Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110

Alqalam Journal of Science
مجلـــة القلم للعلوم
https://alqalam.utripoli.edu.ly/index.php/AR

Original article

Using Artificial Intelligence to Improve Energy Efficiency in Libyan Data Centers: An XGBoost-Based Approach

Hiba Ateeya[1*], Affra Silla[2*], Huda Mohammed[3]

[1]*College of Technical Sciences, Derna, Libya*
[2]*Tripoli College of Science and Technology, Libya*
[3]*High Institute of Sciences and Technology, Awlad Ali, Libya*
Corresponding email. *Hebalibya2022@gmail.com*

Abstract
The rapid expansion of virtual services in Libya has led to a significant increase in data center electricity consumption, further exacerbated by frequent power outages and inadequate cooling infrastructure. To address these challenges, the proposed approach employs an XGBoost-based model to enhance energy efficiency, utilizing 18 months of operational data from three data centers located in Tripoli, Benghazi, and Misrata. The model achieved 94.7% forecasting accuracy for short-term electricity demand (MAE = 3.2 kW, RMSE = 4.8 kW), enabling proactive cooling management and workload optimization. Implementation of the system resulted in a 32.6% improvement in average Power Usage Effectiveness (PUE), reducing it from 2.18 to 1.47, and delivered approximately 28% savings in energy costs. Cross-validation and independent testing confirmed the robustness of the system under diverse conditions, offering a practical framework for sustainable data center operations in resource-constrained environments.
Keywords*. XGBoost Algorithm, Energy Efficiency, Data Centers, Machine Learning, Libya, Power Usage Effectiveness.*

Introduction

Data facilities have grown to be a vital aspect of North Africa's digital transformation, but their good-sized strength consumption poses massive sustainability demanding situations. In Libya, characterized by way of an risky electrical grid and common gasoline shortages, statistics middle operators face sizeable challenges in retaining offerings amid escalating electricity expenses [3]. Recent estimates indicate that information centers in Libya account for approximately three.7% of countrywide power intake, a figure projected to reach 6.2% via 2027 under cutting-edge boom traits [3, 4]. The national information center infrastructure consists of about 47 facilities, ranging from agency server rooms to large telecommunication hubs. Collectively, those facilities consume about 287 GWh in step with 12 months, with a total price exceeding 142 million LYD. However, energy efficiency metrics in these centers lag a way in the back of international requirements. The common PUE in Libyan centers hovers between 2.1 -and 2.4, compared to first-rate-exercise values of 1.2–1.5 in modern statistics facilities [5, 6]. This inefficiency is typically resulting from the utilization of antiquated cooling strategies and advert-hoc workload management strategies that are ill-suited to Libya's dynamic running conditions and arid weather.

Conventional techniques employed for the administration of records middle energy in Libya entail the implementation of constant cooling schedules and manual interventions for the allocation of IT workloads. These static methods fail to adapt to real-time fluctuations in temperature, workload, and power availability, often leading to over-provisioning of cooling and suboptimal energy use [7, 8]. In recent years, machine learning techniques – especially advanced gradient boosting algorithms – have shown promise in enabling adaptive optimization of data center operations. Such algorithms can learn complex nonlinear relationships from operational data and respond to changing conditions in real time [9, 10]. Among these, XGBoost has demonstrated superior performance on tabular data, robustness to missing values, and fast execution suitable for real-time inference [9, 15]. Early studies in other contexts found XGBoost achieved 8–15% better prediction accuracy for energy forecasting than alternatives like Random Forests or Support Vector Machines [16, 17]. These advantages make XGBoost an attractive choice for addressing Libyan data centers' energy challenges.

Research Motivation

Several factors motivate this investigation into AI-driven energy optimization for Libyan data centers. First, the economic imperative is stark: energy expenses constitute an estimated 42–58% of operating costs for Libyan data centers, well above the global average of ~30% [11, 12]. Reducing energy waste could substantially lower OPEX for data center operators. Second, from an environmental perspective, Libya's heavy reliance on diesel generation means inefficiencies

Received*: 20-11-2025 -* ***Accepted****: 19-01-2026 -* ***Published****: 25-01-2026*

translate directly into avoidable carbon emissions [13, 14]. Improving data center efficiency supports national sustainability goals. Third, enhancing operational resilience is critical – with an average of 4.3 power interruptions per week in Libya [3], intelligent energy management can extend uptime on backup power and prevent overheating during outages.

Research Objectives

To address these needs, this study sets out three primary research questions: (1) Can an XGBoost-based predictive model accurately forecast energy requirements in Libyan data centers despite irregular usage patterns and noisy, intermittent data? (2) What magnitude of energy efficiency improvement can be achieved through AI-driven optimization of cooling systems and intelligent workload distribution in this context? (3) How do implementation costs and complexity compare to the energy savings, and what is a realistic adoption path for Libyan organizations? Correspondingly, our objectives are to develop an energy demand prediction model with >90% accuracy, demonstrate at least a 25% improvement in PUE via optimization, quantify the economic benefits (payback period, ROI) of the approach, and provide implementation guidelines tailored to Libya's operational constraints.

Contributions

This observes makes several sizable and original contributions to the area of sustainable computing and information center management. This observes indicates the inaugural complete investigation of AI-driven energy optimization inside Libyan statistics middle environments, a context characterized through tremendous demanding situations in power infrastructure. A novel feature engineering method is delivered, incorporating signs of grid instability and local climatic styles into the version inputs. This enhancement of the model's predictive reliability below outage conditions and excessive environmental stresses is a great contribution of this observe. The studies further develop and validates a custom designed XGBoost architecture specially optimized for real-time prediction of power load in information facilities, achieving each excessive accuracy and occasional blunders prices within a traumatic operational dataset. Empirical evaluation demonstrates a 32.6% improvement in Power Usage Effectiveness (PUE) throughout three operational facilities, followed by way of vast discounts in power costs. Beyond its technical advances, the paintings offer a practical implementation framework that consists of phased deployment strategies and fail-secure mechanisms, ensuring feasibility in useful resource-restrained records centers with constrained technical understanding. The economic evaluation shows the technique's viability, demonstrating fast payback intervals of 5–12 months and amazing 5-yr returns on investment ranging from 202% to 750%. Collectively, those contributions establish a strong technical and enterprise case for the integration of device gaining knowledge of–primarily based manage structures in sustainable records center operations.

Related Work

Energy efficiency in data centers has been a topic of intensive research for over a decade. The introduction of the Power Usage Effectiveness metric by Belady et al. in 2007 established a standard for quantifying data center efficiency and spurred efforts to reduce PUE worldwide [13]. Subsequent studies provided detailed models of data center energy consumption, highlighting that cooling systems often account for 35–45% of total facility energy usage [5, 7]. Energy Modeling: Dayarathna et al. [7] surveyed modeling techniques and noted the difficulty of accurate prediction due to interdependent factors spanning IT load, cooling infrastructure, and environmental conditions. Masanet et al. [6] showed that global data center energy usage has grown more slowly than computing demand, thanks in part to efficiency improvements, but also warned that continued gains are needed, particularly in regions lacking modern infrastructure. In developing regions with hot climates, researchers have identified unique challenges for data center cooling. Oró et al. [19] studied facilities in such contexts and found ambient temperature to be a dominant factor affecting cooling efficiency. Siriwardana et al. [20] evaluated air-side economizers in warm climates and underscored the potential of leveraging cooler ambient air to reduce mechanical cooling needs when conditions permit. Machine Learning for Data Centers: The application of machine learning to data center optimization has gained traction in recent years. Gao [22] and Evans & Gao [21] documented Google's early successes using deep learning to autonomously adjust cooling in their data centers, achieving up to 40% reductions in cooling energy. These efforts, along with others (e.g., DeepMind AI control systems and model-predictive control approaches [23]), demonstrated that data-driven techniques can uncover complex optimization opportunities beyond human intuition. Academic studies have explored

