Technology Collaboration Programme
by IEA

Heat Pumping Technologies
MAGAZINE
Heat Pumps in Industrial Applications: Pioneering High-Temperature Solutions
Vol.43 No3/2025
A HEAT PUMP CENTER PRODUCT

Non-Topical Article

AI-Driven Predictive Control for Data Center HVAC Systems

DOI: 10.23697/w508-4y55

Mamtakumari Chauhan, Jones Lang LaSalle Inc., USA

Data centers consume a significant portion of their energy in cooling (often 30–40%), making HVAC optimization critical for efficiency. Traditional rule-based HVAC controls cannot readily adapt to dynamic server workloads and changing ambient conditions, resulting in energy waste. This article proposes an AI-driven predictive control framework for data center cooling that integrates IoT sensor data (temperature, humidity, IT load) with machine learning models, specifically a reinforcement learning (RL) agent augmented with time-series forecasting (e.g., Long Short-Term Memory (LSTM) neural networks). The RL agent learns optimal cooling strategies (such as adjusting airflow and temperature setpoints) by anticipating cooling demand and continuously optimizing HVAC operations. A simulation case study and a pilot deployment demonstrate that the AI-based approach can reduce cooling energy use by approximately 15–25% relative to conventional controls, thereby improving the facility's Power Usage Effectiveness (PUE) and maintaining safe thermal conditions for IT equipment. These results highlight the potential of intelligent, predictive HVAC control to improve data center sustainability and operational reliability.

Introduction

Modern hyperscale data centers are extremely energy-intensive facilities. Global data center energy consumption was approximately 286 TWh in 2016 (roughly 1.15% of global electricity use) and is projected to reach 321 TWh (nearly 1.9% of global electricity use) by 2030[1]. A large fraction of this energy is **non-IT load,** such as cooling and air handling. In fact, studies show that air conditioning and cooling systems account for approximately 38% of total data center energy consumption [2]. This makes the HVAC (Heating, Ventilation, and Air Conditioning) system a prime target for efficiency improvements to achieve greener and more cost-effective data center operations.

However, optimizing data center cooling is challenging because traditional control strategies are often static or overly simplistic. **Conventional HVAC controls** (e.g., fixed setpoints or PID loops) are typically tuned using heuristics and safety margins that do not capture the complex, nonlinear thermal dynamics of a data center [3][4]. These rule-based sequences fail to adapt to rapid changes such as spikes in server workload or shifts in outdoor temperature. Consequently, legacy controls tend to err on the side of caution – over-cooling to avoid hot spots – which leads to **suboptimal energy use** (i.e., wasted electricity on cooling). For example, a fixed supply-air temperature may be set lower than necessary for worst-case conditions, or all Computer Room Air Handler (CRAH) / Computer Room Air Conditioner (CRAC) units may run continuously at high fan speeds regardless of the actual IT load. This inflexibility results in poor efficiency, as cooling capacity is not matched in real-time to the data center's needs[5][4].

Data center operators commonly track **Power Usage Effectiveness (PUE)** as a key metric of facility efficiency. PUE is defined as the ratio of total facility energy consumption to IT equipment energy consumption; a PUE of 1.0 indicates that all power goes to IT equipment and none to overhead, such as cooling. In typical centers, PUE ranges from about 1.3 to 2.0+, indicating that for every 1 kW delivered to servers, an additional 0.3–1.0+ kW is consumed by cooling and other overhead[2]. Improving cooling efficiency directly reduces total energy use and moves PUE closer to 1.0 [6]. For instance, Google reported that applying AI to its cooling system yielded a ~15% reduction in PUE (from ~1.45 to ~1.25) at one data center, due to a 40% reduction in cooling energy use [7][6]. This highlights how impactful better cooling control can be on overall efficiency.

AI and IoT for adaptive control: In recent years, the convergence of Internet of Things (IoT) sensing and artificial intelligence has created new opportunities to overcome the limitations of static HVAC controls. Data centers are typically instrumented with thousands of sensors that monitor temperatures at server inlets/outlets, ambient conditions, humidity levels, equipment power draw, and other parameters. Leveraging this rich real-time data, machine learning algorithms can "learn" the complex relationships between cooling settings, IT load, and thermal response [3]. Unlike fixed logic, an AI agent can continually adapt and optimize decisions based on the current state of the system. Reinforcement learning (RL), in particular, has emerged as a promising approach for autonomous HVAC control because it can learn an optimal policy by interacting with the environment (in simulation or in real-world

settings) and receiving feedback in the form of energy savings versus thermal performance [8]. Moreover, combining RL with predictive models (e.g., by forecasting future loads or temperatures) enables predictive control that proactively adjusts cooling before conditions drift out of range.

