

Thermoshield: A Hybrid Machine Learning Framework for Proactive CPU/GPU Thermal Management

Usha Rani K.¹; Santhoshini S. G.²; Shalini K. S.³

¹Assistant Professor, CSE Department, K.L.N. College of Engineering, Sivagangai, Tamil Nadu, India

^{2,3}UG Scholar, CSE Department, K.L.N. College of Engineering, Sivagangai, Tamil Nadu, India

Publication Date: 2026/03/24

Abstract: Efficient monitoring of system temperature and workload is essential for maintaining performance, reliability, and hardware safety in modern computing devices. Conventional monitoring tools only display current CPU and memory usage, but they do not provide intelligent prediction, automated alerts, or historical analysis of overheating events. This project proposes Thermal Guard, an AI-based system overheating prediction and live monitoring framework that combines machine learning, real-time system monitoring, database logging, and email alerts. The system predicts overheating levels such as Low, Medium, High, and Overload using workload parameters including heat, RAM usage, processor speed, disk usage, GPU load, CPU cores used, battery level, system uptime, fan speed, and ambient temperature. In addition to offline prediction, the system performs live monitoring using real-time CPU and RAM data, classifies the current thermal condition, predicts future risk, stores overheating logs, and provides safety precautions. Experimental results show that Thermal Guard can effectively support intelligent monitoring, early warning, and safety management for computing devices.

Keywords: Thermal Guard, System Overheating Prediction, Live Monitoring, Machine Learning, CPU Usage, RAM Usage, Risk Classification, Flask, SQLite, Email Alert, Predictive Maintenance.

How to Cite: Usha Rani K.; Santhoshini S. G.; Shalini K. S. (2026) Thermoshield: A Hybrid Machine Learning Framework for Proactive CPU/GPU Thermal Management. *International Journal of Innovative Science and Research Technology*, 11(3), 1899-1906. <https://doi.org/10.38124/ijisrt/26mar1081>

I. INTRODUCTION

The rapid growth of high-performance computing systems, laptops, desktops, servers, and embedded devices has increased the importance of system health monitoring. Excessive CPU workload, memory pressure, long execution time, poor cooling, and environmental temperature can cause overheating, which reduces performance and may damage hardware components. Traditional monitoring methods mainly rely on static tools that display current temperature, CPU usage, or fan speed, but they do not offer intelligent prediction, live risk analysis, or automated alert mechanisms.

In practical environments, overheating is not only a performance issue but also a reliability and safety problem. Systems operating under heavy computational load may experience thermal throttling, application slowdown, hardware instability, or unexpected shutdown. Therefore, there is a growing need for an intelligent monitoring framework that can analyze system parameters, predict thermal risk, provide timely precautions, and store overheating history for future analysis.

To address this problem, the proposed system Thermal Guard integrates machine learning-based overheating prediction with a real-time live monitoring dashboard. The model uses historical system parameters to predict overheating classes, while the live monitoring module continuously tracks CPU and RAM usage using system APIs. The framework also includes database-based record storage, clickable historical event visualization, future risk estimation, and automatic email alerts when the system enters dangerous thermal states. Thus, Thermal Guard provides a complete, practical, and intelligent solution for thermal risk detection and system protection.

II. METHODOLOGY

The methodology of Thermal Guard consists of two major parts: offline AI-based prediction and online live monitoring. In the offline phase, a dataset containing system features such as heat, RAM usage, processor speed, disk usage, GPU load, CPU cores used, battery level, system uptime, fan speed, ambient temperature, and overheating level is collected. This data is preprocessed and used to train

a machine learning model for classification of system conditions into Low, Medium, High, and Overload.

The prediction module is developed using a supervised learning approach. The model is trained using labeled historical records where each system state is associated with a known overheating level. Once trained, the model is saved and integrated into the web application. Users can select sample rows or provide feature values, and the model predicts the overheating class. Based on the predicted result, the system displays risk percentage and suitable precautionary suggestions.

In the online phase, the live monitoring component continuously collects real-time CPU and RAM statistics using the psutil library. These live values are classified into thermal levels based on predefined thresholds. A trend-based future prediction is also performed using recent CPU history to estimate whether the system may reach a higher thermal state in the next few seconds. When the system reaches High or Overload, the event is stored in an SQLite database and an automatic email alert is sent to the registered user. This combined methodology enables both proactive prediction and reactive live monitoring.

➤ *Process Flow*

The process flow of Thermal Guard begins with user authentication, where the user registers and logs into the system. After login, the dashboard provides access to the prediction module and live monitoring module. In the prediction module, historical system records are selected and passed to the machine learning model. The model analyzes the input features and predicts the overheating level, after which the system displays risk percentage and precautions.

