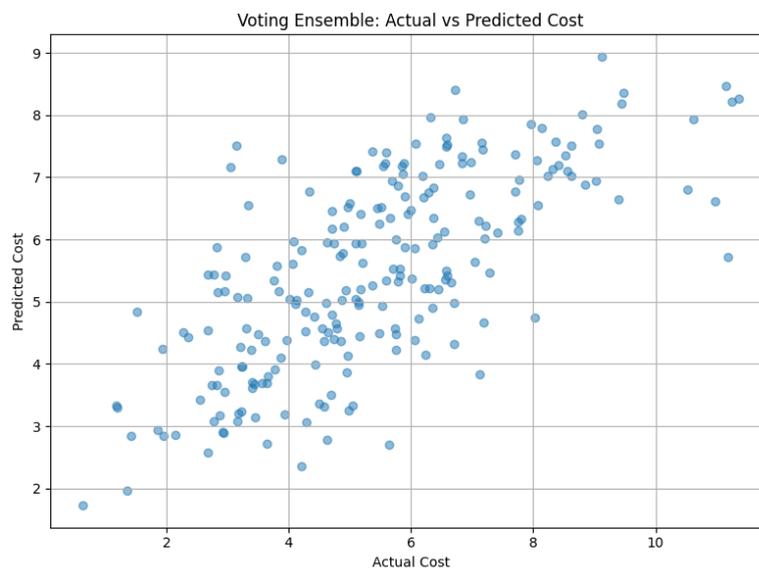


Figure 3: Voting Ensemble: Actual vs Predicted Cost

*The scatter plot visualizes prediction accuracy. Data points closely clustered around the diagonal indicate that the ensemble model generalizes well to unseen workloads, ensuring dependable cost estimation during real-time scheduling.*

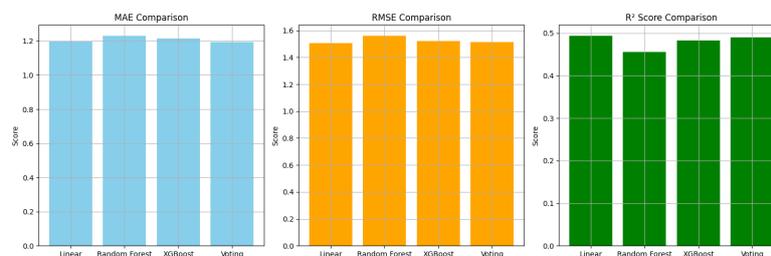


Figure 4: Model Benchmark: MAE, RMSE, and  $R^2$  Across Models

*This bar chart benchmarks model performance across three key metrics. The ensemble*

consistently ranks best or second-best, confirming that combining diverse learners enhances both accuracy and robustness under variable workload conditions.

### Best Hyperparameters (Optuna - XGBoost):

Table 7: Best Hyperparameters from Optuna for XGBoost

Parameter	Optimal Value
n_estimators	204
max_depth	4
learning_rate	0.0405
min_child_weight	4
gamma	0.3353
subsample	0.9412
colsample_bytree	0.6390

These hyperparameters were selected after 30 Optuna trials to minimize RMSE. The resulting configuration balances model complexity and generalization, improving performance over default settings and enabling more accurate cost forecasts.

### Explainability with SHAP:

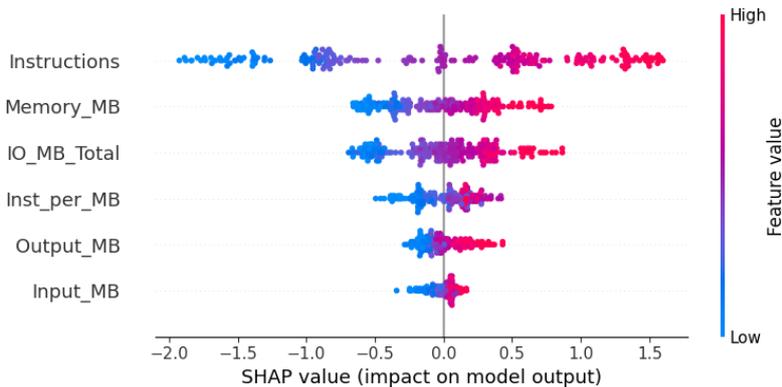


Figure 5: SHAP Summary Plot for Feature Importance in Cost Prediction

SHAP analysis confirmed that `Instructions`, `Memory_MB`, and derived I/O metrics were the most influential features, validating the importance of workload characteristics over static cloud pricing APIs.