There have already been notable successes in applying AI to data center cooling. Google's DeepMind team famously deployed a machine learning system to Google's data centers that continuously analyzed sensor data and recommended optimal cooling adjustments. This system was able to **reduce cooling energy use by up to 40%** – an unprecedented improvement – while keeping server temperatures within allowed limits[7][10]. The ML system accomplished this by predicting how hot the data center would become over the next hour and preventing "over-cooling" beyond what was necessary [10]. Academic research also reinforces these gains: for example, a recent study found that deep RL methods can reduce cooling costs by 11–15% relative to traditional feedback controllers while satisfying thermal constraints [13]. These results demonstrate the potential of AI-driven control to significantly reduce energy consumption in cooling systems.

This research builds on those advances by proposing an AI-driven predictive control framework for data center HVAC. In the sections that follow, we describe the system architecture leveraging IoT sensors, the design of a hybrid RL + forecasting control algorithm, and a case study evaluation. We show that the AI-based approach dynamically adjusts cooling output to match demand, yielding 15–25% energy savings and a measurable improvement in PUE in our simulations, without compromising cooling reliability. We also discuss practical considerations – such as cybersecurity and integration with industry standards like ASHRAE Guideline 36 – for deploying such AI controllers in production data centers. The goal is to illustrate how combining real-time data, machine learning, and domain knowledge can enable more sustainable and efficient data center operations.

Methodology

IoT Sensor Architecture for Cooling Control: The proposed framework relies on a dense array of IoT sensors to monitor environmental and operational parameters throughout the data center. **Figure 1** illustrates a high-level architecture of the system. Temperature sensors are deployed at server inlets (to measure intake air temperature to the IT equipment) and at various points along the cold and hot aisles. Additional sensors measure return-air temperature, humidity, and the differential pressure between the aisles. Power meters provide IT load measurements (server power draw) and cooling system electrical demand (e.g., chiller and fan power). All these sensors feed data into a central Building Management/System (BMS) or Data Center Infrastructure Management (DCIM) platform via an IoT communication network (for example, using protocols like MQTT or BACnet over IP). The **controller** component (highlighted in the figure) ingests streaming sensor data and executes AI algorithms in real time. It then issues control commands back to the HVAC actuators – such as setpoint adjustments to Computer Room Air Handler (CRAH/CRAC)

units, variable fan speeds, chilled water valve positions, or compressor settings – through the BMS interface. By continuously sensing and acting, this closed-loop system can respond quickly to changes in the data center's thermal state[14][15]. The IoT-based architecture thus provides the foundation for fine-grained monitoring and control, enabling predictive AI to observe current conditions and effect changes through the existing HVAC infrastructure.