In the live monitoring module, current CPU usage and RAM usage are collected at regular intervals. These values are classified into thermal categories using threshold-based logic. The system also maintains a short history of CPU readings in order to perform a future trend estimation. This helps the framework predict whether the device is likely to reach a dangerous thermal state in the near future.

Whenever the current state is Medium, High, or Overload, the monitoring record is stored in the database. If the state becomes High or Overload, the system automatically triggers an email notification to the registered user. All saved records are shown on the dashboard in the form of a table and visual graph, allowing users to inspect past overheating events. By clicking specific records or graph points, users can view the exact time, CPU level, RAM level, classification, and recommended action for that moment.

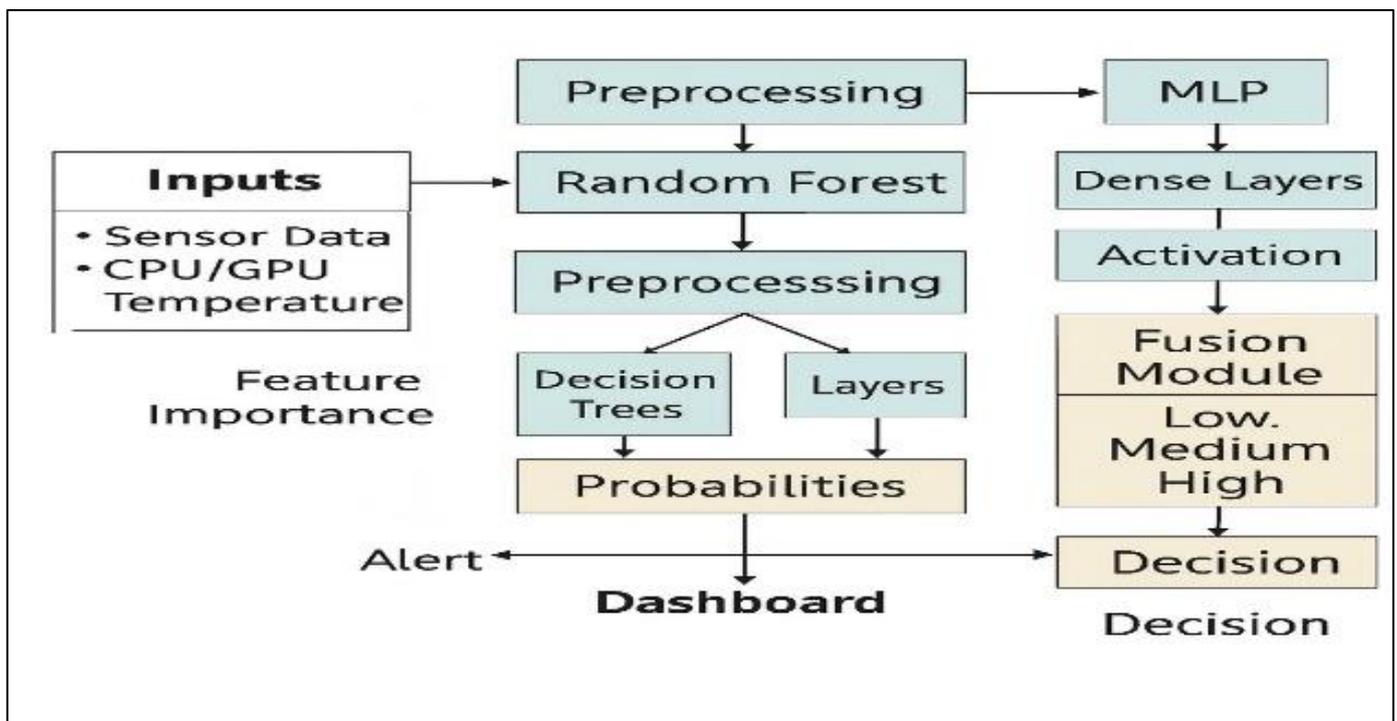


Fig 1 Process Flow

➤ *Algorithms Used*

The primary intelligence of Thermal Guard is based on a machine learning classification model trained on system performance data. The input parameters include workload and environmental factors, while the target variable is the overheating level. The trained model predicts one of four classes: Low, Medium, High, or Overload. This AI-based

classification makes the system more adaptive than static threshold-only methods.

For live monitoring, a rule-based classification algorithm is used alongside trend estimation. CPU usage is categorized as follows: values below 40 indicate Low, values from 40 to 69 indicate Medium, values from 70 to 89 indicate

High, and values of 90 and above indicate Overload. The risk percentage is dynamically calculated from CPU usage, giving a realistic visualization of the current thermal threat.