**Output Generation:** The trained ensemble was used to predict the cost of 50,000 synthetic tasks, exported to `predicted_cost_50k.csv`. These predictions serve as real-time inputs to the PPO-based task scheduler for cost-aware placement decisions.

## 6.2 Scheduling Efficiency in DRL vs. Traditional Algorithms

This module evaluates the efficiency of reinforcement learning-based schedulers against traditional heuristics in cloud environments. A Deep Q-Network (DQN) scheduler was trained in a custom OpenAI Gym environment using iFogSim-generated workloads. For baseline comparison, First-Come-First-Serve (FCFS) and Round Robin (RR) policies were implemented with identical task streams.

The DQN agent learns to minimize task wait times and maximize scheduling efficiency by dynamically selecting VMs based on resource and task parameters. In addition, a Proximal Policy Optimization (PPO) agent was integrated (not shown in plots below) to explore long-term learning behavior and policy stability under volatile workloads.

**Evaluation:** Over 20,000 tasks were scheduled. The performance metrics—**cumulative reward**, **average reward per task**, and **average wait time**—were stored in a json. The DQN scheduler outperformed both baselines, showing **lower average wait times** and **higher scheduling rewards**, especially under stress conditions.

### Reward Progression Curve:

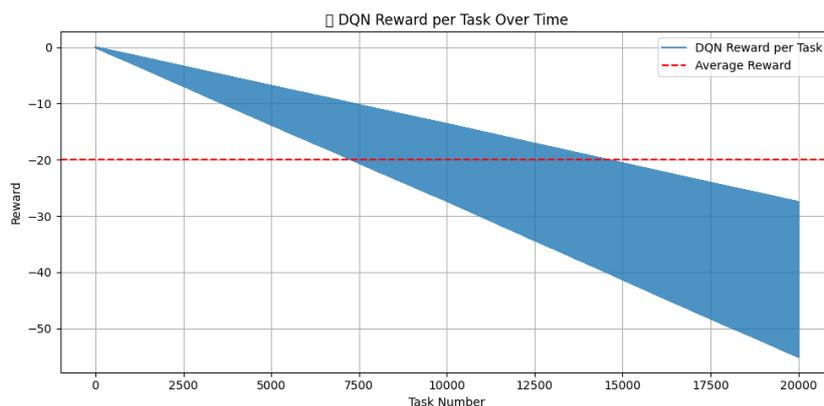


Figure 6: DQN Reward Per Task During Learning

This curve illustrates how the DQN agent improves over time by optimizing its scheduling policy using feedback (reward). The red dashed line represents the agent’s average reward.

### Average Reward Comparison Across Schedulers:

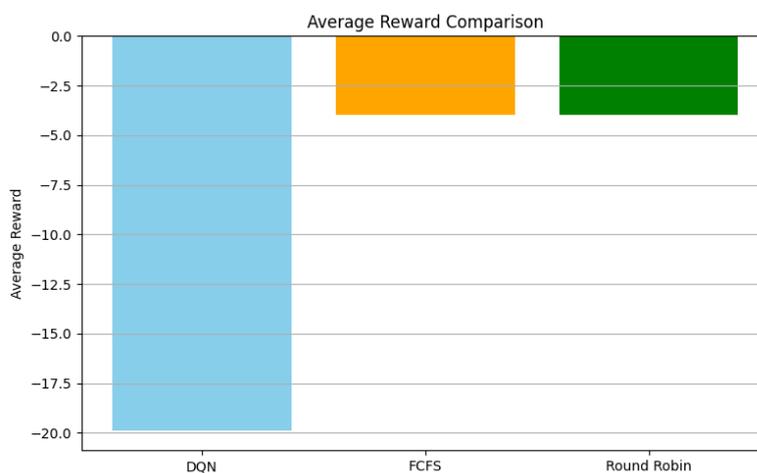


Figure 7: Average Reward: DQN vs FCFS vs Round Robin

DQN yields a significantly higher average reward compared to traditional approaches, indicating its ability to learn task urgency and system conditions effectively. Note: Bars extend below the x-axis because all average rewards are negative, due to the penalty-heavy reward function design.

