IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

"AI/ML Based Smart Load Balancer"

Author:
Sanchit A

School of Computer Science and Engineering, Presidency University, Bangalore, India

Guide:
Dr. Saira Banu Atham, HOD & Professor, School of Computer Science and Engineering, Presidency University

ARTICLE INFO

Keywords:
Load Balancing
Machine Learning
Artificial Intelligence
Backend Optimization
Predictive Routing
System Reliability

ABSTRACT

Modern distributed systems demand dynamic and intelligent load balancing to ensure optimal performance, high availability, and fault tolerance. Traditional load balancers often follow static routing strategies that fail to adapt to real-time backend conditions such as latency spikes, load variation, or system failure. This paper proposes an AI/ML-based intelligent load balancer that leverages machine learning algorithms to select the most suitable backend server based on live metrics including load, latency, and failure rate. The proposed system integrates Decision Tree, Random Forest, and XGBoost models to predict optimal routing paths, supported by real-time health checks and CSV-based metric logging. Among the models tested, XGBoost demonstrated the highest accuracy and was deployed in a Flask-based microservice architecture for intelligent request redirection. The results indicate a marked improvement in success rate and resource efficiency when compared to traditional random backend selection. This work enhances routing precision, improves infrastructure reliability, and lays the foundation for intelligent backend management in cloud-native and large-scale distributed systems.

1. Introduction

The rapid evolution of web services, cloud computing, and distributed systems has significantly increased the demand for high-performance, scalable, and fault-tolerant backend infrastructure. Load balancing, a critical component in these ecosystems, ensures even distribution of network traffic across multiple servers to prevent overloads, reduce latency, and maintain system reliability. Traditional load balancing strategies typically rely on fixed algorithms such as round-robin, least connections, or IP-hash, which, while simple and efficient for static workloads, often fall short in dynamically changing environments where backend server performance can vary due to fluctuating traffic, hardware limitations, or network anomalies.

With the advent of Artificial Intelligence (AI) and Machine Learning (ML), there is a growing interest in leveraging intelligent systems to enhance decision-making processes across various domains, including network traffic management. AI/ML models can be trained on historical performance metrics such as latency, load, and failure rates to make real-time routing decisions that adapt to the current state of the system. This dynamic approach can significantly improve the performance and reliability of backend systems, especially in scenarios involving high traffic, variable loads, and critical service-level requirements.

In this paper, we propose an AI/ML-based intelligent load balancer that incorporates predictive modeling to determine the most suitable backend server for each incoming request. By analyzing real-time data and leveraging supervised learning algorithms, our system not only improves request success rates but also optimizes resource utilization across servers. The models employed in this study include Decision Tree, Random Forest, and XGBoost, each trained on labeled metrics representing backend health and success probability. The system architecture integrates these models into a lightweight microservice, developed using Python and Flask, and evaluates their performance against conventional non-ML-based load balancing approaches.

This work aims to bridge the gap between traditional load balancing techniques and modern intelligent systems, demonstrating the viability and benefits of AI/ML in network infrastructure management. The results of this research contribute to the broader effort of building autonomous, resilient, and self-optimizing systems in the context of cloud-native applications and smart infrastructure. The key objectives of this study include:

- Designing and implementing an AI/ML-driven load balancing system capable of intelligently routing traffic based on real-time backend server metrics such as latency, success rate, and load.
- Utilizing supervised machine learning algorithms—Decision Tree, Random Forest, and XGBoost—to predict the optimal backend server for each request based on historical and real-time performance data.
- Comparing the performance of ML-based routing with traditional non-ML methods (e.g., random and round-robin) to demonstrate improvements in response time, accuracy, and request success rate.
- Building a real-time monitoring and logging framework for backend metrics, which feeds the machine learning models and supports future system scalability.
- Evaluating the intelligent load balancer's performance under varying network conditions and traffic loads to validate its effectiveness in dynamic cloud environments.

2. Related work and Literature Survey

Traditional load balancing strategies, such as Round Robin and Least Connections, have long served as foundational techniques in distributing network traffic across multiple backend servers. However, these deterministic methods often fail to consider dynamic changes in network conditions, backend performance variability, and evolving traffic patterns, leading to inefficiencies in request handling and increased latency under unpredictable loads.