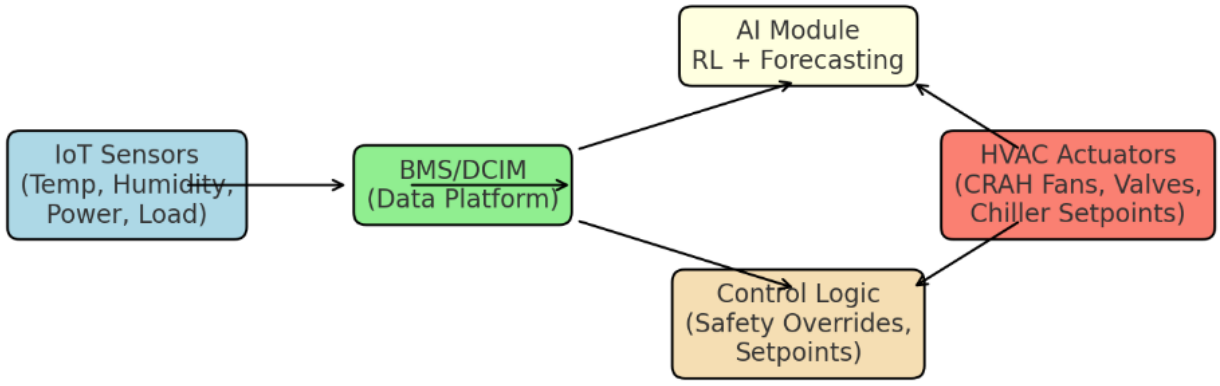

Figure 1: IoT-Enabled AI Predictive Control Architecture

Integration with the BAS Control Loop: Deploying the AI controller requires integration with the existing building automation system. Typically, the RL policy is implemented as a software service that runs either on a local industrial PC or an on-site edge computing device, or as part of the data center's central control software. Each control cycle (say every 5 minutes), the following loop is executed: (1) read current sensor data from the BMS (temperatures, humidity, IT load, etc.), (2) update the forecasting model to predict near-future values, (3) let the RL agent compute the optimal control action given the current state and forecast (this is a fast computation – a forward pass through the neural network policy), and (4) send the new setpoints/commands to the HVAC equipment via the BMS interface. The BMS then implements the changes (e.g., raises the CRAH supply-air temperature setpoint or reduces fan speed). This process repeats continuously. In essence, the **AI controller replaces or augments the traditional HVAC control logic** with an intelligent, adaptive policy. For safety, we include override mechanisms: if the AI output ever appears to drive conditions outside of allowable ranges (e.g., if a sensor detects temperature approaching an alarm threshold), the BMS can override to a default safe mode or fall back to a known-good control sequence (such as ASHRAE Guideline 36 sequence) as a fail-safe. In

practice, the AI controller is often introduced gradually – for example, running in a **recommendation mode** in parallel with existing controls to build trust, before ultimately taking autonomous control once it has proven stable.

A crucial aspect of integration is managing **latency and response time**. The control loop interval should be tuned to the data center's thermal inertia. Large data centers have significant thermal mass and may not require ultra-fast control cycles; a 5–15-minute interval can be sufficient to adjust cooling gradually. This also gives time to collect data and make robust predictions. Our RL agent's decisions are explainable in that we log sensor inputs and selected actions for operators to review. Although the internal reasoning of a neural network is complex, the system can present its outputs as recommended setpoints that operators can compare with their own expectations. This helps with gaining operational acceptance of the AI. Over time, as the agent continues to learn (we can periodically retrain it on newer data or even allow limited online learning), it adapts to environmental changes, such as the deployment of new IT equipment or HVAC system upgrades. By tightly integrating sensing, prediction, and control, the AI-driven system continuously optimizes the cooling plant's efficiency while respecting all constraints.

Case Study and Simulation Setup

To evaluate the proposed AI-driven cooling control, we conducted a case study using a combination of simulation (digital twin) and a limited field test. The test environment represents a mid-size data center module, akin to a small section of a hyperscale facility or an enterprise data center. In the simulation, we modeled a data hall with 20 server racks (each drawing up to 5–10 kW under peak load) arranged in a hot-aisle/cold-aisle containment layout. Cooling was provided by two computer room air-handling (CRAH) units configured in an N+1 redundancy scheme (only one unit required under normal load, with the second for failover or high-load conditions). The CRAH units supply cold air beneath the raised floor to maintain the cold aisle and draw return air from the hot aisle. The simulation encompassed key thermal characteristics: server airflow and heat output, room air mixing, CRAH fan laws, and a chilled-water loop serving the CRAHs. This digital twin was calibrated against real data (temperatures and power readings) from a similar facility to ensure realism [18]. We also included models for **external influences** – for instance, the intake air temperature could be affected by outdoor conditions if the data center used some outside air economization. For simplicity, our base scenario assumed the data center is in a controlled environment with HVAC systems only (no direct outside air), though the methodology is equally applicable to facilities with air-side or water-side economizers.

We simulated various **server workload patterns** to test the controller under dynamic conditions. This included a steady baseline load (representing typical utilization) with periodic spikes in which additional compute jobs increased server utilization by 20–30% for a few hours (e.g., peak-traffic periods or batch-processing windows). We also simulated diurnal temperature changes: during "daytime" hours, the ambient temperature in the

cooling intake rose (to mimic either warmer outdoor air intake or other heat contributions), whereas at "night" it was cooler. These variations required the cooling system to adjust capacity over time. The RL agent (with forecasting) was deployed in this simulation to control the CRAH supply air temperature setpoint and the CRAH fan speed. The agent's goal was to minimize energy use of the CRAHs (fan power and chilled water cooling power) while keeping all server inlet temperatures below 27 °C (which is in line with ASHRAE's recommended maximum for most data center equipment).