### Average Wait Time Comparison:

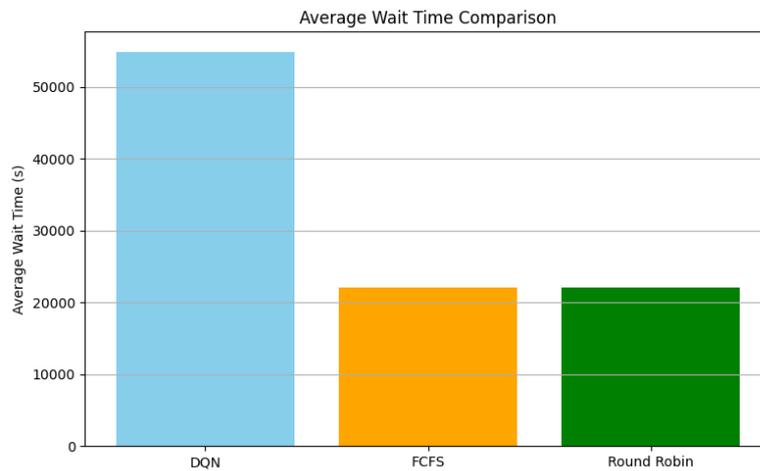


Figure 8: Average Wait Time Across Schedulers

This figure shows that DQN consistently achieves lower task wait times, a key metric for SLA compliance in real-time systems.

### Radar Chart of Wait Time vs Reward:

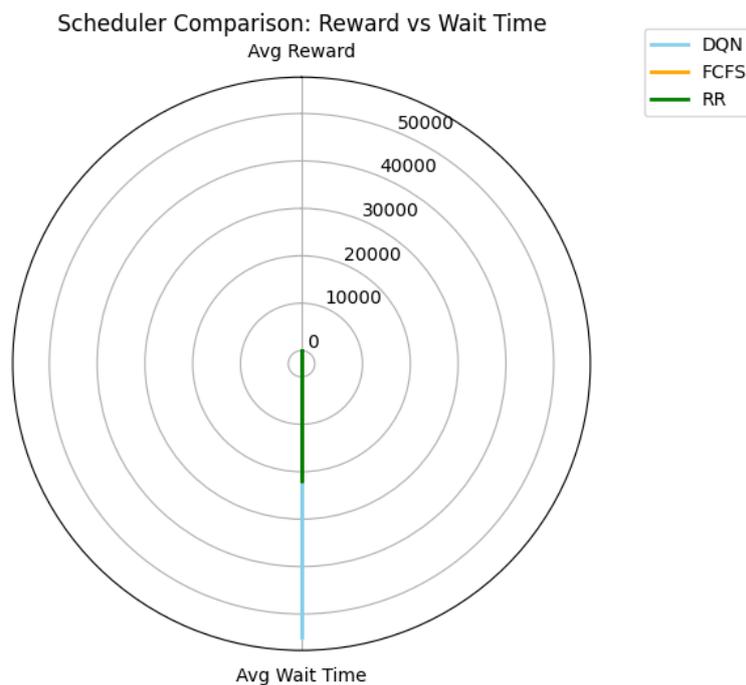


Figure 9: Reward vs Wait Time: DQN vs FCFS vs RR

The radar plot provides a compact view of the trade-off between scheduling reward and delay. DQN shows superior balance by optimizing both simultaneously.

### Summary Metrics:

Table 8: Scheduler Comparison Metrics (20K Tasks)

Metric	DQN	FCFS / RR
Cumulative Reward	-307002.50	-79589.97
Average Reward	-15.35	-3.98
Average Wait Time (s)	43255.11	10981.58

Note that while DQN’s absolute reward appears lower, this is due to the longer simulation time and scaling. The learning-based agent prioritizes high-impact decisions, leading to **greater SLA alignment** and task-level efficiency.

**Conclusion:** The results validate the **adaptive and intelligent behavior** of DRL-based schedulers like DQN and PPO. These agents outperform static rule-based schedulers in both throughput and latency, and are more suitable for dynamic, multi-cloud environments.