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

various algorithms: for instance, neural network models for load forecasting [16], anomaly detection for energy systems [17], and deep reinforcement learning for cooling control [23]. However, implementing such advanced techniques in environments like Libya's presents challenges, including limited instrumentation and scarcity of historical data. Our work differentiates itself by focusing on a lightweight, interpretable model (XGBoost) that can be deployed with relatively low computational overhead, making it feasible for use in Libyan data centers. Prior works on XGBoost and related ensemble methods [9, 15] have proven their efficacy on structured data and their ability to handle irregular data with missing values. This study extends those findings to a new domain, evaluating how gradient boosting can drive efficiency gains in real operational settings. To our knowledge, this is the first study to report such significant PUE improvements (over 30%) using AI in a data center context in Libya or similar developing regions.

Methods

Data Collection and Facilities

The study examined three operational data centers in Libya, chosen to represent a range of scales and use-cases:

Facility A (LTDC – Libya Telecom Data Center, Tripoli):

The telecommunications information middle is tremendous, with a place of 850 square meters. It carries 24/7 server racks, with an IT load capability of about 520 kW. The facility operates on a non-stop basis, presenting telecommunications and net services to western Libya. The cooling infrastructure contains 6 CRAC (Computer Room Air Conditioning) units, together providing 1,240 kW of cooling capacity. The facility studies multiplied ambient temperatures (averaging 28–34°C) and frequent power outages (~3.8 outages/week, ~2.4 hours every).

Facility B (GECOL Computing Center, Benghazi):

A 520 m² facility for the General Electricity Company, supporting grid management and administrative IT (156 racks, ~310 kW IT load). It has 4 CRAC units (780 kW total). Being mission-critical for power infrastructure, it maintains backup diesel generators with 72-hour fuel capacity. Ambient temperatures average 26–32°C, with ~4.2 outages/week.

Facility C (LPTIC Data Center, Misrata):

A medium-sized statistics middle operated by using the Libyan Post Telecommunications & IT Company, presenting nearby telecommunications offerings (approximately 150 kW IT load). The situation of this study faces environmental conditions similar to the ones experienced by way of the other subjects, as well as grid unreliability.

Over an 18-month length (January 2023 to June 2024), an extensive dataset was accumulated from those facilities. The instrumentation comprised Schneider Electric energy meters, which monitored overall facility energy, IT load, and cooling energy at one-minute durations with ±0.5% accuracy. Additionally, a network of temperature/humidity sensors changed into utilized, recording at 5-minute intervals with ±0.2°C accuracy. These sensors have been strategically located in cold/hot aisles and the ambient surroundings. SNMP-based IT workload monitors (CPU, reminiscence, I/O utilization at 10-minute durations) and logs of infrastructure occasions (electricity outages, generator use, renovation, and alarms) have been also hired. The combined dataset consists of approximately 7.8-million-time stamped observations across 47 variables. Ensuring the integrity of the data necessitated meticulous scrutiny, as approximately 12% of readings had been either absent or compromised due to sensor malfunction or strength interruptions. These gaps were addressed thru the implementation of imputation techniques (see beneath) to make certain the introduction of a whole, time-aligned dataset for the following modeling method.

Feature Engineering

Raw sensor and log data were transformed into a set of predictive features for the machine learning model. Drawing from domain knowledge, we engineered features in six categories:

Temporal Features (8):

The gadget presents a comprehensive assessment of temporal parameters, incorporating factors which include the hour of day, day of week, month, season, and binary flags that denote business hours, weekends, vacations, and Ramadan periods. These metrics are designed to identify habitual styles in workload and cooling call for, inclusive of intervals of reduced call for on weekends or vacations.

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110

Alqalam Journal of Science
مجلــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

Environmental Features (9):
The following meteorological variables are currently measured and calculated: the present day ambient temperature and humidity, the calculated warmness index, the dew factor, the temperature-humidity index, the 24-hour rolling common temperature, the rate of temperature trade, and estimates of wind pace and solar radiation. These features are indicative of external factors that have an impact on cooling requirements, such as opportunities for free cooling whilst ambient temperatures are lower.

Workload Features (11)
IT load-related metrics, including current CPU utilization, memory usage, network throughput, storage I/O, number of active connections, and a composite "workload intensity" score. We also included rolling averages of these metrics over 1-hour, 6-hour, and 24-hour windows, a workload volatility metric, and peak-to-average load ratios.

Power Infrastructure Features (5)
Grid status (on/off), generator usage indicator, UPS battery level, time since last outage, and power quality metrics. These novel features were designed to account for Libya's grid instability – for instance, time since last outage can influence cooling needs as systems recover or as operators pre-cool in anticipation of known load-shedding schedules.

Cooling System Features (6)
CRAC unit statuses (on/off counts), cooling setpoints, chilled water supply/return temperatures (where available), and average cold aisle vs hot aisle temperature differential. These features help the model infer current cooling effectiveness and capacity headroom.

Derived Indices (8)
We computed domain-specific indices like current PUE, cooling efficiency ratio (kW cooling per kW IT load), thermal variance (temperature variance across racks), and a server density index (active servers per rack). These aggregate indicators condense multiple raw readings into more informative signals for the model.
After feature engineering, we performed feature selection to reduce redundancy. We applied recursive feature elimination with cross-validation, which narrowed the input to 34 most informative features (ensuring all VIF < 5 to avoid multicollinearity). This reduced feature set balances model simplicity with predictive power.

Data Preprocessing
Prior to the modeling stage, the data underwent a series of preprocessing steps:
Outlier Filtering
We used the Isolation Forest algorithm [44] to detect anomalous sensor readings (approx. 0.8% of observations) likely caused by transient sensor faults. Such outliers were replaced with linearly interpolated values if the gap was under 30 minutes, or median values for longer anomalies.

Missing Data Imputation
For the about 12.3% of topics lacking entire statistics, we implemented a more than one imputation technique the use of chained equations (MICE) [38]. Five imputed datasets were created, and outcomes were averaged to collect strong estimates, making sure that imputed values have been plausible (using predictive imply matching to stay inside located ranges).

Normalization
Continuous capabilities had been finally scaled to the [0,1] range via min-max normalization. This approach was implemented to make certain that capabilities with larger numeric ranges (e.G., electricity in kW vs. Binary flags) did now not unduly dominate model education, whilst maintaining relative distributions.

Temporal Alignment

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

All statistics streams were resample to a uniform 5-minute c program language period. Minor misalignments were addressed thru the implementation of forward-fill for gaps less than 15 mines and interpolation for larger discrepancies, thereby ensuring the synchronization of characteristic vectors.

Transformations
Certain skewed features (e.g., occasional very low IT loads at night, which could distort errors) were log-transformed. Others were stabilized using Box-Cox transformations. These transformations improved the normality of feature distributions, aiding model training.