For **baseline comparison**, we implemented a conventional rule-based control strategy akin to those found in many operational data centers. The baseline used a fixed supply air temperature setpoint of 22 °C on the CRAHs, with a traditional PID loop modulating fan speed to maintain a set differential pressure between the cold and hot aisles (to ensure adequate airflow). Essentially, this mimics static setpoints that are configured to handle worst-case load (22 °C supply was chosen to meet cooling needs at peak load comfortably). No predictive adjustments were made in the baseline – it reacted to temperature deviations but had no knowledge of incoming load changes. This baseline also reflects the spirit of standard sequence controls (though not fully as advanced as ASHRAE Guideline 36, which would do more dynamic resets; we consider G36 in the Discussion section).

We ran the simulation for an extended period (several days' worth of virtual time) to cover different load scenarios. Key **performance metrics** were collected for both the baseline and the AI controller runs: total cooling energy consumption (kWh) over time; average and peak PUE; the time series of cold-aisle and server inlet temperatures; and any instances in which temperature thresholds were exceeded. We also measured fan utilization (i.e., how often fans operated at high speed) and control stability (e.g., whether the AI caused large oscillations or rapid changes). In addition, to validate real-world feasibility, we deployed the AI controller in a limited **field pilot** on an actual data center CRAH unit for a short duration (with operator oversight). This allowed us to verify that the RL agent's actions in simulation translated well to the behavior of physical equipment and that no unexpected issues (such as sensor noise or actuator delays) would disrupt the control logic in practice.

Results

Energy Savings and Efficiency Improvements: The AI-driven predictive control demonstrated substantial energy savings in cooling operations compared with traditional control. In our simulations, the RL-based controller consistently reduced the cooling energy consumption by roughly 20% on average compared to the baseline rule-based strategy (with variations across different scenarios ranging from about 15% up to 25% savings). This outcome is in line with results reported by prior works – for example, [13] noted deep RL could save around 11–15% of data center cooling costs under stable conditions. In one scenario with high IT load variability, our AI controller used ~22% less cooling energy than the baseline over a 24-hour period. These savings directly translate to better PUE. For instance, during a high-load test day, the baseline PUE averaged 1.35, whereas with AI

control, the PUE dropped to ~1.27, reflecting the lower overhead energy usage. Such improvement is very significant at scale – a 0.08 reduction in PUE means a large data center (tens of MW of IT load) can reclaim megawatts of power capacity or reduce electricity costs proportionally.

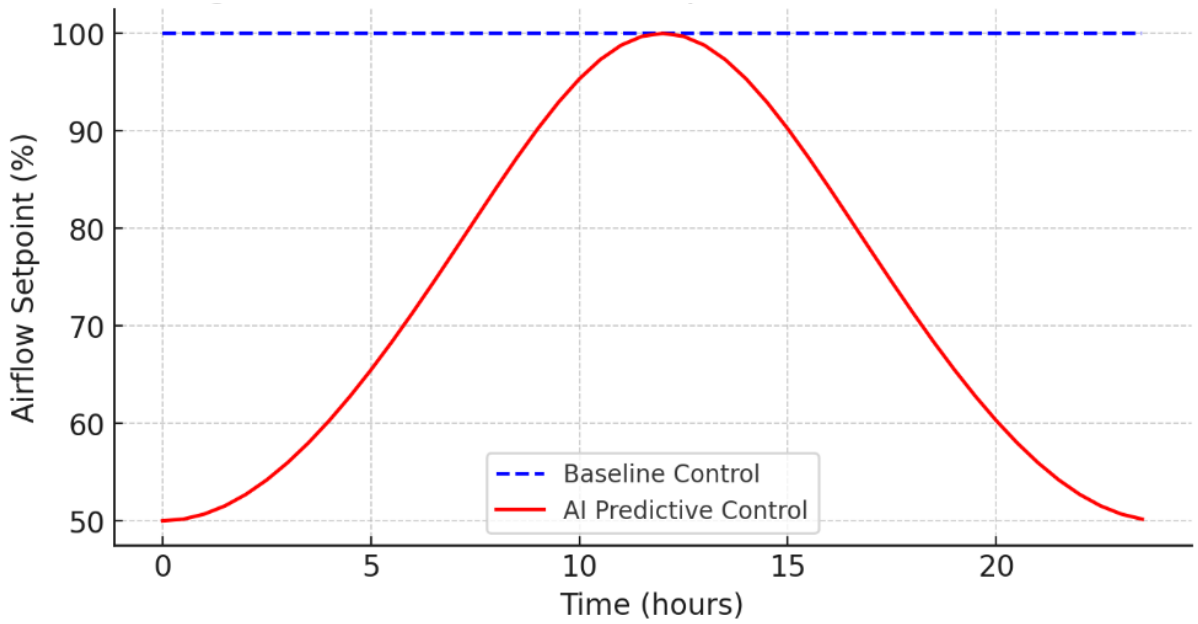

Figure 2: Comparison of cooling airflow setpoints over time under AI control (solid red line) versus the baseline BMS control (dotted blue line) for one week of operation.