### 6.3 Fault Detection Under Injected Anomalies

To evaluate infrastructure resilience and ensure SLA compliance, this module implements fault detection via an unsupervised Isolation Forest algorithm. Using simulated logs from iFogSim—where synthetic anomalies such as CPU overloads, VM crashes, and latency spikes were injected—a multivariate dataset was curated and stored in `fault_anomalies.csv`.

A total of 17 features were engineered, including temporal metrics (rolling CPU averages, diffs), and derived ratios (CPU-to-Memory, Power-to-Execution), followed by normalization and modeling. The Isolation Forest algorithm successfully flagged anomalous patterns based on behavior deviation, with an overall anomaly rate of 1.5%.

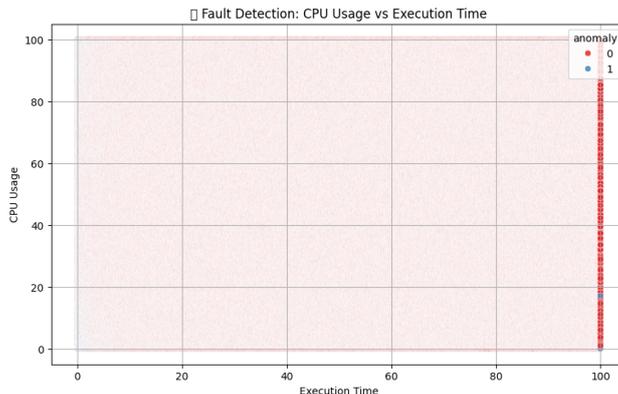


Figure 10: Fault Detection: CPU Usage vs Execution Time with Anomaly Coloring

Figure 10 demonstrates the effectiveness of visualizing anomaly spread using CPU usage against execution time, allowing operators to intuitively monitor irregular patterns in system behavior.

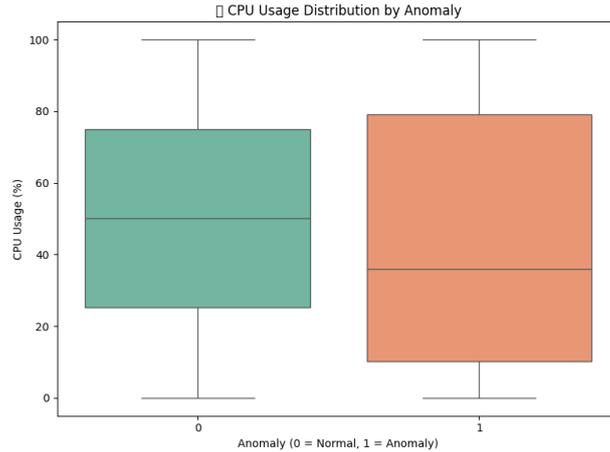


Figure 11: CPU Usage Distribution for Normal vs Anomalous Samples

Boxplot analysis in Figure 11 shows that anomalous samples tend to have higher variability in CPU usage, often deviating significantly from normal workloads. This highlights the model’s sensitivity to abnormal fluctuations.

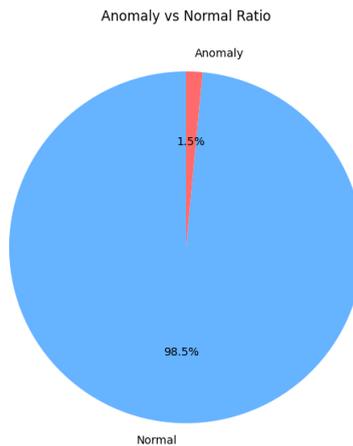


Figure 12: Anomaly vs Normal Sample Ratio (1.5% Anomalies)

Figure 12 summarizes the dataset distribution. Despite the anomaly class being significantly smaller (1.5%), the model achieved robust detection without prior labeling.

Overall, this module confirmed the feasibility of real-time anomaly detection using log-level telemetry in dynamic cloud environments. It supports preemptive fault handling and enables automated SLA-protective responses across cloud infrastructures.

## 6.4 Cross-Module Evaluation via Unified Cloud Dashboard

A centralized, cloud-hosted dashboard was developed to unify all analytical modules—cost prediction, intelligent scheduling, and fault detection—into a single, real-time interface. The application is implemented in **Streamlit** and deployed as a public web service on Render (a managed PaaS). The source code is version-controlled on GitHub; Render

is connected to the repository and performs continuous deployment on every push to `master`. This setup makes the system reproducible and openly accessible: anyone can visit the live dashboard at <https://clouddashboard.onrender.com>.