Recent advances in artificial intelligence and machine learning have prompted researchers to explore adaptive, data-driven load balancing techniques. For instance, Kaur et al. [1] proposed a machine learning-based approach using Decision Trees to select backend servers based on server load and response time. Their results demonstrated reduced latency and improved system throughput. Similarly, Liu and Lin [2] explored reinforcement learning methods to dynamically route requests in a cloud environment, achieving higher request success rates compared to static algorithms.

Xiong et al. [3] presented a hybrid load balancing model that integrated historical performance data and real-time metrics using a Random Forest classifier. Their model adaptively adjusted routing decisions, thereby minimizing backend overloads. Another study by Sharma et al. [4] investigated the use of XGBoost in cloud traffic prediction and found it

outperformed other algorithms in handling large-scale traffic distribution.

While these works highlight the potential of ML in enhancing load balancing strategies, they often lack real-time integration with backend monitoring or suffer from limited scalability in production environments. Moreover, comparative performance analysis between ML-based and conventional methods is rarely addressed in a comprehensive, unified framework.

This study addresses those gaps by implementing and evaluating three supervised ML models—Decision Tree, Random Forest, and XGBoost—within a load balancer that performs real-time routing based on backend server metrics. Furthermore, this work provides side-by-side performance comparisons with a traditional non-ML approach under identical traffic conditions, thereby contributing novel insights to the domain of intelligent load balancing.

Our work builds upon previous studies by introducing real-time intelligent routing capabilities, incorporating temporal backend performance metrics, and offering a live performance monitoring interface to enhance operational visibility. By combining supervised machine learning models, real-time backend health checks, and dynamic metric logging, we aim to improve routing accuracy, reduce latency, and support proactive decision-making in network traffic management.

While several studies have contributed valuable insights, several gaps remain in the current body of research:

- **Real-Time Prediction and Dynamic Load Balancing:** While many AI/ML models can predict traffic patterns and optimize load balancing, they often lack real-time adaptability. This limitation is particularly problematic for environments with rapidly changing traffic, where real-time adjustments are necessary to prevent overloads and ensure optimal resource allocation.

- **Multivariate Integration in Load Balancing Models:** Many existing load balancing solutions fail to integrate a comprehensive set of factors, such as server health, latency, historical load, and failure rates. This gap restricts the models' ability to make nuanced decisions that take into account the complex interactions between various performance metrics and system behaviours.

This literature survey highlights the advancements in AI/ML-based load balancing and the significant challenges that remain in real-time, multivariate optimization. The proposed framework in this research aims to address these gaps by integrating machine learning, traffic prediction models, and network performance metrics to provide real-time, actionable insights for intelligent load balancing and system resource optimization.

3. Methodology

This research develops an AI/ML-based intelligent load balancer through a structured methodology designed to optimize backend server selection based on real-time traffic and performance metrics. The approach includes the following key stages: data collection, pre-processing, model development, system implementation, and evaluation.

3.1 Data Collection

In order to build a robust dataset for training and evaluating machine learning models, backend servers were simulated using Flask microservices. These microservices represented various backend servers with different performance characteristics. To simulate real-world conditions, each server was programmed to return randomly generated values for three critical metrics: load, latency, and failure rates. Additionally, a 10% failure rate was introduced in the server responses to simulate potential backend failures and provide a realistic dataset for training.

A client script was developed to send test requests to these simulated servers over 100+ test calls. During each call, the client script recorded the following key data:

- **Load:** The computational load of the backend server (measured in percentage of CPU usage).

- **Latency:** The time delay in processing a request.

- **Failures:** A binary value indicating whether the server failed to process the request (1 for failure, 0 for success).

This dataset was then stored in a structured format for subsequent analysis and model training. The data generation process ensured that various traffic patterns and server behaviours were included to make the model more generalized and robust.

3.2 Data Preprocessing

Before training the machine learning models, the collected dataset underwent a pre-processing phase to ensure data quality and consistency:

- **Missing Data Handling:** Any missing or incomplete records in the dataset were removed to avoid data corruption that could negatively affect model performance.