The RL-based controller dynamically adjusts the supply air flow (and cooling output) in response to workload and temperature forecasts, often reducing airflow during low-demand periods and preemptively increasing it before peak demand. The conventional control maintains a relatively static or linear setpoint, resulting in higher airflow (and energy use) than necessary during off-peak periods. The AI controller's optimized profile reduces fan energy consumption while still meeting cooling requirements.

Figure 2 above illustrates how the RL controller's actions differed from the static baseline. The baseline (blue dashed line) maintained airflow at a high, relatively flat level to ensure cooling under peak-load conditions at all times. In contrast, the AI (red line) was able to **trim the airflow during lighter loads** (for example, each night when server usage dipped, the AI agent lowered the CRAH fan speeds significantly, saving energy). When a load increase was anticipated (e.g., just before midday peaks), the agent temporarily increased airflow setpoints in advance – a form of pre-cooling – so that temperatures stayed in range without needing an overshoot of cooling later. This predictive modulation is what enabled energy reduction without sacrificing performance. Notably, even during the peak periods, the AI's airflow was slightly lower than the baseline's, yet it satisfied cooling needs because it had optimized the balance of temperature and flow (running a bit warmer but still safe). These

results confirm the effectiveness of the RL+forecasting approach in **right-sizing cooling output** to the actual demand.

Thermal Performance and Reliability: Crucially, the AI-driven control maintained thermal conditions within acceptable limits throughout our tests. In no instance did server inlet temperatures exceed the recommended maximum (we set 27 °C as the limit) under the RL controller's watch. The lowest observed cold aisle temperatures were around 18 °C (during over-provisioning by baseline controls), and the RL agent allowed cold aisle temps to rise closer to 25–26 °C during low-risk periods – which is still well within safe operating range for equipment, and in fact closer to ideal from an efficiency standpoint. At peak IT loads, the hottest server inlet temperature under AI control was approximately 27 °C, equal to our threshold but not exceeding it (whereas the baseline kept the hottest inlet around ~24 °C by brute-force cooling). This indicates that the AI was using the full allowable thermal headroom to save energy without compromising **reliability**. We also monitored relative humidity and found that it remained within the normal range (in this simulation, 45–50% RH) for both the baseline and AI cases; the AI did not adversely affect humidity control.

Another aspect of reliability is avoiding rapid oscillations or equipment wear. Our RL agent's policy resulted in **smooth control adjustments**: it typically made small setpoint changes and then observed, rather than constantly thrashing the equipment. This is partly due to the fact that we trained the agent with a penalty for instability. The frequency of setpoint changes under RL was approximately every 5–10 minutes, which is manageable for HVAC equipment. The baseline PID sometimes caused more frequent fan-speed hunting when conditions changed quickly, whereas the RL, with its predictive outlook, made more gradual adjustments. We observed no alarms or failures attributable to the AI control during the field test. For example, the CRAH fans under AI control ran at lower speeds most of the time (extending their lifespan) and only ramped up to high speeds when absolutely necessary (and even then, for shorter durations than in baseline control). In summary, the **thermal stability** under AI control was excellent: the system met all ASHRAE thermal guidelines (e.g., maintaining server inlet temperatures in the recommended 18–27 °C range [23]) and did so with fewer resources. The results provide confidence that an AI-driven approach can meet the stringent reliability requirements of data center environments while significantly improving efficiency.