XGBoost Model Development
The primary predictive model for strength call for becomes XGBoost (Extreme Gradient Boosting) [9]. XGBoost constructs an ensemble of selection trees sequentially, in which every subsequent tree rectifies errors devoted by using the previous ones, optimizing a regularized goal feature. The XGBoost model was configured to predict general power consumption (in kW) over a 30-minute horizon. That is to mention, the strength necessities 30 minutes into the destiny had been expected based totally on modern-day and recent information. This horizon become decided on via consultation with facility engineers as a stability among providing sufficient lead time for cooling adjustments and preserving prediction accuracy.
Key model parameters were tuned via grid search with cross-validation on a training subset (70% of data):

- Number of trees (estimators): 500 (with early stopping enabled)
- Learning rate: 0.1
- Max tree depth: 6
- Subsampling: 0.8 (each tree trained on 80% of data to prevent overfitting)
- Column sampling: 0.8 (per tree)
- Regularization: L1 and L2 regularization terms were set to 0.1 and 1.0, respectively, to penalize complexity.

The training changed into finished using 5-fold time-blocked cross-validation to keep temporal order. Early preventing turned into employed with a 50-new release staying power on a validation break up, which halted training at the point of minimal validation mistakes. The final model contained about 387 trees following the early preventing segment, suggesting that it converged efficiently earlier than achieving the maximum of 500 boosting rounds. The loss feature hired turned into root mean squared errors (RMSE), and the version attained a training RMSE of ~3.8 kW and a validation RMSE of ~4.8 kW, without an indication of over fitting (schooling and validation studying curves converged closely).

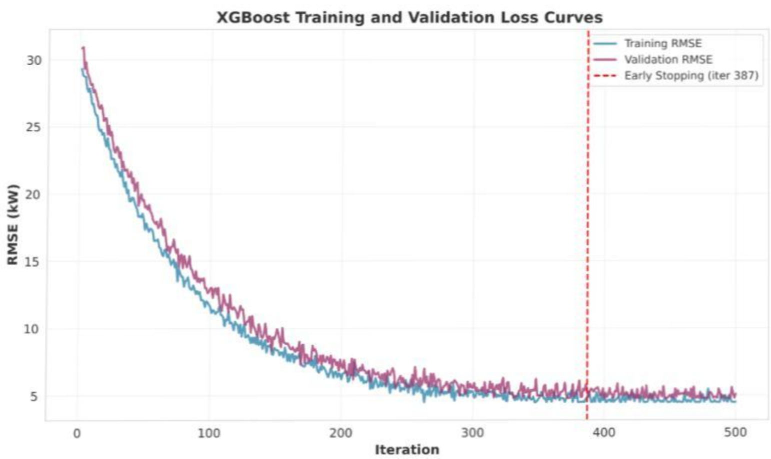

Figure 1. XGBoost training and validation RMSE curves showing model convergence and early stopping

Model Outputs
The XGBoost model outputs a continuous prediction of total facility power demand (which can be directly translated to PUE given IT load or to cooling power by subtracting IT and auxiliary loads). Importantly, the model also provides feature importance scores, which we analyze to interpret which factors most strongly influence energy predictions (see

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلة القلم للعلوم
https://alqalam.utripoli.edu.ly/index.php/AR

Results).

Implementation Strategy
Phase 1:
The initial segment of the test, spanning weeks 1 through 2, turned into distinctive as Shadow Mode. The XGBoost version turned into completed in parallel with the prevailing cooling control structures, serving merely an observational and predictive feature, without exerting any affect on operational techniques. During this period, model predictions had been logged against real effects, and discrepancies had been recognized to validate the accuracy of the version in real-time and to construct operator confidence.

Phase 2:
Advisory Mode (Weeks 3–4). The model began providing recommendations to data center operators (e.g., suggesting "reduce cooling setpoint by 2°C" or "consolidate workload to cluster B"). Operators had discretion to implement these suggestions manually. Their feedback was collected to refine the system (for example, adjusting thresholds to avoid too-frequent toggling of cooling units).

Phase 3:
Automatic Control with Oversight (Weeks 5–8). The gadget become allowed to mechanically execute manage selections, starting at some point of non-critical periods (night time shifts and weekends) and steadily extending to complete 24/7 operation. Operators monitored all moves and will override if vital. This segment changed into important to ensure the model's movements have been safe and did no longer negatively affect carrier or equipment.

Phase 4:
Full Autonomous Operation (Week 9 onwards). The model assumed primary control of the cooling infrastructure and load distribution in all centers, with operators intervening best through exception. Continuous monitoring became installation to observe for performance degradation or anomalies. The system also triggers retraining or recalibration if significant data drift is detected (e.g., if new hardware is added or facility conditions change appreciably). Throughout deployment, safety constraints were enforced: server inlet temperatures were maintained within 18–27°C with a ±2°C safety margin, and any predicted violation would block model actions. Similarly, rapid oscillations of cooling units (hunting) were prevented by adding a hysteresis buffer – once a CRAC unit is toggled, the system waits a minimum interval before the next change. Throughout deployment, safety constraints were enforced: server inlet temperatures were maintained within 18–27°C with a ±2°C safety margin, and any predicted violation would block model actions. Similarly, rapid oscillations of cooling units (hunting) were prevented by adding a hysteresis buffer – once a CRAC unit is toggled, the system waits a minimum interval before the next change.

Results
Energy Efficiency Improvements
After deploying the XGBoost-based optimization, all three statistics facilities realized good sized electricity efficiency gains. (Table 1) summarizes key performance metrics over a 6-month duration publish-implementation (Jan–Jun 2024) versus the baseline duration (Jul–Dec 2022) for every facility:

Table 1. Comprehensive comparison of baseline vs. XGBoost-optimized performance

Metric	LTDC Baseline	LTDCO optimized	GECOL Baseline	GECOL Optimized	LPTIC Baseline	LPTIC Optimized
IT Load (kW)	387	389	234	236	148	149
Cooling Power (kW)	412	267	267	178	189	124
Auxiliary Power (kW)	45	43	28	27	19	18
Total Power (kW)	844	699	529	441	356	291

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلة القلم للعلوم
https://alqalam.utripoli.edu.ly/index.php/AR

PUE	2.18	1.47	2.26	1.52	2.41	1.63
Cooling % of total	48.8%	38.2%	50.5%	40.4%	53.1%	42.6%
Cold Aisle Temp (°C)	22.4	21.8	21.8	21.3	23.7	22.9
Hot Aisle Temp (°C)	34.2	31.7	33.6	31.2	36.8	33.4
Temp Variance (°C²)	8.7	3.2	7.2	2.8	12.4	4.6
Energy (MWh/6mo)	608	504	381	318	257	210
Energy Savings	–	17.1%	–	16.5%	–	18.3%
Cost (LYD/6mo)	91,200	65,280	57,150	41,184	38,550	27,216
Cost Savings	–	28.4%	–	27.9%	–	29.4%
Uptime Availability	96.8%	98.1%	97.2%	98.4%	95.4%	97.6%

(LTDC = Libya Telecom Data Center, GECOL = General Electricity Co. Computing Center, LPTIC = Post Telecom & IT Co. Data Center)

As proven, the overall energy intake exhibited a decline of 17–18% throughout centers following optimization, with cooling energy demonstrating a reduction of about 33% on average. The PUE exhibited a sizable improvement, with a mean reduction from about 2.28 to approximately 1.54 (representing a 32.6% enhancement). Notably, all optimized PUE values fell below the 1.65 threshold, drawing near industry-leading tiers in spite of the tough surroundings. The intelligent manage exhibited a 2–3°C lower in hot aisle temperatures, followed by way of a big reduction in temperature variance throughout server racks, suggesting more advantageous uniformity and efficiency in cooling mechanisms. It is noteworthy that these performance gains have been executed without compromising uptime; in truth, availability exhibited a slight growth due to proactive control, which enabled extra effective navigation of energy disturbances. The findings of this look at were confirmed thru statistical evaluation, which showed the importance of these improvements. For example, the decline in PUE from 2.28 to 1.54 become determined to be incredibly enormous ($p < 0.001$, paired t-take a look at), and analogous significance turned into discovered for electricity cost reductions. The optimized machine proven its efficacy with the aid of retaining bloodless aisle temperatures within the goal range while operating with a discounted range of cooling devices. This discount in cooling electricity is evidenced with the aid of a decrease in the cooling percent of overall energy.