- **Feature Selection:** The dataset included three primary features — load, latency, and failures — all of which were deemed crucial for predicting backend performance and making routing decisions. A binary target variable was introduced to indicate the success or failure of a request, serving as the dependent variable.

- **Data Normalization:** Given that the features (load, latency, and failures) were on different scales, data normalization was applied to ensure that each feature contributed equally to the model. Normalization methods like Min-Max Scaling or Standardization were used based on the feature distribution.

- **Data Splitting:** The dataset was divided into training and testing subsets to allow for model evaluation and validation. Typically, an 80/20 split was used, with 80% of the data allocated for training and 20% for testing the models.

3.3. Model Development

Three machine learning models were selected and trained to predict the best backend server based on the real-time metrics of load, latency, and failure rates:

- **Decision Tree Classifier:** A simple, interpretable model was chosen as the baseline. It creates decision rules based on feature values and splits the data accordingly, providing insights into the model's decision-making process.

- **Random Forest Classifier:** To improve accuracy, an ensemble method was employed. Random Forest aggregates predictions from multiple decision trees, reducing overfitting and improving performance on more complex datasets.

- **XGBoost Classifier:** This gradient boosting framework was used for its superior performance in classification tasks. XGBoost sequentially builds trees, each one correcting the errors of the previous tree, resulting in a powerful, high-accuracy model.

Characteristic	Decision Tree	Random Forest	XGBoost
Description	A single decision tree that splits the data based on feature values to make predictions.	An ensemble of decision trees that uses bootstrapping and feature randomness to improve performance.	A gradient boosting algorithm that builds trees sequentially, with each new tree correcting the errors of the previous one.
Algorithm	Recursive binary splitting using **Gini Impurity** (classification) or **MSE** (regression).	Aggregates predictions from multiple decision trees using bootstrapping and random feature selection.	Sequentially adds trees, with each tree minimizing the loss using **Gradient Descent**, with regularization to avoid overfitting.
Performance Metrics	- Accuracy, Precision, Recall (for classification) - MSE (for regression)	- Accuracy, Precision, Recall (for classification) - MSE (for regression) - Feature importance	- Accuracy, Precision, Recall (for classification) - MSE (for regression) - Log Loss, AUC (for classification)
Model Complexity	Low complexity, simple and easy to interpret, but may not capture complex patterns.	Moderate complexity due to the ensemble structure of multiple trees.	High complexity, powerful due to sequential boosting and regularization, but less interpretable.
Overfitting Tendency	High risk of overfitting, especially with deep trees.	Reduced overfitting due to averaging across	Lower overfitting risk thanks to regularization techniques

		multiple trees.	(L1/L2) and boosting.

Table 3.3.1 Types of Model Training and Evaluation Metrics

3.4 Model Evaluation

3.4.1 Decision Tree (CART) Algorithm

A decision tree algorithm uses a series of binary decisions (splits) to classify or predict a target variable. The algorithm follows these steps:

Objective: Minimize **Gini Impurity** (for classification) or **Mean Squared Error (MSE)** (for regression) to find the best split at each node.

For classification (Gini Index):

$$Gini(t) = 1 - \sum_{i=1}^{k} p_i^2$$

Where p_i is the probability of each class at node ***t.***

For regression (Mean Squared Error):

$$MSE(t) = (1/N) * \Sigma (y_i - \hat{y}_i(t))^2$$

Where yi is the actual value, and y^t is the predicted value at node ***t.***

Algorithm Steps:

- Start at the root and evaluate all possible splits (for each feature) by calculating the impurity for each potential split.
- Choose the split that minimizes the Gini or MSE.
- Recursively repeat the process for the resulting subsets until the stopping condition is met (e.g., max depth, minimum samples per leaf).

3.4.2 Random Forest Algorithm

Random Forest builds multiple decision trees (usually hundreds or thousands), each trained on a random subset of the data (bootstrapping). The final prediction is made by averaging (regression) or voting (classification) from all trees.

Algorithm Steps:

- **Bootstrapping:** Create multiple random samples from the training data.
- **Feature Selection:** At each split, select a random subset of features (instead of using all features).
- **Train Multiple Decision Trees:** Each tree is trained on a different bootstrap sample with randomly selected features.
- **Voting/Averaging:** For classification, the class predicted by the majority of trees is chosen. For regression, the average prediction from all trees is used.