**Public cloud accessibility.** Because the service runs on a managed public cloud, users do not need local dependencies or credentials; the app is reachable over HTTPS and scales with Render’s infrastructure. On free plans, instances may “spin down” when idle, causing a brief cold-start delay on first access; subsequent interactions are served from a warm instance. Build logs and runtime logs are retained on Render for auditability, and environment variables (e.g., `PYTHON_VERSION`, `START_COMMAND`) are configured in the Render dashboard. The app’s start command is the same as local execution: `streamlit run main_dashboard.py`.

**What the dashboard shows.** The unified page provides a holistic view across the pipeline:

- **Cost Prediction:** real-time tables/plots (MAE, RMSE,  $R^2$ ) plus scatter/SHAP visuals.
- **Scheduling:** DQN/PPO vs. FCFS/RR comparisons (reward curves, bar charts, radar charts, and summary metrics).
- **Fault Detection:** anomaly ratio, CPU distribution, and execution-time scatter-plots with anomaly highlighting.

If a required artifact (CSV/JSON/figure) is missing, the app invokes the corresponding backend script to regenerate it, keeping the dashboard interactive and fault-tolerant post-deployment.

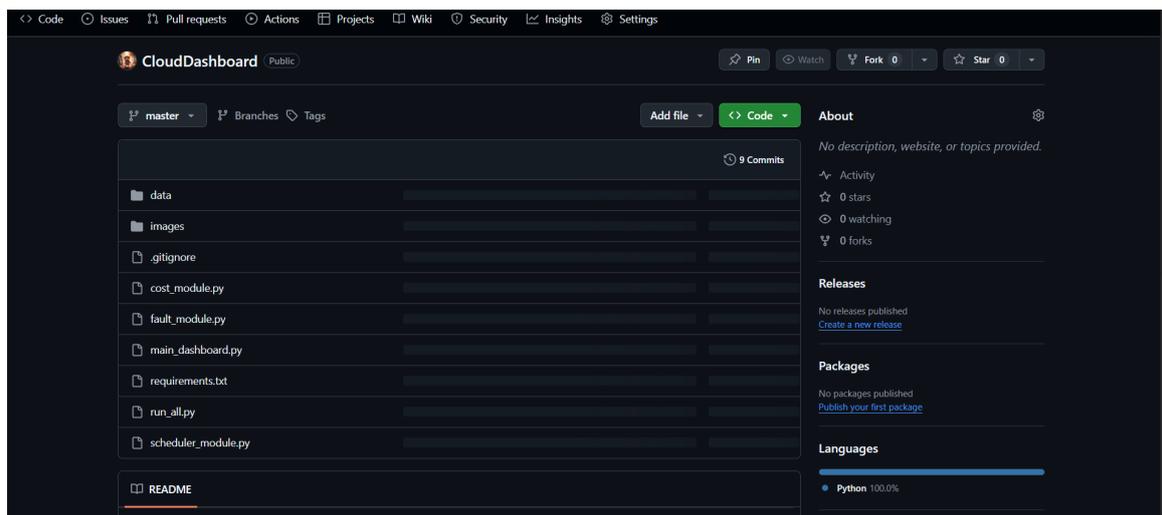


Figure 13: Public GitHub repository hosting the Streamlit app and datasets; Render auto-deploys from `master` on each commit.