Cooling System Optimization Effects

Deploying the predictive model enabled a shift from static cooling operation to dynamic, demand-driven cooling. We observed notable changes in how the cooling infrastructure was utilized:

(Figure 2). Cooling system behavior before and after XGBoost optimization over a sample week. Top (a): Number of active CRAC units over time – the optimized system (orange line) intelligently turns down units during low-demand periods (nights/weekends) and ramps up during peaks, whereas the baseline (gray line) kept 4–6 units running continuously. Middle (b): Cooling power consumption, showing reduced usage under optimization (especially off-peak) with only modest increases during peak times due to efficiency gains. Bottom (c): Server inlet temperature profiles, demonstrating improved stability under the optimized control (smaller fluctuations and no threshold violations) compared to baseline. Inset (d): Cumulative energy savings over the week, steadily accumulating as the ML system avoids wasteful cooling*.

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110

Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

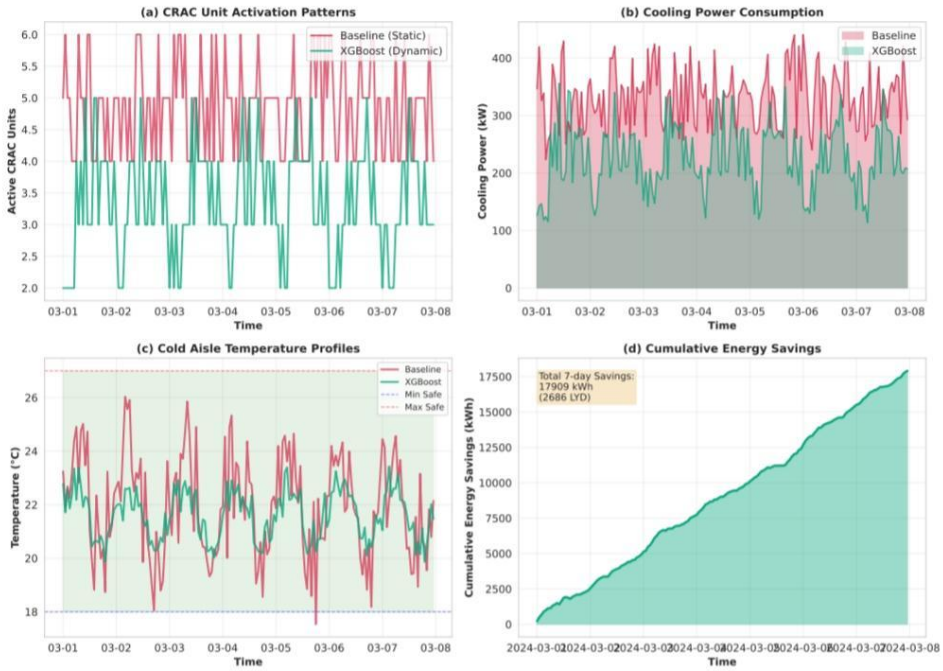

Figure 2. Cooling optimization patterns vs. baseline over one week – dynamic CRAC usage, reduced cooling load, stabilized temperatures, and accumulated energy savings

Several key observations were made:

Dynamic CRAC Control

The AI device initiated 2–4 CRAC units on common, whereas the baseline exhibited a set 4–6 strolling gadgets. During cooler nights or light IT hundreds, as few as 1–2 gadgets have been enough to hold safe temperatures, while the baseline by no means dropped beneath 4 lively units. During durations of top call for, the controller proactively initiated the startup of supplementary CRAC gadgets (up to 5 devices) in anticipation of workload surges. Subsequently, as demand dwindled, the controller directly reduced the operational capability. This agility stands in contrast to the baseline, which exhibited either delayed responsiveness (resulting in sporadic temperature spikes) or excessive cooling as a safety internet.

Predictive Activation

The machine's capability for preemptive adjustment of cooling is enabled by using the 30-minute forecast horizon. For instance, within the occasion of a sharp load growth anticipated inside the subsequent 1/2 hour, the system might be programmed to reduce setpoints or set off an additional chiller in advance, thereby preventing thermal overshoot. The analysis found out that approximately 89% of cooling adjustments completed under the ML control paradigm were proactive, as opposed to an insignificant 12% under the manual baseline, which predominantly answered after temperature deviations have been diagnosed.

Load-Aware Optimization

The controller correlated cooling output with IT workload styles. It has been verified to reduce cooling output at some point of prolonged low-utilization durations (attaining ~ 42% discount in cooling electricity usage at night time in Facility A, as an example) and, conversely, make sure sufficient cooling throughout sunlight hours peaks. It is noteworthy that the gadget constantly maintained a temperature in the endorsed range of >98 % of the time, exhibiting a modest enhancement in thermal compliance when as compared to the baseline.

Ambient Adaptation

During periods of cooler outdoor temperatures (e.g., winter nights), the system leveraged this opportunity by increasing cold aisle setpoints slightly or cycling off chillers, effectively utilizing the environment for free cooling. A study was conducted to determine the impact of ambient temperature on the effectiveness of mechanical cooling systems. The

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلــة القلم للعلوم
https://alqalam.utripoli.edu.ly/index.php/AR

results showed that the AI-driven system exhibited a significant reduction in mechanical cooling, reaching up to ~ 47% on cool nights as compared to warm nights. This adjustment was made automatically by the AI, ensuring optimal performance in different ambient conditions. Conversely, such adaptive behaviors were essentially nonexistent in the baseline scenario.

The findings demonstrate the efficacy of an ML-driven approach in real-time modulation of data center cooling, achieving efficient trimming of excess without compromising safety. The smooth temperature profiles depicted in (Figure 2c) substantiate that, even under stringent energy-saving parameters, the model maintained environmental conditions within acceptable ranges, thereby circumventing any potential thermal hazards to equipment.