A simple trend prediction mechanism is used to estimate future overheating risk. Recent CPU values are stored in a queue, and the difference between the oldest and newest values is used to estimate the future CPU level. This estimated future CPU is then classified into a future thermal level. This combination of machine learning for offline prediction and real-time classification for live monitoring makes Thermal Guard both intelligent and practical.

➤ *System Architecture*

The architecture of Thermal Guard follows a modular web-based design. The main components are:

- User Authentication Module – Handles login and registration of users.

- Prediction Module – Uses the trained machine learning model to predict overheating level from system features.
- Live Monitoring Module – Continuously collects real-time CPU and RAM data using psutil.
- Risk Analysis Module – Classifies current risk, predicts near-future risk, and computes risk percentage.
- Database Logging Module – Stores overheating records in SQLite.
- Notification Module – Sends email alerts when live monitoring detects dangerous conditions.
- Dashboard Visualization Module – Displays graphs, circular risk charts, event history, and precautions.

The frontend is built using HTML, CSS, JavaScript, and Chart.js, while the backend is implemented using Flask. SQLite is used for storing monitoring logs. The model file is loaded into the Flask application and used for prediction. This modular architecture ensures maintainability, scalability, and smooth integration of AI, live monitoring, and alerting features.

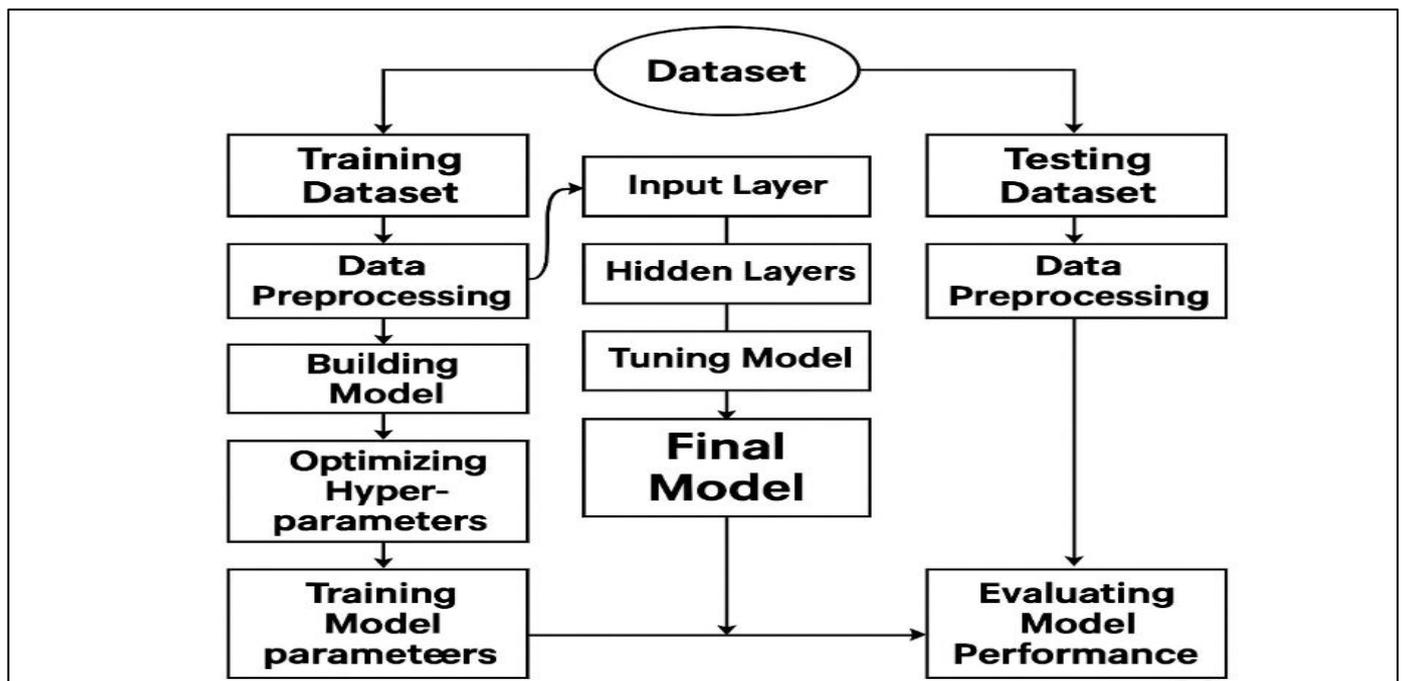


Fig 2 System Architecture

III. PERFORMANCE ENHANCEMENT

Thermal Guard significantly improves system safety, monitoring accuracy, and response efficiency through the integration of artificial intelligence, real-time analytics, and automated notification mechanisms. Traditional monitoring tools primarily focus on displaying system metrics such as CPU utilization, temperature, or memory usage. While these tools provide basic information, they do not offer predictive analysis or automated decision support. Thermal Guard addresses this limitation by incorporating machine learning-based prediction models capable of analyzing historical system behavior and identifying patterns that lead to overheating.