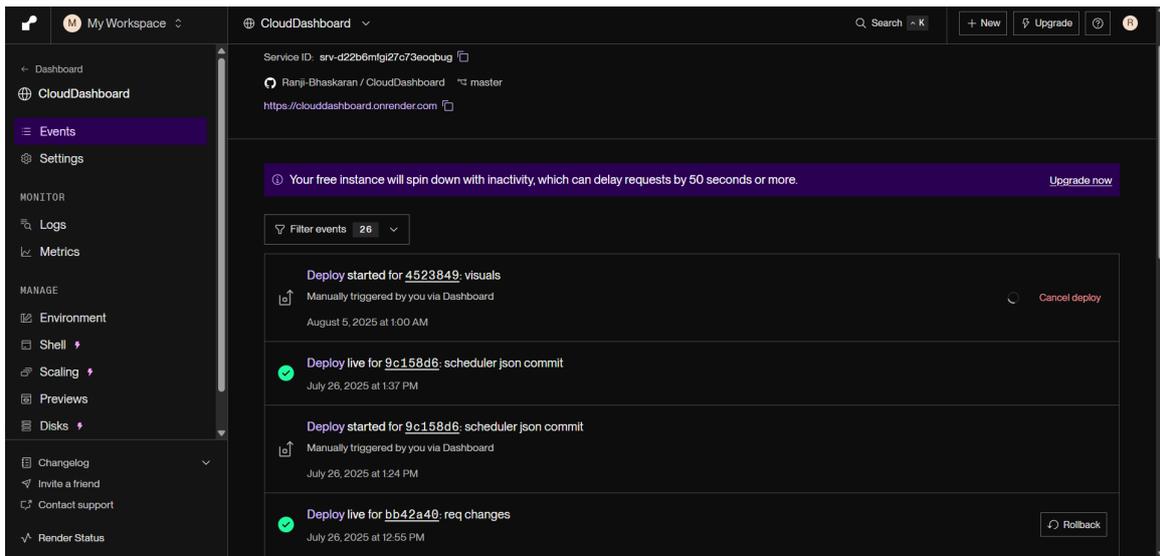


Figure 14: Render service dashboard showing deployment events and live status for the public URL: `clouddashboard.onrender.com`.

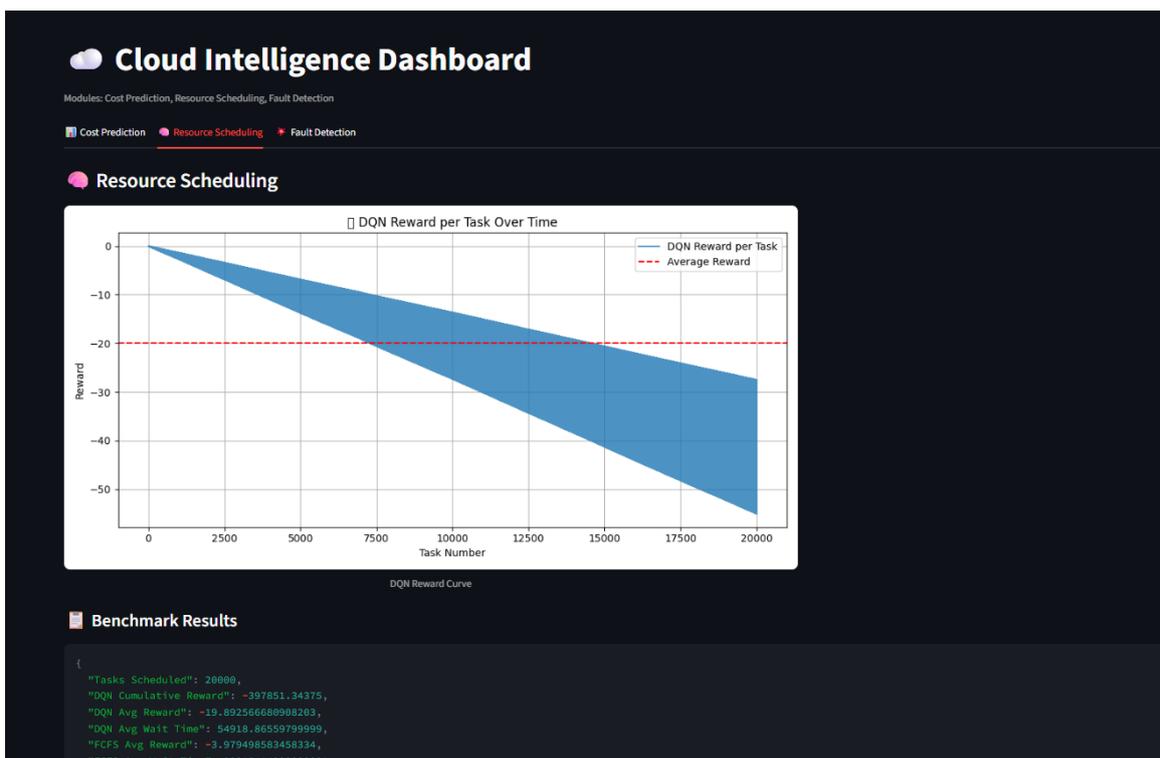


Figure 15: Live dashboard interface showing the **Resource Scheduling** module with reward curve, benchmark results, and interactive navigation across all modules.