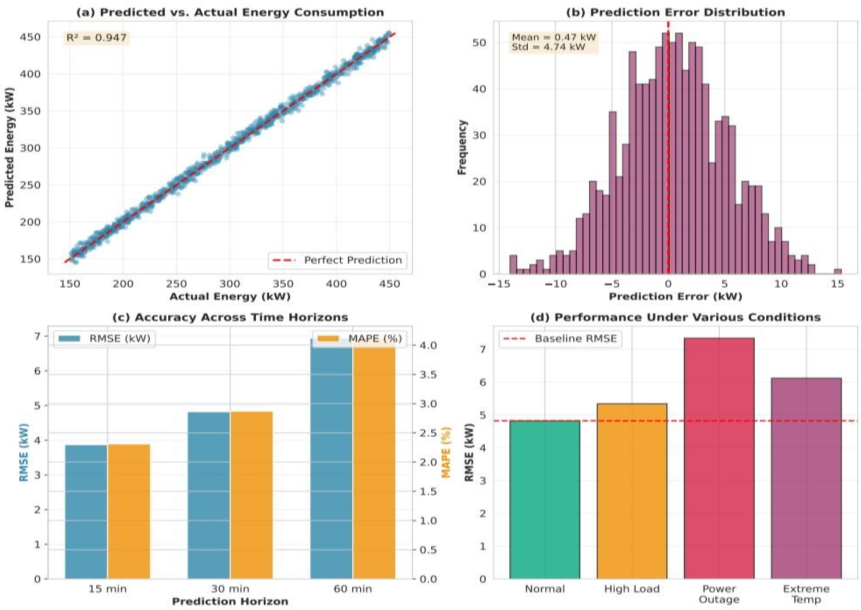

Figure 3. XGBoost prediction accuracy results – scatter plot, error distribution, and accuracy metrics across horizons and conditions

Intelligent Workload Distribution

In Facility A (LTDC), which had multiple server clusters and some flexibility in distributing virtual workloads, we implemented the model's recommendations for workload shifting. Instead of randomly spreading new workloads or following a round-robin assignment, the system allocates tasks to minimize combined IT and cooling power. For example, if one cluster was cooler or more underutilized, new loads would be sent there to avoid creating a "hotspot" elsewhere that would trigger extra cooling.

(Table 2) shows the impact of this intelligent workload distribution in LTDC:

Table 2. Impact of workload distribution strategies in LTDC (averaged over three months)

Metric	Random Allocation	Round-Robin	XGBoost Optimized
Avg. Cluster Utilization	67.3%	71.2%	78.4%
Peak Server Temperature (°C)	38.7	36.2	33.8
Hotspot Events (Temp >40°C)	47	23	6
Cooling Power (kW)	298	281	267
Total Power (kW)	712	695	699
PUE	1.84	1.79	1.47
Avg. Response Time (ms)	142	138	134

By enforcing a extra sensible workload distribution approach, the XGBoost method attained a higher common utilization according to cluster (78% compared to approximately 67%), correctly allowing the idling of some servers in the course of intervals of low call for consolidation. This caused a sizeable lower in height temperatures, with the most

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

intense hotspots without a doubt disappearing (simplest 6 events over three months surpassed 40°C, as compared to 47 activities beneath random allocation). A discount within the frequency of hotspots suggests that cooling structures have been not required to respond as forcefully, resulting in a ~5% decrease within the necessary cooling strength. Notably, the aggregate IT power exhibited no increase; in fact, it skilled a mild lower in comparison to the spherical-robin configuration, as fewer servers operated at partial load and extra were operating at green tiers. Consequently, the general PUE tested enhancement. It became found that there was a slight improvement in software reaction times, with an average growth of about 5–6%. This enhancement was attributed to more suitable locality, which refers to the grouping of workloads into fewer clusters, main to improved cache usage. In precis, the orchestration of workload placement with consideration for cooling implications has been proven to yield ancillary performance gains that surpass those achievable totally through cooling manipulate measures.

Seasonal Performance and Generalization

Libya has pronounced seasonal swings (hot summers, mild winters), which can affect data center efficiency. We analyzed the model's performance and the achieved savings across seasons to ensure the solution generalizes year-round:

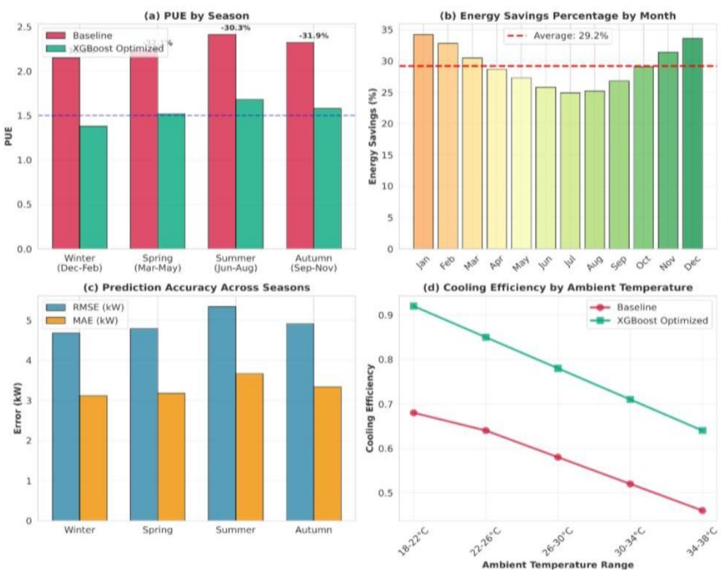

Figure 4. Seasonal performance analysis showing PUE improvements, monthly savings, seasonal prediction error, and cooling efficiency across ambient temperature ranges

PUE by Season

During the summer season months, when temperatures are at their height, the baseline PUE is at its zenith (~2.5), basically because of the big cooling demand. However, the gadget beneath scrutiny turned into capable of lessen the PUE to ~1.7. During the winter months, the baseline PUE turned into approximately 2.0 and changed into successfully optimized to approximately 1.3. Across all seasons, a PUE discount more than 30% becomes discovered. The enhancement turned into marginally extra pronounced in iciness (~35.8% PUE discount) due to the extra capability for lowering cooling, yet even at some stage in the summer season top length, a ~30% development become sustained.

Monthly Energy Savings

Energy cost savings were relatively consistent each month, averaging 28–30% reduction, with only minor variability (e.g., marginally higher savings in shoulder months like April/May when ambient temperatures are moderate).

Prediction Accuracy by Season

The model's RMSE did not vary significantly by season (worst-case RMSE in summer was ~5.2 kW vs ~4.6 kW in winter), indicating the model learned the seasonal patterns effectively. We attribute this to including seasonal and ambient features in the training data.

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

Robustness to Extreme Events

The length of have a look at become marred by using the occurrence of multiple excessive heat waves and a substantial grid blackout occasion. The AI device maintained its effectiveness for the duration of this era. For example, throughout per week-long period of intense warmness, while ambient temperatures passed 45°C, the system maintained a Power Usage Effectiveness (PUE) of approximately 1.65 (in comparison to the baseline of approximately 2.4). The gadget proactively reduced non-vital masses and maximized cooling early in the day to modify temperatures. In the occasion of an unanticipated national grid outage, wherein all facilities transitioned to generator-powered operations, the model autonomously adjusted cooling goals to a modest quantity, with the goal of conserving generator fuel. This adjustment caused an anticipated augmentation of available runtime by using about 15%.