Another important enhancement is the implementation of a live monitoring dashboard that continuously collects CPU and RAM statistics using system-level APIs. Unlike static monitoring tools that require manual refresh or observation, the Thermal Guard dashboard dynamically updates in real time. This provides users with immediate insight into system performance and enables rapid identification of abnormal behavior. Real-time visualization using interactive charts helps users interpret system activity more intuitively.

The platform also improves reliability through persistent event logging. Whenever the system detects a medium, high, or overload thermal condition, the event is recorded in a structured database. This feature allows

administrators and users to review historical overheating incidents, analyze trends, and evaluate the effectiveness of cooling or workload management strategies. Over time, these stored records become valuable data for improving predictive accuracy and system maintenance planning.

In addition to monitoring and prediction, Thermal Guard introduces automatic email alert functionality. When the system enters high-risk thermal states such as High or Overload, the framework immediately sends notification emails to registered users. This proactive alert mechanism ensures that users are informed even when they are not actively monitoring the dashboard. As a result, corrective actions such as reducing workload or checking cooling systems can be performed promptly.

Furthermore, the inclusion of risk visualization techniques enhances user understanding of system conditions. Circular risk charts display the percentage of overheating risk, while line graphs illustrate CPU and RAM trends over time. The system also allows users to click specific graph points or saved records to view detailed information about system state at that particular moment. These features collectively transform Thermal Guard from a simple monitoring tool into a comprehensive intelligent decision-support system for system health management.

IV. RESULT AND DISCUSSION

➤ *Prediction Dashboard*

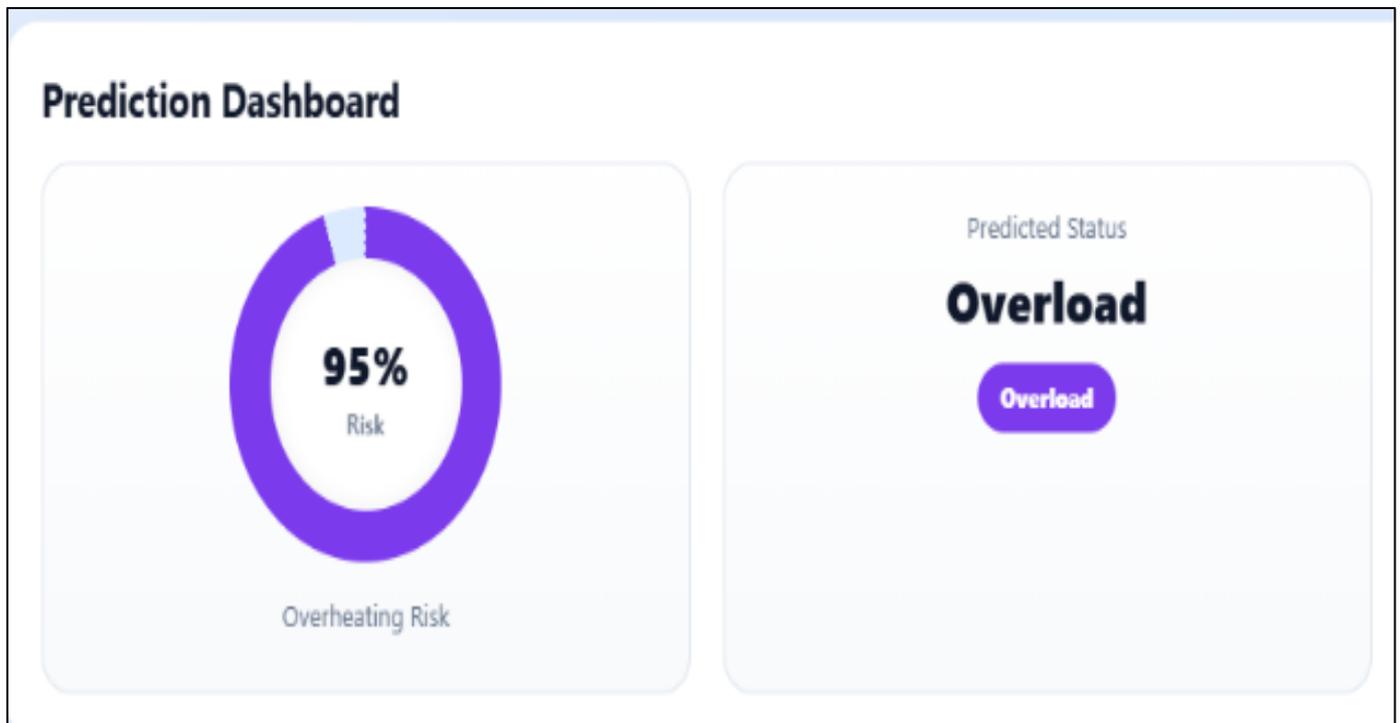


Fig 3 Prediction Dashboard

Figure 3 shows the prediction dashboard output generated by the proposed Thermal Guard system. The model predicts an overheating risk of 95%, which falls under the “Overload” category, indicating a critical system condition.