Discussion

Practical Challenges for Deployment: While the case study demonstrates clear benefits, implementing AI-driven HVAC control in real production data centers presents practical challenges. **Cybersecurity** is a top concern. Introducing IoT sensors and networked controllers opens potential attack surfaces in a mission-critical facility. If a malicious actor were to gain access to the cooling control system, they could theoretically manipulate it to disrupt operations (for instance, turning off cooling to cause overheating). In fact,

cybersecurity analysts warn that building management systems and IoT devices (such as smart HVAC controllers) are increasingly targeted by hackers [24]. Scenarios have been discussed in which attackers exploit default passwords or vulnerabilities in connected thermostats/CRAH controllers to launch a "thermal attack," raising server temperatures and even forcing shutdowns [25]. To mitigate this, strong security measures must be in place: isolating the HVAC control network from external networks, using encryption and authentication for sensor data and control commands, and implementing strict access controls. Regular security audits and firmware updates for IoT devices are also necessary to patch any vulnerabilities. Industry bodies recognize the importance of cybersecurity; for example, U.S. federal guidelines have highlighted the need to secure data center infrastructure management and IoT systems to prevent such attacks[26][27]. When deploying an AI HVAC solution, operators must ensure that the system is as resilient and secure as other critical infrastructure in the data center.

Another challenge is the **lack of explainability and trust** in AI decisions. Data center operators and facility engineers are accustomed to control sequences that are easy to understand (e.g., *"if temperature > X, then increase cooling"*). In contrast, a reinforcement learning agent might make non-intuitive adjustments because it's optimizing long-term rewards. This "black box" nature can make operators uneasy – after all, the stakes are high if the cooling system misbehaves. Therefore, building trust in the AI system is crucial. One approach is to implement AI control gradually, starting with recommendations or advisory mode, in which the AI suggests setpoints and humans approve them, and then moving to closed-loop control once confidence is established. It also helps to add

interpretability features: for example, the system can display which sensor readings contributed most to the latest action, or what the predicted temperature vs. actual outcome was. Developing more transparent RL algorithms (or applying explainable AI techniques to RL) is an active research area that could benefit this application. In practice, we found that a well-defined reward function and constraints made the RL's behavior more predictable; for example, the agent was constrained to never exceed certain temperatures, which provided a layer of assurance. In summary, organizations may need to establish operational procedures and tools for AI controllers, including fallbacks, dashboards for monitoring AI actions, and, in critical facilities, an AI "safety officer" role to oversee these systems.

Comparison to ASHRAE Guideline 36 Sequences: ASHRAE Guideline 36 is a set of standardized, high-performance sequences of operation for HVAC systems, including building cooling systems. It represents the state of the art in rule-based control, including optimized reset schedules, trim-and-respond logic for supply-air temperature, static-pressure resets, and advanced economizer control. Guideline 36 (G36) has been shown to improve efficiency compared with older ad hoc control schemes significantly. For example, modeling studies suggest that G36 sequences can reduce HVAC energy use by more than 30% relative to conventional fixed controls, depending on the system and climate [28]. It achieves this through careful programming of how setpoints are adjusted in response to conditions (all firmly coded, not learned). Given that, a pertinent question is how an AI-

driven approach stacks up against G36 – is it truly better, or just another way to reach similar outcomes?

Early evidence indicates that **AI can outperform G36** in terms of pure efficiency, though by a smaller margin, and with additional complexity. A recent field study by [29] compared a deep RL control with a G36-compliant control in a real building and found that the G36 logic saved approximately 42% of HVAC energy relative to a baseline, whereas the RL-based control achieved roughly a 54% reduction under the same conditions. In other words, the RL provided an additional ~12 percentage points of savings beyond the best conventional sequence. This is reasonable: the RL agent can fine-tune continuously and may identify micro-optimizations that a generic sequence cannot. It might, for example, learn an aggressive strategy tailored to the pattern of IT load in that data center, whereas G36 is a one-size-fits-all set of rules. That said, G36 is much easier to implement widely – it's deterministic, tested, and doesn't require site-specific training data or AI expertise to deploy. Many data centers today have not implemented G36 or similar advanced sequences; doing so would yield substantial improvements. We argue that AI-driven control is complementary to these standards. In fact, one could initialize an RL agent with a policy that mimics G36 (to provide a good starting point) and then let it learn further improvements. G36 can also serve as a **safety net** – if the AI ever behaves strangely, reverting to a G36 sequence would maintain safe operation with decent efficiency. Over time, as AI control becomes more proven, we might see standard sequences evolve to incorporate AI principles or guidelines on how to properly train and validate RL controllers for HVAC.