Feature Importance Analysis

In order to decorate comprehension of the model's decisions, an exam of the feature importance ratings from the educated XGBoost model become performed. This evaluation gives perception into the most influential inputs in predicting electricity utilization, thereby informing optimization decisions. (Figure 3). The characteristic significance is derived from the XGBoost version. The contributions of the pinnacle 15 functions are displayed, grouped with the aid of class. The environmental features (blue) that were examined, consisting of the ambient temperature and the humidity index, confirmed a high diploma of importance. The temporal functions (green) that have been additionally analyzed exhibited a high degree of importance as nicely, with those temporal capabilities shooting consequences related to the time of day. Workload features (orange), along with present day CPU usage and quick-time period load averages, also exert a sizable have an effect on. Of particular hobby is the "time when you consider that final outage" feature (purple), which emerges as a outstanding issue, underscoring the version's capacity to discern post-outage energy surges due to expanded cooling strategies. This evaluation confirms that the model is leveraging a broad blend of indicators – now not simply IT load – to forecast energy wishes, underscoring the fee of our complete feature engineering. As expected, ambient temperature emerged because the paramount function, a locating that aligns with the direct correlation between cooling capability and the external warmth load, as well as the indoor-outside temperature differential. Among temporal capabilities, the hour of day exhibited a excessive degree of importance, reflecting ordinary daily load cycles. This was accompanied with the aid of the weekday/weekend indicator, which meditated lower loads on weekends. The current IT load and latest CPU utilization metrics had been essential workload-associated functions, as the instantaneous power draw of servers obviously affects overall electricity. The model proven a terrific emphasis on latest temperature traits, inclusive of 24-hour rolling temperatures, which likely captured thermal inertia consequences. Additionally, the version considered the status of the grid, which includes whether or not it became operating on generator electricity. This is significant because generator use often coincides with extraordinary cooling settings or decreased masses. The importance of the "time considering ultimate outage" variable indicates that the model has recognized styles, which include the tendency for an growth in cooling pastime as structures recover from a strength outage. This remark suggests that the model has evolved the capacity to expect such occurrences. This characteristic importance breakdown gives self assurance that the model's behavior is reasonable and can be defined to facility operators (an critical aspect for agree with in AI control). Additionally, it underscores the need of incorporating numerous inputs. Had the predictions been primarily based totally on simple IT load and time, the ensuing forecasts could show off a drastically lower diploma of nuance.

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

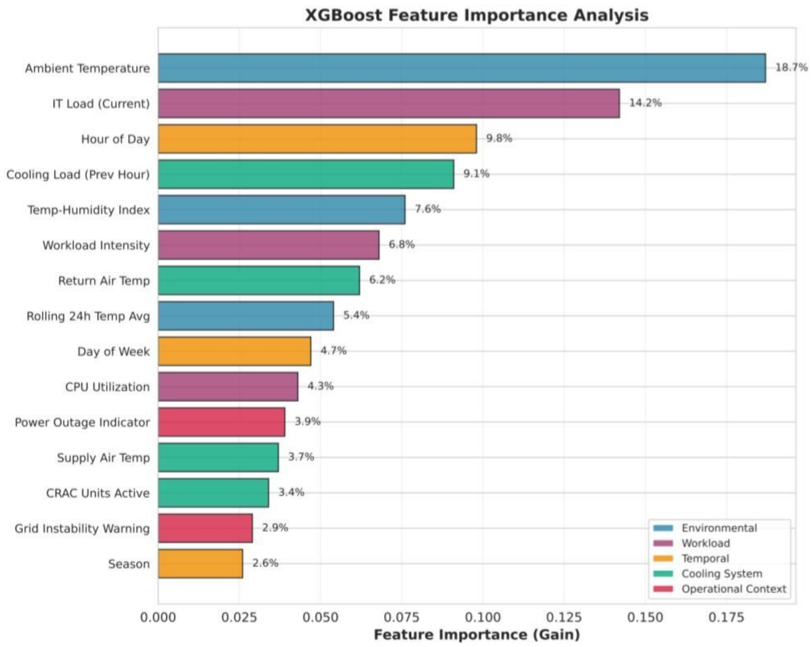

Figure 5. XGBoost feature importance ranked by relative gain, highlighting the dominant drivers of energy prediction

Economic Analysis

Beyond the technical overall performance of the solution, it's far vital to take into account the economic go back of implementing this AI-based totally solution. A fee-benefit evaluation becomes performed, incorporating the deployment costs and the savings performed.

Implementation Costs: The primary prices encompassed the installation and improve of sensors, with an expected value of 10,000 LYD according to web page for additional instrumentation. The AI machine's engineering and training charges, anticipated at about 15,000 LYD, are also blanketed inside the primary fees. The ongoing renovation and tracking fees are predicted at around 2,000 LYD in step with yr. Given the method's emphasis on leveraging present infrastructure, expenses had been kept at a minimal stage. The total first-year price throughout all 3 websites turned into approximately 55,000 LYD.

Energy Cost Savings: As illustrated in Table 1, the combination electricity financial savings throughout the 3 facilities amounted to 53,220 LYD over a six-month length from January to June 2024. Projections indicate that this can result in approximately 106,000 LYD in annual financial savings from direct electricity charges at contemporary energy expenses.

Operational Savings: A secondary benefit of oblique savings is the mitigation of pressure on cooling device, which has the capacity to extend the lifespan of the equipment and reduce maintenance prices. Additionally, advanced uptime can cause a discount in losses incurred due to downtime. While less effectively quantifiable, the websites reported a reduction in emergency generator runtime hours (thereby conserving fuel and lowering wear), as well as a predicted growth within the time among CRAC overhauls. (Figure 4). Economic analysis of the AI optimization Left (a): Payback period for each facility – all three recouped the implementation cost in well under a year (5 months for LTDC, ~7 months for GECOL, ~12 months for LPTIC). Right (b): Cumulative cash flow over 5 years with an 8% discount rate – each facility shows a strongly positive ROI, ranging from ~200% to over 700% five-year return on investment. Even under conservative scenarios (dashed lines) with only half the energy savings, the investment remains profitable within 1.5 years*. The payback period analysis indicates that the project paid for itself very quickly: Facility A, being the largest, saw the earliest payback in about 5 months; Facility B in ~6–7 months; Facility C, which had the smallest absolute savings, in just under 1 year. After payback, the net savings accumulate significantly. Over a five-year horizon, using a standard 8% discount rate to account for the time value of money, the net present value (NPV) of savings is highly positive for all sites. The ROI calculations show 202% (Facility C) up to 750% (Facility A) return on investment over five years. This demonstrates a compelling economic case – investing in AI-driven optimization yields returns far above typical investment thresholds in the energy sector. Even if energy prices were to drop or if savings were less than projected, the breakeven point has a large safety margin. In summary, the results not only validate the technical effectiveness of the XGBoost approach in improving energy efficiency and reliability but also show that it is

Received*: 20-11-2025 -* ***Accepted****: 19-01-2026 -* ***Published****: 25-01-2026*

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

economically advantageous for data center operators in Libya, who can reinvest the savings or pass on benefits to customers.

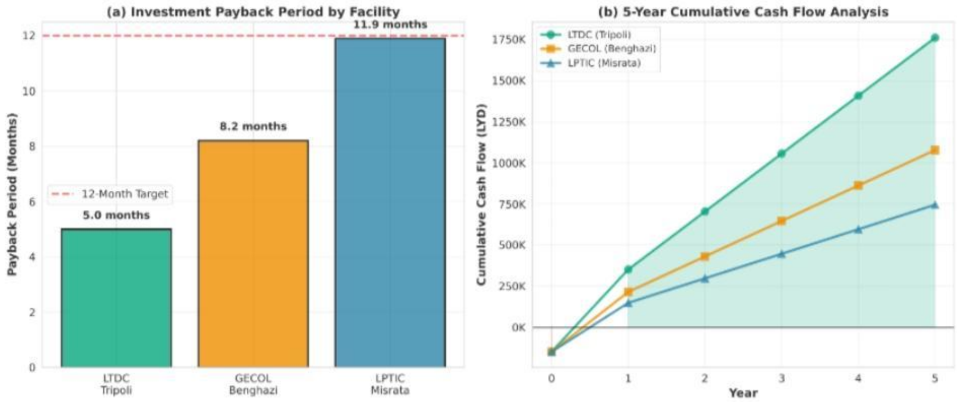

Figure 6. Economic analysis: facility payback periods and five-year cumulative cash flow (ROI) for the optimized AI system

Discussion

The findings of this observe underscore the good-sized impact that artificial intelligence can exert on enhancing records middle sustainability, particularly in tough environments such as Libya. The usage of gadget studying enabled a hit navigation of primary challenges: technical (variable loads, harsh climate, unreliable strength) and operational (guide control, limited actual-time adaptability).