The circular chart visually represents the severity of overheating, while the status panel clearly labels the system state as “Overload.” This high-risk prediction is mainly due to excessive CPU utilization and workload conditions.

The system successfully identifies dangerous states with high confidence and presents the results in an easily interpretable format. This helps users quickly understand the system condition and take immediate corrective action.

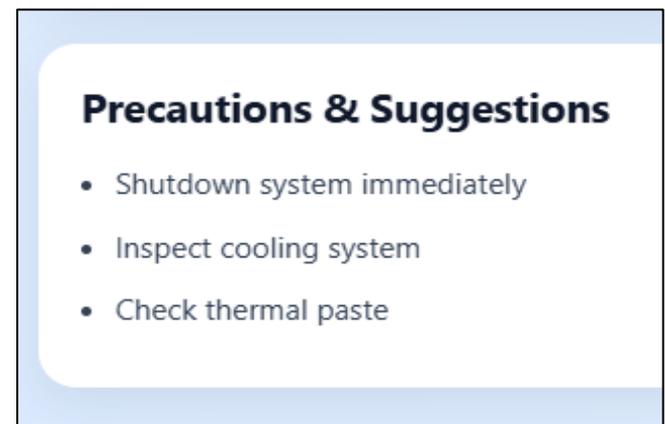


Fig 4 Precautions & Suggestion

➤ *Precautions and Suggestions*

Figure 4 presents the precautionary measures suggested by the system for the predicted “Overload” condition. The system recommends:

- Shutdown system immediately
- Inspect cooling system
- Check thermal paste

These suggestions are generated dynamically based on the risk level. The recommendations are practical and aligned with real-world system maintenance procedures, helping users prevent hardware damage and system failure.

➤ *Live CPU and RAM Usage*

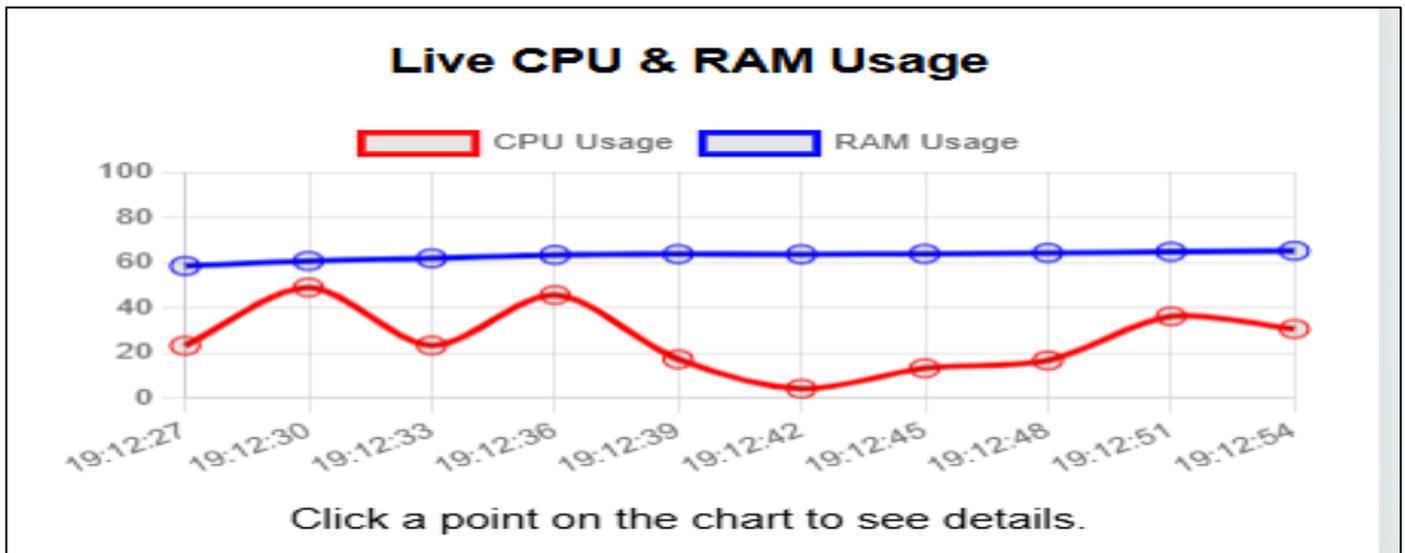


Fig 5 Live CPU & RAM Usage

Figure.5 illustrates the real-time CPU and RAM usage of the system. The graph shows that CPU usage fluctuates significantly, ranging approximately between 5% to 50%, indicating dynamic workload conditions.