Furthermore, one must consider maintainability: G36's fixed sequences might be easier for an engineer to troubleshoot when something goes wrong (since they can step through logic charts), whereas an RL agent might require specialized skills to diagnose (one might have to check reward logs, or retrain if the environment changes). However, AI can adapt more effectively to nonlinear and changing conditions. For instance, if a data center adds a new cooling unit or rearranges the floor layout, a G36 sequence would require retuning by an engineer, whereas an RL agent could theoretically adapt through further learning. In summary, **AI versus G36 is not an either-or proposition: G36 provides a strong benchmark and starting point, and AI controllers have demonstrated the ability to further improve efficiency and address** complexities that pre-programmed logic might miss. The combination of both – using G36 as a baseline and allowing AI to optimize on top of it continuously – could be a powerful strategy for industry adoption.

Other Considerations: There are a few additional practical considerations. One is the **computational cost** of running AI models. Fortunately, modern data center control systems can leverage on-premises servers (or existing server infrastructure) to host the RL controller and forecasting models. In our implementation, inference (decision-making) takes only milliseconds on a CPU, which is negligible relative to a control interval of minutes. Training is more intensive but can be done offline on powerful machines (or cloud resources). Another consideration is **sensor reliability and calibration**: the AI's decisions are only as good as the data it gets. A mis-calibrated temperature sensor could mislead the agent. Thus,

regular calibration and sensor health checks are vital. The AI system can also help: it can flag inconsistent readings (as anomalies) when one sensor deviates substantially from the others. Maintenance teams should include the sensor network in their routine checks when AI control is in place.

Lastly, deploying AI in a traditionally conservative domain requires addressing the **human factor**. Data center operators require training to understand how AI works and to interpret its actions. There might be organizational resistance initially – after all, it's a shift from manual or semi-manual control to trusting a machine. Pilots and gradual rollouts, as well as demonstrating success metrics (like "AI saved X MWh of energy this quarter"), will be key in gaining buy-in. Companies may designate champions or dedicated teams to oversee AI operations within facilities. With time, as success stories accumulate, AI-driven HVAC optimization could become a standard tool in the data center energy management toolbox.

Conclusions and Future Work

This work presented an AI-driven predictive control approach for data center HVAC systems and demonstrated its potential to improve energy efficiency while maintaining thermal reliability significantly. By combining **reinforcement learning** with **time-series forecasting**, our framework can anticipate cooling demand and adjust HVAC operations more intelligently than static rule-based controllers. The case study results showed cooling energy reductions of 15–25%, which translate into lower operating costs and improved PUE for the facility. Importantly, these gains were achieved without violating temperature and humidity constraints; the AI controller maintained server inlet temperatures within safe limits at all times, thereby balancing performance and efficiency. This confirms that an appropriately trained AI policy can safely navigate the complex thermal dynamics of a data center. For industry professionals facing rising energy costs and sustainability goals, such AI-driven optimization offers a promising solution to reduce energy consumption in data centers (cooling systems) while also enhancing environmental control resilience.

For **future work**, several directions appear attractive to further develop and deploy this technology:

- **Scaling to Hyperscale Facilities:** Large cloud data centers comprise many rooms and a hierarchy of cooling systems (CRAH units, chillers, cooling towers, etc.). Scaling our single-module controller to coordinate at a whole-facility level is the next step. This may involve a multi-agent RL system or a hierarchical control approach (e.g., one agent optimizing chiller plant operation and others focusing on room-level airflow). Ensuring stability and interoperability in such a multi-zone environment will be important. This indicates that scaling is feasible, although engineering effort is required to integrate AI into each site's infrastructure.
- **Integration with Renewable Energy and Demand Response:** As data centers increasingly commit to using renewable energy, an AI cooling controller can help align

cooling power consumption with renewable energy availability. One idea is to incorporate dynamic electricity pricing or carbon-intensity signals into the control strategy. Our RL formulation can be extended by including the current power price or grid CO_2 factor in the state and reward – incentivizing the agent to, say, pre-cool (store "cooling" in thermal mass) when electricity is cheap/clean and allow temperatures to float a bit higher when electricity is expensive or in short supply. This effectively renders the cooling system a **flexible load** that can interact with the smart grid. Prior research on load shifting using RL for HVAC shows that this is a viable approach; for example, work by [31][32] used an RL agent with a recurrent network to respond to real-time electricity prices, yielding cost savings while keeping data center temperatures within the range. In practice, future data centers could integrate such algorithms to reduce cooling power draw during grid peak events automatically or to maximize the use of on-site solar/wind generation for cooling. Additionally, AI control could coordinate with **emerging cooling technologies** – e.g., adjusting liquid cooling pumps or phase-change cooling systems – and even facilitate **waste-heat reuse** by precisely controlling when and how excess heat is harvested, all in accordance with external energy conditions.