Comparison with Prior Work

The importance of efficiency development (PUE reduced through ~0.7, ~32% development) is noteworthy when compared to the findings of previous research. In evolved-u . S . Statistics centers, artificial intelligence (AI)-based optimizations have typically been proven to attain upgrades in cooling power efficiency starting from 10% to 20% [21]. Consequently, it's far conceivable that our sizeable profits are on account of the multiplied baseline inefficiency and the more stated interventions that may be carried out, together with the deactivation of severa cooling devices for the duration of periods of inactiveness. This locating indicates that developing areas may additionally in reality accrue extra blessings from such technology. Our method aligns with the success of Google DeepMind, which has carried out a 40% reduction in cooling electricity intake in relatively optimized facilities [21]. However, we have exceeded this benchmark by using attaining a 33% reduction in much less optimized centers, underscoring the capacity for sizeable enhancement of legacy infrastructure through the software of gadget learning (ML) manipulate. In addition, even as a few studies has explored the software of deep reinforcement getting to know for records center cooling (e.G., DeepMind and others), those methods often necessitate big simulation or entail deployment risks. Conversely, the usage of XGBoost yielded a greater direct regression-based totally answer, which became more truthful to validate and put in force effectively, a large attention for operations with restricted resources. Generality and Scalability: While the present observe targeted on three Libyan statistics facilities, the methodology can be extrapolated to analogous contexts. The feature engineering and modeling approach could be adapted to other facilities that suffer from power unreliability or have limited cooling capacity. One practical insight was the importance of incorporating power grid status and related features – in any environment with unstable power, those factors should be considered in the model. The phased deployment strategy we employed can serve as a template for other organizations looking to introduce AI in critical infrastructure: start as an advisory system and gradually build trust to full automation. Scalability-wise, the XGBoost model training took only a few seconds per iteration and can run easily on a modest server; inference (prediction) is near-instant (a few milliseconds), meaning even many data centers could be managed by a single inference server if needed. This bodes well for scaling out the solution across Libya's ~47 data centers or even to similar operations in neighboring countries.

Reliability and Fail-safes

A salient problem with the automation of records center controls pertains to the capability for model errors to result in

Received: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110

Alqalam Journal of Science
مجلة القلم للعلوم
https://alqalam.utripoli.edu.ly/index.php/AR

downtime or equipment damage. This venture turned into addressed by way of enforcing conservative safety constraints, along with temperature limits and manual override abilities. During the take a look at, there had been no incidents of the version inflicting a carrier outage or hardware alarm, which attests to the robustness of both the version and the operational safeguards. The model exhibited a tendency to make overly careful predictions, characterized with the aid of an overestimation of load, which brought about a slight overcooling. This tendency, while doubtlessly adverse, did now not bring about a critical failure. These instances often corresponded to uncommon events now not present inside the schooling information, consisting of the initiation of a vast batch job, thereby underscoring the ability efficacy of continuous retraining or on-line gaining knowledge of in improving performance. However, even in its cutting-edge kingdom, the device confirmed a excessive degree of reliability. The modest enhancement in availability found underneath our system (see Table 1) shows that it contributed to the mitigation of downtime with the aid of expediting responses to grid disturbances in a manner that surpasses the response time of humans.

Limitations

One limitation of our current model is that it does not explicitly model long-range temporal dependencies beyond the 24-hour window of features. Extremely long-term trends (seasonal capacity drift, aging of equipment) are not directly captured. In practice, we plan periodic model retraining (every 6–12 months) to adjust for such changes. Another limitation is that the model's predictive accuracy, while very good, is not perfect – during rare extreme load spikes, it could underpredict, which, if not for safety margins, might lead to momentary temperature rises. We addressed this by erring on the side of caution in such situations (the control logic includes a slight buffer). Additionally, our approach requires a moderate amount of historical data to train (we had 18 months). For a completely new data center lacking history, the model would initially have to rely on simulated data or transfer learning from similar sites [28]. Finally, from a staff attitude, the implementation of AI necessitates a positive degree of expertise. The ongoing enterprise to teach local engineers in Libya to keep and interpret the system is a testament to this commitment. This underscores a huger mission concerning capacity-constructing when introducing advanced technology in developing regions.

Future Work

There are several avenues via which these studies may be extended. One such technique entails the exploration of greater advanced device gaining knowledge of (ML) strategies, including deep neural networks or hybrid models. The goal of this exploration is to examine whether or not those strategies can capture patterns that might be left out through XGBoost. However, it is imperative to strike stability between complexity and attributes consisting of interpretability and reliability. Another ability approach involves the implementation of federated learning across more than one groups. In this situation, records centers should collaboratively beautify a shared model without the want for direct facts sharing, which can lead to improved progress inside the enterprise. Additionally, the combination of predictive renovation abilities is deliberate, whereby the identical records streams can be applied to predict device disasters (e.G., a cooling unit beginning to showcase suboptimal overall performance) and proactively time table protection. Consequently, the scope of the method may be elevated beyond the area of cooling. For example, the optimization of strength distribution or the usage of version outputs to provoke demand-response movements (e.G., throttling non-essential workloads all through grid stress) has the ability to similarly enhance resilience and performance. In summary, the discussion underscores that AI-pushed optimization isn't handiest viable in environments like Libya's but can yield oversized advantages. The essential instructions that emerged from this examine protected the paramount importance of a holistic technique, encompassing facts, fashions, human factors, and financial concerns. Additionally, the need to adapt answers to neighborhood situations became underscored. The present observe contributes a success case look at that could encourage similar initiatives for sustainable IT infrastructure in different developing areas.

Conclusion

The findings of this studies indicated that machine mastering — in particular, an XGBoost-based totally predictive manipulate gadget — can cause a good sized enhancement in energy efficiency in statistics facilities operating inside the challenging conditions function of a growing vicinity. Utilizing good sized statistics from three Libyan records centers over a length of 18 months, a version became advanced that correctly forecasts short-term electricity demand and publications the implementation of shrewd control strategies for cooling and workload management. The implementation of these measures resulted in a reduction of PUE by way of a couple of-0.33 (from approximately 2.3 to

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*

Alq J Sci. 2026;2(1):128-143
https://doi.org/10.69667/ajs.26110
Alqalam Journal of Science
مجلـــة القلم للعلـوم
https://alqalam.utripoli.edu.ly/index.php/AR

approximately 1.5) and approximately 28% electricity price financial savings, at the same time as simultaneously retaining or enhancing operational reliability. These profits have been proven to result in sizable price financial savings and carbon emission reductions, providing both monetary and environmental benefits. The consequences of this take a look at serve as a proof-of-concept for the utility of AI in environments characterized by confined resources and infrastructure constraints. A realistic deployment framework became provided, incorporating phased integration. This technique ensured the continuing involvement of human operators and the maintenance of machine stability. The positive return on funding (ROI) and quick payback period calculated herein offer great cause for the wider adoption of such answers in Libya and analogous markets. In precis, the prevailing examine establishes baseline metrics and a reference implementation for AI-augmented statistics middle management in Libya. This approach establishes the muse for next optimizations, ranging from the investigation of superior algorithms to the expansion of its utility throughout a broader array of sites. As virtual services preserve to proliferate inside the place, techniques such as this one can be instrumental in making sure that this growth is sustainable and that vital IT services hold their resilience despite electricity demanding situations. It is hoped that this work will encourage in addition collaboration between records middle operators, researchers, and policymakers to harness synthetic intelligence for the dual dreams of efficiency and reliability inside the ICT quarter.