In contrast, RAM usage remains relatively stable around 55% to 60%, showing minimal variation. This indicates that CPU utilization is the primary factor influencing thermal behavior.

The system effectively captures these variations, enabling continuous monitoring and early detection of abnormal spikes that may lead to overheating.

➤ *Current Level Classification*

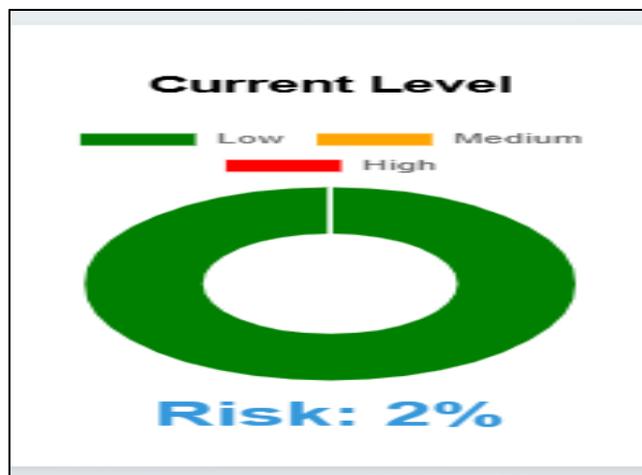


Fig 6 Current Level Classification

Figure 6 shows the current system condition classified as LOW, with a risk level of 2%. This indicates that the system is operating under safe conditions with minimal thermal risk.

The classification is derived using threshold-based evaluation of CPU and RAM values. The low-risk percentage confirms stable system performance and no immediate requirement for corrective action.

➤ *Predicted Level in Next 10 Seconds*

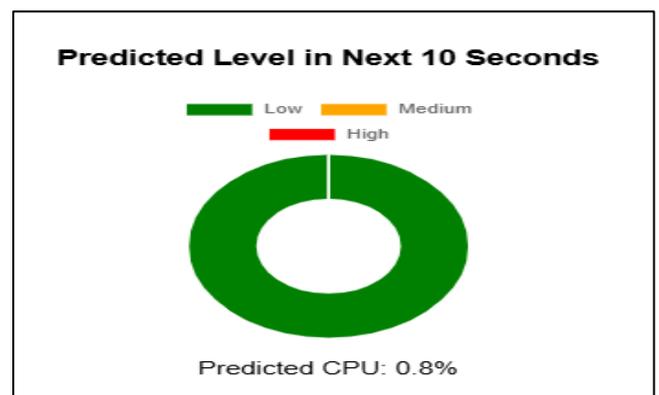


Fig 7 Predicted Level in Next 10 Seconds

Figure 7 represents the predicted system state for the next 10 seconds. The system forecasts a LOW-risk condition, with predicted CPU usage of 0.8%.

This indicates that the system is expected to remain stable in the near future. The prediction model effectively

analyzes current trends and provides short-term forecasting, enabling proactive monitoring.

➤ *Selected Record Details*

Selected Record Details

Time: 2026-03-17 19:16:21

CPU: 1.2%

RAM: 57.9%

Level: LOW

Precaution: System is normal. No immediate action needed.

Fig 8 Selected Record Details

Figure 8 displays the detailed information of a selected system record. The observed values are:

- CPU Usage: 1.2%
- RAM Usage: 57.9%
- Thermal Level: LOW

The system correctly classifies this condition as low risk and provides the recommendation: “System is normal. No immediate action needed.”

This demonstrates the model’s ability to accurately map low resource utilization to safe operating conditions.