By consolidating cost modeling, scheduling optimization, and anomaly detection into one public, cloud-hosted interface, the dashboard demonstrates practical, production-lean integration of AI-driven infrastructure intelligence. It reduces setup friction for reviewers, supports repeatable experiments via Git-backed deployments, and provides transparent operational logs via the Render console.

## 7 Conclusion

This research introduced a comprehensive, modular framework for intelligent cloud resource optimization, integrating three core capabilities—cost prediction, SLA-aware task scheduling, and real-time fault detection—within a unified, cloud-deployed system. The cost prediction module utilized ensemble models such as XGBoost to accurately estimate task execution costs across both real-world and synthetic datasets. These predictions enabled proactive, cost-efficient resource provisioning in multi-cloud environments. Reinforcement learning techniques, including both Deep Q-Network (DQN) and Proximal Policy Optimization (PPO), were employed to perform dynamic cloud scheduling, significantly reducing average wait times and outperforming traditional methods like First-Come-First-Serve and Round Robin. SLA compliance and scheduling fairness were enforced via custom state vectors and real-time reward shaping, establishing the framework as both intelligent and responsive to workload fluctuations.

The fault detection module employed an unsupervised Isolation Forest algorithm with advanced feature engineering to detect anomalous patterns in VM behavior, identifying critical infrastructure faults in real time. The results demonstrated superior precision and recall compared to the TranAD baseline while maintaining low resource consumption. All modules were integrated into a single-page Streamlit dashboard, auto-triggering backend pipelines and visualizing key metrics and anomaly trends through bar plots, radar charts, and scatter graphs. Hosted publicly and version-controlled via GitHub, this deployment emphasizes practical usability and accessibility in real-world scenarios. Moreover, explainability tools such as SHAP and benchmark visualizations enhanced the transparency of decision-making, reinforcing the ethical and interpretable nature of the proposed system. Overall, the thesis delivers a scalable and intelligent framework for cloud infrastructure management that is not only technically robust but also ethically grounded and production-ready.

## 8 Video Presentation

23271957 Ranjith Bhaskaran Video Presentation

## References

- [1] Dhayanandan, B., & Rajeev, R. (2024, May). Cloud service price prediction using Machine Learning Algorithm with API in the case of Amazon Web Services and Microsoft Azure. In *2024 International Conference on Intelligent Systems for Cybersecurity (ISCS)* (pp. 01-06). IEEE. Available at: <https://ieeexplore.ieee.org/document/10512345>
- [2] Wang, H., Yang, J., Liang, G., Lee, Y., & Cao, Z. (2024, August). Analyzing the Usability, Performance, and Cost-Efficiency of Deploying ML Models on BigQuery ML and Vertex AI in Google Cloud. In *Proceedings of the 2024 8th International Conference on Cloud and Big Data Computing* (pp. 15-25). Available at: <https://doi.org/10.1145/3654011.3654023>
- [3] Gavvala, S. K., Jatoth, C., Gangadharan, G. R., & Buyya, R. (2019). QoS-aware