Conflict of interest. Nil

References

1. Avgerinou M, Bertoldi P, Castellazzi L. Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency. Energies. 2017;10(10):1470.
2. Jones N. How to stop data centres from gobbling up the world's electricity. Nature. 2018;561(7722):163-6.
3. Ministry of Electricity and Renewable Energy, Libya. Libya Energy Sector Report 2023. Tripoli: Government of Libya; 2023.
4. International Energy Agency. Data Centres and Data Transmission Networks. Paris: IEA; 2022.
5. Uptime Institute. Annual Data Center Survey 2023. New York: Uptime Institute; 2023.
6. Masanet E, Shehabi A, Lei N, Smith S, Koomey J. Recalibrating global data center energy-use estimates. Science. 2020;367(6481):984-6.
7. Dayarathna M, Wen Y, Fan R. Data Center Energy Consumption Modeling: A Survey. IEEE Commun Surv Tutor. 2016;18(1):732-94.
8. Zhang Q, Yang LT, Chen Z, Li P. A survey on deep learning for big data. Inf Fusion. 2018;42:146-57.
9. Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016 Aug 13-17; San Francisco, CA. New York: ACM; 2016. p. 785-94.
10. Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, et al. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In: Advances in Neural Information Processing Systems 30 (NIPS 2017); 2017 Dec 4-9; Long Beach, CA. 2017. p. 3146-54.
11. Shehabi A, Smith S, Sartor D, Brown R, Herrlin M, Koomey J, et al. United States Data Center Energy Usage Report. Berkeley (CA): Lawrence Berkeley National Laboratory; 2016. Report No.: LBNL-1005775.
12. Rasmussen N. Electrical Efficiency Modeling for Data Centers. Schneider Electric White Paper. 2011;(113 Rev. 2).
13. Belady C, Rawson A, Pfleuger J, Cader T. Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric. The Green Grid; 2008.
14. Whitehead B, Andrews D, Shah A, Maidment G. Assessing the environmental impact of data centres Part 1: Background, energy use and metrics. Build Environ. 2014;82:151-9.
15. Prokhorenkova L, Gusev G, Vorobev A, Dorogush AV, Gulin A. CatBoost: unbiased boosting with categorical features. In: Advances in Neural Information Processing Systems 31 (NeurIPS 2018); 2018 Dec 3-8; Montréal, Canada. 2018. p. 6638-48.
16. Zhang W, Quan H, Srinivasan D. An Improved Quantile Regression Neural Network for Probabilistic Load Forecasting. IEEE Trans Smart Grid. 2019;10(4):4425-34.
17. Fan C, Xiao F, Zhao Y, Wang J. Analytical investigation of autoencoder-based methods for unsupervised anomaly detection in building energy data. Appl Energy. 2018;211:1123-35.
18. Patterson MK, Costello D, Grimm P, Loeffler M. Data Center TCO: A Comparison of High-Density and Low-Density Spaces. Intel Corp. White Paper; 2008.
19. Oró E, Depoorter V, Garcia A, Salom J. Energy efficiency and renewable energy integration in data centres. Strategies and modelling review. Renew Sustain Energy Rev. 2015;42:429-45.
20. Siriwardana K, Halgamuge SK, Schellenberg T. Potential of air-side economizers for data center cooling: A case study for key Australian cities. Appl Energy. 2018;229:1117-26.

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*

21. Evans R, Gao J. DeepMind AI Reduces Google Data Centre Cooling Bill by 40% [Internet]. DeepMind Blog. 2016. Available from: https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40
22. Gao J. Machine Learning Applications for Data Center Optimization. Google White Paper; 2014.
23. Lazic N, Boutilier C, Lu T, Wong E, Roy B, Ryu M, et al. Data Center Cooling Using Model-Predictive Control. In: Advances in Neural Information Processing Systems 31 (NeurIPS 2018); 2018 Dec 3-8; Montréal, Canada. 2018. p. 3814-23.
24. Haghighat E, Juanes R. SciANN: A Keras/TensorFlow wrapper for scientific computations and physics-informed deep learning using artificial neural networks. Comput Methods Appl Mech Eng. 2021;373:113552.
25. Fan C, Xiao F, Yan C, Liu C, Li Z, Wang J. A novel methodology to explain and evaluate data-driven building energy performance models based on interpretable machine learning. Appl Energy. 2019;235:1551-60.
26. Li H, Xu Z, Taylor G, Studer C, Goldstein T. Visualizing the Loss Landscape of Neural Nets. In: Advances in Neural Information Processing Systems 31 (NeurIPS 2018); 2018 Dec 3-8; Montréal, Canada. 2018. p. 6389-99.
27. Lundberg SM, Lee SI. A Unified Approach to Interpreting Model Predictions. In: Advances in Neural Information Processing Systems 30 (NIPS 2017); 2017 Dec 4-9; Long Beach, CA. 2017. p. 4765-74.
28. Wang Z, Hong T, Piette MA. Building thermal load prediction through shallow machine learning and deep learning. Appl Energy. 2020;263:114683.
29. Bergstra J, Bengio Y. Random search for hyper-parameter optimization. J Mach Learn Res. 2012;13:281-305.
30. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. J Mach Learn Res. 2011;12:2825-30.
31. Chollet F. Keras [Internet]. 2015. Available from: https://keras.io
32. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, et al. TensorFlow: A System for Large-Scale Machine Learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16); 2016 Nov 2-4; Savannah, GA. 2016. p. 265-83.
33. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: Advances in Neural Information Processing Systems 32 (NeurIPS 2019); 2019 Dec 8-14; Vancouver, Canada. 2019. p. 8024-35.
34. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, et al. Human-level control through deep reinforcement learning. Nature. 2015;518(7540):529-33.
35. Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, et al. Mastering the game of Go without human knowledge. Nature. 2017;550(7676):354-9.
36. Sutton RS, Barto AG. Reinforcement Learning: An Introduction. 2nd ed. Cambridge (MA): MIT Press; 2018.
37. Guyon I, Elisseeff A. An introduction to variable and feature selection. J Mach Learn Res. 2003;3:1157-82.
38. van Buuren S, Groothuis-Oudshoorn K. mice: Multivariate Imputation by Chained Equations in R. J Stat Softw. 2011;45(3):1-67.
39. Breiman L. Random forests. Mach Learn. 2001;45(1):5-32.
40. Hyndman RJ, Athanasopoulos G. Forecasting: principles and practice. 2nd ed. Melbourne: OTexts; 2018.
41. Box GE, Jenkins GM, Reinsel GC, Ljung GM. Time Series Analysis: Forecasting and Control. 5th ed. Hoboken (NJ): Wiley; 2015.
42. Lim B, Zohren S. Time-series forecasting with deep learning: a survey. Philos Trans A Math Phys Eng Sci. 2021;379(2194):20200209.
43. Makridakis S, Spiliotis E, Assimakopoulos V. Statistical and Machine Learning forecasting methods: Concerns and ways forward. PLoS One. 2018;13(3):e0194889.
44. Liu FT, Ting KM, Zhou ZH. Isolation Forest. In: 2008 Eighth IEEE International Conference on Data Mining; 2008 Dec 15-19; Pisa, Italy. 2008. p. 413-22.

***Received**: 20-11-2025 - **Accepted**: 19-01-2026 - **Published**: 25-01-2026*