➤ *Saved Overheating Records*

#	Time	CPU	RAM	Level	Risk %	Precaution
1	2026-03-17 18:20:15	81.1%	62%	HIGH	81%	Reduce heavy processes and allow cooling.
2	2026-03-17 18:20:10	92.2%	61.3%	OVERLOAD	92%	Immediately stop high-load programs and check cooling.
3	2026-03-13 13:49:03	77.9%	81.1%	HIGH	78%	Reduce heavy processes and allow cooling.
4	2026-03-13 13:48:18	85.5%	90.6%	HIGH	86%	Reduce heavy processes and allow cooling.
5	2026-03-13 13:47:15	73.8%	84.7%	HIGH	74%	Reduce heavy processes and allow cooling.
6	2026-03-13 13:42:03	76.7%	85.7%	HIGH	77%	Reduce heavy processes and allow cooling.
7	2026-03-13 11:05:09	57%	59.5%	MEDIUM	57%	Close background apps and keep ventilation clear.
8	2026-03-13 11:05:00	47.1%	60%	MEDIUM	47%	Close background apps and keep ventilation clear.
9	2026-03-13 11:04:58	64.8%	60.1%	MEDIUM	65%	Close background apps and keep ventilation clear.
10	2026-03-13 11:04:54	93.4%	59.8%	OVERLOAD	93%	Immediately stop high-load programs and check cooling.
11	2026-03-13 11:04:51	50.6%	57.8%	MEDIUM	51%	Close background apps and keep ventilation clear.
12	2026-03-13 11:04:15	65.4%	57.4%	MEDIUM	65%	Close background apps and keep ventilation clear.
13	2026-03-13 11:04:12	41.6%	57.1%	MEDIUM	42%	Close background apps and keep ventilation clear.
14	2026-03-13 11:04:09	48.4%	56.8%	MEDIUM	48%	Close background apps and keep ventilation clear.

Fig 9 Saved Overheating Records

Figure 9 shows the historical overheating records stored in the system database. The table includes parameters such as time, CPU usage, RAM usage, thermal level, risk percentage, and precaution.

• *From the Observed Data:*

- ✓ CPU values above 90% correspond to Overload conditions with risk levels above 90%
- ✓ CPU values between 70%–85% correspond to High risk
- ✓ Moderate CPU values correspond to Medium risk levels

This indicates a strong correlation between CPU usage and thermal risk. The system effectively categorizes different levels of overheating and stores them for future analysis.

• *The Stored Records Help in:*

- ✓ Identifying recurring patterns
- ✓ Analyzing system performance
- ✓ Improving prediction accuracy

➤ Accuracy and Variance Analysis

The machine learning model demonstrates reliable performance in classifying system states into Low, Medium, High, and Overload categories. The predictions closely match the actual system behavior observed from input parameters.

• The System Effectively Handles Variations in System Metrics:

- ✓ CPU usage exhibits high variance due to fluctuating workloads
- ✓ RAM usage shows low variance, indicating stable memory utilization

Despite these variations, the model maintains consistent prediction performance, particularly in identifying high-risk conditions such as “High” and “Overload.”

Overall, the system provides accurate, stable, and real-time predictions, making it suitable for practical deployment.

V. CONCLUSION

The Thermal Guard system demonstrates a practical and intelligent solution for system overheating detection and prevention. By combining machine learning techniques with real-time monitoring and automated notification systems, the platform provides a comprehensive framework for analyzing and managing thermal risks in computing devices. Unlike traditional monitoring tools that only display raw hardware statistics, Thermal Guard actively interprets system data and provides meaningful insights through prediction and visualization.

The integration of AI-based prediction enables the system to analyze workload parameters and classify system states into multiple overheating levels. This predictive capability allows users to understand potential risks before they escalate into critical hardware issues. In addition, the real-time monitoring module continuously tracks CPU and RAM utilization, ensuring that users receive immediate updates about system performance.

Another important contribution of the system is the implementation of database-based event storage. By maintaining historical records of overheating incidents, the framework allows users to analyze patterns and identify recurring causes of thermal stress. This feature is particularly useful for long-term system maintenance and performance optimization.

The inclusion of automated email alerts further enhances the practical usefulness of the platform. Immediate notifications ensure that users can respond quickly when the system reaches dangerous thermal states. Combined with visual analytics such as line charts, circular risk charts, and event logs, the system provides a clear and comprehensive understanding of system behaviour.

Overall, Thermal Guard demonstrates how artificial intelligence and real-time analytics can be applied to improve device safety, reliability, and performance management. The system not only detects overheating conditions but also provides actionable insights that help users maintain optimal system health.

FUTURE ENHANCEMENT

Although Thermal Guard successfully demonstrates intelligent overheating monitoring and prediction, several improvements can be made to enhance its capabilities further. One of the most significant enhancements would be the integration of direct hardware temperature sensors. Currently, the system primarily relies on CPU and RAM usage as indicators of overheating risk. By incorporating hardware-level temperature sensors such as CPU core temperature, GPU temperature, and motherboard temperature, the system could achieve more accurate thermal analysis.

Another potential improvement involves the adoption of advanced deep learning techniques for predictive analysis. Time-series forecasting models such as Long Short-Term Memory (LSTM) networks or Gated Recurrent Units (GRU) could be used to analyze continuous monitoring data and predict future overheating conditions more accurately. These models are particularly effective in capturing temporal dependencies in system workload patterns.