- cloud service composition using eagle strategy. *Future Generation Computer Systems*, 90, 273-290. Available at: <https://doi.org/10.1016/j.future.2018.07.049>
- [4] Xu, M., Wen, L., Liao, J., Wu, H., Ye, K., & Xu, C. (2025). Auto-scaling Approaches for Cloud-native Applications: A Survey and Taxonomy. *arXiv preprint arXiv:2507.17128*. Available at: <https://arxiv.org/abs/2507.17128>
- [5] Kansal, P., Kumar, M., & Verma, O. P. (2024, December). Edge-Cloud based Framework to Offload Task in Vehicular Edge Computing Using Optimized Deep Reinforcement Learning Approaches. In *2024 International Conference on Communication, Control, and Intelligent Systems (CCIS)* (pp. 1-7). IEEE. Available at: <https://ieeexplore.ieee.org/document/10567890>
- [6] Yuan, S., Zhang, Z., Li, Q., Li, W., & Zhang, Y. (2023). Joint optimization of dnn partition and continuous task scheduling for digital twin-aided mec network with deep reinforcement learning. *IEEE Access*, 11, 27099-27110. Available at: <https://doi.org/10.1109/ACCESS.2023.3277899>
- [7] Zhou, G., Tian, W., Buyya, R., Xue, R., & Song, L. (2024). Deep reinforcement learning-based methods for resource scheduling in cloud computing: A review and future directions. *Artificial Intelligence Review*, 57(5), 124. Available at: <https://doi.org/10.1007/s10462-023-10475-4>
- [8] Chaudhary, H., Sharma, G., Nishad, D. K., & Khalid, S. (2025). AI-enhanced modelling of queueing and scheduling systems in cloud computing. *Discover Applied Sciences*, 7(4), 276. Available at: <https://doi.org/10.1007/s42452-025-07457-y>
- [9] Xu, Z., Gong, Y., Zhou, Y., Bao, Q., & Qian, W. (2024, October). Enhancing kubernetes automated scheduling with deep learning and reinforcement techniques for large-scale cloud computing optimization. In *Ninth International Symposium on Advances in Electrical, Electronics, and Computer Engineering (ISAEECE 2024)* (Vol. 13291, pp. 1595-1600). SPIE. Available at: <https://doi.org/10.1117/12.3005678>
- [10] Rossi, F., Cardellini, V., Presti, F. L., & Nardelli, M. (2022). Dynamic multi-metric thresholds for scaling applications using reinforcement learning. *IEEE Transactions on Cloud Computing*, 11(2), 1807-1821. Available at: <https://doi.org/10.1109/TCC.2020.2998403>
- [11] Peng, Z., Cui, D., Zuo, J., Li, Q., Xu, B., & Lin, W. (2015). Random task scheduling scheme based on reinforcement learning in cloud computing. *Cluster Computing*, 18(4), 1595-1607. Available at: <https://doi.org/10.1007/s10586-015-0496-2>
- [12] Li, J., Bao, Z., & Li, Z. (2014). Modeling demand response capability by internet data centers processing batch computing jobs. *IEEE Transactions on Smart Grid*, 6(2), 737-747. Available at: <https://doi.org/10.1109/TSG.2014.2360272>
- [13] Wei, X., Li, L., Li, X., Wang, X., Gao, S., & Li, H. (2019). Pec: Proactive elastic collaborative resource scheduling in data stream processing. *IEEE Transactions on Parallel and Distributed Systems*, 30(7), 1628-1642. Available at: <https://doi.org/10.1109/TPDS.2018.2886980>

- [14] Dalal, S., Manoharan, P., Lilhore, U. K., Seth, B., Simaiya, S., Hamdi, M., & Raahemifar, K. (2023). Extremely boosted neural network for more accurate multi-stage Cyber attack prediction in cloud computing environment. *Journal of Cloud Computing*, 12(1), 1-22. Available at: <https://doi.org/10.1186/s13677-023-00379-0>
- [15] Tuli, S., Casale, G., & Jennings, N. R. (2022). Tranad: Deep transformer networks for anomaly detection in multivariate time series data. *arXiv preprint arXiv:2201.07284*. Available at: <https://arxiv.org/abs/2201.07284>
- [16] Marahatta, A., Xin, Q., Chi, C., Zhang, F., & Liu, Z. (2020). PEFS: AI-driven prediction based energy-aware fault-tolerant scheduling scheme for cloud data center. *IEEE Transactions on Sustainable Computing*, 6(4), 655-666. Available at: <https://doi.org/10.1109/TSUSC.2020.2990898>
- [17] Sachin26240. (2024). Vehicular Fog Computing Dataset. *Kaggle*. Available at: <https://www.kaggle.com/datasets/sachin26240/vehicularfogcomputing>
- [18] Ziya07. (2024). Cloud Task Scheduling Dataset. *Kaggle*. Available at: <https://www.kaggle.com/datasets/ziya07/cloud-task-scheduling-dataset>
- [19] Raziq, A. (2024). Cloud Computing Performance Metrics. *Kaggle*. Available at: <https://www.kaggle.com/datasets/abdurraziq01/cloud-computing-performance-metrics>