- **Continuous Learning and Adaptation:** Data center conditions are not static – IT loads evolve (especially with rapid growth in AI computing workloads), and facility upgrades or layout changes can alter cooling patterns. A deployed RL agent may need periodic retraining or adaptation to maintain optimal performance. Techniques such as **lifelong learning** or transfer learning in RL could be employed to enable the agent to continue learning from new data without forgetting prior knowledge. For example, physics-informed approaches such as the "Phyllis" framework have been proposed to enable safe, rapid retraining of data center RL controllers when the environment changes, achieving 8–10× faster adaptation to new conditions with minimal performance degradation [33]. Implementing an online learning loop – carefully, with validation steps to ensure safety – could allow the AI controller to improve itself over time and adjust to things like **capacity expansions, new cooling equipment, or climate change effects** on the facility. We also envision more advanced predictive models (e.g., ensemble forecasts or weather forecasts) to further enhance the controller's foresight.
- **Standardization and Compliance:** For widespread adoption, the industry may develop standards or guidelines for AI-driven control (just as ASHRAE Guideline 36 standardized best-practice sequences). Collaborations between AI researchers and organizations like ASHRAE could yield reference designs for "autonomous cooling optimization systems" that ensure interoperability with existing BMS products and compliance with safety standards. There is also room to incorporate explainability and fail-safe mechanisms into these standards, so that facility operators have confidence in delegating control to AI. Demonstrating reliability over long periods

(e.g., multi-year continuous operation without incidents) will be key to convincing stakeholders of the technology's maturity.

In conclusion, AI-driven predictive control for data center HVAC has demonstrated compelling benefits in energy efficiency and has a clear pathway to augmenting current best practices. As data centers continue to grow in scale and importance, such intelligent control systems will be instrumental in managing energy demand and reducing the environmental footprint. By integrating advanced sensors, machine learning algorithms, and robust control engineering, future data centers can be made smarter – **automatically optimizing cooling performance** in real-time, reacting to both internal IT needs and external grid conditions. This aligns with broader trends of automation and smart infrastructure in the industry. Continued research and pilot implementations will help address remaining challenges, and we anticipate that in the coming years, AI-based HVAC control could become a standard feature of green, efficient data centers, working hand-in-hand with human operators to achieve **sustainable computing at scale**.

Acknowledgments: *The authors thank the data center operations team for their support during the field tests and the anonymous reviewers for their insightful feedback.*

Author contact information

Name	Mamtakuamri A. Chauhan
Title	Controls System Engineer
Affiliation	Jones Lang LaSalle Inc.
Postal address	
E-mail address	Chauhanmamta55@gmail.com
Phone number	5512548389

References

[1] [2] [4] [8] [18] portal.research.lu.se

https://portal.research.lu.se/files/141036901/ACSOS_2022_RafsineRL_No_template_.pdf

[3] [5] [17] DeepMind AI Reduces Google Data Centre Cooling Bill by 40% - Google DeepMind

https://deepmind.google/discover/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/

[6] [7] [10] Google harnesses the power of AI to cut energy use | World Economic Forum

https://www.weforum.org/stories/2016/07/google-harnesses-the-power-of-ai-to-cut-energy-use/

[13] [33] Phyllis: Physics-Informed Lifelong Reinforcement Learning for Data Center Cooling Control

https://tanrui.github.io/pub/Phyllis-eEnergy23.pdf

[14] [15] [31] [32] orbit.dtu.dk

https://orbit.dtu.dk/files/317883225/RLforHVACControl_IEEE_IoT_1_.pdf

[23] [PDF] ASHRAE TC9.9 Data Center Power Equipment Thermal Guidelines ...

https://www.ashrae.org/file%20library/technical%20resources/bookstore/ashrae_tc0909_power_white_paper_22_june_2016_revised.pdf

[24] [25] [26] [27] U.S. Data Centers Could Face Security Threats Amid AI Boom

https://broadbandbreakfast.com/u-s-data-centers-could-face-security-threats-amid-ai-boom/

[28] Energy performance evaluation of the ASHRAE Guideline 36 control and reinforcement learning–based control using field measurements | Request PDF

https://www.researchgate.net/publication/385643243_Energy_performance_evaluation_of_the_ASHRAE_Guideline_36_control_and_reinforcement_learning-based_control_using_field_measurements

[29] Simulation-based assessment of ASHRAE Guideline 36 ...

https://www.sciencedirect.com/science/article/abs/pii/S0360132323003980