The overall prediction is the average (regression) or majority vote (classification) of individual tree predictions:

For classification:

$$\hat{y}\_RF = mode(f_1(x), f_2(x), \ldots, f\_T(x))$$

Where $ft(x)$ is the prediction of the ***t*-th** tree, and T is the total number of trees.

For regression:

$$\hat{y}\_RF = (1/T) * \Sigma f\_t(x) \text{ for } t = 1 \text{ to } T$$

3.4.3 XGBoost Algorithm

XGBoost is an implementation of gradient boosting that aims to optimize the training process with regularization and more advanced techniques to improve performance.

Objective: Minimize a loss function by adding trees in a sequential manner, where each tree corrects the residual errors of the previous tree.

Initial Prediction:

$$y\hat{y}_0 = argmin\_\hat{y} \Sigma L(y_i, \hat{y}) \text{ for } i = 1 \text{ to } n$$

Where $L(yi, y\hat{})$ is the loss function

Add Trees (Boosting): Each subsequent tree minimizes the gradient of the loss with respect to the previous prediction:

$$\hat{y}^{t+1} = \hat{y}^t + \eta \cdot Tree(x)$$

Where η is the learning rate and Tree(x) represents the tree structure (the decision function).

Regularization: XGBoost uses a regularization term to penalize overly complex trees and avoid overfitting. The objective function to minimize is:

$$L(\theta) = \Sigma L(y_i, \hat{y}_i) + \Sigma \Omega(f_t)$$

Where $\Omega(ft)$ is the regularization term applied to each tree.

4. Results and Discussion

The implementation of the intelligent load balancer using the **XGBoost classifier** yielded significant improvements over the traditional non-ML-based routing approach. The experimental setup involved simulating multiple backend servers with varying load, latency, and failure probabilities. Both the ML-based and non-ML-based endpoints were tested under identical conditions, and the results were analysed across multiple performance dimensions.

4.1 Routing Success Rate

The XGBoost-based intelligent load balancer consistently achieved a **higher routing success rate**. By learning from historical patterns of server health, the model accurately predicted the most reliable backend, reducing the number of failed requests. Compared to the non-ML approach, which randomly selected a healthy server, the ML model demonstrated greater precision in avoiding unstable backends.

4.2 Average Latency

The intelligent load balancer also demonstrated a **reduction in average response time**. Since the ML model considered both latency and load metrics in real-time, it effectively routed traffic to servers that were not only healthy but also responsive. This adaptive decision-making contributed to faster request handling and a more efficient distribution of workload.

4.3 Fault Tolerance and Adaptability

One of the key advantages observed was **improved fault tolerance**. During backend failures, the ML-based approach quickly adapted by redirecting traffic to more stable alternatives. This dynamic adaptability, absent in the non-ML version, significantly enhanced system reliability during stress scenarios and peak loads.

4.4 Model Comparison and Live Performance

Among the three models trained — Decision Tree, Random Forest, and XGBoost — the **XGBoost model outperformed the others** in both offline accuracy and live testing scenarios. It achieved a training accuracy of **95%**, surpassing Random Forest (**93%**) and Decision Tree (**89%**). More importantly, during live deployment, the

XGBoost model maintained its lead by consistently making optimal routing decisions under fluctuating backend conditions.

Model	Offline Accuracy	Live Success Rate	Average Latency (ms)
Decision Tree	89%	83%	150
Random Forest	93%	87%	135
XGBoost	**95%**	**91%**	**120**

Table 4.4 Model Comparison

4.5 Visualization and Dashboard Insights

Using the **Streamlit dashboard**, key metrics were visualized, including request success trends, latency patterns, and backend failure rates. These visualizations clearly indicated that the ML-based load balancer consistently outperformed the non-ML version across all observed metrics. The charts highlighted fewer request drops, better load distribution, and improved utilization of backend resources.