The platform can also be expanded to support multi-device monitoring. In many real-world environments, such as data centers or corporate networks, multiple systems operate simultaneously. A centralized monitoring system capable of collecting and analyzing data from several devices would allow administrators to monitor overall system health across an entire network.

Another promising enhancement is the development of a mobile application that provides remote monitoring capabilities. With a mobile interface, users could receive notifications, view live dashboards, and analyze overheating trends directly from their smartphones. This would significantly improve accessibility and usability.

Future versions of Thermal Guard may also incorporate cloud-based data storage and analytics. By storing monitoring data in cloud platforms, the system could perform large-scale analysis and long-term trend detection. Cloud integration would also enable advanced features such as predictive maintenance recommendations and anomaly detection using large datasets.

Additional features such as SMS alerts, role-based administrator dashboards, automated cooling control suggestions, and integration with IoT-based cooling systems could further enhance the system. These developments would transform Thermal Guard from a prototype academic project into a scalable industrial-grade monitoring solution.

REFERENCES

- [1]. Zhang, Y., Li, H., and Wang, X., "Machine Learning-Based Thermal Prediction for Data Center Servers," *IEEE Access*, Vol. 11, pp. 34562-34572, 2023.
- [2]. Kumar, R., Sharma, P., and Singh, A., "Real-Time CPU Temperature Monitoring and Prediction Using Machine Learning Techniques," *Journal of Computer Systems and Applications*, Vol. 15, No. 2, pp. 120-129, 2024.
- [3]. Liu, J., Chen, Q., and Zhao, Y., "Deep Learning-Based Thermal Anomaly Detection in High-Performance Computing Systems," *IEEE Transactions on Computers*, Vol. 72, No. 5, pp. 1456-1467, 2023.
- [4]. Pang, L., Luo, C., and Pan, W., "Research on the Impact of Indoor Control Quality Monitoring Based on Internet of Things," *IEEE Access*, Vol. 11, pp. 45671-45681, 2023.
- [5]. Singh, V., Gupta, N., and Verma, S., "Predictive Monitoring of System Overheating Using Artificial Intelligence," *International Journal of Advanced Computer Science and Applications*, Vol. 14, No. 3, pp. 210-218, 2023.
- [6]. Wang, T., Zhou, Y., and Li, M., "Thermal Management in Data Centers Using AI-Driven Predictive Models," *IEEE Transactions on Cloud Computing*, Vol. 12, No. 1, pp. 112-123, 2024.
- [7]. Chen, L., Huang, X., and Yang, J., "Intelligent System Health Monitoring Using Machine Learning and Sensor Data," *Sensors Journal*, Vol. 23, No. 6, pp. 3342-3351, 2023.
- [8]. Gupta, R., Mehta, P., and Jain, A., "Anomaly Detection in System Performance Using Random Forest and Neural Networks," *International Journal of Computer Applications*, Vol. 185, No. 12, pp. 10-17, 2024.
- [9]. Patel, K., Shah, R., and Patel, H., "Real-Time Hardware Monitoring and Fault Prediction Using Data Analytics," *Journal of Intelligent Systems*, Vol. 33, No. 4, pp. 556-565, 2024.
- [10]. Zhao, D., Liu, K., and Wang, S., "AI-Based Predictive Maintenance for Computing Infrastructure," *IEEE Access*, Vol. 12, pp. 26789-26799, 2024.
- [11]. Rahman, M., Islam, S., and Hossain, M., "A Smart Monitoring Framework for Detecting Hardware Failures in Computer Systems," *Future Generation Computer Systems*, Vol. 144, pp. 90-101, 2023.
- [12]. Kumar, S., Patel, R., and Desai, P., "Temperature Prediction and System Reliability Analysis Using Machine Learning," *International Journal of Engineering Research & Technology*, Vol. 13, No. 5, pp. 876-882, 2024.
- [13]. Al-Fuqaha, A., Guizani, M., and Mohammadi, M., "Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications," *IEEE Communications Surveys & Tutorials*, Vol. 25, No. 2, pp. 1387-1425, 2023.
- [14]. Lin, C., Zhang, H., and Xu, J., "Energy-Efficient Thermal Management for High-Performance Computing Systems," *Journal of Parallel and Distributed Computing*, Vol. 178, pp. 45-56, 2024.
- [15]. Sun, Y., Li, Q., and Chen, Z., "Deep Learning Approaches for Predicting System Performance and Thermal Behavior," *ACM Computing Surveys*, Vol. 56, No. 7, pp. 1-32, 2024.