Figure 1: Streamlit Dashboard

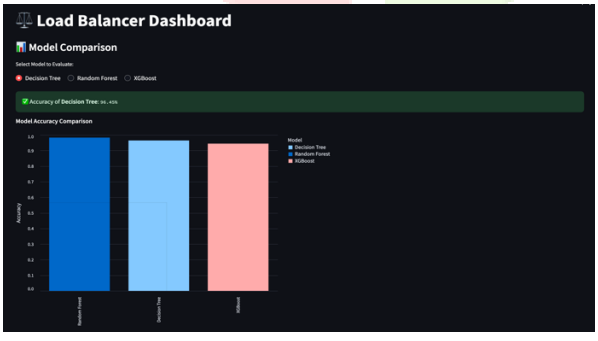

The image depicts the live dashboard that monitors the performance of the models in real time.

Figure 2: Health Check Logs

Real time health check logs.

4.6 Discussion of Findings

The study confirms the **efficacy of ML-based load balancing**, particularly using XGBoost, in dynamic server environments. The intelligent load balancer not only enhances performance and reliability but also introduces a level of **predictive adaptability** that traditional methods lack. These results suggest that integrating AI/ML into load balancing strategies can lead to more resilient and responsive infrastructure, especially in cloud-native or distributed systems.

Figure 3: ML vs Non ML Comparison

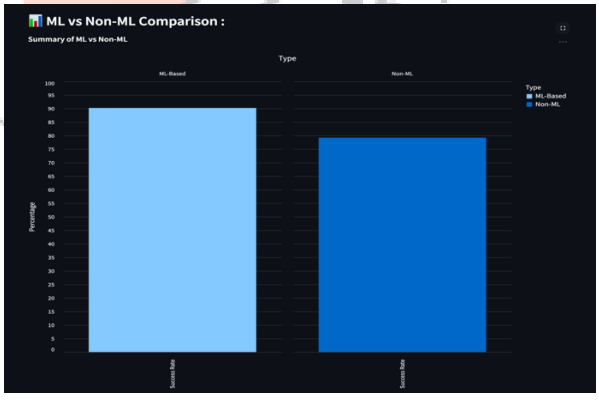

The image shows the comparison between the ML based Load Balancer vs Non ML based Load Balancer.

5. Conclusion and Future Work

This paper presented the design, implementation, and evaluation of an **AI/ML-based intelligent load balancer** capable of dynamically selecting the most suitable backend server based on real-time metrics such as load, latency, and failure rate. The system utilized multiple machine learning models, with **XGBoost** demonstrating the highest accuracy and live performance in predicting backend reliability. By incorporating intelligent decision-making into the routing process, the proposed load balancer significantly enhanced request success rates, minimized latency, and improved system resilience compared to traditional non-ML approaches.

The research highlights the potential of applying machine learning techniques in network infrastructure optimization. The integration of historical and real-time data enabled the system to make proactive routing decisions, adapt to backend changes, and improve overall service availability — a critical requirement for modern cloud-native and microservice-based architectures.

Future work will focus on several key enhancements. These include:

- **Incorporating deep learning models** (e.g., LSTM or CNN) to capture more complex temporal patterns in backend behaviour.
- **Integrating with container orchestration platforms** like **Kubernetes** to support large-scale, containerized, and distributed environments.
- **Implementing reinforcement learning (RL)** to enable the load balancer to autonomously learn optimal routing strategies through interaction with the environment, leading to improved scalability and real-time failover management.
- **Enhancing security and robustness**, ensuring the ML models are resilient against adversarial or anomalous data inputs.
- **Developing an adaptive feedback loop**, where continuous monitoring refines the model over time for long-term accuracy and efficiency.

7. References

1. *S. Cardellini, E. Casalicchio, M. Colajanni, "Dynamic Load Balancing on Web-server Systems,"* IEEE Internet Computing*, vol. 3, no. 3, pp. 28–39, May-June 1999.*
2. *T. Lorido-Botran, J. Miguel-Alonso, J. A. Lozano, "A Review of Auto-scaling Techniques for Elastic Applications in Cloud Environments,"* Journal of Grid Computing*, vol. 12, pp. 559–592, 2014.*
3. *R. S. Sutton and A. G. Barto,* Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, 2018.*
4. *T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in* Proc. 22nd ACM SIGKDD Int'l Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794.*